IMPROVING WORD-FINDING IN ASSISTIVE COMMUNICATION TOOLS: A MIXED-INITIATIVE APPROACH

SONYA NIKOLOVA

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
COMPUTER SCIENCE
Advisor: Perry R. Cook

September 2010

Abstract

Navigating a vocabulary consisting of thousands of entries in order to select appropriate words for building communication is challenging for individuals with lexical access disorders like those caused by aphasia. Most existing assistive communication 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 confusing searches; others are simply a list of arbitrary categories which causes excessive scrolling and a sense of disorganization. Ineffective vocabulary organization and navigation hurt the usability and adoption of assistive communication tools and ultimately fail to help users build functional communication.

We argue that to provide effective word-finding, an assistive vocabulary needs to adapt to individual user's word usage patterns and to the semantic associations present in a speaker's mental lexicon, where words are stored and organized in ways that allow efficient access and retrieval. To test our thesis, we employ a mixed-initiative approach to the design of the Visual Vocabulary for Aphasia (ViVA). ViVA is adaptive in that it automatically models an individual user's mental lexicon according to psycholinguistic theories that propose a semantic network structure of the lexicon and spreading activation as supported by semantic priming. The vocabulary is also adaptable and can be customized to reflect user preferences. ViVA compensates for some of the impaired links in a user's mental lexicon by building a dynamic network where words are linked based on semantic association measures, human judgments of semantic similarity and past vocabulary usage. Thus, the tool tailors the vocabulary organization according to both user-specific information and general knowledge of human semantic memory.

We evaluated how our system performs compared to a widely used vocabulary access system in which words are organized hierarchically into common categories and subcategories. The results indicate that word retrieval is significantly better with ViVA. In addition, we present results from a longitudinal single-case study with an aphasic participant which illustrates the importance of personal associations for creating an effective assistive vocabulary such as ViVA.

Contents

List of Figures

List of Tables

Acknowledgments

I would first like to thank the participants in my user studies who have all contributed to this work. I am grateful to the staff at the Adler Aphasia Center for their generous support and to Pamela Woronoff who welcomed me to the Princeton Lingraphica User's group.

I have had the opportunity to work with and have been advised by a number of extraordinary people. I am grateful to Dr. Maria Klawe for introducing me to the field of human-computer interaction and for encouraging me to pursue it. My Ph.D. advisor Professor Perry Cook showed me what it means to be passionate about one's work and his support enabled me to do the research that excited me the most. I am particularly grateful to Professor Marilyn Tremaine for her constant guidance, support and invaluable feedback throughout the years. I am also grateful for having worked with Dr. Christiane Fellbaum whose door was always open when I needed advice or encouragement. I thank Professor Szymon Rusinkiewicz and Professor Adam Finkelstein for their thoughtful comments and guidance during the most important milestones of graduate school. My colleagues Xiaojuan Ma and Jordan Boyd-Graber have both directly contributed to this dissertation by helping me shape ideas, and design and conduct user studies.

I shared a few enjoyable moments at Princeton with Najwa Aaraj, Miloš Ilak, Ronny Luss and Nebojša Stanković who made graduate school life significantly better. From further away, Nevena Nancheva knew how to cheer me up whenever I needed it and Nyoko Muvangua provided consistent encouragement. Slaveya Boncheva always welcomed me back home and I could not wish for a more caring friend. Jelena Zurovac made the hectic last months at Princeton better and Maria Ivanova provided me with a

roof whenever I needed to escape to New York City. My host parents, Sara Compton and Terry Peterman, gave me home away from home. I thank them all for their friendship.

Through the years, I have often drawn strength from my family's faith in my capabilities and my potential to succeed. My parents, Maria and Stoyan, nurtured my ambitions even though this kept me far away from home. However, despite the distance, I never felt alone because of the unconditional love and care that my sister, Nina, surrounded me with. Finally, I could not imagine being able to complete this journey without David; I thank him for the boundless support and love he has given me.

The research in this dissertation has been supported in part by Microsoft Research (Intelligent Systems for Assisted Cognition grant), the National Science Foundation (under grants 0101247, 0509447 and 9984087), the Kimberly and Frank H. Moss '71 Research Innovation Fund of the Princeton School of Engineering and Applied Science, Interval Research and the Arial Foundation. I would also like to thank Andrew Gomory, CEO of Lingraphicare America, Inc., for giving me access to the Lingraphica vocabulary and other company resources. All Lingraphica images in this dissertation are copyrighted by Lingraphicare America, Inc. and used with permission.

На майка ми и на другите изключителни жени които ме подкрепяха по този път.

To my mother and to the other exceptional women who supported me along the way.

Chapter 1

Introduction

Augmentative and Alternative Communication (AAC) tools assist people with communication exchanges. In order to be effective for functional communication, AAC tools need to provide a rich and expressive vocabulary. To meet this requirement, assistive vocabularies offer extensive collections of words which are cumbersome to navigate by users with language and cognitive impairments such as those present in aphasia.

Aphasia refers to a family of acquired communication disorders that impact an individual's language abilities. It affects close to one million people in the United States (National Aphasia Association 2010) and is typically acquired as the result of a stroke, brain tumor or other brain injuries. Depending on the area or degree of damage to the brain, the resulting impairments to the ability to understand and produce language vary. Rehabilitation can improve people's abilities; however, a significant number of people with aphasia are left with life-long chronic impairments which impact their everyday life significantly. The most common, chronic, impairment that is observed across individuals with aphasia is anomia, the inability to access and retrieve words from the mental lexicon (Goodglass & Wingfield 1997).

People with aphasia cope with their inability to communicate by using different low-tech strategies such as drawing pictures, writing notes, pointing, mimicking and gesturing as well as by relying on their caregiver. There have also been commercial and research efforts to build AAC technological tools to help these individuals communicate independently and thus, regain some of their social life. AAC tools for people with aphasia provide multi-modal vocabularies consisting of text, speech-audio and picture-based representations of concepts. A typical AAC vocabulary contains thousands of words which are organized in a hierarchy of concepts (e.g., *food–dessert–cake*) or a list of categories (e.g., *clothes* category). Due to anomia, many people with aphasia cannot retrieve from their mind the words they need to communicate. Instead, to build communications, users browse through the vocabulary for the desired concepts (e.g., *shoes*, *new*, *buy*) and once found, link them to form a phrase or a sentence (e.g., *buy new shoes*). It is challenging to minimize the complexity of navigating a large vocabulary and enable the user to intuitively and quickly find words. There has not been sufficient research investigating what type of vocabulary organization optimizes word-finding, but previous work studying the use of AAC tools has shown that existing vocabularies are challenging to navigate (Beukelman & Mirenda 2006, Boyd-Graber et al. 2006).

We argue that to provide effective word-finding, an assistive vocabulary needs to adapt to individual user's word usage patterns and the semantic associations present in the user's mental lexicon (where words are stored and organized in ways that allow efficient access and retrieval). To test our thesis, we employ a mixed-initiative (Horvitz 1999) approach to the design of the Visual Vocabulary for Aphasia (ViVA). ViVA is adaptive in that it automatically models an individual user's mental lexicon according to psycholinguistic theories that propose a semantic network structure of the lexicon and spreading activation as supported by semantic priming (Collins & Loftus 1975, Swinney 1979). The vocabulary is also adaptable and can be customized to reflect user preferences.

In addition to ViVA's design, this dissertation presents ViVA's evaluation. We compared our approach to word-finding to searching for words in the hierarchical organization of a commercial AAC tool built specifically for people with aphasia. We asked non-aphasic and aphasic people to search for words in two vocabularies implementing the two different organizations. Through a longitudinal case study with one aphasic user, we also investigated how ViVA can be personalized to assist the user better and how the vocabulary evolves with usage.

1.1 Motivation

Designing technology that satisfies the needs and expectations of the intended user is a fundamental challenge in the field of human-computer interaction. This is particularly challenging when designing technology for people with aphasia due to the variability of the resulting impairments. Aphasia can affect speaking, language comprehension and writing to varying degrees and in any combination in an individual as well as across individuals. For example, some people may speak fluently but have impaired auditory comprehension while others may have impaired speech but good reading comprehension. Given the importance of language communication in all aspects of daily life, it is not surprising that most individuals with aphasia experience a reduction in their ability to participate in everyday activities with the result that social isolation and depression are relatively common (Kauhanen et al. 2000, Martin et al. 2002).

There have been consistent efforts in improving the lives of those with aphasia through technology, but existing assistive communication tools fail to address the problems arising from the heterogeneity of the user population. This shortcoming has stimulated additional research efforts that show it is essential to seek flexible and customizable solutions when building AAC tools. Researchers as well as speech-language therapy experts who treat people with aphasia have advocated designing AAC tools that take into consideration individual user's language abilities and idiosyncrasies (Moffatt et al. 2004, van de Sandt-Koenderman 2007, Waller et al. 1997).

An assistive communication tool that attempts to meet the needs of individual users demands an expressive and flexible vocabulary. Vocabulary expressiveness and word organization, however, often compromise the effectiveness of existing communication tools (Beukelman & Mirenda 2006). Although initial vocabulary sets can be formed from words frequently needed by the target population, no packaged system has the depth or breadth to meet the requirements of every individual. In addition, most of the existing assistive 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 cumbersome searches; others are simply a list of arbitrary categories which causes excessive scrolling and create a sense of disorganization. To address these issues, it is important to build a well-structured computerized vocabulary that offers efficient navigation and search capabilities.

As a solution, we propose a flexible and customizable assistive vocabulary which consists of semantic networks tailored to an individual user's profile. The Visual Vocabulary for Aphasia (ViVA) can be adapted by the user as well as automatically suggests word associations that guide word-finding. We base ViVA's design on experimental evidence suggesting that words are organized in a speaker's mental lexicon by various similarity relations, in particular, phonological and semantic similarity. Some of these relationships are lost or impaired due to aphasia, making it challenging for people to retrieve words expressing desired concepts. Speech-therapy researchers treat word retrieval by guiding a patient in identifying important semantic features of the target word. It has been shown that such exercises restore and strengthen some of the damaged associations. ViVA provides support for word-finding by incorporating meaningful semantic features expressed by human judgments of semantic relatedness, predicted semantic association, word frequency usage statistics and user preference. Thus, the vocabulary reflects general knowledge of semantic relatedness and models a specific user's mental lexicon.

1.2 Research Questions

The high-level goal of this dissertation work is to improve word-finding in assistive vocabularies for people with lexical-access impairments such as anomia. To accomplish this goal, we built a vocabulary of semantic networks that model a user's mental lexicon by incorporating word frequency usage statistics, known and predicted semantic word associations and a large collection of human judgments of semantic relatedness that we collected. More specifically, this thesis addresses the following research questions:

1. How can we improve word-finding in assistive communication vocabularies?
   a. Can we improve word-finding by incorporating semantic word associations retrieved from existing large-scale lexical databases such as WordNet (Fellbaum 1998, Miller 1990)?

   b. Can we improve word-finding by incorporating human judgments of semantic relatedness?

2. How can we tailor word-finding to a specific user?

   a. Do word frequency usage statistics improve word-finding?

   b. Based on word-usage and semantic-word-association data, can we predict associations between words in the vocabulary that improve word-finding?

3. How does a static hierarchical vocabulary organization compare to a dynamic (adaptive) organization?

   a. How does an adaptive vocabulary organization affect user performance (e.g., time to find a word)?

   b. How does an adaptive vocabulary organization affect user satisfaction?

1.3 Thesis Overview

In this section, we outline the key elements of the dissertation and summarize our contributions.

1.3.1 Design of the Visual Vocabulary for Aphasia

The Visual Vocabulary for Aphasia (ViVA) relies on adaptable and adaptive functionality to build semantic networks that the user can navigate effectively when searching for words. We define an ***adaptable*** tool to be one that can be reconfigured by the user, whereas an ***adaptive*** tool is one that tailors itself automatically to a user's profile (based on usage characteristics or other factors). Users, especially ones with cognitive impairments, tend to rely on consistency and stability within an interface. Thus, we believe that it is necessary to explore a ***mixed-initiative*** design, an effective blend of automation and direct manipulation (Horvitz 1999). This 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.

Our goal is to let the user directly influence the organization of the vocabulary and simultaneously have the structure change to better suit usage needs. However, we still need an initial organization to allow the user to successfully use ViVA from day one. We constructed ViVA's initial vocabulary set such that it is a collection of commonly used words as well as ones relevant to our target population, people who have aphasia. This was achieved by mining words from two sources: the "core" WordNet (Fellbaum 1998, Miller 1990), consisting of frequent and salient words, and Lingraphica's visual vocabulary (Lingraphicare Inc. 2010).

Lingraphica is an assistive tool designed specifically for people with aphasia. It provides a vocabulary hierarchy that has been carefully designed and proven useful via its application in a commercial assistive device widely used across the U.S.A. Lingraphica's vocabulary consists of a set of common words used in daily communication which has evolved over a number of years incorporating feedback from users with aphasia, their caregivers and speech-language pathologists who prescribe the device and use it to treat patients. The words in Lingraphica's vocabulary are organized according to shared contexts that are common in daily life. If you need to find *milk*, for example, you select the icon for *kitchen*, then the *refrigerator* category and then you find *dairy products*. The icon for *milk* is in *dairy products*. Lingraphica's hierarchy tends to be deep and often requires multiple clicks to find even simple words. Thus, users readily get lost while browsing for a word. Even if the user keeps track of the path they have taken, the task of navigating the vocabulary often becomes cumbersome because a number of hierarchy branches have to be traversed when composing a simple phrase.

ViVA takes advantage of Lingraphica's established vocabulary, but to improve word-finding, it enriches its core hierarchy with additional associations between words which aim to compensate for some of the impaired links in a user's mental lexicon. This is achieved by making use of models of human semantic memory that build on the notion of semantic similarity and relatedness. Such models, constructed on evidence gained from psycholinguistic experiments, form the basis of the large lexical database WordNet (Fellbaum 1998, Miller 1990). Since WordNet connects meaningfully related words to one another, it guided our initial design. However, incorporating only WordNet associations in ViVA was insufficient because noun–verb and noun–adjective connections are sparse; combining words from different parts of speech is integral to functional communication. To address this problem, we collected ratings from thousands of English speakers about the strength with which a given word evokes another, (e.g. *eat–hungry*).

1.3.1.1 Collecting Evocation Ratings from Online Annotators

To augment Lingraphica's vocabulary hierarchy with meaningful links between words, we concentrated on the measure of ***evocation***, i.e., a rating of how much one word brings to mind another word. 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 collected scores for strength of evocation for 100,000 word pairs through a large-scale online experiment that asked untrained annotators to provide ratings. The data, which correlated well with ratings gathered from trained annotators, was then used to augment the core vocabulary hierarchy with links reflecting human judgments of semantic relatedness. This created densely connected semantic networks which provide new paths between words that assist the user with linking concepts to shape a sentence. We have made the evocation data available to other researchers for use in assistive communication for other user populations and for improving on our dataset.

1.3.2 Simulated Sentence Construction with ViVA

Our first evaluation of ViVA simulated vocabulary usage and sentence construction using sentences collected from blogs of elderly people.

The evocation data was used to augment Lingraphica's hierarchy with links reflecting how strongly people associate two words. We also implemented a prediction algorithm that added additional links between words based on prior usage. Prior usage of the vocabulary was simulated by gathering sentences from blogs of elderly people who write about their daily life on topics such as cooking, gardening and travelling. The sentences were parsed to extract nouns, verbs and adjectives. Links between adjacent words in a sentence were introduced in the vocabulary. The experiment investigated whether the paths between words in a new sentence from the same blogger shortened after augmenting the vocabulary with evocation, usage links and predicted associations.

The results revealed that, compared to the paths available in the basic hierarchy, ViVA shortened approximately 52% of the browsing paths between words used in a sentence.

1.3.3 Guiding Word-Finding with Semantic Associations for Non-Aphasic Individuals

We first investigated whether users will take advantage of the shorter paths possible due to the semantic association we incorporated in the vocabulary with an able user population. We asked the participants to find, as quickly as possible, the missing words in a number of phrases, using two vocabularies. The first vocabulary, LG, implemented Lingraphica's hierarchical organization and the second one, ViVA, inherited the same hierarchy but was augmented with semantic associations which translated into related words that participants could see once they had selected a word.

The results of the experiment showed that participants took significantly less time to find words with the augmented vocabulary, taking shorter paths guided by the provided semantic associations. All participants agreed that having related words helped them find words faster and most thought that finding words in ViVA was less confusing than searching in the basic hierarchy. Using the augmented vocabulary, people tended to naturally search for words via related-word links instead of concentrating on remembering what category a target word belongs to.

1.3.4 Guiding Word-Finding with Semantic Associations for Aphasic Individuals

To show that semantic word associations can also guide people with aphasia in finding words more effectively, we adapted the experiment we ran with a healthy population, and ran it with aphasic participants. We again asked users to search for words in a hierarchical vocabulary, LG, and in one augmented with associations based on the evocation ratings we gathered, ViVA. To activate the semantic networks around a target concept, the task provided additional context through a scenario comprised of an image and a sentence related to what was portrayed in the image.

Participants found words with ViVA faster, using more effective paths to connect words in completing the stimulus phrases. They also voted ViVA less confusing to navigate and agreed that the associations that it provided helped them find words faster. The study demonstrated ViVA's potential to improve vocabulary navigation and searching for words for users with lexical access impairments as those present in aphasia.

1.3.5 Single-Case Study on Personalizing ViVA

It is challenging to evaluate a tool that adapts according to usage statistics and long-term user preferences in controlled experiments. In order to investigate how ViVA should be personalized for a specific user to reflect his idiosyncrasies and needs, we conducted a three-phase case study with an aphasic individual over an eight-week period.

In the first phase of the study, we collected information on the participants' social and personal life (e.g., family and hobbies) through ethnographic interviews. The data was mined for specific conversation topics and frequently used words that could be used to customize and evaluate ViVA. The collected data was also informative of the language abilities of the participant, his compensatory communication strategies and his use of assistive technology.

In phase two, we mined the collected conversation data for word association specific to the participant. We constructed words pairs reflecting those associations and the participant was then asked to confirm the strength of evocation between some of the resulting pairs. The second phase also revealed the need for alternative and personalized access points to the vocabulary that provide shortcuts to words and topics frequently needed by the user.

In the final phase of the study, we customized ViVA to incorporate the semantic associations extracted from our conversations with the participants, and we introduced personalized access points to the vocabulary. Through a set of structured task, we evaluated the participant's ability to find words using ViVA. The collected data and the evaluation results show the importance of semantic associations, personal and general, in guiding the user. There was also evidence that the personalized access points help the user find words that appear in both familiar and unfamiliar contexts faster.

1.4 Summary of Contributions

The main contributions of this dissertation work are:

1. A novel approach to vocabulary organization in assistive communication tools that enables effective word-finding.
2. A dataset of human judgments of semantic relatedness that can be used to model a user's mental lexicon.
3. Empirical evidence that enhancing a basic vocabulary hierarchy with semantic word-associations improves word-finding compared to existing alternatives.
4. Guidelines for personalizing an assistive vocabulary to fit users' needs and improve word-finding.

In addition to these contributions, we expect that the lessons learned during the design and evaluation of ViVA will have some broader applications. Our approach to building customized semantic networks can be applied to the design of communication tools for people with other cognitive impairments and of educational tools for children and foreign-language learners. Our investigation into the design of an adaptive vocabulary for people with language impairments also demonstrates the potential of using adaptive tools to assist users with other communication and cognitive impairments.

1.5 Thesis Outline

In Chapter 2, we summarize some related work. Chapter 3 describes the design of the Visual Vocabulary for Aphasia and elaborates on the online experiment we used to collect evocation ratings, its results and their implications. Chapter 4 describes evaluating ViVA by simulating usage of the vocabulary with data mined from blogs of elderly people. Chapter 5 and Chapter 6 turn to evaluating ViVA with users without and with lexical-access impairments respectively. Chapter 7 details a case study that was used to investigate how ViVA adapts to a specific user's profile. In Chapter 8, we summarize our contributions and discuss areas for future work.

Chapter 2

Related Work

In this chapter, we review the literature that has informed and influenced our work. We begin with a description of a debilitating language disorder known as aphasia. We then turn to a discussion of some available technological tools created to assist people with aphasia with communication. Because our work is motivated by the belief that the effectiveness of such tools can be enhanced by our knowledge of human semantic memory, we also review some of the language theories that have guided our research. We conclude with a discussion of the techniques we have used to address some of the challenges inherent to the design and evaluation of assistive communication tools.

2.1 Aphasia

Aphasia refers to a family of acquired communication disorders that impact an individual's language abilities. It affects close to one million people in the United States (National Aphasia Association 2010) and is often acquired as the result of a stroke, brain tumor or other brain injuries. Depending on the area or degree of damage to the brain, the resulting impairments to the ability to understand and produce language vary; for example, some people may speak fluently, but have impaired auditory comprehension while others may have impaired speech, but good reading comprehension. Substantial variations in the nature and degree of severity of impairments can be observed in an individual as well as across individuals with aphasia.

There are a number of classifications for aphasia based on characteristics such as the part of the brain that has been damaged or the residual language abilities (Goodglass 1993). Our work does not target a specific type of aphasia but a collection of symptoms associated with the difficulty of accessing and retrieving words from the mental lexicon. Impaired word-finding or anomia is a defining feature that is present across aphasia types (Goodglass & Wingfield 1997). Even though rehabilitation can improve people's abilities, a significant number of aphasic individuals are left with life-long chronic impairments among which anomia persists (Davis 2000, Goodglass & Wingfield 1997).

2.2 Assistive Communication Tools for People with Aphasia

People with aphasia cope with their inability to communicate using different low-tech strategies such as drawing pictures, writing notes, pointing to objects, mimicking and gesturing. Traditionally, assistive tools for people with aphasia have focused on essential therapeutic efforts and the recovery of basic language functions such as spelling and naming. There have also been commercial and research efforts to build technology that helps this user population communicate and thus, regain some of their independence. High-tech assistive communication tools have a number of advantages: 1) they are good for repetitive tasks such as practicing spelling, naming and pronunciation; 2) they provide consistent and unemotional feedback; 3) they have evolved to be mobile and unobtrusive and thus, can assist users in different contexts outside of their home. This last advantage addresses the stigma issues arising from using a bulky computer to speak for the user.

Most efforts have been focused on developing technologies that provide augmentative and alternative communication (AAC) and thus help people with aphasia with communication exchanges. Existing AAC tools for people with aphasia make use of picture-based representations of concepts due to research showing that aphasic individuals retain abilities that can be used to improve communication via visual prompts (Thorburn et al. 1997). AAC tools provide users with a collection of multi-modal icons that combine images, text and speech audio in order to compensate for the heterogeneity in language impairments across individuals, e.g. if a user could not read, they could hear the sound associated with an icon. Such tools support phrase and sentence composition by enabling the user to assemble words in a linear fashion. For example, if the user wants to communicate the phrase *I am hungry*, s/he needs to browse the collection of icons, find the icon illustrating the pronoun *I*, the verb *am* and the adjective *hungry*, and arrange them in the correct order. The resulting communications can also be stored and reused in the future.

TouchSpeak is a commercial example of an AAC tool for people with aphasia (TouchSpeak 2010). TouchSpeak offers a hierarchical vocabulary supplemented with speech audio, phonetic cueing, typing and sketching. It can be customized for each user with the help of a speech therapist and the vocabulary can incorporate graphical symbols or icons as well as images from the user's personal collection or ones that have been taken with the device's camera. Lingraphica (Lingraphicare Inc. 2010), a dedicated device designed specifically for people with aphasia, is another commercial example. We elaborate on its design and functionality in the next section.

Some other assistive devices, marketed for broader user populations, include DynaVox (Dynavox Mayer-Johnson 2010), Proloquo2Go (AssistiveWare 2010) and Vantage Lite by PRC (Vantage 2010). They also provide multi-modal vocabularies whose elements can be pieced together to build communications.

The TalksBac system (Waller et al. 1998) and a related system, PROSE (Waller & Newell 1997), targeted the ability of some higher-functioning aphasic individuals to recognize familiar words and short sentences. TalksBac did not provide multi-modal representations of concepts, but enabled the user to store textual information about their regular conversation partners, favorite subjects and other utterances meaningful to him/her. In a longitudinal study, TalksBac was shown to be helpful for some individuals, although its reliance on caregivers to maintain the system was problematic. Such a heavy reliance on either the user or a caregiver (or therapist) to import and manage the contents of the assistive system (e.g., images and audio) reduces greatly the usability of the tool.

The user was responsible for collecting and annotating the content in the system created by Hine, Arnott, and Smith (Hine et al. 2003). They developed a multi-media communication tool for story-telling. It provided people with communication difficulties means to relate stories that consisted of video and audio clips, and static images. Stories were organized in topics and subtopics. Once a topic was selected, five stories associated with the topic were retrieved. Daemen et al. also recognized the unmet need of aphasics to share their daily feelings and built a free style multi-media interface for storytelling (Daemen et al. 2007). In addition to images, audio and text annotation, the tool provided icons for expressing emotions which could be used to augment the content of a story.

There have also been research efforts in harnessing computer technology to enable aphasic users to manage independently daily tasks such as appointment scheduling (Moffatt et al. 2004), following recipes (Tee et. al 2005) and file and photo management (Allen et al. 2008, Davies et al. 2004).

2.3 Lingraphica

Lingraphica (Lingraphicare Inc. 2010) is a communication device built specifically for people with aphasia. It provides the user with a tool for basic communication and for practicing natural speech. Lingraphica is based on the C-VIC and VIC systems. VIC, which stands for Visual Communication, is an icon-based language developed in the eighties (Gardner et al. 1976) for patients with aphasia. It consists of a limited set of black and white symbols drawn on index cards. Each symbol is a meaningful unit that was chosen over a gestural language and aims to circumvent the visual memory deficits caused by aphasia. VIC was subsequently enhanced by a research team of specialists in neurology, speech language pathology, linguistics, psychology and computer science, and became a computer based interactive system known as C-VIC (Computerized Visual Communication System) (Steele et al. 1989). Both VIC and C-VIC were found to require extensive training in order to be usable (Steele et al. 1989).

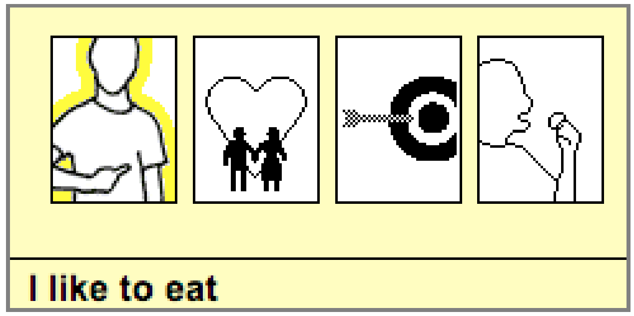

Figure 2.1. A Lingraphica storyboard for the phrase *I like to eat*.

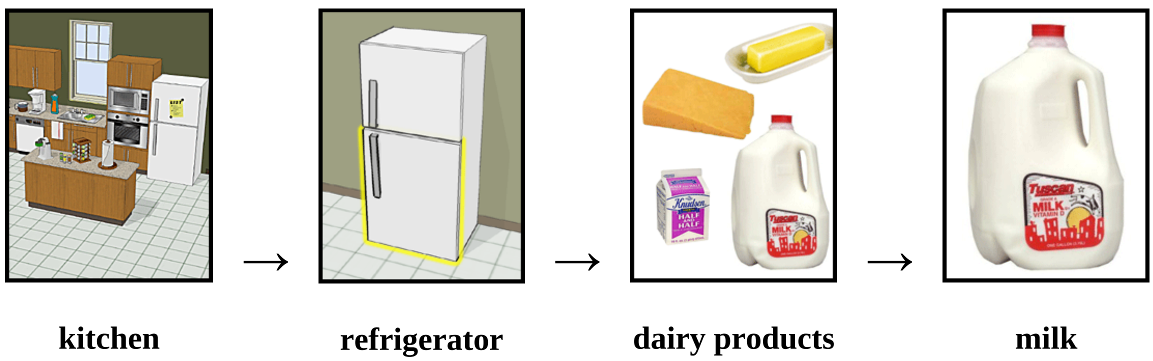

Figure 2.2. In Lingraphica, *milk* is found in the *kitchen*, in the *fridge*, among other *dairy products*. Each word is represented by an icon consisting of a picture, text and speech audio.

Lingraphica extends C-VIC with animation and additional icons to create a vocabulary of over 3000 words (base forms, excluding conjugations and verb tense forms) represented by triplets of image, text and speech audio. Nouns are represented by static images while verbs are represented by an animation depicting the action. The central metaphor of the Lingraphica system is a storyboard (see Figure 2.1), a collection of visual symbols arranged to approximate English syntax. Users of Lingraphica perform hierarchical searches for pictorial representations of the concepts they want to express. They move these graphical icons to the storyboard portion of the screen, combining them to create phrases. Lingraphica's vocabulary attempts to mimic real-life contexts by grouping words according to shared contexts. If you need to find *milk*, for example, you select the icon for *kitchen*, then the *fridge* category and then you find *dairy products*. The icon for *milk* is in *dairy products* (see Figure 2.2). This organization is not necessarily obvious for all users, 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*. At an additional effort of the user, this problem is addressed by enabling him/her to customize the vocabulary categories and icons, and include personal images.

Although Lingraphica is a successful commercial product (approval by Medicare since 2002 has made it accessible to many aphasia patients in the US (Lingraphicare Inc. 2010)), there is concern that its use is predominately limited to the home. Speech-language pathologists who prescribe Lingraphica to their patients and work with it on daily basis have noted that users tend to use the device to practice phrases that have been created for them rather than compose their own sentences and use the device to assist them in conversation at home (for example, talking to a friend) or away from home (for example, at the doctor's office) (Boyd-Graber & Nikolova et al. 2006). One reason for the device's lack of use outside the home may be its form factor. The Lingraphica software comes pre-installed on a dedicated-use Apple iBook and there are several reasons why this form factor may be limiting use: 1) the size and weight of a laptop makes it cumbersome and inconvenient to transport, especially when aphasia is accompanied by motor impairments; 2) the time required to take out the laptop and start up the software, and to find the appropriate communication can reduce the timeliness of the support it provides; 3) the obtrusiveness of a laptop may undesirably interfere with social interactions. It has been noted that assistive technologies that draw attention to the disability often create a stigma hindering adoption (Hirsch et al. 2000). Lingraphicare Inc. has begun addressing these issues with software called SmallTalk (SmallTalk 2009). SmallTalk is Lingraphica's mobile extension (currently with limited vocabulary and functionality) running on an iPod Touch (Apple Inc. 2010).

2.4 Desktop-PDA Assistive System

To address the mobility and stigma issues that have been consistently confirmed by speech-language pathologists, caregivers, and individuals with aphasia, we designed a hybrid Desktop-PDA assistive system. The system is composed of a desktop and a mobile component. The mobile component named ESI Planner II is based on ESI Planner (Moffatt et al. 2004). It attempts to harness the advantages of a mobile personal digital assistant (PDA) shown to benefit one user with aphasia who incorporated the PDA into his daily communication strategies (Davies et al. 2004). It also overcomes the difficulties encountered by Moffatt et al. in evaluating ESI Planner. Many individuals struggled with input on the PDA because of the limited screen size and interaction capabilities afforded by the device. We overcame these limitations by incorporating the mobile component with a desktop component similar to Lingraphica (Boyd-Graber & Nikolova et al. 2006).

The Desktop-PDA system uses the lexical elements of Lingraphica for the task of composing appointments, reminders, phrases and checklists on a desktop computer. These are then transferred to ESI Planner II for portable use (see Figure 2.3). We envisioned that aphasic users would use our system independently to create phrases at home with Lingraphica for a future activity, associate these phrases with a particular date and time, and then automatically transfer these phrases and appointments to the PDA for use outside the home.

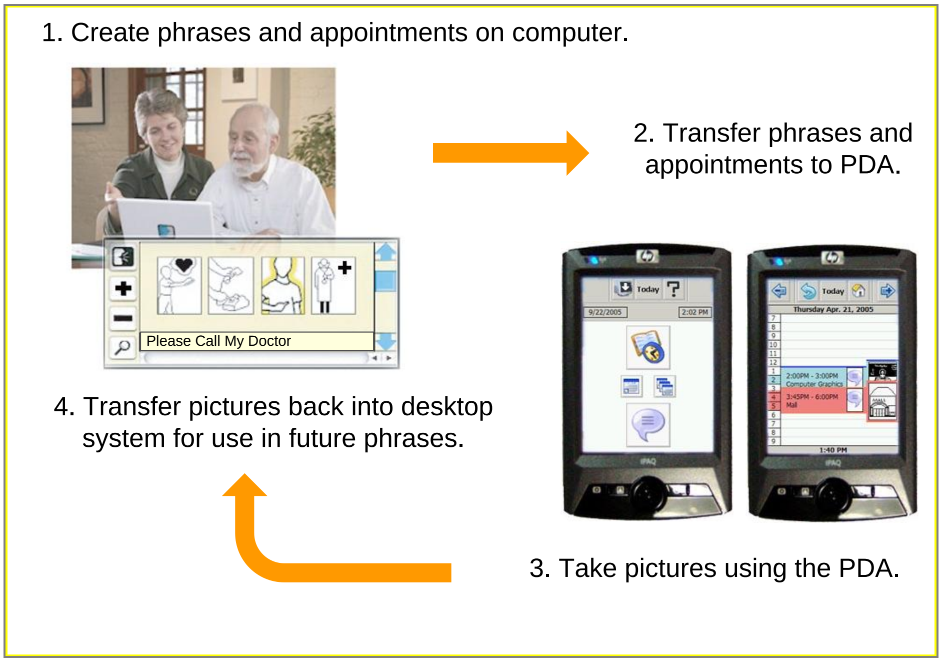

Figure 2.3. Desktop-PDA assistive system that helps people with aphasia with appointment-related communications.

During a discussion with speech language pathologists about the strengths and weaknesses of Lingraphica, it was revealed that it has a steep learning curve which often hinders adoption. To be able to evaluate ESI Planner II, we built a simpler interface to serve as the desktop component of the system that we called LgLite. The design of LgLite minimizes the complexity of Lingraphica by eliminating many of the hierarchical searches required for navigating Lingraphica's extensive vocabulary. Replacing Lingraphica with LgLite reduced the scope of the system's vocabulary to that which is pertinent to appointment creation and it allowed us to incorporate a feature that we were interested in exploring but that had not been possible with Lingraphica at the time of the initial design – the introduction of custom photos into the image library. A medium fidelity prototype of the system was evaluated with aphasic individuals from the Adler Aphasia Center (Adler Aphasia Center 2010). The participants seemed to enjoy and were able to use the system, but a need for a more flexible and a customizable interface was revealed (Boyd-Graber & Nikolova et al. 2006). For example, some users wished concepts to be represented only with text and sound while others found only the image-sound pair useful. All participants took advantage of the ability to take pictures with the PDA and the majority used them as means for communication. A few participants took multiple pictures to complement the limited vocabulary we provided which highlighted the need for a more expressive vocabulary. This in turn led us back to the problem of users not being able to navigate effectively a more sizable vocabulary which can be a very cumbersome task when using vocabulary hierarchies such as Lingraphica's.

2.5 Addressing the Need for Flexibility and Customization

The findings from the evaluation of our Desktop-PDA prototype (Boyd-Graber & Nikolova et al. 2006) confirmed the need for flexibility and customization in assistive technology for people with aphasia which is consistent with previous results (e.g. Moffatt et al. 2004, Sutcliffe et al. 2003, van de Sandt-Koenderman 2004). Enabling users to add their own photographs to the image library of the Desktop- PDA system was only a step towards satisfying this need and providing a customizable vocabulary. Other tools such as Lingraphica (Lingraphicare Inc. 2010) and TouchSpeak (TouchSpeak & TypeSpeak 2010) have also enabled users to customize their vocabulary by incorporating personal multi-media such as videos and by creating personal categories of words such as family-related items. While customization empowers the user in many cases, it also introduces an additional burden of managing the introduced changes, which could be especially challenging for people with cognitive impairments.

All users tend to rely on consistency and stability within an interface. It is not surprising that this dependency is even more pronounced among users with cognitive impairments. In order to address this concern and still create a flexible vocabulary interface, we explore a ***mixed-initiative approach*** to customization, an effective blend of automation and direct manipulation (Horvitz 1999). We chose a mixed-initiative design, because it has been recognized as the approach that has the greatest potential of bringing balance between user and system control, providing for a reliable and satisfactory interaction. The resulting vocabulary solution is both ***adaptable***, able to be customized by the user, and ***adaptive***, able to dynamically change to better suit past actions and future needs.

Early research efforts in applying adaptive techniques to improving interaction with computer interfaces date back to 1985 and the first rigorous study of adaptation (Greenberg & Witten 1985). More recent examples include investigating the benefits and drawbacks of adaptive menus (Findlater 2009) and automatically generating user interfaces to fit user preferences, specific devices, usage and user's motor abilities (Gajos 2008). For a detailed review of adaptive human-computer interaction we direct the reader to Benyon (Benyon 1993), and Norcio and Stanley (Norcio & Stanley 1989). However, only since the late 1990s researchers have started introducing a measure of intelligence in systems designed to support elderly people and people with cognitive disabilities (Fink et al. 1998). Much of that effort has been concentrated on web accessibility, and scheduling and prompting systems. For example, Fink, Kobsa and Nill (Fink et al. 1998) proposed to automatically modify HTML code in response to adaptive user profiles, with the primary purpose of making web pages more accessible to the disabled and elderly. Through visual and audio clues and prompts, PEAT (Levinson 1997) provides mobile support for plan execution for patients with traumatic brain injury. Both PEAT and Autominder (a cognitive orthotic system helping elderly people adapt to cognitive decline) (Pollack et al. 2003) build and maintain a detailed model of the client's plan and step through its execution. Autominder is more ambitious in that instead of only providing automatic reminders, it reasons about when and whether it is the appropriate time to issue a prompt. The Assisted Cognition Project (Kautz et al. 2002) develops systems that compensate for the loss of some of the memory and problem-solving abilities observed in Alzheimer's patients through the use of ubiquitous computing and artificial intelligence. Other projects that attempt to provide intelligent aids for cognitive support include MAPS (Carmien 2002), i.l.s.a. (Haigh 2006) and AVANTI (Stephanidis 1998).

Investigation into adaptive AAC tools for people with aphasia is very limited even though it was first highlighted as an alternative approach for some people with aphasia with the design of the predictive text-based TalksBac system (Waller et al. 1995, Waller & Newell 1996). TalksBac enabled users to create and store communication content and associated it with different topics and conversation partners. The tool assisted the user by having him/her specify a topic or a conversation partner and letting the system suggest relevant communications. Lingraphica, on the other hand, can assist the user with word prediction if s/he is able to type the first few letter of the word they are looking for.

The work presented in this dissertation contributes to this sparsely explored research area (in its broader applications, often termed intelligent systems for assisted cognition (Kautz et al. 2002)) in that we employ adaptive and adaptable techniques to improve word-finding in assistive communication tools.

2.6 The Difficulties of Navigating Assistive Vocabularies

An essential component of a communication system that attempts to be flexible, extensible, and expressive is the vocabulary that it offers to the user. To be expressive, assistive communication 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 (Beukelman & Mirenda 2006). It is challenging to provide functional communication assistance by enabling the user to select words quickly and effortlessly. Many people with aphasia cannot always type the desired word in a search box, but instead have to browse through an extensive vocabulary of pictorial representations of concepts until they find the one that expresses their current communication needs. Minimizing the complexity of navigating the vocabulary and supporting efficient word-finding is essential to the usability of AAC tools.

The ease of word-finding depends on how the words are organized. Most existing assistive visual vocabularies have a lexical organization scheme based on a simple list of words, a list of categories of words or a hierarchy of categories and subcategories. While there is no consensus on what the best way to organize an assistive vocabulary is, speech-language pathologists suggest that fewer categories are less confusing and easier to navigate (Boyd-Graber & Nikolova et al. 2006). However, populating a category with a sufficient number of words introduces the problem of excessive scrolling. The VocaSpace, featured in an augmentative and alternative communication software product called Proloquo2Go (AssistiveWare 2010), takes this category-centered approach. The words in VocaSpace are organized in functional categories, e.g. greetings and questions, and common word categories such as colors, places, and clothes. Unfortunately, no published evaluation on its effectiveness is available. On the other hand, commercial tools such as TouchSpeak (Touchspeak & TypeSpeak 2010) and (Lingraphicare Inc. 2010) offer categories of words that are organized in a hierarchy. For example, the *dinner* category is a child of the *meals* category, which is in turn a child of the *food* category. There have also been efforts to improve on traditional linear syntax and word-finding by providing phrase starters and semantic rather than syntactic schemas (e.g., AssistiveWare 2010, Lingraphicare Inc. 2010, Patel et al. 2004). Despite this assistance, the user still has to search through the vocabulary to find most of the concepts s/he wishes to express.

Each organization has some disadvantages. A hierarchy can help a user build, over time, a mental model of a vocabulary but it often leads to deep and confusing searches. Well-populated common word categories such as a collection of food items could be obvious as an organization but can result in excessive scrolling. More abstract categories, on the other hand, may introduce a sense of disorganization especially since speech-language pathologists suggest that the majority of people with aphasia have difficulties abstracting a word to its superordinate (e.g., finding *sad* in a *feelings* category). 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 and that having to perform multiple clicks to reach a single word is a disincentive for the user. These informal studies, feedback from speech-language pathologists and previous research work (e.g., Beukelman & Mirenda 2006, Boyd-Graber & Nikolova et al. 2006) underlined the need to investigate effective approaches to vocabulary organization and word-finding.

2.7 Language Organization and Retrieval Theories

To address the problem of cumbersome vocabulary navigation, we propose a vocabulary structure that is based on theories that explain how the human mind organizes words. We 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 occasional inconvenience of being unable to retrieve a certain word (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. While for people without language impairments this inconvenience is caused by temporary impaired semantic connections in the mental lexicon, for people with anomic aphasia the problem is more severe and persistent.

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* (Moss & Older 1996). Semantic priming (Swinney 1979), 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 (Collins & Loftus 1975) 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 (Fellbaum 1998, Miller 1990) is a large-scale lexical database inspired by network theories of semantic memory that accommodate the spreading activation paradigm among related words and concepts. Taking advantage of the knowledge encoded in WordNet, we attempt to build a system that can compensate for some of the missing semantic connections in a user's mental lexicon. WordNet groups synonymous words together into sets, called "synsets." Depending on the part of speech, synsets are interlinked according to specific meaningful relations. For example, *eat* (to take in solid food) is linked, among other verbs, to *lunch*, *breakfast* and *picnic*. The noun *lunch* is associated with other nouns such as *dinner*, *feast* and *brunch*. We hypothesize that such relationships could guide users in searching for words and make word-finding more efficient and the interaction with the vocabulary more satisfactory.

2.8 Semantic Therapy

Researchers in the field of speech-therapy have recognized the potential of semantic associations in helping people with aphasia rebuild some of the impaired links in their mental lexicon. As part of a treatment, Webster et al., for example, employed a training strategy in which the participant was asked to generate words that are associated with a target verb (Webster et al. 2005). This strategy, termed *Semantic Feature Analysis* (SFA), is an approach mostly applied to treating the retrieval of nouns and verbs and has shown great potential (Boyle & Coelho 1995, Edmonds et al. 2009). SFA is used to guide the patient in identifying important semantic features of the target word. Given the word *apple*, for example, the patient is asked to identify the properties of an apple, its use, where you find it, what actions apply to it, any superordinate categories and other associations (Boyle 2004). This exercise helps activate the semantic network surrounding the target word, apple, and consequently aids in its retrieval (Boyle & Coelho 1995). In the process of identifying features of the target item, non-targeted and semantically related words may also benefit from the treatment, because they share features that are being accessed or retrained. We build on these positive findings and create semantic networks that the user can navigate in search of a specific word. These networks model the organization of words in a person's mental lexicon and can compensate for some impaired connections, successfully guiding the user to their desired word.

Guided by the theories described above and encouraged by the positive results of semantic feature analysis in the treatment of people with anomia, we incorporate semantic associations in an assistive vocabulary, providing a novel approach to vocabulary navigation. The vocabulary incorporates human judgments of semantic relatedness and it dynamically reassigns the links between words based on user preferences, word frequency usage and predicted associations.

2.9 Designing Tools for People with Aphasia

There have been a few challenges, specific to designing for our target user population, which we had to address while completing this dissertation work. Here we name a few of them and discuss the approaches we have taken to overcome them.

The first challenge is inherent to our target user population – it is difficult to investigate user needs, and gather user requirements and feedback from individuals with communication difficulties. We addressed this problem by procuring the feedback of domain experts such as speech-language pathologists at the early stages of the design process and reaching out to users at the later stages. We also developed collaborations with local support groups for people with aphasia in order to be able to attend meetings, observe their interactions and communication strategies, and incorporate the insights into our work.

The second challenge is recruiting study participants from a limited and isolated pool to evaluate ViVA. Social isolation and depression, result of the inability to communicate, are common among people with aphasia (Kauhanen et al. 2000 and Martin et al. 2002). This made it difficult to find and engage people with aphasia in our projects. We addressed this problem by strategically using feedback from domain experts to supplement and generalize the feedback from target users which has been shown to be a successful design approach. Other researchers have also examined the use of non-target individuals in designing tool for people with communication difficulties. Both (Moffatt 2004) and (Davies 2004) also suggested using advocate users, that is, aphasic individuals who are, for one reason or another, better able to contribute in a participatory design setting. In addition, we have fostered a connection with a local aphasia support center and a support group for Lingraphica users. This has enabled us to connect with people with aphasia, conduct ethnographic observations of their use of assistive technology, and recruit them as participants in some of our studies.

ViVA's adaptive component introduces an additional challenge in terms of user evaluation. ViVA's goal is to adapt over time to an individual user's profile and to help him/her find the words specific to their needs faster. To evaluate its potential, one would need to train and test it on a large scale corpus of transcribed communications generated by users of assistive communication tools. The corpus would also need to cover a variety of topics. Unfortunately, as highlighted by other researchers, no such corpus exists (Trnka & McCoy 2007). A potential source of similar data is available from repositories of switchboard exchanges (Switchboard 2007). However, there would not have been an easy way to filter the data such that they are appropriately representative of our target user population. Instead, as a basis for our vocabulary, we chose Lingraphica's collection of words. It provided us with a well-structured vocabulary that has evolved over the years with the goal of meeting the needs of the user population. The vocabulary also incorporates feedback from professionals such as speech-language pathologists. In addition, we also created our own corpus from text describing everyday life experiences of elderly bloggers. Even though healthy elderly people lead more active lives than people with aphasia, they have similar social interactions and needs, which makes their communications relevant to the type of sentences our system will be processing.

Finally, in addition to controlled usability studies, we chose to conduct a longitudinal case study with a single participant. The study enabled us to understand how ViVA can be personalized to meet our user's needs and how it adapts to the user's profile. Conducting such an investigation on a larger scale would have been unmanageable in the time span of a graduate tenure. While the case study does not provide generalizable solutions to the problem of vocabulary organizing and effective word-finding, it reveals considerations for further investigation and guidelines for working with our target user population.

Chapter 3

Design of the Visual Vocabulary for Aphasia

In this chapter, we elaborate on our mixed-initiative approach to creating a vocabulary that is both customizable by the user and automatically adapts to suit users' past actions and future needs. We describe how semantic similarities play a role in ViVA's design and report results from an online experiment that we conducted to collect a large set of evocation ratings.

3.1 Vocabulary Design Components

All users tend to rely on consistency and stability within a user interface. This dependency is even more pronounced among users with cognitive impairments. In order to address this concern and still achieve our goal of creating a flexible vocabulary, we built a tool that relies on a mixed-initiative approach to customization (Horvitz 1999). The Visual Vocabulary for Aphasia (ViVA) is ***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 approach enables the user to be in control by making changes to the vocabulary and anticipating ones that have been initiated by the system. Thus, ViVA provides stability of the vocabulary structure while automatically predicting associations between words that could help the user navigate to the target concept. ViVA has two modules – the user preference module and the active learning module (Figure 3.1).

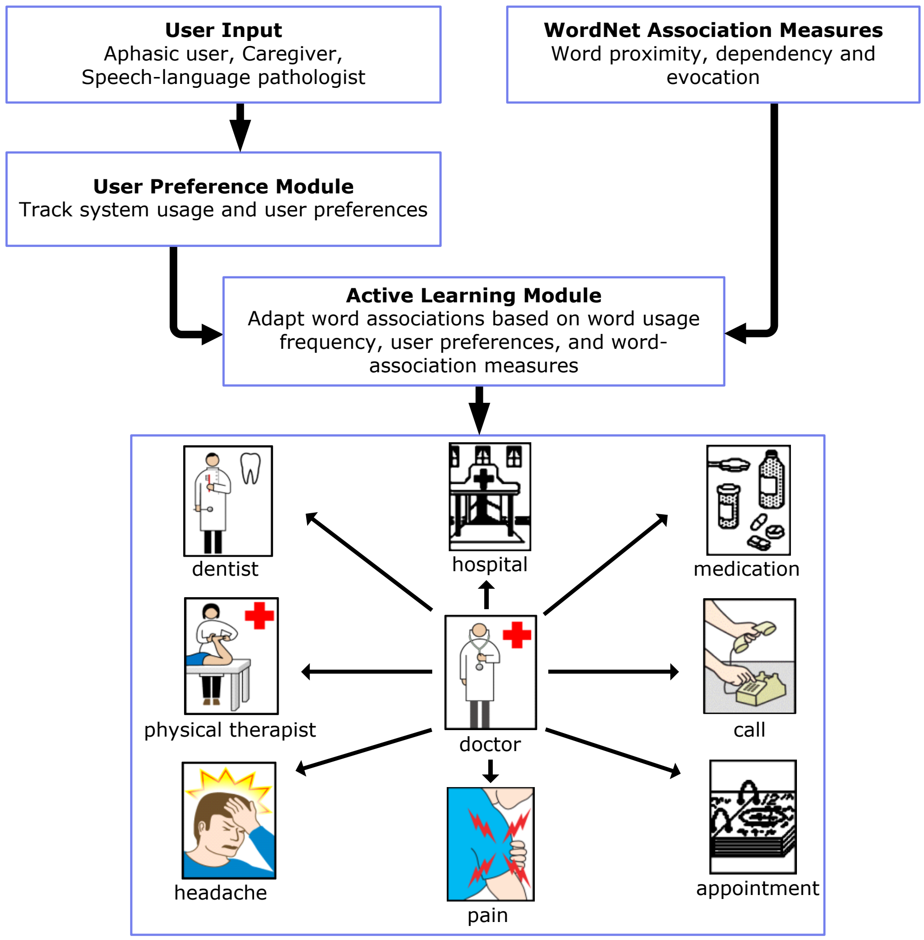

Figure 3.1. Schematic of components of the Visual Vocabulary for Aphasia.

The ***user preference module*** implements the adaptable component of our mixed-initiative vocabulary. It keeps track of changes initiated by the users such as adding and removing vocabulary items, grouping words in personalized categories (for example a *Hobbies* folder or words related to *Family*), and enhancing words with personal images and sounds. In addition to practical concerns of having sufficient vocabulary terms to express the needed concepts, the ability to adapt a system invests in the user a sense of ownership and empowerment. This attachment to the system, brought about by a sense of accomplishment, is an important aspect of the rehabilitation process (Allen et al. 2008).

ViVA's adaptive component is implemented by the ***active learning module***. We explain its role with an example: if the user wishes to compose the phrase *need appointment with doctor* and s/he searches for *doctor* first, the vocabulary network centered on *doctor* may look as the one shown in Figure 3.1. The links between the words may exist because the user has previously composed sentences using *doctor* and *medication* or using *doctor* and *appointment, hospital* and *doctor,* for example, may be linked because of a prediction based on known word association measures and usage statistics. We explain in more detail the role of semantic association in the next section.

ViVA's goal is to seamlessly incorporate user preferences and word frequency usage, and to adapt the organizational structure of the vocabulary to better suit usage needs. However, we still need an initial organization to allow the user to successfully use ViVA from day one. We derive this scaffold from the body of work investigating how the human mind organizes concepts into a mental lexicon.

3.2 WordNet and Evocation

WordNet is a large-scale lexical database with a rich structure that connects its component synonym sets (synsets) according to a variety of relationships. Noun synsets are interlinked by means of hyponymy, the *super-subordinate* or *is-a relation*, as exemplified by the pair [poodle]-[dog][1]. Meronymy, the *part-whole* or *has-a* relation, links noun synsets like [tire] and [car] (Miller, 1990). Verb synsets are connected by a variety of lexical entailment pointers that express manner elaborations [walk]-[limp], temporal relations [compete]-[win], and causation [show]-[see] (Fellbaum 1998). The links among the synsets structure the noun and verb lexicons into hierarchies, with noun hierarchies being considerably deeper than those for verbs.

We exploit the structure of WordNet to help "find" intended concepts and words by navigating along the paths connecting WordNet 's synsets. However, the internal density of WordNet is insufficient – there are too few connections among the synsets. Boyd-Graber et al. (Boyd-Graber et al. 2006) attempted to create thousands of new links that go beyond the relations specified in WordNet. The authors introduced the measure of ***evocation***, a measure of how much one concept brings to mind another.

In summary, evocation (Boyd-Graber et al. 2006) aims to add cross-part-of-speech links, connecting nouns to verbs and adjectives. Such syntagmatic relations 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 such concepts as [eat], [hungry] and [food] should be encoded in WordNet. This

[1] We follow the convention of using a single word enclosed in square brackets to denote a synset. Thus, [dog] refers not just to the word dog but to the set – when rendered in its entirety – consisting of {dog, domestic dog, canis familaris}.

work (Boyd-Graber et al. 2006) also addressed another shortcoming of WordNet, namely the absence of weights that indicate the semantic distance between the members of related pairs. These human judgments of evocation were collected via a laborious, expensive method. Undergraduate students were put through a training and vetting process to consistently rate pairs of synsets through a specially designed interface. Because the pairs of synsets were randomly selected, many of the ratings, as expected, were zero. We originally hoped that these initial ratings, collected over the course of a year and with a significant outlay of time and money, would allow us to automatically label the rest of WordNet with directed, weighted links. However, machine learning techniques could not reliably replicate human ratings. In Section 3.4, we describe our method of collecting additional empirical similarity ratings using a far less expensive annotation strategy. First, we describe how we built ViVA's core vocabulary.

3.3 Core Vocabulary

Customization of the vocabulary can be a powerful feature, but it is necessary to have an initial set of words to allow the user to successfully use ViVA from day one. We selected ViVA's initial vocabulary set such that it is a collection of commonly used words as well as ones relevant to our target population, people who have aphasia. ViVA's core vocabulary was mined from two sources: the "core" WordNet consisting of frequent and salient words selected for collecting evocation data (Boyd-Graber et al. 2006) and the visual vocabulary of an assistive device for people with aphasia created by Lingraphicare (Lingraphicare Inc. 2010), described in the background section of this dissertation.

The core WordNet consists of two sets of 1000 and 5000 salient words. The sets were constructed by collecting the most frequently used words from the British National

Corpus (The British National Corpus 2007). Each word from the resulting list was then assigned its most salient meaning available in WordNet (Boyd-Graber et al. 2006). We used all synsets from the core 1000 synsets used by (Boyd-Graber et al. 2006), all verbs in Lingraphicare's vocabulary, and all nouns and adjectives in both Lingraphica's vocabulary and the core 5000 synsets (WordNet Research @ Princeton 2010). We chose to use Lingraphica’s word collection as ViVA’s core vocabulary, because it consists of words commonly used and needed by our target user group. In addition, the device has also been sold and used for a number of years and thus, has been vetted by both users and domain experts such as speech language pathologists. Finally, it provided us with a well-developed multi-modal vocabulary so we could concentrate on word organization and not have to collect or design visual representations of each concept in the vocabulary.

Figure 3.2. Different Lingraphica icons for *glass* enable sense disambiguation.

Lingraphica's multi-modal vocabulary consists of icons that combine text, a pictorial representation of the concept and speech output of the text. We used the pictorial representation to perform a form of coarse disambiguation. For each concept in Lingraphica's vocabulary, based on the pictorial representation, we selected the corresponding WordNet sense to create a single, unified representation of the vocabulary. For example, Lingraphica has images for the following definitions of *glass*: a) drinking glass, and b) a brittle transparent solid (see Figure 3.2). Thus, both senses were represented and were matched to the appropriate icon. On the other hand, the only pictorial representation for *bitter* is for the sense related to taste as opposed to acrimonious, resentful, for example. The resulting vocabulary consists of approximately 1300 words which were organized according to Lingraphica's hierarchy.

Even though Lingraphica's hierarchy was chosen as the basic organization, our goal was to augment it and form a vocabulary network that incorporates links between words which reflect usage frequencies and semantic association. As illustrated in Figure 3.1, such a network would make finding related concepts faster by connecting them directly and compensating for impaired links in a person's mental lexicon. In the following section, we elaborate on our experiment of collecting semantic associations to supplement the basic hierarchy.

3.4 Collecting Evocation Ratings from Online Annotators

Many natural language processing tasks such as determining evocation require human annotation that is expensive and time-consuming on a large scale. Snow et al. (Snow et al. 2008) demonstrated the potential of Amazon's Mechanical Turk (Amazon.com, Inc. 2009) as a method for collecting a large number of inexpensive annotations quickly from a broad pool of human contributors. Their experiment illustrated that labels acquired through Amazon Mechanical Turk (AMT) from non-expert annotators are in high agreement with gold standard annotations from experts (Snow et al. 2008).

Amazon Mechanical Turk (AMT) is an online community that provides access to a vast pool of individuals, called workers, who complete tasks for a monetary reward. The people who post work on AMT, called requesters, design a task and rules for completing it, determine the reward, and release it in the pool of available work. Workers browse the work pool and select the tasks they wish to do. Once a task has been completed, its requester can approve or reject the results. The approval rate, which provides a minimal level of quality control, is part of a worker's profile visible to all requesters. It represents how often the individual's work has been found satisfactory. The positive results reported by Snow et al. (Snow et al. 2008) motivated us to collect evocation ratings to be used in the visual vocabulary for aphasia through an AMT experiment that we describe next.

3.5 Experimental Methodology

3.5.1 Building the Dataset

We used a machine learning algorithm to form the synset pairs to be rated via Amazon Mechanical Turk annotators. We used many of the features found to be predictive of evocation including those based on WordNet connectivity (Jiang & Conrath 1997), pointwise mutual information based on words appearing in the same sentence, and context similarity. We duplicated high evocation pairs (having a median rating greater than 15, on the scale of 0 to 100) to create a high-recall training set and trained a classifier using AdaBoost (Schapire 2003). The pairs selected to be rated via AMT were the subset of all pairs of synsets labeled as having high evocation predicted by the machine learning algorithm.

3.5.2 Task

The final set of synset pairs was split into 200 tasks consisting of 50 pairs each. The design of the template we posted on AMT was closely modeled after the computer program used by Boyd-Graber et al. (Boyd-Graber et al. 2006) to collect ratings from undergraduate annotators. Anchor points on a scale from 0 to 100 were available to rate evocation (Figure 3.3).

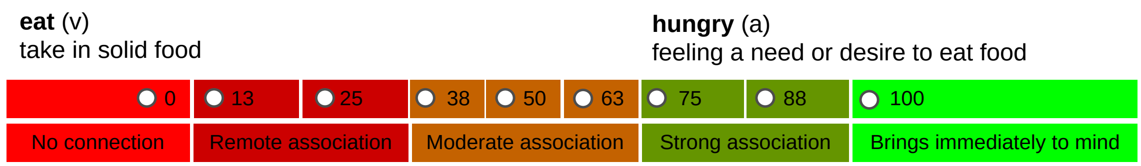

Figure 3.3. In the AMT experiment subjects are asked to rank how much the first word brings to mind the second word.

Raters were first presented with the following set of instructions:

1. Rate how much the first word brings to mind the second word using the provided scale.
2. The relationship between the two words is not necessarily symmetrical. For example, *dollar* may evoke *green* more than the reverse.
3. Pay attention to the definition of the words given on the second line; words can have more than one meaning. For example *dog* (the animal) would not bring to mind *bun* (the piece of bread you serve with a hot dog).
4. The letter in parenthesis signifies whether the word is ***a***: an adjective, ***n***: a noun or ***v***: a verb.
5. Don't use information from your personal life. For example, if you had a dog named bog you personally would associate bog and dog, but the average person wouldn't.

6. Don't use the spelling of words to make your decisions. For example, even though bog and dog rhyme, they are not associated.
7. We cannot offer you a big reward for your time, but we greatly appreciate your sincere effort. There are a few pairs with known average ratings embedded in the task. If your ratings for those pairs do not fall in a generously set acceptance bound, we will have to reject your responses.

The last instruction was included to forewarn annotators that sloppy contributions such as clicking all zeros will not be rewarded. We embedded five checks, unknown to the annotators, in each task which were later used to determine the validity of the gathered results. Annotators were paid $0.05 to complete a task.

We did not gather any demographic information from the AMT workers, but only ones with approval ratings of 97% or higher were recruited to complete the tasks.

3.6 Results

We collected ratings for 107,550 synset pairs over a period of three months. The average time to complete the task of rating 50 pairs was 3.6 minutes, resulting in an average pay of $0.83 per hour. To ensure the quality of the ratings, we compared the results against the ones provided by Boyd-Graber et al. (Boyd-Graber et al. 2006) by using the five embedded checks in each task of 50 pairs to allow us to filter the unreliable results. The ratings for four of those checks were collected from the Boyd-Graber et al. dataset. The fifth check required annotators to rank a pair consisting of the same synset, for example [help] and [help]. We ran three different reliability tests depending on the number of checks we wanted satisfied. If the annotator's rating for the fifth check was 100 and a number of the remaining checks were met within certain acceptance bounds, the annotations were considered reliable. The acceptance bounds were defined as follows.

The scale of 0 to 100 was split into 5 intervals, [0-10), [10-30), [30-70), [70-90), [90-100]. If an annotator's rating fell within the same interval as the corresponding check or in the upper half of the immediately lower interval or the lower half of the immediately higher interval, the rating was considered reliable. The first reliability test required ***all*** checks to be met. For this set, 43.4% of the pairs were rated as having no association and 2.7% fell in the category immediately brings to mind. The second reliability test required ***most***, three or more, checks to be met in addition to satisfying the complete-evocation check. The final and most relaxed reliability test required ***some***, two or more, checks to be met in addition to the complete-evocation check.

Table 3.1. Correlation of the mean and median of the AMT ratings against evocation annotations collected by trained undergraduate annotators.

Filtering Method	**Correlation with Mean**	**Correlation with Median**	**Number of Pairs**
All Checks	0.54	0.54	46900
Most Checks (4 or more)	0.45	0.43	55400
Some Checks (3 or more)	0.37	0.34	56850

Table 3.1 shows the number of synset pairs for each of the reliability levels, and Table 3.2 has explicit examples of mean evocation ratings for the three levels. Finally, Table 3.1 also shows mean and median correlation of the three reliability sets against the ratings provided by undergraduate students in (Boyd-Graber et al. 2006). As expected, there is more variability when fewer checks are applied. The set of synsets where all checks were met resulted in the highest correlation to the original evocation data. This correlation on a very difficult task is sufficient to show that with good quality control, gathering ratings through AMT was a valid approach. While AMT annotators seemed to rate, on average, evocation lower than the trained annotators (Figure 3.4), ratings from the untrained online annotators correlated well (0.54) with those collected by trained undergraduate annotators.

Table 3.2. Examples of mean evocation ratings, on the scale of 0 to 100, given three different methods to ensure rater reliability. For comparison, evocation ratings from trained annotators are also shown.

Untrained Annotators			**Trained Annotators**	**Rated Pair**	
All Checks	***Most Checks***	***Some Checks***		**Synset 1**	**Synset 2**
50	10	61	88	trust.v.01	responsible.a.01
39	44	41	44	surgeon.n.01	responsible.a.01
25	18	22	42	deservingness.n.01	exceed.v.02
29	30	30	20	television receiver.n.01	performance.n.02
46	57	62	19	log.n.01	leaf.n.01
34	33	31	16	diligence.n.02	craft.n.04
25	20	27	16	abundant.a.01	harmony.n.02
23	19	18	0	eyelid.n.01	wrist.n.01
25	28	26	0	reason.n.02	reference point.n.01
4	5	9	0	spread.n.05	pill.n.02

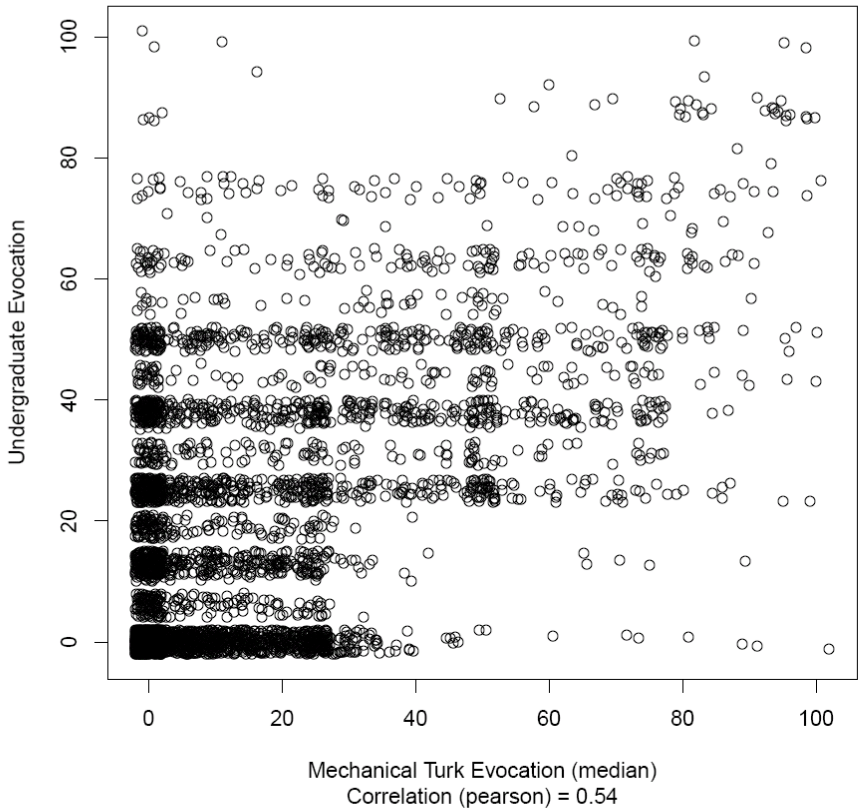

Figure 3.4. Ratings from untrained online annotators correlated well (0.54) with those collected by trained undergraduate annotators.

3.7 Discussion of Results

When interpreting the results, it is important to bear in mind the difficulty of the task. First, the nature of the task was such that we asked the participants to actively *produce* a rating, rather than to agree or disagree with a pre-set judgment or to select one from a few pre-defined options. Second, the ratings were to be expressed using anchored points on a scale from 0 to 100, thus allowing for – and in fact, encouraging – subtle judgments that permitted significant disagreement. Third, while we controlled for intra-rater reliability, we did not know who our raters were in terms of educational level, literacy, and familiarity with the words and concepts that were presented. Indeed, we had no way to ascertain that the raters were native or near-native speakers of English. Finally, the raters might have received insufficient training given the cognitive demands of the task.

The results must be compared to those obtained in the carefully controlled study reported by Boyd-Graber et al. (Boyd-Graber et al. 2006). At the outset of that experiment, it was unknown whether any reasonable reliability could be obtained at all, as the authors were well aware of the difficulty of the task, for which no precedent existed, and they considered the results encouraging. The raters used by Boyd-Graber et al. (Boyd-Graber et al. 2006) came from a small, homogeneous pool – Princeton undergraduate students – whose identity the authors knew and whom they trained carefully and provided with personal feedback. In light of the different methods of data collection, the results of the study described above are comparable.

The inherent noise in the collected evocation data should also be viewed as a positive result. We asked participants to try to abstract their personal experience and background from their answers, but the reality is that an individual's mental lexicon evolves throughout a person's life reflecting his own experiences and shaping semantic association that sometimes cannot be easily generalized. Thus, the variability in the data reflects idiosyncrasies in world knowledge; only by accepting this reality and incorporating it into assistive technologies can we hope to build devices that can truly help a heterogeneous target population. ViVA's goal is to provide users with a vocabulary that can meet common user needs, but also that adapts over time to incorporate these idiosyncrasies and thus assists each individual user better.

Even though our strictest reliability test invalidated half of the collected data, using Amazon Mechanical Turk to gather evocation is still more efficient and economical compared to using a small group of trained annotators. This resulting collection of reliable evocation ratings adds on to the scaffolding of our assistive vocabulary by providing meaningful links between words. Such links will compensate for impaired access to the user's “mental lexicon” and assist him/her in communicating. A network of words whose organization reflects human semantic memory has the potential to help users with anomic aphasia navigate the vocabulary effectively and thus find the concepts they are trying to express faster.

Finally, the AMT approach to collecting semantic data also introduces an exciting avenue for interface design by allowing hundreds or thousands of individuals with diverse backgrounds to help inform the design of assistive technology. It ensures that no single person has unreasonably shaped the interface and that the design reflects a broad spectrum of society.

Chapter 4

Simulated Sentence Construction with ViVA

In Chapter 3, we presented the Visual Vocabulary for Aphasia which is designed to enable effective word-finding. To compensate for impaired semantic links in a user's mental lexicon, ViVA builds a rich vocabulary network by incorporating word frequency usage, user preferences and semantic word associations.

ViVA's goal is to adapt to a user's needs and communication patterns in order to better assist him/her in finding words to build functional communication. It would be challenging to evaluate an adaptive tool – one that is supposed to assist people with everyday tasks – through controlled user studies because not enough time is available to carry out the adaptation. On the other hand, running longer studies that ask users to incorporate an early and possibly unstable prototype in their everyday life could cause significant frustration, especially when the individual is experiencing cognitive difficulties. Thus, as a first step to evaluating ViVA, we concentrated on assessing its backend rather than user-initiated adaptive functionality.

In the field of augmentative and alternative communication, adaptive functionality has been gaining prominence mainly for word prediction in communication tools designed for people who use letter-by-letter spelling, but have no cognitive or language impairments (Trnka et al. 2009). The convention for testing such a tool is to first train it on a dataset representative of the system's intended usage. Once the system has been trained, its performance is assessed using a second test set of representative data. The optimal training and testing datasets would be built from transcribed communications produced by users of assistive communication tools and covering a variety of topics. Unfortunately no such corpus exists (Trnka & McCoy 2007). No corpus of the recorded communications of people with language impairments such as aphasia exists either. Compiling an aphasia-specific corpus would be challenging due to the different modes of communication people with aphasia employ to compensate for their language difficulties (e.g., drawing pictures and using multi-modal icons). A recent effort to collect multimedia interactions for the study of communication in aphasia is underway and, among other contributions, could result in such a corpus (AphasiaBank 2010).

To address the lack of available data that would allow us to train and test ViVA's adaptive module, we collected a small set of transcribed communications. We gathered text from blogs of elderly people and used the basic semantic information contained in it to simulate usage of the vocabulary. After training the vocabulary with the simulated usage data, we examined how the system performed with new sentences composed by the same blogger. The test program searched automatically for related words in two vocabularies, ViVA and Lingraphica's vocabulary hierarchy, and calculated the corresponding browsing distances between words. We hypothesized that by enhancing the vocabulary with known and predicted semantic associations, ViVA would be able to make it easier to find related words in the process of composing a phrase because meaningful user-specific connections discovered from usage statistics would be available.

We used the vocabulary hierarchy of Lingraphica as our core organization and as the baseline for comparison in the simulation experiment. To our knowledge, Lingraphica is the only commercial device designed specifically for people with aphasia and widely used in the United States of America. Its design has evolved through a number of years based on the input of speech language pathologists and users with aphasia; thus, it presents a realistic and practical standard to evaluate against.

4.1 Vocabulary Usage Profiles

To test our idea that building a vocabulary network that incorporates usage statistics and word association measures would make word-finding more effective, we started with an experiment that circumvented the involvement of actual users at an early stage of the evaluation. We created five usage profiles from data collected from five bloggers. The blogs were chosen from the Ageless Project (Ageless Project 2010). All bloggers were in their sixties with the exception of one younger blogger who we selected because she was writing about her experience after acquiring aphasia. The rest of the blog postings covered everyday-life topics such as cooking, gardening, health and family.

Text from the most recent postings of each blogger was gathered and broken into sentences. Using the NLTK part-of-speech tagger (Natural Language Toolkit 2010), we tagged and extracted all nouns, verbs, and adjectives from the bloggers' sentences. We then formed word pairs to represent vocabulary usage by linking $word_n$ in a sentence to $word_{n+1}$. For example, processing the sentence *I checked my credit balance and called the dentist*. resulted in the following set of pairs: *checked–credit*, *credit–balance*, *balance–*

called, and *called–dentist*. The data was also filtered so that it contained only words that are part of the Lingraphica vocabulary. We did this to be able to use Lingraphica's hierarchical organization as a baseline for comparison. Because Lingraphica and WordNet both represent different senses for homographs, we assumed that all words were the most frequently applied sense for that word. Table 4.1 shows a summary of the bloggers' profiles in terms of age, number of sentences from each blogger and the size of the resulting word collections.

Table 4.1. The sentences extracted from senior bloggers which covered topics such as cooking, family visits, traveling and gardening.

User	Age	Topics	Example Sentences	# Sentences
User_1	66	Everyday stories, cooking	*I am also going to have to learn how to bite and chew with dentures – especially biting.* *I use a little garden bench so I don't have to bend over quite so much.*	1150
User_2	61	Everyday stories, ruminations, cooking	*I cleaned out the cupboard which had held all the bread and made a potato salad.* *Cover with the cabbage add the cheddar and feta cheese sprinkle the top with paprika.*	1208
User_3	67	Everyday stories, cooking, family	*I piled a couple of pillows next to the window above my head put a fuzzy blanket over it and moved my own pillows down a bit.* *Most zucchini bread recipes have a cup of oil in them and I have always found them terrible greasy.*	1269
User_4 *aphasic*	39	Journal of life after aphasia	*Everything I used to do before is a question mark in my mind.* *The screen comes up I move the mouse around click on some things that is also OK I can navigate the several buttons and drop down menus.*	887
User_5	62	Everyday stories, travel, gardening	*If I could convince a student that he/she can do a thing and motivate him/her to do that thing then I've accomplished something.* *That's when I got in the mood for organizing the kitchen counters and cabinets as well.*	1037

4.2 Word-finding with a Static and an Adaptive Vocabulary

We started off with two basic lexical inventories. One was Lingraphica's hierarchy of words and the other one was ViVA's organization which used Lingraphica's hierarchy as a core organization. To create a semantic network with meaningful links among words in the vocabulary, ViVA augmented the basic hierarchy with usage data and semantic word associations. In the experiment, we compared word-finding in Lingraphica to word-finding in the enhanced hierarchy, ViVA.

4.2.1 Creating a Vocabulary Network

We augmented ViVA's core hierarchy with links between words based on the evocation data that we collected (discussed in Chapter 3). If the mean evocation rating between two words was ranked to be higher than 30 (moderate or higher evocation on the scale of 0 to 100), a link was introduced. The evocation subset used to create these links was the intersection of the whole data set and the words contained in the Lingraphica vocabulary. We constrained the data to words available in Lingraphica's vocabulary in order to be able to draw a fair comparison of the two approaches to vocabulary organization. This resulted in a vocabulary of Lingraphica's size, approximately 3000 words (counting base verb forms and singular noun forms only).

4.2.2 Simulating Usage of the Vocabulary

To simulate usage of the vocabulary for sentence composition, we used the word pairs we created based on the usage profiles described above. We randomly selected 80% of the word pairs from each usage set. These subsets were used to further enhance ViVA with direct links between words simulating past usage of the system. In other words, if a usage set contained the pair *bake–cake*, for example, we introduced a link between *bake* and *cake* in ViVA (*cake* was considered part of *bake's* semantic network). We matched the simulated usage links with an equal number of links based on the evocation data to form a training set for the vocabulary. The remaining 20% of the usage data were used as a testing set.

To predict new links between words in the vocabulary, we ran a logistic regression using as input the training set. The features of the input vectors for the logistic regression were: 1) a usage score (0 or 1); 2) an evocation score (0 or the corresponding mean from the ratings collected through the Amazon Mechanical Turk experiment described in Chapter 3); 3) a score based on semantic distance introduced by (Jiang & Conrath 1997); 4) one last score based on semantic relatedness as computed by (Lin 1998). For the purposes of training, a link between $word_x$ and $word_y$ was assigned if:

- $word_x$ and $word_y$ were connected due to usage (the pair was used in a sentence from a usage set)
- $word_x$ and $word_y$ were connected due to predicted association based on the evocation data, usage links and the two additional association measures from the literature.

The (Jiang & Conrath 1997) and (Lin 1998) scores were used to balance the prediction and avoid having the evocation and usage scores completely guide the outcome of the logistic regression. If both of these primary scores were zero, the higher of the semantic distance and semantic relatedness scores determined whether the two tested words should be linked or not. The predicted links were incorporated in ViVA and we evaluated whether the paths between the words in the testing set were shorter in

ViVA compared to the ones supported by Lingraphica's hierarchy. We present the results in the following section.

4.3 Results

The experiment of creating a vocabulary network based on links created due to the collected evocation data, simulated usage data and predicted connections showed improvement over the original Lingraphica hierarchy, that is, the paths between the words in the tests sentences were shorter in ViVA than in Lingraphica. Adding evocation and simulated usage data links alone resulted in shortening the distances between approximately 44% of the words that appeared next to each other in a sentence from the usage sets (the remaining 56% of the paths were the same length as in Lingraphica). Predicted links due to logistic regression improved the results by 8% on average. On average, 22% of the paths became shorter by two or more steps (see Table 4.2 for a summary of the results). Table 4.3 illustrates specific examples of shorter ViVA paths.

Table 4.2. Statistics for decreased browsing distance between related words resulting from augmenting the core vocabulary with links between words based on usage, word association measures and predicted associations.

Usage Profile	Total Usage Pairs	Intersection with Lingraphica	Percentage of Shorter Paths after Evocation & Usage	Increase in Percentage of Shorter Paths after Logistic Regression Prediction	Percentage of Paths Found to be Shorter by two or more Steps	Performance Improvement with Random Links Added to Baseline Network
User 1	5844	2539	43.9%	8.1%	27.9%	1.6%
User 2	6164	2914	42.4%	5.8%	25.3%	0.6%
User 3	3497	1537	44.6%	6.9%	15.6%	0.3%
User 4	4500	2077	46.5%	9.7%	21.8%	1.2%
User 5	4910	1865	42.3%	7.6%	19.2%	0.7%

Table 4.3. Examples of browsing paths between related words in ViVA and in Lingraphica. Often the only way to reach the second word, having found the first, is to default to the *home* (hierarchy root) category in Lingraphica.

$word_x$	$word_y$	ViVA Path	Lingraphica Path
rice	cheese	rice→cheese	rice→home→dictionary→things→house→ kitchen→refrigerator→dairy products→cheese
get	ticket	get→buy→ticket	get→home→dictionary→things→leisure→ outings→movies→ticket
baby	brother	baby→brother	baby→more people→people→family→ family relations→brother
hard	try	hard→teach→try	hard→home→dictionary→actions→ communicating→thinking→try
table	drink	table→glass→drink	table→home→dictionary→things→food→drink

We constrained our working vocabulary only to words available in Lingraphica so that we can draw a clear comparison with a practical baseline. As seen in Table 4.2, even with this constraint, our network improved approximately 45% of the search paths for all of the simulated usage data sets. We also used only parts of the word pairs from the evocation data set; 43% were excluded for this experiment, because one or both of the words in a pair were not part of the Lingraphica vocabulary. Constraining the data to match words in Lingraphica eliminated a number of links that could have shortened paths between related concepts even further.

We performed an additional “naïve” baseline test to show that our improvement in the distances between usage-related words due to link prediction cannot be achieved simply with a random increase in the density of the vocabulary network. We contrasted adding to the initial Lingraphica vocabulary links predicted using logistic regression and adding the same number of randomly chosen links. As shown in Table 4.2, there was a 0.88% average improvement on path distances, but it was marginal compared to the improvement due to the predicted links.

4.4 Discussion

We considered two different sources of data to evaluate ViVA’ adaptive functionality. Since we envisioned a tool that will help users find words to build phrases for communication, we thought about using data from repositories of switchboard exchanges (Switchboard 2007). However, it would have been very difficult to filter this data such that they appropriately represent our target user population. Instead, we decided to use informal text written by elderly people, describing their daily life and routines. We created a corpus of sentences from text collected from elderly bloggers. All bloggers except for one aphasic blogger 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 these people probably lead more active lives than people with aphasia, we argue that they have similar social interactions and needs, which makes their communications and thus, word associations relevant to the type of data our tool will be handling.

We constrained our vocabulary to words available in Lingraphica's vocabulary so that we could compare ViVA's performance to a practical baseline. In addition, Lingraphica's multimodal vocabulary has been validated by aphasic users and speech-language pathologists. Using Lingraphica icons made it possible for us to later test our approach to vocabulary organization with actual users who have aphasia without having to collect and evaluate our own image representations. However, one could also imagine a more powerful solution with ViVA's organization overlaid on a rich image library such as ImageNet (Deng 2009). ImageNet populates WordNet with hundreds of clean, high resolution images matched to the appropriate synsets. Such a combination is likely to create a rich visual vocabulary that, combined with good navigation support, will be able to help people with language impairments communicate.

We used a logistic regression to predict additional links for enhancing the vocabulary network. This allowed us to include words that our limited blogger associations and the evocation data did not address. The prediction was informed by simulated usage of the vocabulary, our evocation data and two additional semantic association measures. The simulation produced results indicating that ViVA significantly reduces path lengths between related words in the vocabulary. Our prediction algorithm was relatively simplistic. Thus, it is certainly worth exploring better predictive algorithms to enable possible shorter paths between related words and thus, improve word-finding even further. However, it is important, first, to explore what challenges are present in assisting a user with an adaptive solution. A dynamic organization is flexible and connects words faster, in theory, but it can also create an additional cognitive load on the user which can compromise its effectiveness. We continue with an investigation of how non-aphasic users search for words with the support of ViVA.

Chapter 5

Guiding Word-Finding with Semantic Associations for Non-Aphasic Individuals

It is challenging to navigate a dictionary consisting of thousands of entries in order to select appropriate words for building written communication. This is true both for people trying to communicate in a foreign language who have not developed a full vocabulary, for school children learning to write, for authors who wish to be more precise and expressive, and also for people with lexical access disorders who cannot retrieve words to express a concept in mind. We make vocabulary navigation and word-finding easier by augmenting a basic vocabulary with links between words based on human judgments of semantic similarity. In this chapter, we present the results from an evaluation of how our system, ViVA, performs compared to a commercial vocabulary access system in which words are organized hierarchically into common categories and subcategories.

Aphasia results in a number of language impairments which vary in nature and severity across individuals. The variability of impairments and the challenge of communicating with subjects who have aphasia are major hurdles to overcome when evaluating assistive technologies designed for this user group. Thus, as a stepping stone to evaluating ViVA with aphasic individuals, we first recruited people without language impairment to test our vocabulary organization. As other researchers have done, we could have used elderly people as a substitute to subjects with aphasia. However, ViVA incorporates links between words that model a speaker's mental lexicon independent of cognitive decline, so we chose able subjects. This way we could test our initial assumption of the role of evocation-based links in guiding word-finding by gathering faster and richer feedback from users without communication impairments. This also enabled us to form hypotheses about the broader application of our approach to vocabulary organization.

As discussed in Chapter 2, there is experimental evidence that words are organized in a speaker's mental lexicon by various similarity relations, including semantic relatedness, that facilitate evocation (e.g., *doctor* – *nurse*). The goal of the study described in this chapter is to determine whether providing users with an organization that adapts based on semantic word associations will enable them to find words faster than searching through hierarchically organized word collections. If users take shorter search paths when provided with a related words organization, we argue that this search path efficiency will generate less cognitive load on the user and enable him / her to write faster and more effectively. Participants were asked to find the missing words in a number of phrases using two vocabularies (hierarchical and semantically related) that they could navigate through a multi-modal user interface.

5.1 Word-Finding with ViVA and Lingraphica

We begin by describing the two vocabularies we used for the experiment as well as the user interfaces we built to access them. The first vocabulary was based on Lingraphica's hierarchy. As we stressed in the previous section, we chose Lingraphica as a baseline for comparison, because it is an assistive communication device built specifically for people with aphasia. The device is sold in the U.S.A. and through the years, it has been tested by and incorporates feedback from users and speech-language pathologists. The second vocabulary, ViVA, implemented Lingraphica's hierarchy as the basic organization but was enhanced with the evocation ratings we collected via the Amazon Mechanical Turk experiment described in Chapter 3. This enhancement introduced additional links between words in the vocabulary which we expect would guide users in navigating the vocabulary more effectively.

5.1.1 LG Vocabulary Condition and User Interface

We used a random subset of Lingraphica's vocabulary, which we call LG, to form the basic vocabulary for the experiment. The subset was chosen such that it provided all possible paths to the words missing in the phrases presented in the task. We describe how the phrases were constructed further down in this section. The resulting vocabulary consisted of 270 words (Lingraphica's vocabulary consists of approximately 3,000 base-form words). The maximum depth of a path to a missing word was 7; for example, *dictionary* → *things* → *house* → *kitchen* → *refrigerator* → *vegetables* → *garlic*.

The user interface is a simplified version of the Lingraphica interface which was first evaluated in a pilot study (Song 2009). It enables the user to browse the vocabulary, searching for missing words in the provided phrase. Figure 5.1 illustrates the home screen for the LG vocabulary interface. The vocabulary is accessed through a dictionary icon on the home screen. Clicking on the down arrow button leads the user to the subcategories of a specific concept. For example, if the user is trying to complete the phrase *tea or coffee for breakfast*, to find *tea*, s/he would have to start by clicking on the *dictionary* icon which results in the view in Figure 5.2. To get to *tea,* the user needs to traverse the portion of the hierarchy shown in Figure 5.3. Once the user has found *tea*, s/he can click on the plus button icon and the choice will be reflected in the lower portion of the screen where the phrase and the missing words, indicated by a question-mark icon, are displayed (see Figure 5.4).

The ***home*** button seen in Figure 5.4 in the upper left corner of the interface, takes the user to the home page that, through the dictionary icon, gives access to the highest level of the vocabulary hierarchy. The ***back*** and ***forward*** buttons enable users to move forward and backwards through their browsing history, similar to the navigation buttons on a web browser. The ***arrow button*** on the left of the storyboard allows for skipping phrases.

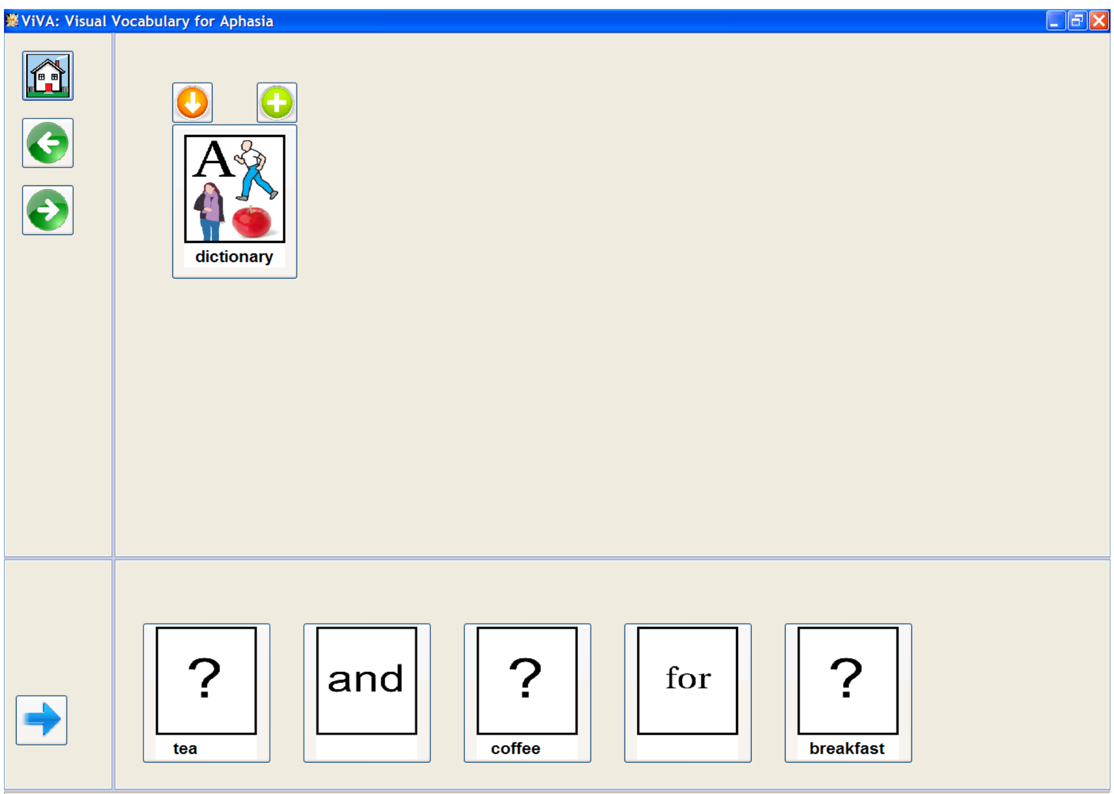

Figure 5.1. Home screen of the LG interface with the phrase to be completed displayed at the bottom of the screen. Target words are indentified by a question-mark icon.

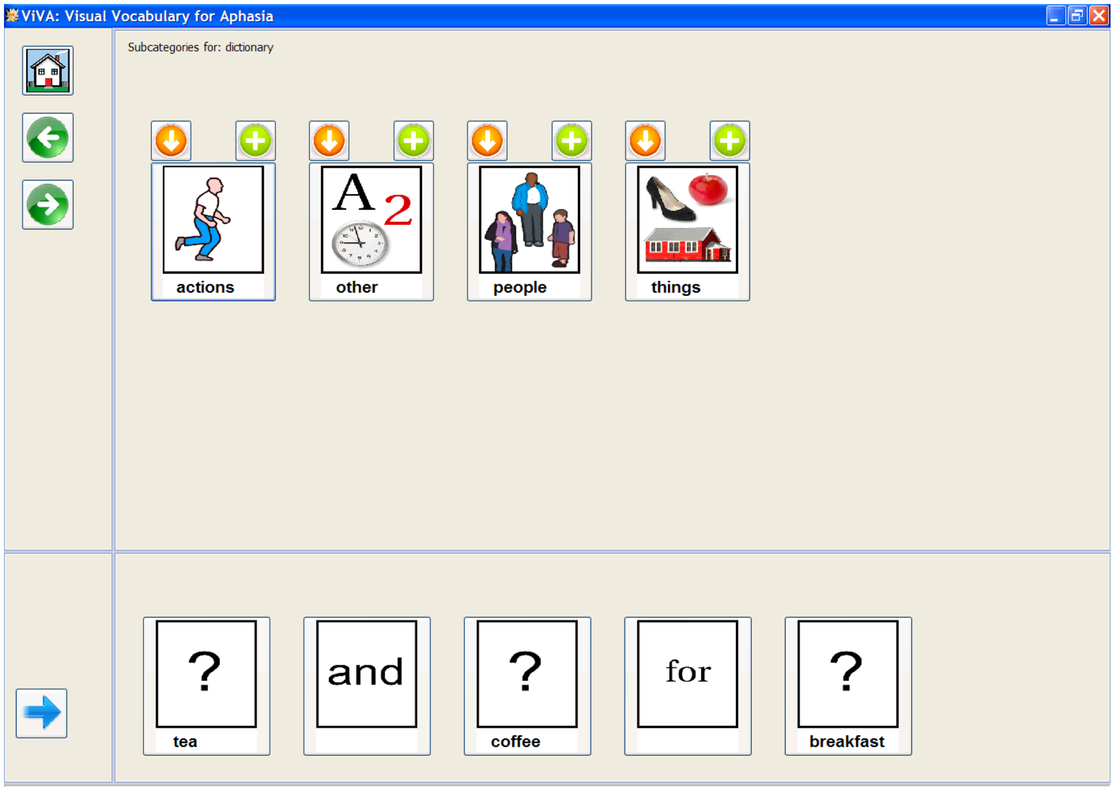

Figure 5.2. ***Action*, *other*, *people* and *things* are children of the root of the core hierarchy.**

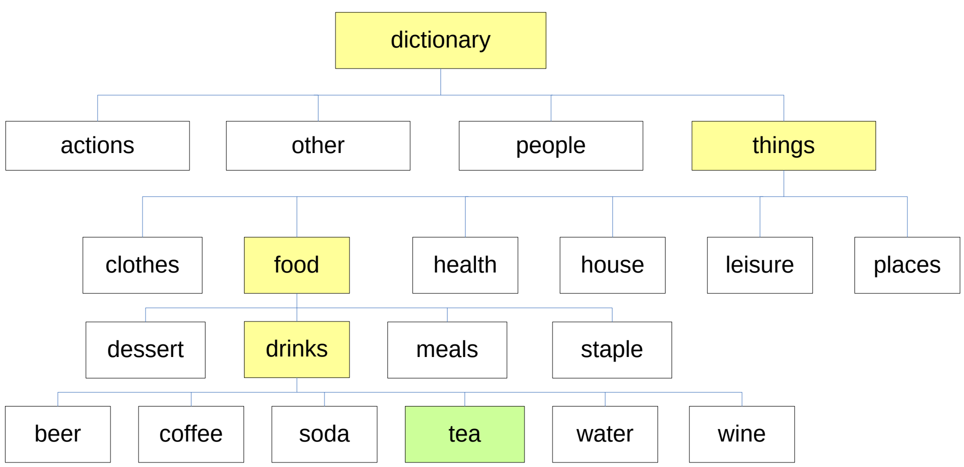

Figure 5.3. Hierarchical search leading to *tea*.

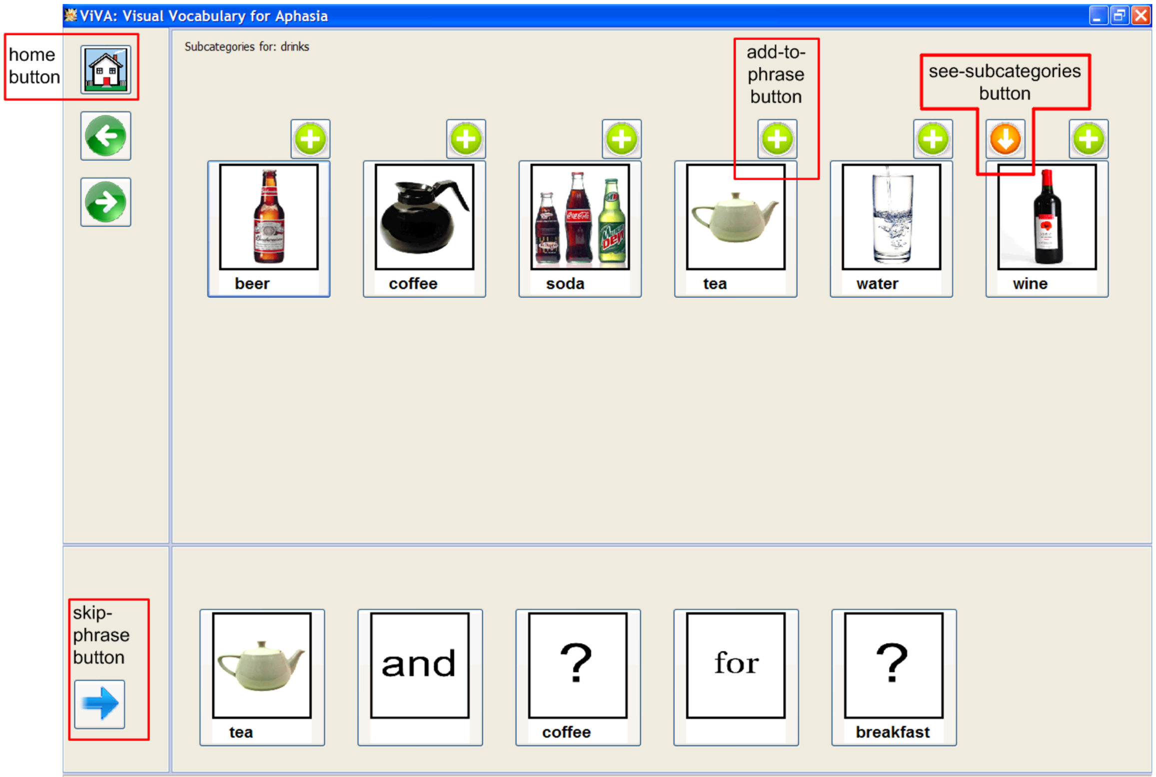

Figure 5.4. Clicking on the plus button above *tea* replaces the corresponding question-mark icon in the phrase (used to indicate that *tea* is a missing word) with the icon for *tea*.

5.1.2 ViVA Organization and User Interface

The augmented vocabulary, ViVA, implemented the LG hierarchy but also provided links between words based on our evocation data. No additional words were added to the base of 270 words. Moderate to strong evocation was considered sufficient for creating a link between two words.

The ViVA vocabulary was accessed and navigated in the same way as in the LG interface. The difference was that throughout the navigation words related to the concept the user clicked on were displayed in a related-words panel in the upper part of the vocabulary panel. If the user is trying to find *tea* in ViVA, they can still traverse the hierarchy shown in Figure 5.3. However, as shown on Figure 5.5, when the user has clicked on food, s/he can see *tea* as a related word, as well as *tasty* and *rice*. In this example, the path to *tea* is shortened by one click only, because *tea* is in the *drinks* category. However, if the user were to connect *food* and *cook* using the hierarchy, it will take five clicks through: *home* → *dictionary* → *actions* → *daily* activities → *cook* as opposed to using the direct link between: *food* → *cook*. The path from *food* to *hungry* is even less straightforward: *home* → *dictionary* → *other* → *feelings* → *physical feelings* → *hungry*.

Figure 5.5. Related words in the ViVA interface are displayed in a panel above the hierarchy subcategories' panel.

5.1.3 Stimulus Phrases

We chose frequently used words (frequency was judged by the author) that were present in the Lingraphica vocabulary to construct phrases for the experimental task. The overall number of phrases was 78 (39 for each vocabulary condition). The missing words in the ViVA phrases were constructed such that 50% of the words were directly associated. In other words, $word_{n+1}$, missing in a phrase, could be found in the related words panel of $word_n$. The relationship was not necessarily bidirectional as evocation is not always a symmetric relationship. For example, *dollar* may evoke *green* more than the reverse. The strength of association for the rest of the missing words was as follows – 36% of the words were strongly associated or there was a path between them through a common related word and 14% of the words were moderately associated or there was a path between them through two levels of common related words. This distribution allowed us to study people's satisfaction of having the word that they are looking for appear immediately as well as observe the alternative paths they may take when navigating through the words presented in the related words panel.

The phrases for the LG vocabulary were constructed so that the optimal paths between missing words was consistent with the optimal paths between missing words in the ViVA condition. The number of missing words per phrase was also kept consistent. There were a number of missing words that repeated across vocabulary conditions and as expected, because the vocabulary was not very rich and in words and the hierarchy was not very complex, there was a level of learning which we discuss below in the results section. In both sets of phrases, between 2 and 3 words per phrase were missing. All stimulus phrases are listed in Appendix A.

5.2 Experimental Methodology

5.2.1 Design

A within-subjects design was used. Thus, each participant used both vocabularies to search for words. Order of presentation was counterbalanced and participants were randomly assigned to conditions. The independent variable was the type of vocabulary used to retrieve words. The dependent variables measured length of search path, time to find the words requested, user's organization method preference and the number of words skipped (not found).

5.2.2 Task and Procedure

Participants were asked to find the missing words in 78 phrases using two different vocabulary access systems, ViVA and LG. They were first given a tutorial on how to use the first interface (randomly assigned) and on how to complete a phrase. After completing another example phrase, they proceeded to completing the remaining 37 phrases using the first vocabulary. To avoid some of the expected learning effect (mainly due to limited size of our experimental vocabulary), we asked participants to find the missing words as quickly as possible. Once a phrase was completed, the system would notify the user with a sound and automatically move on to the next phrase. Participants were allowed to skip words and whole phrases if they felt like they had tried hard enough to find the missing words and the task was becoming frustrating. After completing the first set of phrases, participants were presented with the second interface and the differences were highlighted. At the end of the study the investigator asked a number of background questions as well as questions related to the participant's experience with the two vocabularies. The study took approximately 60 minutes and participants were compensated $10.

5.2.3 Quantitative and Qualitative Measures

We logged automatically each interaction with the interface, including whether words were selected from the main panel or the related words panel, which button (down arrow, plus, home, back, forward, or next phrase) was clicked on, as well as the time when the interaction happened. This allowed us to keep track of the paths that people took to find each word, if they used the related words in ViVA, as well as how long it took participants to complete each phrase.

In addition, the investigator kept notes of what words were skipped, what words seemed particularly difficult to find, and of participants' comments during the study. Using a questionnaire at the end of the study, we also collected information on participants' demographic information, language background and feedback on their experience with the two vocabularies.

5.2.4 Apparatus

We ran the experiment on a 3.0GHz Intel Pentium 4 desktop computer with 1.0 GB of RAM and Microsoft Windows XP Professional. We used a 20" LCD monitor with 1280x1024 resolution; the experiment was coded in Java 1.6.

5.2.5 Participants

Sixteen people (seven female) were recruited via campus email lists and website postings, advertising the experiment. Their average age was 24 (SD = 9.83). One of the participants was a professor, six were graduate students in computer science, and nine were undergraduate students from different majors. Seven people were non-native English speakers, but all participants spoke fluent English.

5.3 Results

5.3.1 Completion Time, Skipped Words, and Learning Effect

Table 5.1 shows the overall completion time of the tasks in the two vocabulary conditions, LG and ViVA, according to the order the vocabularies were presented to the participants (first or second). The completion time in both conditions accepts homogeneity of variance. Analysis of variance (ANOVA) results 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 ($F(1, 31) = 33.70$, $p < .01$), which in turn signaled a learning effect. This was anticipated because both vocabularies rely on the same core hierarchy. Once participants became familiar with the vocabulary structure and locations of common categories (e.g. *food*), finding words became faster.

Table 5.1. Experiment completion time (in minutes).

Presentation Order	**LG**		**ViVA**	
	Mean	SD	Mean	SD
First	33.88	5.96	25.50	3.59
Second	25.75	6.11	14.25	1.83

ViVA benefited more from the learning effect in terms of time, because most of the missing words were linked through related words (1 to 3 clicks apart) and thus most of the time was spent on locating the first word. Table 5.2 shows the average completion time of each phrase with both LG and ViVA. Although the phrases with the same index may not be the same, the paths for the missing words were consistent across the two vocabulary conditions.

Table 5.2. Average time (in seconds) to complete a phrase.

Phrase # : # of missing words	**LG time**	**ViVA time**	**Phrase # : # of missing words**	**LG time**	**ViVA time**
0:2	36.69	30.16	17:2	25.25	18.95
1:2	98.83	25.16	18:2	52.77	22.52
2:2	47.67	22.85	19:2	20.98	30.7
3:2	67.02	30.76	20:2	56.42	10.95
4:3	32.85	28.14	21:3	23.09	22.53
5:3	20.71	16.9	22:2	27.07	16.65
6:3	74	25.32	23:2	106.93	40.6
7:2	32.86	15.34	24:3	34.43	26.55
8:3	35.08	31.54	25:3	25.87	27.25
9:2	45.91	34.96	26:2	38.21	38.49
10:2	29.81	23.77	27:3	30.91	28.24
11:2	21.54	23.37	28:3	24.47	22.93
12:2	30.85	16.53	29:2	83.08	17.57
13:2	17.67	36	30:2	31.47	19.74
14:3	22.1	50.89	31:2	52.95	18.7
15:3	42.05	17.34	32:3	34.54	75.02
16:3	62.22	32.15	33:2	20.67	12.7

Table 5.3. Number of words skipped.

Order	LG		ViVA	
	Mean	SD	Mean	SD
First	6.38	2.77	1.75	1.58
Second	2.63	2.33	1.38	1.41

Table 5.4. Words skipped and the number of times they were skipped.

<table>
<tr><th colspan="2">LG</th><th colspan="2">ViVA</th></tr>
<tr><th>Times</th><th>Words</th><th>Times</th><th>Words</th></tr>
<tr><td>15</td><td>cigarette</td><td>8</td><td>cigarette, mow</td></tr>
<tr><td>8</td><td>sleep</td><td rowspan="2">5</td><td rowspan="2">poison</td></tr>
<tr><td>6</td><td>poison</td></tr>
<tr><td>5</td><td>salt, approve</td><td rowspan="5">1</td><td rowspan="5">salt, explain, grass, smell</td></tr>
<tr><td>3</td><td>milk, vegetables, broccoli</td></tr>
<tr><td>2</td><td>speed, van, health, team</td></tr>
<tr><td>1</td><td>cheese, honey, sugar, increase, hot</td></tr>
</table>

LG benefited significantly more from the order effect ($F(1,15) = 8.58$, $p < .05$); less words were skipped when it was tested second (decrease to a mean of 2.63 skipped words from a mean of 6.38) (Tables 5.3 and 5.4). With LG, participants could find a word only by traversing the hierarchy. This was challenging for words located in uncommon categories, e.g. *sleep* is under *dictionary* → *actions* → *daily routines* → *inhale* → *sleep*. On the other hand, ViVA provided alternative paths to words and people did not have to know the exact category. For example, many people did not know *milk* was under *house* → *kitchen* → *refrigerator* → *diary*, but could still find it through related words such as *coffee* and *tea*. As a result, less words per sentence were skipped with ViVA ($F(1,31) = 15.70$, $p < .01$).

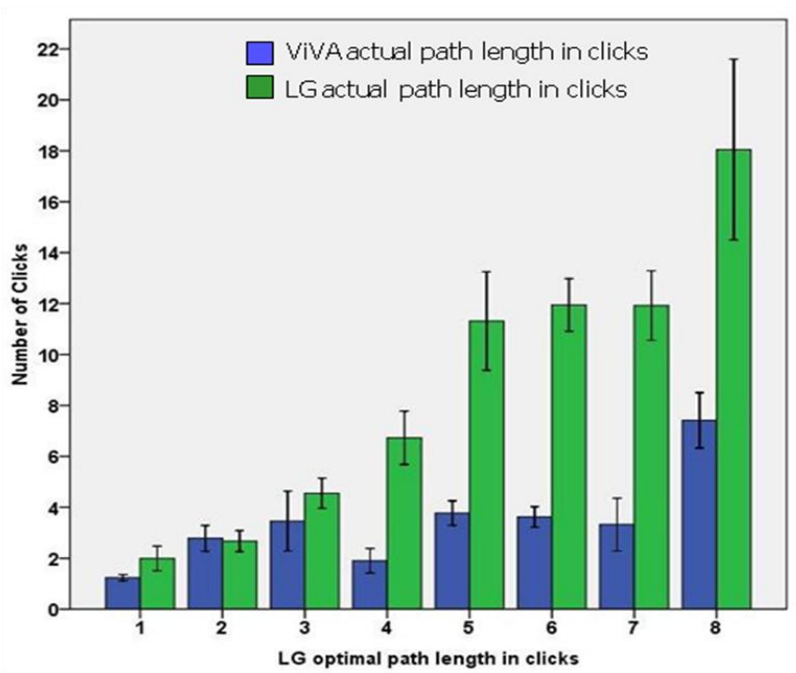

Figure 5.6. Participants took significantly shorter paths to find words in ViVA than in LG.

5.3.2 Comparison of Path Length between LG and ViVA

ANOVA results showed that on average significantly fewer clicks were required to find a word in ViVA than in LG (Figure 5.6: $F(1, 83) = 33.7013$, $p < .01$). The results suggest that ViVA significantly shortens distances between words. This is especially valuable for words in different parts of speech (e.g. *food* (noun), *eat* (verb), and *hungry* (adjective)). In LG, the actual number of clicks used in the study increased as the optimal path length in LG increased ($F(1, 83) = 10.8161$, $p < .01$). However, the interaction between condition and LG optimal path length was significant ($F(1, 83) = 5.9907$, $p < .01$), which suggests that the actual path length in ViVA was not affected by the LG optimal path length as in LG. This indicated that people did take advantage of the provided related words.

Table 5.5. Comparison of actual number of clicks in ViVA for word pairs in the two possible search orders.

Original order	ViVA average path in clicks	Reversed order	ViVA average path in clicks
smell, food	6.00	food, smell	24.89
mow, grass	1.00	grass, mow	16.33
explained, student	3.30	student, explained	15.80
drive, motorcycle	2.62	motorcycle, drive	12.67
swallow, pill	1.00	pill, swallow	9.29
butter, bread	1.00	bread, butter	7.33
eat, hungry	1.58	hungry, eat	7.83
grandfather, lonely	1.92	lonely, grandfather	8.00
brown, rice	1.00	rice, brown	7.00
plants, animals	1.00	animals, plants	5.75
cook, bacon	4.20	bacon, cook	1.00

beer, dinner	1.00	dinner, beer	4.00
pudding, dessert	8.00	dessert, pudding	5.00
bacon, tasty	5.70	tasty, bacon	8.67
travel, buses	2.25	buses, travel	5.00
meat, brown	9.50	brown, meat	7.00
drive, buses	2.60	buses, drive	5.00
tasty, breakfast	3.40	breakfast, tasty	1.20
sister, lonely	3.00	lonely, sister	5.00
uncle, daughter	2.83	daughter, uncle	1.00
cereal, breakfast	3.25	breakfast, cereal	1.67
lunch, cook	3.71	cook, lunch	5.25
rice, eggs	5.33	eggs, rice	4.00
call, write	1.00	write, call	2.00
coffee, milk	1.00	milk, coffee	2.00
university, class	1.00	class, university	2.00
meat, steak	1.80	steak, meat	1.00
garlic, food	1.69	food, garlic	1.00
rice, meat	4.20	meat, rice	3.60
tea, coffee	1.00	coffee, tea	1.56
bacon, lunch	4.50	lunch, bacon	5.00
room, kitchen	1.38	kitchen, room	1.00
drive, travel	1.22	travel, drive	1.00
aunt, uncle	1.00	uncle, aunt	1.00
chair, cushion	5.75	cushion, chair	5.75
sausage, breakfast	1.00	breakfast, sausage	1.00
smell, garlic	1.00	garlic, smell	1.00
tea, milk	1.00	milk, tea	1.00
breakfast, bread	1.00	bread, breakfast	1.00

5.3.3 Impact of Distance to Related Words

In ViVA, for each phrase, the missing words were linked through a path of related words consisting of 1 to 3 clicks. This distance had significant impact on the length of the actual path participants took to find the target words in ViVA ($F(1, 31) = 32.4812$, $p < .01$). Figure 5.7 illustrates that if the second missing word was a member of the related words of the first word participants found, people usually picked it up within one to two clicks. However, if the two missing words were not directly related, people had to explore more to see if it surfaced eventually. If the target word was not found in five clicks, people were likely to give up searching through the related words and fall back on the main hierarchy.

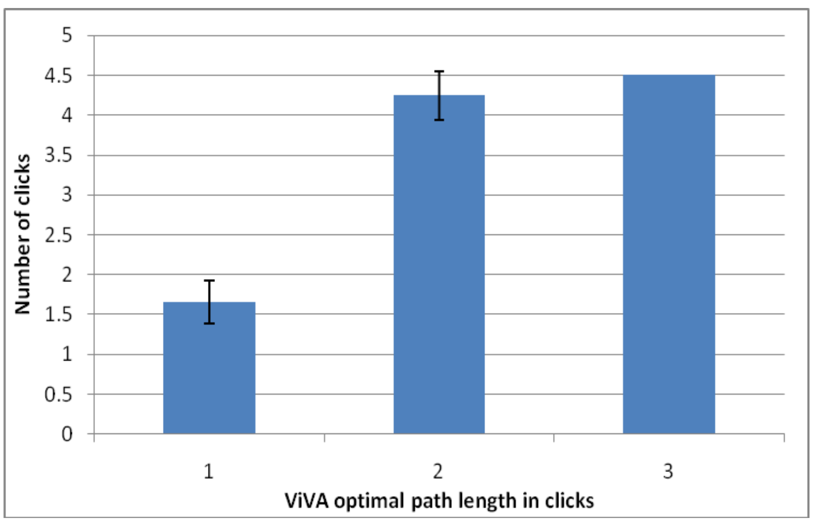

Figure 5.7. Actual path length in clicks using ViVA compared to the optimal path utilizing the provided related words.

As discussed earlier, evocation, which determined what related words were displayed, has a direction. Thus, which target word was found first had impact on the number of clicks needed to reach the second word in a pair. Table 5 lists the path length in clicks between word pairs found in different orders. For the first ten pairs (from *smell–food* to *plants–animals*), the second word was reachable from the first word within three clicks through related words (e.g. *smell*, *plants*) to the second word (e.g. *food*, *animals*) but not the other way around. Thus, the path length varied greatly, highlighting the advantages of having semantic associations in the vocabulary.

5.3.4 Subjective Feedback: Easiness and Confusion

We gathered participants' subjective feedback using a questionnaire at the end of the experiment. Overall (see Table 5.6), 12 people found it easy to use ViVA to search for words; only one person found word-finding easy with LG. Six people thought the LG vocabulary organization was confusing to them versus one person who thought ViVA organization was confusing. All people agreed that having related words helped them to find words faster and nine people considered putting words in categories (context) helped them to find words faster (Table 5.7).

Table 5.6. Participants indicated on a 5-point Likert scale how easy it was to find words with both vocabularies and whether the organizations were confusing. Here we present a count of the participants' responses.

Ease	LG	ViVA	Confusion	LG	ViVA
Very easy	0	3	***Strongly disagree***	0	3
Easy	1	9	***Disagree***	3	11
Neutral	10	3	***Neutral***	7	1
Hard	3	1	***Agree***	4	1
Very hard	2	0	***Strongly agree***	2	0

Table 5.7. Count of participants' feedback on whether having words in context and providing related words helped them find words faster.

	Context/categories	**Related Words**
Strongly disagree	1	0
Disagree	4	0
Neutral	2	0
Agree	3	5
Strongly agree	6	11

Nine people stated that all of the associations between given words and their related words make sense to them, two people said most of the time, three people said sometimes (e.g. *grandfather* and *lonely* should not be linked, *beef* and *sausage* should be linked). Five people indicated that they did not analyze why words were related and merely utilized what was given. With ViVA, especially when presented first, participants tended to search for words via related words instead of trying to locate what category the word should belong to. For example, two participants always got to food by going to *eat* and used its related words, instead of going to *things* and then *food.* Thirteen people indicated that they would prefer to organize the words in categories themselves, suggesting that idiosyncratic links are really important and valued. The other three people considered building an entire vocabulary from scratch as a challenging task, but indicated that if given a basic vocabulary structure, they would like to be able to change it and group words according to their preferences. Four people felt strongly that the vocabulary should be a network of words with associations crossing different categories and linking words in different parts of speech (as opposed to a static tree structure).

5.4 Discussion and Future Work

ViVA was significantly better in terms of search time, shorter paths taken, user preference and least words skipped than the hierarchical organization. Even though the number of missing words as well as the optimal path to find them was consistent across conditions, it was unreasonable to do a pair wise comparison of the average time participants took to complete each phrase. This is explained by the fundamental challenge present in organizing words in categories. Different users expect to find the same thing in different categories depending on factors such a life style and life experiences. A few participants indicated that they would like to organize the words in categories themselves. This response might have been facilitated by the small vocabulary they had to work it. Once the vocabulary grows to be sufficiently expressive for functional communication (in an assistive tool for communicating in a foreign language, for example), having the user manage it would become cumbersome. Thus, it is important to build a vocabulary with an organization which reflects user preferences and user input in addition to word association measure which can still serve as a framework for the vocabulary and assist with adapting its structure.

The simplicity of the phrases and the limited vocabulary allowed us to design a fun, manageable task for our participants. This was also a drawback, because to really understand how people navigate the vocabulary and search for words, we need to present them with more challenging tasks. This is especially important in exploring the broader implications of this work. Aside from searching for words in assistive vocabularies that help people with language impairments to communicate, searching for words in electronic dictionaries occurs in a variety of contexts. Examples include searching for words to communicate in a foreign language, using a thesaurus to find the best word to express a concept when writing, and searching for words as part of a game, e.g., crossword puzzles. It would be interesting to apply our approach to vocabulary organization to the design of assistive communication tools for other user populations, foreign language learners, for example.

Showing significant difference in user performance and satisfaction with fluent English speakers who have their semantic links intact would suggest aphasic people can benefit even more from a tool that aims to compensate for impaired semantic associations in a speaker's mental lexicon. Any benefits would be conditional on the level of assistance the vocabulary can provide and users' acceptance and expectations of its adaptive functionality. As the results show, for example, users tended to abandon their reliance on related words to find the target once they had exceeded five clicks. They would then fall back on the basic hierarchical organization. This illustrates the advantages of the mixed-initiative approach we employed which gives the user choices, but also highlights the need for further investigation into adaptivity and its precision. Next, we build on evaluating ViVA with actual users and describe a study with people with aphasia, our target user group.

Chapter 6

Guiding Word-Finding with Semantic Associations for Aphasic Individuals

Our first evaluation of ViVA simulated vocabulary usage and sentence construction using sentences collected from blogs of elderly people. The experiment found that, compared to the paths available in Lingraphica, ViVA shortened approximately 52% of the browsing paths between words in a sentence (Chapter 4). To investigate whether able people would take advantage of these shorter paths, we conducted an experiment evaluating ViVA with participants without language impairments. We asked the participants in our study to find the missing words in a number of phrases as quickly as possible, using two vocabularies. The first vocabulary that we call LG has Lingraphica's hierarchical organization and the second one, ViVA, inherits the same hierarchy but is augmented with semantic associations which translated into related words that participants could see once they had clicked on an icon.

The results of the experiment showed that participants took significantly less time to find words with ViVA, taking shorter paths guided by the provided semantic associations (we define a browsing path to be the number of words retrieved between $word_n$ and $word_{n+1}$ in a sentence). All participants agreed that having related words helped them find words faster and most thought that finding words in ViVA was less confusing than searching in LG (Chapter 5). Using ViVA, people tended to search for words via related-word links instead of trying to locate what category the word should belong to.

In this chapter, we describe how we adapted the experimental task used in the study with non-aphasic participants to evaluate ViVA with people who have aphasia. We demonstrate that a vocabulary hierarchy augmented with associations that reflect human judgments of semantic relatedness enables people with aphasia to find words in the context of a sentence faster and that the associations guide their search effectively. We present the results from a study comparing word-finding using a commercial hierarchical vocabulary access system and using ViVA's organization. We discuss the paths people take in searching for target words, the associations they expected to lead them to the target words, users' feedback on interacting with the two organizations and the challenge of searching for the first word to initiate a phrase.

6.1 Experimental Task Design

The experimental task was adapted from the study with people without language impairments (described in Chapter 5) such that it provided support for aphasic participants. We offered additional context to the task of searching for target words in the vocabulary, reduced the number of phrases we gave to the participants and adjusted the interaction with the interface.

6.1.1 Stimulus Construction

We followed three high-level guidelines in redesigning the task for aphasic participants. The first guideline is avoiding fatigue. People with aphasia tend to get tired fast when presented with a linguistic task. Thus, to keep the experiment within the limits of an hour, we reduced the number of phrases they had to complete as part of the task. The second guideline is eliminating any factors that may intimidate the participants, e.g., making the participants aware that we were measuring speed or error rate. For this study, instead of asking people to find the words as fast as possible, we encouraged participants to search for words at their own pace and we did not specify a minimum (or a maximum) number of target words they had to find. The third guideline is making the experience rewarding by ensuring that participants can successfully complete most of the task. To address this consideration, we simplified the phrases and provided additional context so that it was easier for people to guess the missing words.

We first simplified the task by targeting only frequently used words provided by the British National Corpus (The British National Corpus 2007). We paired nouns with verbs and nouns with adjectives to be used as the target words in a phrase. Presenting only a phrase to the participant places the target word in a limited context. To emphasize the context and stimulate the part of the mental lexicon where the word could be retrieved from, we added an additional step to the study protocol. We created a set of scenarios consisting of images related to the target word pairs. The images were selected by using Google Images (Google Inc. 2010). This selection was done by typing the verb/noun or noun/adjective pairs in the search box and selecting the best image from the first page of search results. Results were subjectively filtered based on image quality and how well they represented the desired terms. For example, the image judged best for *travel* and *bus* is the image shown in the upper half of the stimulus shown on Figure 6.1. We then constructed simple sentences that related to the image and included the targeted words.

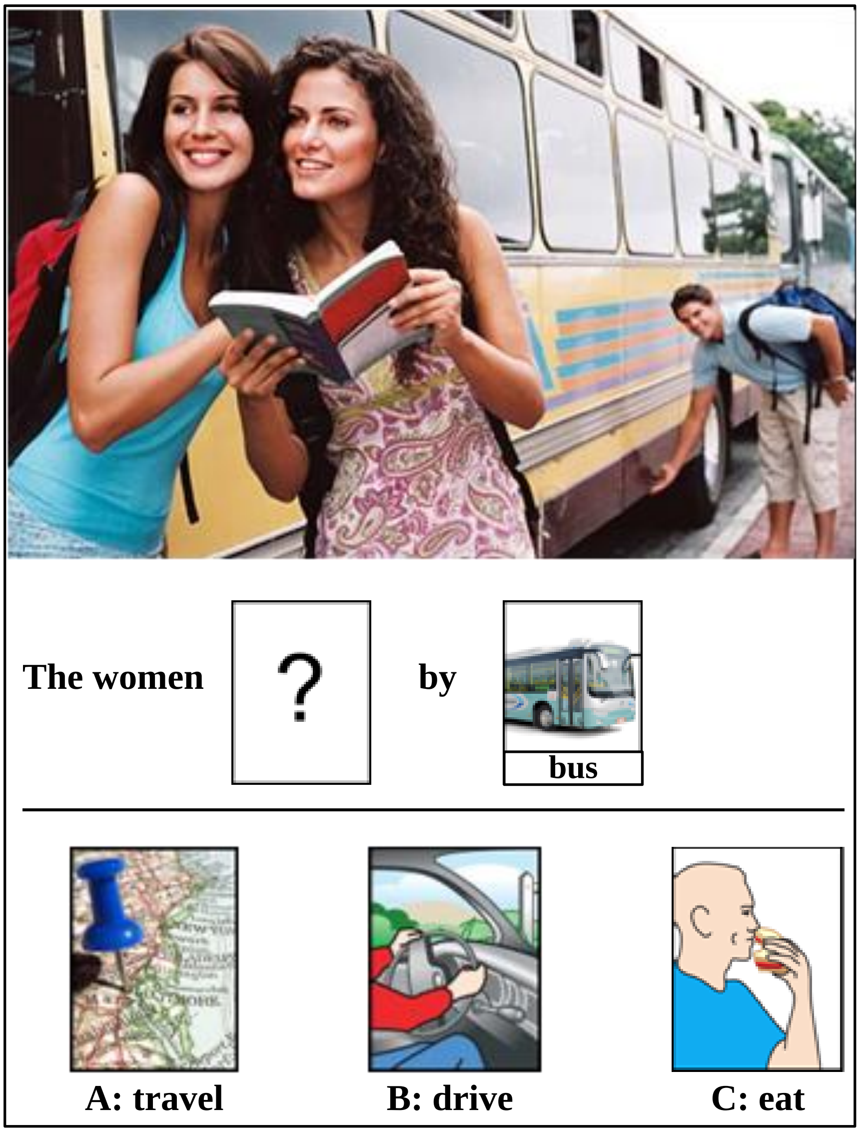

Figure 6.1. Example stimulus: based on the image and sentence context, the participant selects the missing word.

The resulting collection of 30 images and the corresponding sentences was presented for evaluation to a linguist and two computer scientists, all with experience in designing assistive technology for people with aphasia. They were asked to view each image, read the sentence following it, and based on the picture and the context of the sentence, guess the missing word. In case of ambiguity, the reviewers listed alternative words that the image accompanied by the sentence context may evoke.

Based on this first round of evaluation, some images were discarded or replaced and the sentences were adapted to be less ambiguous. The adjusted scenarios were then presented to two people with aphasia. They were first shown the image. Then the researcher read the sentence out loud and finally asked the person to guess the missing word. This exercise showed that there was still room for ambiguity among the images and the participants were sometimes confused. In addition, they felt under pressure to produce a “right” answer and wanted to confirm every guess with the researcher. Thus, we decided to constrain the task further and instead of asking people to produce a guess, we provided a multiple choice format.

Each scenario was made to fit on an A4 sheet of paper (see Figure 6.1 for an example); we showed the image first, followed by a short sentence. One of the target words in the sentence was represented by a question-mark icon and the other by the corresponding Lingraphica icon. The example shown in Figure 6.1 presents an image of two women with bag packs who seem to have just gotten off a bus. They are holding an open book that could be a tourist guide. The image is accompanied by the sentence *The women ? by bus*. The user is given three choices for the missing word. One of the choices, *travel*, is the correct answer; one could be used in the sentence, *drive*, but not in the context implied by the image and also not in the provided sentence structure. The last choice, *eat*, does not fit in the overall context. By providing context through an image and a sentence, we aimed to assist the retrieval of the correct word, but also to stimulate semantic associations surrounding that word.

6.1.2 Pilot Testing

The scenarios and the interface for accessing and browsing the vocabulary (described in section 6.2.1.2) was evaluated by two staff members at a local support center for people with aphasia where we planned to recruit participants. One staff member is a research speech language pathologist and the other one is a senior computer coach. No changes except for an improvement on the interaction with the interface were recommended. We provide more details on that improvement in section 6.2.1.4.

Two people with aphasia were recruited to pilot the study. They were given thirty minutes to find the missing words in 15 scenarios using with each one of two vocabularies. They were able to complete only five sentences per vocabulary condition. As a result, the final experiment presented only seven scenarios. Participants were easily discouraged when they could not find the target words; to simplify the search, we eliminated all words that were more than seven clicks away from the top level of the core vocabulary hierarchy. We also eliminated words that were in categories we judged "less intuitive" (e.g., the word *broken* is reached by navigating through *dictionary* → *other* → *modifiers* → *state* → *broken*). Next, we describe the experiment in more detail and discuss the results.

6.2 Word-finding with ViVA and LG

The purpose of the study was to compare word-finding in a static vocabulary hierarchy and in a vocabulary hierarchy adapted with links between words reflecting semantic relatedness. Our goal was to investigate whether people take advantage of the shortcuts possible through the provided semantic associations and how evocation guides their search. Based on ViVA's previous evaluations described in Chapters 4 and 5, we formed the following hypotheses:

H1. *The paths participants take to find words with ViVA will be significantly shorter than the paths in LG.* We expected that people will take advantage of the related words provided by ViVA which will guide them quicker to the desired word.

H2. Study participants will rate ViVA's organization less confusing than LG.

H3. Study participants will find it easier to locate words in ViVA than in LG.

We hypothesized that the related words available in ViVA will speed up the search and also provide more satisfactory vocabulary navigation for the participants.

6.2.1 Methodology

6.2.1.1 Vocabulary Conditions

The first vocabulary, **LG**, provided a hierarchical organization. LG's vocabulary is a subset of 200 words of Lingraphica's vocabulary (Lingraphica's vocabulary consists of approximately 3,000 words). The subset was chosen such that it provided paths to all target words. The maximum depth of a path to a missing word was seven.

The interface to the vocabulary was similar to the one used in the experiment described in Chapter 5. It enabled the user to search for words by browsing the hierarchy. Figure 6.2 illustrates the home screen for the LG vocabulary interface. The vocabulary is accessed through the dictionary icon on the home screen. Clicking on the down arrow button leads the user to the subcategories of a specific concept which in turn takes the user deeper into the hierarchy. For example, if the user is trying to complete the phrase *travel by bus*, to find *travel*, she would have to traverse the portion of the hierarchy shown in Figure 6.3. Once the user has found *travel*, she can click on the plus button above the travel icon and the choice will be reflected in the lower portion of the screen where the target words, indicated by a grey icon and a questions mark on top, are displayed (see Figure 6.4). To draw the participants' attention on the target words, the interface displayed a phrase comprised mainly of the target words extracted from the original sentence. The arrow button to the left of the phrase enabled the participant to skip phrases.

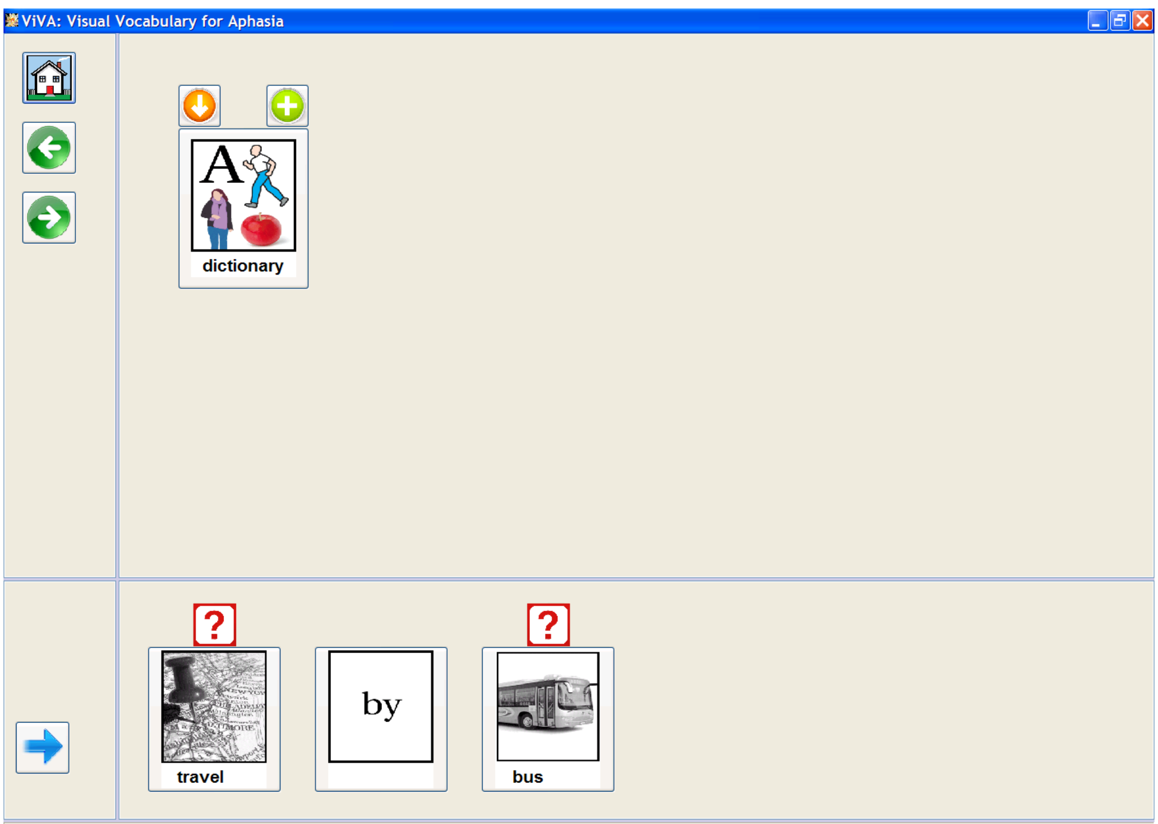

Figure 6.2. Home screen of the LG vocabulary interface with the phrase to be completed displayed at the bottom.

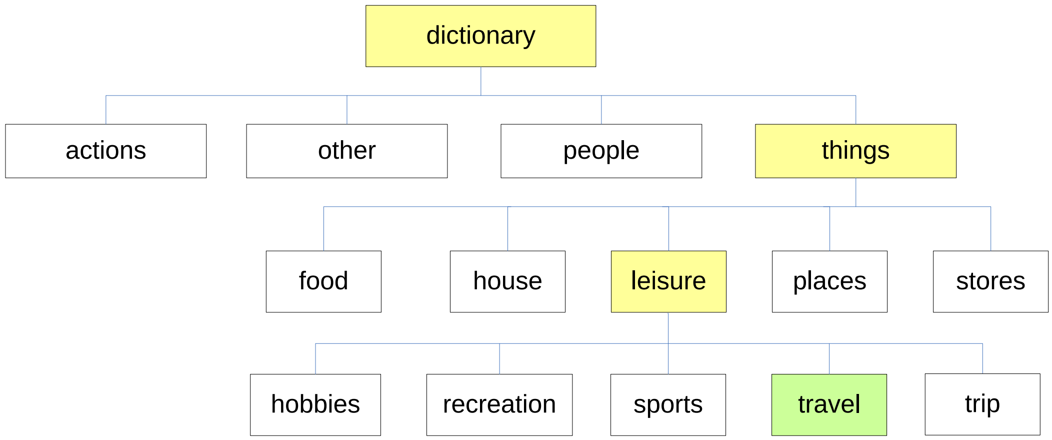

Figure 6.3. Hierarchical search leading to *travel*.

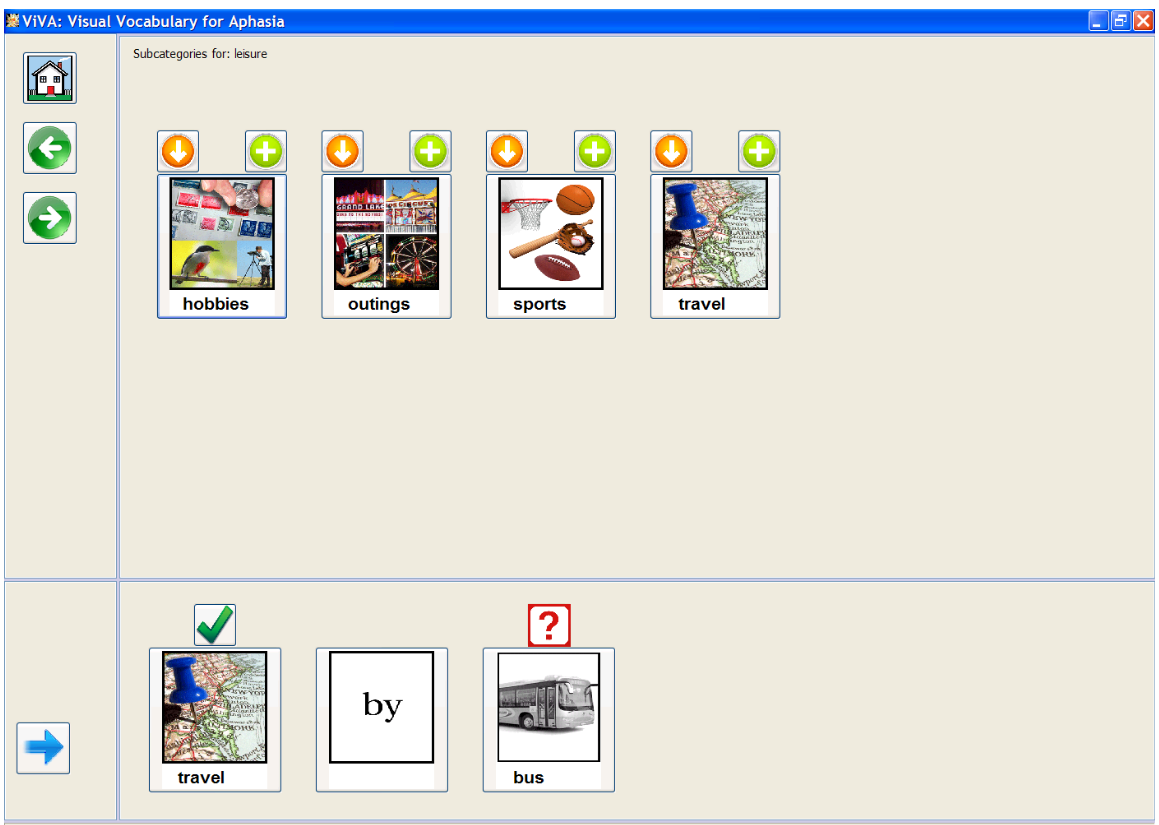

Figure 6.4. *Travel* is checked as found after the user clicks on the plus button above the travel icon.

Figure 6.5. ViVA provides words related to the one the user has clicked on, e.g., *drive* and *plane* are related to *travel*.

ViVA, the second vocabulary, implemented the LG hierarchy but also provided links between words based on the evocation data. No additional words were added to the core 200 words. Moderate to strong evocation was considered sufficient for creating a link between two words.

The interface for browsing ViVA had the same layout and functionality as the LG interface with one exception. When browsing the vocabulary, words related to the concept the user clicked on were displayed in a related-words panel in the upper part of the navigation screen. Figure 6.5 shows what the user sees after clicking on *travel*; *bus* and *drive* are related to *travel*. In addition, the target word *car* can be reached through *drive* → *car*. The maximum number of related words displayed was limited to five to avoid browsing.

6.2.1.2 Task

Participants had to find a number of words using two different vocabularies. The task consisted of fourteen scenarios (similar to the one shown in Figure 6.1) comprised of an image followed by a sentence related to the image. Participants were first shown the image and the sentence was read out loud. They were then asked to identify, from a choice of three words, the word that was missing in the sentence. Once they had chosen the correct word, participants had to find in the vocabulary the missing word and the word it is associated with (*travel* and *bus* in Figure 6.1). A sheet of paper displaying only the image remained in front of them to remind them of the context.

There were seven sentences per vocabulary condition. In the ViVA condition, 60% of the pairs were directly associated (the second word was displayed in the first word's related words panel). The rest of the words were associated through a common related word. The scenarios were distributed across the two conditions so that the depth of target words and the optimal path between target words in a sentence was balanced. The sentences used as a stimulus are listed in Table 6.1 and Appendix B contains the image scenarios for the two vocabulary conditions.

Table 6.1. Stimulus phrases for the two vocabulary conditions. The targets words are shown in bold.

LG Stimulus	ViVA Stimulus
The woman ***travel*** by ***bus***.	They ***drive*** by ***car***.
People ***fly*** to Europe by ***plane***.	***Coffee*** with ***cream***.
Tea with milk and ***honey***.	The men ***play checkers***.
The boys ***play baseball***.	The woman ***bakes*** fresh ***cookies***.
The boy ***listens*** to ***music***.	The boy ***smells*** the ***flower***.
People ***eat dinner*** in a restaurant.	The ***baby cries***.
They ***shop*** in the ***mall***.	The girl ***writes*** a ***letter***.

6.2.1.3 Procedure

The experiment was designed to last approximately 60 minutes and the time was divided into two 30-minute slots. Participants were first introduced to the task using an aphasia friendly consent form describing the components of the study (the consent form can be found in Appendix B). When all questions were answered and their signature was collected, they were given a brief tutorial on how to use the first interface which was randomly assigned (ViVA or LG). The scenarios were printed on A4 sheets of paper and bound in a booklet. The investigator would show the scenario drawing attention to the image first, then read out loud the sentence, omitting the missing word, and let the participant make a choice. Once the participant had chosen the correct answer, their attention was redirected to the computer screen and the vocabulary interface where they had to find both of the target words. Participants were encouraged to try the first scenario as an example before completing the remaining six. They were allowed to skip words or whole phrases if they felt like they had tried hard enough and the task was becoming frustrating.

After completing the first set of phrases or after the first half hour had passed, participants were asked questions about their experience with the vocabulary they had interacted with and given five minutes rest. They were then the presented with the second interface and the differences were highlighted. After completing an example with the second vocabulary, they completed the last six scenarios. At the end of the study the investigator asked a number of background questions as well as questions related to the participant's experience with the second vocabulary. Participants were compensated $10 which they could keep or donate to the center where the study was run.

6.2.1.4 Apparatus

We ran the experiment on two laptop computers connected to external monitors. One was a 1.86 GHz Intel Pentium M laptop computer with 1.50 GB of RAM, running Microsoft Windows XP. The second one was a 2.40 GHz Intel Core2 Duo laptop computer with 3 GB of RAM, running Microsoft Windows 7 Professional. We used two 17" LCD monitors; one with a 1280x1024 resolution and the other one with a 1280x800 resolution. The experiment was coded in Java 1.6. The inconsistency in computer performance and screen resolution was irrelevant to the experimental results because we were not interested in task completion speed or accuracy in interacting with the interface (e.g., target acquisition accuracy) across subjects.

The computer coach at the center where participants were recruited suggested making the interaction with the interface easier for all participants (especially those with motor impairments that prevented them from using a mouse) by having them point to the interface instead of using a mouse to navigate the vocabulary. To have the participants feel in control of the interface and the task (and due our inability to find a touch screen and adapt the experiment at a short notice), we simulated a touch screen experience. The investigator navigated the mouse tracking the participant's pointing and clicked when the participant applied pressure to the screen. The simulation worked well in that a few participants did not realize that the investigator was controlling the interaction. Those that eventually did, continued applying pressure to the screen since that translated into a mouse click.

6.2.1.5 Participants

We recruited 20 participants from a local support center for people with aphasia. Two of them helped us pilot the study and two did not finish it so their data was excluded from the analysis. We also excluded the data for two additional participants, randomly chosen, to counterbalance the order of presentation of the vocabularies. Participants met two selection criteria: 1) impaired speaking abilities and 2) good comprehension abilities. Participants were informally evaluated by staff at the center and were all medium- to high-functioning. In addition, we recorded their self-reported communication abilities. On average, participants found it easy to understand what is said in a conversation ($M = 3.8$ on the scale of 1: very difficult to 5: very easy) and somewhat difficult to express their thoughts, wishes and needs ($M = 2.7$). Most participants had aphasia due to a stroke, between 3 and 11 years post onset. Two participants have had aphasia less than a year, but their performance did not differ significantly from the rest. The age range for participants was between 40 and 89 with the majority of people in the 50-59 range.

6.2.1.6 Qualitative and Quantitative Measures

We logged each interaction with the interface automatically – which word was selected (from the main dictionary or from the related words) and which button (down arrow, add, home, back, forward, or next phrase) was clicked on. This was done to track what paths people took to find each target word and whether they took advantage of the associations in ViVA.

During the study, the investigators also kept notes of what words were skipped or were particularly difficult to find and recorded participants' comments during task completion. At the end of the experiment, participants were asked to fill out a questionnaire collecting their demographic information and feedback on their experience with the two vocabularies.

6.2.1.7 Design

A within-subjects design was used. Thus, each participant used both vocabularies to search for words. The order of presentation was counterbalanced. While participants were randomly assigned to conditions, we ensured that their cognitive abilities (medium or high) were balanced across order of presentation. The independent variable was the type of vocabulary used to retrieve words and the dependent variable was length of path taken to reach the target words. We also collected subjective feedback on the participants' experience with the vocabularies.

6.3 Results

6.3.1 Quantitative Analysis

A 2x2 (vocabulary x presentation order) repeated measures ANOVA on path length to the target words revealed that there was no significant main or interaction effects of order of presentation of the vocabularies and thus no evidence of a learning effect. This lack of order impact contradicts the results from the study with non-aphasic participants which we described in Chapter 5. We suspect that this is due to the limited number of words that the aphasic participants had to search for (14 as opposed to more than 80 per vocabulary condition in the previous study). We may have observed a learning effect if our aphasic participants had more time to explore the vocabularies.

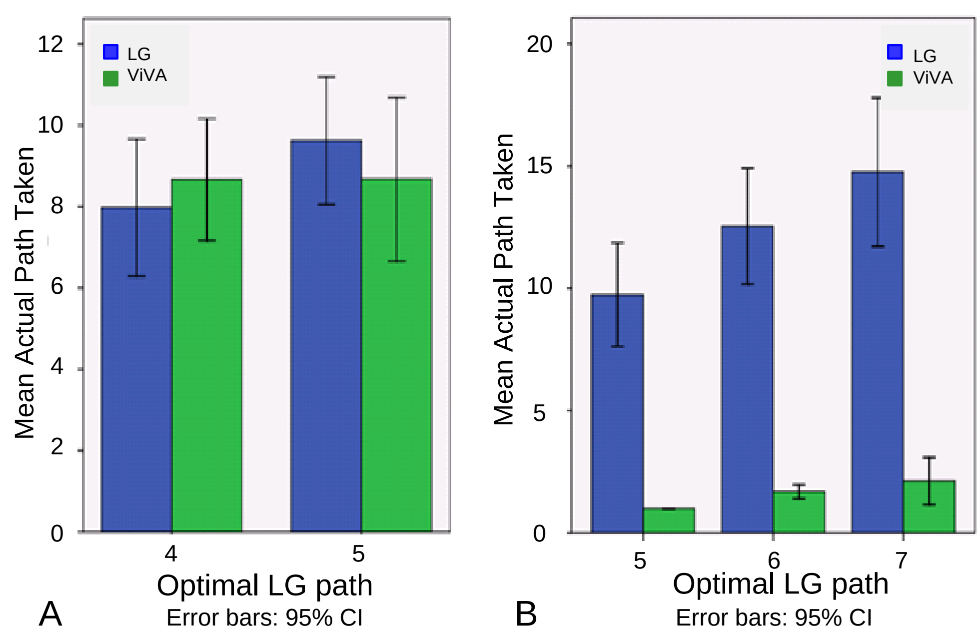

Figure 6.6a & 6.6b: Comparison of the optimal LG path length, measured in clicks, and the path length that participants took when using ViVA and LG. Graph A shows the average path taken to find the first word participants searched for; graph B shows the average path taken to find the second word.

As seen in Figure 6.6b, participants took advantage of the provided associations and once they found one of the words in a phrase, fewer clicks were required to find the second word and complete the task. The average path users took to connect concepts in ViVA was significantly shorter than in LG ($F(1, 13) = 60.58$, $p < 0.01$), supporting our first hypothesis. However, this did not hold for finding the first word in a sentence (Figure 6.6a), highlighting the difficulty of locating which category words belong to. Even though Lingraphica's hierarchy mostly consists of common categories such as food and clothes, some categories such as modifiers were difficult to decipher. Thus, participants had a hard time finding words that are hidden under unusual categories, for example *sleep* is under *dictionary → actions → daily routines → inhale → sleep*. ViVA offers the advantage of finding a word without having to know which category it belongs to. For example, many people did not know *milk* was under: *house → kitchen → refrigerator → dairy*, but could still find it through related words such as *coffee* and *tea*.

6.3.2 Subjective Feedback and Expectations

In order to gather feedback on participants' experience with the two vocabularies, we asked three five-point Likert scale questions after people had completed the tasks with each one of the two vocabularies, ViVA and LG. In support of our second hypothesis, Friedman's test revealed ($N = 14$, $X^2(1) = 7.00$, $p = .008$) that ViVA's organization is less confusing ($M = 2.7$, $SD = 0.9$) than LG's ($M = 3.4$, $SD = 0.9$). Even though both vocabularies rely on the same basic hierarchy, the fact that participants found ViVA less confusing to navigate suggests that having the associations resulted in a more satisfactory experience with the vocabulary.

Our third hypotheses was also supported (based on Friedman's test, N = 14, $X^2(1)$ = 5.33, p = .021), because participants felt that it was easier to find words with ViVA (M = 2.5, SD = 1.0) than with LG (M = 3.2, SD = 0.8). Finally, participants agreed that having related words helped them find words faster (M = 3.9, SD = 0.7). While this claim is only partially supported by the data (all participants found it challenging to navigate to the first word they targeted), the positive feedback suggest that based on the associations, ViVA was perceived generally faster.

Table 6.2. Some direct associations were expected but missing in the vocabulary. For example, 5 people clicked on *leisure* expecting to find *fly*; 4 expected the association *food-milk*. The second table provides additional examples where no exact count was possible.

word_1	word_2	#
leisure	fly	5
food	milk	4
family	shop	4
food	shop	4

word_1	word_2
recreation	bake
family	play
backyard	baseball
time	fly

Participants were not actively encouraged to talk out loud while completing the tasks (to avoid overwhelming them with additional responsibilities and because speaking is challenging for most of them), but some gave us feedback throughout the experiment. This revealed certain expectations they had about associations between words that we had not provided (e.g., *"the* ***mom*** *could* ***make tea*** *... or the* ***dad****"* while searching for *tea* while in the *family* category). These expectations were further revealed in analyzing the paths people took to find certain target words. Some of the associations that were most frequently expected are listed in Table 6.2 along with a few additional examples.

Based on user feedback and the investigators' observations during the experiment, a number of approaches in interacting with the vocabulary emerged. Some participants

explored the vocabulary by clicking on icons without any evident plan of action. Others ***memorized the organization*** while browsing to help them find words faster in subsequent tasks. A few participants were ***guided by the associations*** and based their next move solely on what was displayed on the current screen. Some participants ***formed associations*** and expectations in advance and let their intuition about where a word should be found guide them. It would be interesting to explore these approaches and user profiles further in order to understand how to design better vocabularies and provide better user-vocabulary interaction.

6.4 Discussion and Future Work

The study results revealed that participants found words with ViVA significantly faster because they took advantage of the semantic associations used to augment the basic vocabulary hierarchy. However, the results also highlighted the problem of using preset categories to organize the words in the vocabulary. It was difficult to find the first word in a sentence both for ViVA and LG, because they relied on the same basic hierarchy of categories. To facilitate searching for the first word, people could create their own categories or access points to the vocabulary, but this would add additional burden on the user. When asked whether they would prefer organizing the vocabulary in categories themselves, participants' response on a five-point Likert scale was neutral (M = 3.0, SD = 1.3). Further investigation in what the higher levels of the hierarchy should be is required. One possibility of addressing this problem is to provide custom access points to the vocabulary that reflect the person's profile, but also branch out to more general words in the vocabulary.

Even though the number of missing words as well as the optimal path to find them was consistent across conditions, we thought it was unreasonable to do a pair wise comparison of the average path length participants took to complete each phrase. While the user can find *flower* and *cookies* in the same number of clicks, the path to *cookies* is easier to predict *things* → *food* → *dessert* → cookies. The word *flower*, on the other hand, is reached along the path *things* → *house* → *backyard* → *flower*. Naturally, different users expect to find different things in different categories depending on factors such as life style and personal experiences. Having the words stored in preset categories provides a stable vocabulary organization which facilitates learning over time, but it also often does not make sense to users who tend to get easily discouraged by unsuccessful attempts to find a word.

There is compelling evidence that while performing a task, people's behavior and their decision on what to do next is strongly influenced by the current context (Suchman 1987). This notion known as situated action (Suchman 1987) has been of great importance to interface design in shaping the idea that people take advantage of contextual queues when manipulating an interface to accomplish a task. At each stage of the process, users check the environment to make a decision on what action to take next instead of forming a detailed plan in advance. Applying situated action to the problem of vocabulary organization supports our hypothesis that instead of expecting the users to memorize the vocabulary organization in order to be able to navigate it effectively, we should create the right context at each step of the navigation. The reported results suggest that creating this context by mimicking how words are organized in a speaker's mental lexicon has the potential to assist users with aphasia with word-finding. All users took advantage of the provided associations and a few of them expected certain associations that were not available to lead them to the target word. Thus, user preferences and vocabulary usage patterns should also be taken into consideration in organizing the vocabulary. We explore the question of personalization in the next chapter which describes a longitudinal single-subject case study with an aphasic person.

Chapter 7

Single-Case Study on Personalizing ViVA

After showing that ViVA can improve word-finding by simulating usage of the vocabulary (described in Chapter 4), and in controlled experiments with able users (described in Chapters 5) and with aphasic users (described in Chapter 6), we turned to investigating how ViVA can be adapted to a specific person to reflect his conversation needs. We conducted a six-week single-case study with an aphasic participant. The study enabled us to examine what role personal associations can play in creating ViVA's semantic networks and gave us insights into how the structure can be further adapted according to the user's needs and idiosyncrasies.

The study comprised three phases and was a mix of ethnographic interviews, hypothesis building, and evaluation through controlled tasks. The first phase was of an exploratory nature and let us familiarize ourselves with the participants, find out conversation topics of interest to him, and collect semantic associations that are specific to his daily life and experiences. During the second phase, we evaluated the semantic associations we compiled in phase one, personalized the vocabulary and introduced ViVA to the participant. In this phase, it also became obvious that the vocabulary structure could be improved by introducing additional access points to the vocabulary. Using a number of controlled tasks, we evaluated this idea in the final third phase.

We begin by discussing our choice of methodology and continue with a description of the participant. After elaborating on the three phases, we conclude with a discussion of the results and their implications.

7.1 Choice of Methodology

To build a good understanding of how ViVA can be personalized for and adapted to a specific user's needs, we chose to conduct a case study over the period of six weeks. This choice enabled us to combine ethnography with an experimental evaluation of personalizing ViVA, providing for a more qualitative and comprehensive investigation. The case study is a research methodology that came to prominence in exploring cases where each situation was unique and could not be easily controlled or generalized. It has been well-utilized in the social sciences (e.g., in sociology, business and anthropology). Its advantages have been discussed by a number of researchers among whom (Eisenhardt 1989), (Flyvbjerg 2006), (Noor 2008) and more prominently (Yin 2003). According to Yin's classification of case studies (explanatory, descriptive and exploratory) (Yin 2003), our work fits in the framework of an exploratory study that aimed to investigate how ViVA could be personalized using semantic associations so that the vocabulary adapts to the profile and mental lexicon of a specific user.

Single-case studies and single-subject experiments are also common in the field of aphasiology. They are mainly applied to studying the effects of treatments, because they recognize and account for the variability in impairments among aphasic individuals. As patients' language abilities tend to vary greatly, the responses to treatments can be different. Thus, averaging of results, which is required in group experimental studies, can dilute signs of the treatment's effectiveness for specific patients (Beeson 2006). This is often taken under consideration particularly during early stages of experimentation when the effects of a certain treatment may be unknown. On the other hand, a more in-depth and longitudinal analysis of the response of a single individual may give better insights into the treatment's effectiveness – including results that can be generalizable. Once the treatment effects have been established in individual subjects, it is appropriate to move forward with a study that involves a larger number of participants (Beeson 2006).

While we did not plan to expose our participant to a treatment, we were interested in investigating closely how ViVA can be personalized according to his background, experiences and communication needs, which led us to consider and choose the case study as a research methodology. Case studies are not as common in the field of human-computer interaction as in the social sciences, but they are sometimes applied to interdisciplinary research. Examples include a comparative study on devices for users with severe physical disabilities (Lau & O'Leary 1993), an ethnographic case study to inform the design of an assistive tool for a person with Alzheimer's disease (Cohene et al. 2005), and an early study on a relational database system interface by (Mantei & Cattell 1982). For us, conducting a case study was also advantageous, because, it enabled us to avoid the costs (e.g., substantial time investment) associated with conducting exploratory research with multiple subjects over a long period of time.

7.2 Participant

The participant in our case study was recruited from a local support group for people with aphasia. In order to preserve his anonymity, we will refer to him with the fictional name Sam. Sam is male and 63 years old. He acquired aphasia as a result of a stroke approximately seven years ago. Before his stroke, Sam owned a company, which he has since sold; he no longer works. He leads a relatively independent life, because despite right side hemiparesis (weakness on the right side of the body), he can still drive. He attends the meetings and activities of two support groups for people with aphasia. He is also involved in a number of language rehabilitation and research projects in universities and hospitals in the area.

To help us acquire a better understanding of Sam's language abilities, during phase one of the study, we asked Sam to describe three videos that we selected from YouTube (YouTube 2010). We chose not to use any of the more traditional evaluation stimuli used in aphasiology (e.g., the cookie-theft image (Adams et al. 2006)), because Sam indicated that he has worked with a few of them multiple times and that the task would be boring. All three videos were less than a minute long, showing a brief but clear sequence of events. We played the videos without sound so that the participant can concentrate on the visual details. The first video we picked shows a woman taking her boyfriend to his 30th birthday surprise party (Video 1 2010). She ties his eyes and leads him to a room full of his friends and decorated for the event. The second video shows three people at a table in a café who are paying the waiter (Video 2 2010). One of them drops a few coins on the ground and when they all reach for the change, the empty glasses are knocked off the table and shatter. The third video (Video 3 2010) is a prank on unsuspecting picnickers in a park. A man walks up to them, picks up their fruit tray and walks off with it. He is later shown sitting under a tree and eating the fruit. We used Sam's description of the videos to gain a better understanding of his language abilities which we describe next.

Sam has nonfluent aphasia; he can sometimes produce very short meaningful phrases but is often limited to the use of single words when he wants to communicate something.

His condition is chronic and his abilities have leveled off which is typical for a person who has had aphasia for more than a year. Aphasic patients go through intensive rehabilitation in the first six months post onset. While treatment can help improve residual communicative ability, the progress generally plateaus after a year (Cherney 2002) leaving individuals with chronic aphasia in a relatively stable state. Sam also has apraxia, a deficit in the motor programming stage of speech (Duffy 1995). This often prevents him from articulating words even when has found the word he wants to communicate. Sam compensates for his inability to speak by writing on paper. He often writes a word or even just the first few letters of a word and depends on his communication partner to guess out loud what he is trying to convey. Figure 7.1, for example, shows the notes Sam wrote while talking about some of his trips. He then confirms or rejects the guess. In the case of rejection, he attempts to elaborate by providing a helper word or more letters. He evaluates his auditory comprehension as high (5 on a 5-point Likert scale), his ability to communicate verbally as moderately good or 3, his ability to write – 3, and to read – 3. His reading and writing have always been impaired due to dyslexia. Research suggests that dyslexic readers have impaired access to words in their lexicon (German 1986). Studies with dyslexic children also show that they are poor readers and that they are slow and sometimes inaccurate on tests of rapid automatic naming (German 1986). Thus, Sam’s problems with word-finding due to aphasia could be exacerbated by dyslexia.

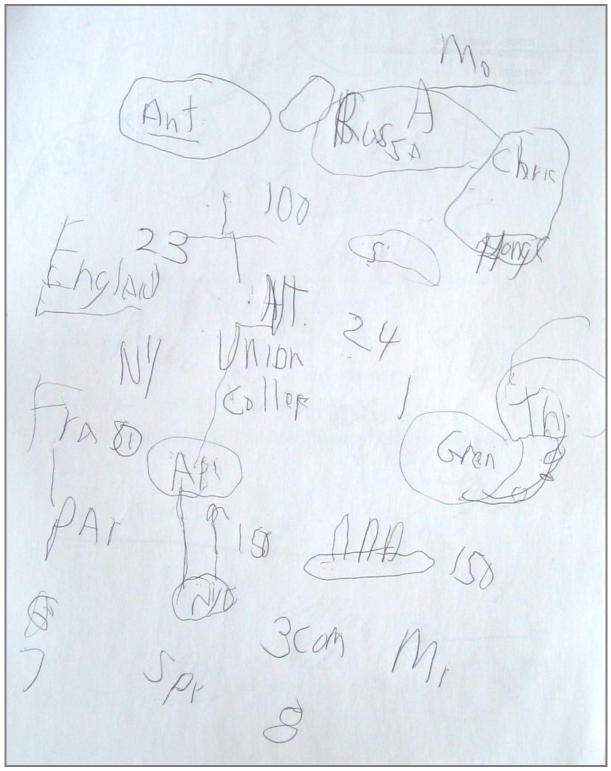

Figure 7.1. Sam compensates for his inability to communicate verbally by writing hints for his conversation partner. *Fra* and *Par* above stand for France and Paris, for example. Here he also drew (in the upper right corner) Russia, China and Hong Kong to explain where his son had travelled for work.

During the first phase of the study, we conducted a set of ethnographic interviews (Spradley 1979) to gathered information on Sam's background, family, daily activities and hobbies. The conversations we had also helped us gain a better understanding of his language abilities. Sam's word-finding difficulties extend to all words, including proper nouns such as the names of members of his family. When unable to verbalize specific words, Sam often tries to help his conversation partner guess the word by providing semantically-related words (e.g., place for food is the dining room, place where you sleep is the bedroom). He often uses empty or indefinite words as fillers, e.g. "that stuff", "good stuff," and "people." To give himself time to form the word while alerting the listener that he is trying to name something specific, he begins with “is called.” For example, when we asked Sam about his children’s occupations, the conversation went as follows (the names of his children are fictional to preserve anonymity):

Researcher: “So what do they do? Tom is …”

Sam: “Tom is a doctor. And her, him and her, ahm …”

Researcher: “His wife?”

Sam: “His wife the same”

Researcher: “They are both doctors?”

Sam, pointing to the name of his other son which he wrote down earlier: “Yes. And him, ahm, is called, it’s called, ahm.” He writes down the letter M. “is called, ahm, what, ahm.”

The conversation continues for approximately two minutes and the researcher eventually guesses that both his son and daughter work in marketing.

When Sam’s communication partner has found a word that is similar or close to the one he is targeting, he often encourages him/her by saying “no, but one more.” Sam has difficulties with numbers as well. He can write down numbers easily (e.g., price, mileage, year, phone number), but he cannot verbalize them. He only attempts to say numbers between 1 and 10 and reaches the desired number by counting out loud from one up. For example, talking about travelling by train in Europe:

Sam: “Is called … that place.” He writes the letters Eu.

Researcher: “Europe?

Sam: “Yes”

Researcher: “Aha. Did you take the train there?”

Sam: "Yes!"

Researcher: "From where to where?"

Sam: "Ahm, but…" He waves his hands signifying that he took the train more than once.

Researcher: "Many …"

Sam: "Yes!"

Researcher: "Oh, ok, that's good. So … London–Paris, for example?"

Sam: "And, yeah. And Greece and …" *Writes down It.*

Researcher: "Italy?"

Sam: "Yes."

Researcher: "Oh, in Italy. Just in Italy?"

Sam: "No!"

Researcher: "Italy to …" Sam is writing.

Sam: "One place, ahm. Oh, yes." Researcher guesses England.

Researcher: "That must have been a long trip."

Sam: "Yes, but the train, this train is the … then, the train one, two, three, four, five days."

Researcher: "It took five days?"

Sam: "Yes."

Sam is very comfortable with computers. He has an undergraduate degree in computer engineering. After working for a number of big companies such as AT&T, he founded his own; his company provided computer network services. At the start of the case study, Sam had just acquired a Lingraphica device (Lingraphica Inc. 2010) and the

SmallTalk software (SmallTalk 2009). This made him a particularly interesting subject for our study because he was just starting to use an assistive communication tool which gave us the opportunity to see how his use evolves and whether the tool meets his needs. His experience with computers and his motivation to use an assistive device also spared us the effort involved in introducing an elderly person to a technological tool and needing to provide extensive training.

We met Sam twice a week over the period of six weeks. The meetings were one hour long and were conducted at his house at his convenience. Sam was compensated $10 for each meeting.

7.3 Phase 1: Getting to Know the Participant

The first phase of the case study consisted of six ethnographic interviews. The purpose of this phase was to familiarize ourselves with the participant, i.e., learn about his daily life, his family, his past and present. More specifically, we wanted to understand what he likes to talk about and to compile conversation topics and a list of words that he uses often. We needed this information to personalize ViVA and design exercise to evaluate it with Sam. We began with questions on how he had been using his assistive communication tools (Lingraphica and SmallTalk).

The support staff at Lingraphicare Inc. composed sentences for Sam that he transferred to his SmallTalk system, Lingraphica's mobile accessory (see Figure 7.2). SmallTalk runs on an iPod Touch which he preferred to use because of its form factor and mobility. Sam had stored sentences about his family, about his work before he had a stroke, and about his family's vacation destinations. For example, he has a few sentences listing where he has worked over the year in the format "I worked for X number of year in Y company." In Lingraphica, he also has phrases and sentences that he could use when talking to his wife (e.g., "Are you hungry?", "How was your day?", and "Want to go to the mall?"). In addition, Sam had a folder with phrases describing a recent car accident he was in. Sam had not tried putting sentences together by himself at this time and he was dependent on the support of the Lingraphicare team. When asked about what he likes to do with the two assistive tools he had just acquired, he expressed preference for prerecorded phrases that he can use later on the SmallTalk as opposed to finding single words on Lingraphica in a practice exercise.

Figure 7.2. Lingraphica's mobile companion, SmallTalk, runs on an iPod Touch and provides access to a limited subset of Lingraphica's vocabulary.

We proceeded with a set of ethnographic interviews which were a combination of general exploratory questions such as "Tell me about your family?" and "Tell me about your life before aphasia?" and more specific and targeted questions such as "What does your son do?" and "How often do you visit your children?" Once we narrowed down a topic that Sam liked to discuss, we would ask him to elaborate on it as much as possible.

For example, he used to travel often, because of work as well as for pleasure, with his family. He told us that he took a trip to Greece so we asked him what places in Greece he visited, what kind of food he tried there, and what he liked and disliked about his trip. Over the course of our first six meetings, we found out about Sam's family, his time in college, the trips he has been on and his hobbies. All meetings were voice-recorded. As a result of these ethnographic interviews, five topics emerged as most important to Sam. Below we summarize the topics and give excerpts from our conversations to familiarize the reader with the participant.

Car: Sam likes cars. He used to have a Porsche which was his favorite car, but he had to sell it recently. The car was manual which made it hard for him to drive after he had a stroke which resulted in right side hemiparesis. Recently, after he damaged his other car in an accident, he decided to buy a new BMW. He visited different dealerships in the area until he chose the car he liked. He had it fitted with a steering wheel knob and a special pedal to make it easier to drive. Sam likes to drive fast and used to take his Porsche to a race track. Below we include a discussion of Sam's decision to buy a new car after his most recent car was badly damaged due to a car accident. He also talks about selling his Porsche.

Sam: "Yes, but now my car …"

Researcher: "Yes, show me …" Sam navigates to a picture of the car he wants to buy which he stored in his Lingraphica.

Sam: "Ahm, think … I think, this one"

Researcher: "So that's the one you picked?"

Sam: "Yes, yes."

Researcher: “BMW 335 xDrive?”

Sam: “No, ahm, no, one more, no, no. My car here …”

Researcher: “The Porsche?”

Sam: “One week more.”

Researcher: “And then it’s gone? Who’s buying it?”

Sam: “Ahm, ahm. One person, I don’t know the name. And Julie that stuff” (*Julie is the name we will use for Sam’s wife.)*

Researcher: “She is dealing with selling it?”

Sam: “Yes, ‘cause me, this money” *Sam writes down 22,000.*

Researcher: “22?”

Sam: “Yes, but her …” He writes down 25,000.

Researcher: “She managed 25?”

Sam: “Yea, is good. So what …”

Researcher: “The car is gonna be gone? Well, so then, you can get the new car.”

Sam: “Yes.”

Researcher: “You said it’s gonna be a used car or you haven’t decided?”

Sam: “Me ... Julie called my car two years old … Me, nothing.”

Researcher: “You want a new one?”

Sam: “Yes, yes!”

Family: Sam has two sons and a daughter. His wife is a teacher, one of his sons is a doctor and he lives on the west coast, and his other son and his daughter work in marketing and live in New York City. Both of his sons have children. Sam and his wife visit their children when they can. Recently, his daughter left her dog for Sam and his wife to take care of her, because she could no longer keep it in her apartment in the city. Below is an excerpt of a conversation with Sam where he talks about visiting his children in New York:

Researcher: “Do they usually come to visit or do you go?”

Sam: “Yes, yes.”

Researcher: “Both?”

Sam: “Yes, yes. Sometimes here, sometimes there.”

Researcher: “Sometimes there. Do you like going to New York?” … “What do you like to do when you go to visit?”

Sam: “Well, sometime theater. Sometimes, them go the houses and the two things … kids” *Sam explains where his children live in New York city.*

Researcher: “How do you decide whom to stay with?”

Sam: “Ahm, the same. But, but him, ahm, go the baby.”

Researcher: “Oh, they now have a baby? A small baby?” Sam writes down the baby’s age - 1½.

Researcher: “For example, when you were there last time, what did you do? … Did you just go out for dinner or did you do something else”

Sam: “The park or one … this here, the … ahm, this …two them … ahm, museums” *Sam draws the location of the Metropolitan museum*. “and the thing for people.”

Researcher: “history …”

Sam: “Yes, yes!”

Hobbies: Sam collects coins. He has an extensive collection of old gold and silver US coins. He receives a newsletter for coin collectors which helps him decide what coins to buy next. He also has a collection of Pelikan and Waterman pens and fountain pens. He used to collect stamps when he was in his 20s. Sam also likes to swim. He used to be on the swim team in high school and in college as well as a lifeguard during a few summers. Sam is showing his coin collection:

Researcher: “These look more familiar - 1992”

Sam: “Yes, ‘cause, yes. And that stuff … gold and silver.”

Researcher: “These are silver, except for that one.”

Sam: “Yea.”

Researcher: “How do you find out about the special editions?”

Sam: “Ahm.” Sam walks to a cabinet and pulls out a drawer. “Here.”

Researcher: “Oh, you get a newsletter? I see.”

Sam: “Yes.”

Travel: Sam loves to travel. He used to travel often when he was working and had his own company. He has also travelled with his wife and with his children. Currently, he mainly travels to visit his children, friends in Vermont and to the Cayman Islands where his family has a timeshare. His favorite trip was to Greece and Turkey. He went there with his wife and his daughter to visit one of his sons who was there on a semester-abroad program. He especially enjoyed a boat trip they took around a number of Greek and Turkish islands. He also liked the food and the markets there. Sam talks about his favorite trip:

Researcher: “What would you say was your favorite trip?”

Sam: “I don’t know.”

Researcher: “So think about everywhere you’ve been, because you’ve travelled a lot.”

Sam: “Yes.”

Researcher: “Which one do you think was your favorite? It could be more than one, but just think about why you liked it, what you liked about it.”

Sam: “No, ‘cause my...”

Researcher: “It’s hard to decide?”

Sam: “Yea, ‘cause same thing for me, same.”

Researcher: “You liked all of them?”

Sam: “Yea, ‘cause me do two, ahm... is called, the places.”

Researcher: “Mmm, so let’s see. What was the most different experience? Because, for example, when you travel in Europe, sometimes cities may seem similar”

Sam: “Oh, well, yes.”

Researcher: “They are all beautiful, but you could say this city looks like the other city. What was the most different, newest, something you hadn’t seen so it was impressive in that way?”

Sam: “Ahm.” *He writes Gre.* “That trip.”

Researcher: “That’s the place? Greece?”

Sam: “Yes, and Turkey.”

Researcher: “and Turkey. So that was the most different?”

Sam: “Well, yes.”

Researcher: “Why was that?”

Sam: “’cause them.”

Researcher: “Let’s start with Greece? Why was Greece so different than anything else you had seen before?”

Sam: “’cause the …” *followed by speech that is hard to decode.* “Turkey and Greece, the islands.”

Researcher: “the islands?”

Sam: “Yes, and... so, one, two, three times there.”

Researcher: “You went three times there?”

Sam: “Yes.”

Work: Sam used to be a computer engineer. He worked for a number of big companies such as AT&T until he decided to start his own. His company provided network services and he employed approximately ten people. He sold the company after he had a stroke. In the following excerpts, Sam talks about running his company from his home.

Researcher: “Were you running your business from home?”

Sam: “Yes.”

Researcher: “Oh, ok, that’s why you had to make sure that everything is well connected.”

Sam: “Yes, yes.”

Researcher: “Did you have other people working for you?”

Sam: “Yes.” He writes down the number of his employees on a piece of paper.

Researcher: “Six?”

Sam: “Yes, but them go to the, ahm, the people.”

Researcher: “To the people? Six people were working for you?”

Sam: “Yes, yes.”

Researcher: “They were going to visit people?”

Sam: “Well, and …”

Researcher: “Were they sales?”

Sam: “No, my sales.”

Researcher: “They were selling for you?”

Sam: “No! Them, the, the people there, ahm, engineers.”

Researcher: “Engineers?”

Sam: “Yes.”

Researcher: “Oh, so, these six people were engineers?”

Sam: “Yes.”

Researcher: “So were they working from their home?”

Sam: “No!”

Researcher: “In an office?”

Sam: “The people, ahm.”

Researcher: “People’s houses?”

Sam: “No”

Researcher: “Company?” Sam writes down a note. “At AT&T?”

Sam: “Yes, and that people.”

Researcher: "Was AT&T your client?"

Sam: "Yes!"

Researcher: "So they worked at the client?"

Sam: Laughs in approval.

This first phase resulted in an understanding of the participant's life, daily interactions and language abilities. We compiled a list of Sam's favorite topics for conversation and a list of frequently used words that relate to these topics. We used the information we collected to personalize ViVA. We describe that process next.

7.4 Phase 2: Personalizing ViVA

Phase two was dedicated to personalizing ViVA, understanding how Sam takes advantage of semantic associations and introducing him to the idea of using them in navigating the vocabulary. Phase two was completed in three meetings.

7.4.1 Mining Personal Associations

Based on the recordings of our conversations with Sam, we compiled a list of 100 frequently used words related to the topics we discussed above. We constrained our choice of words to verbs and nouns which carried most of the meaning in our discussions. Since our conversations with Sam compiled to approximately six hours total, our definition of a frequently-used word was a noun or a verb that he used to describe a person, object or an event two or more times during phase one. For example, in talking about his son, Sam explained that he is a doctor and that he has a wife with whom he has two children. Because he used the words more than twice, *son*, *doctor*, *wife* and *children* were considered frequently used. Another example consists of the words *travel*, *Greece*, *Turkey*, *boat*, and *island* referring to his trip to this Mediterranean region, which was his favorite vacation. Table 7.1 lists all hundred words, their frequency of use by Sam and their frequency of use according to the British National Corpus (The British National Corpus 2007). The BNC frequencies show that the collected set of words presents a good balance between words frequently used in general (e.g., *call*, *doctor* and *family*) and words that are specific to Sam' lexicon (e.g., *BMW*, *lifeguard* and *Cayman Islands*). A comparison of the two columns also suggests that to be able to effectively represent the mental lexicon of a specific user, it is not sufficient to account only for associations reflecting general knowledge of semantic relatedness.

The frequently-used words were then paired according to their use within the same sentence (or general context where a complete sentence was not formed). We paired words that were either used by Sam or uttered by the author, but confirmed as the correct target word by Sam. Within the context of Sam's Mediterranean trip, *Greece* and *Turkey* were both paired with *travel and island*, for example; within the family context, *son* was paired with *doctor*, *wife* and *children*. The resulting list consisted of 160 pairs (see Table 7.1). We expected that these pairs of words would be strongly associated for Sam and would evoke each other, e.g. *son* would evoke *doctor* and *doctor* would evoke *son*. We chose a subset of the pairs to evaluate with Sam (evaluating the whole set would have been too time-consuming and tedious for the participant). We randomly picked fifty pairs and asked Sam to produce a rating representing the strength of bi-directional evocation between the two words.

Table 7.1. List of words frequently used by the participant to describe past experiences and current events. Note that the BNC frequencies are not associated with a specific meaning of the word. For comparison, frequently used word such as *eat* and *drink* have a BNC count of 7264 and 7674 respectively.

Word	Number of times used	BNC Frequency	Word	Number of times used	BNC Frequency
accident	3	6316	lifeguard	3	60
apartment	4	1233	like	4	147744
baby	2	8605	magazine	3	4649
backyard	4	125	market	3	29902
beach	6	3724	money	6	36551
BMW	7	519	mountain	3	3824
boat	4	5266	museum	3	6771
business	2	35512	new	6	124308
buy	7	12295	New York City	7	206
call	5	18766	old	4	52716
camp	3	3362	own	3	68878
car	10	26886	park	2	11333
Cayman Islands	5	34	pen	5	2132
children	6	45724	people	20	121711
city	4	22873	play	4	21143
coin	6	1204	police	2	27056
collect	3	2765	pool	4	4478
college	7	10257	Porsche	6	259
company	7	39945	price	2	18324
compete	2	1913	race	2	7855
computer	12	13810	read	3	17125
crash	2	2427	rent	2	3440
dance	3	4324	restaurant	5	3438
daughter	4	9194	roof	2	4017
doctor	5	10224	rowing	2	413
dog	4	7872	San Francisco	4	621
drink	3	7674	sell	6	7588
drive	5	8800	silver	2	4882
employee	2	3150	ski	3	875
engineer	3	2262	soccer	3	1322
Europe	5	18087	son	7	12644
exercise	4	8582	stop	2	14553
expensive	3	5743	swim	11	1396
family	7	33761	teacher	3	8757

fast	5	7349	team	2	18786
find	3	40906	theatre	2	5788
fish	4	10750	town	2	17881
food	3	18681	toy	3	945
fraternity	5	228	train	5	7861
friend	7	16374	travel	10	7221
fun	3	5004	trip	5	4496
gold	4	7558	Turkey	4	2132
grandchildren	4	492	university	4	23507
Greece	7	1655	vacation	6	290
hit	2	9656	Vermont	4	68
hobby	3	607	visit	5	12322
home	2	52702	watch	2	9119
house	5	49424	wedding	4	3238
invite	3	1209	wife	7	16497
island	7	6541	win	2	10427
kitchen	3	7633	work	10	89823

We showed Sam a list of word pairs and asked him to indicate how strongly each pair is associated. We emphasized that he should concentrate on evaluating the strength of evocation according to his personal experiences, not just based on general usage. We asked the participant to indicate his answer on a 5-point Likert scale similar to the one shown in Figure 7.3. He rated more than 90% of the associations (see Figure 7.4) as immediate associations. Examples include *wife–teacher*: 5, *travel–work*: 5, *fun–swim*: 5. The evaluation indicated that we had collected associations personal to Sam.

1	**2**	**3**	**4**	**5**
No Association	Remote Association	Moderate Association	Strong Association	Immediate Associations

Figure 7.3. Scale used by participant to indicate how strongly two words are associated.

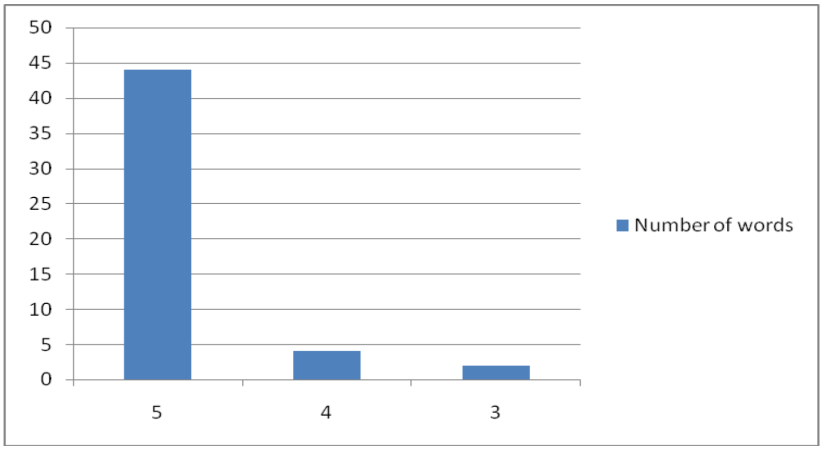

Figure 7.4. The majority of a random sample of the associations extracted from the conversations with Sam was strongly related.

To adapt ViVA to Sam's lexicon, we incorporated his personal associations into the vocabulary. ViVA's core vocabulary consisted of Lingraphica's core hierarchy – all Lingraphica words in their base forms, which compiled to approximately 3000 words. The basic hierarchy was augmented with semantic associations extracted from the evocation data we collected through an online experiment (see Chapter 3 for more details). For each word in the core hierarchy, we included the seven highest rated associations. The number was chosen so that the participant does not need to scroll through the related words when using the interface to the vocabulary. This added approximately 2000 links between pairs of words in the core vocabulary. The interface to the vocabulary was the same as the one we used in the study with aphasic participants described in Chapter 6. The categories that comprise Lingraphica's hierarchy were displayed in the lower portion of the main panel and the related words were displayed in the upper portion (see Figure 7.5). Once the user clicked on a word, his personal associations were displayed in alphabetical order, followed by association due to evocation displayed in decreasing order of strength of evocation.

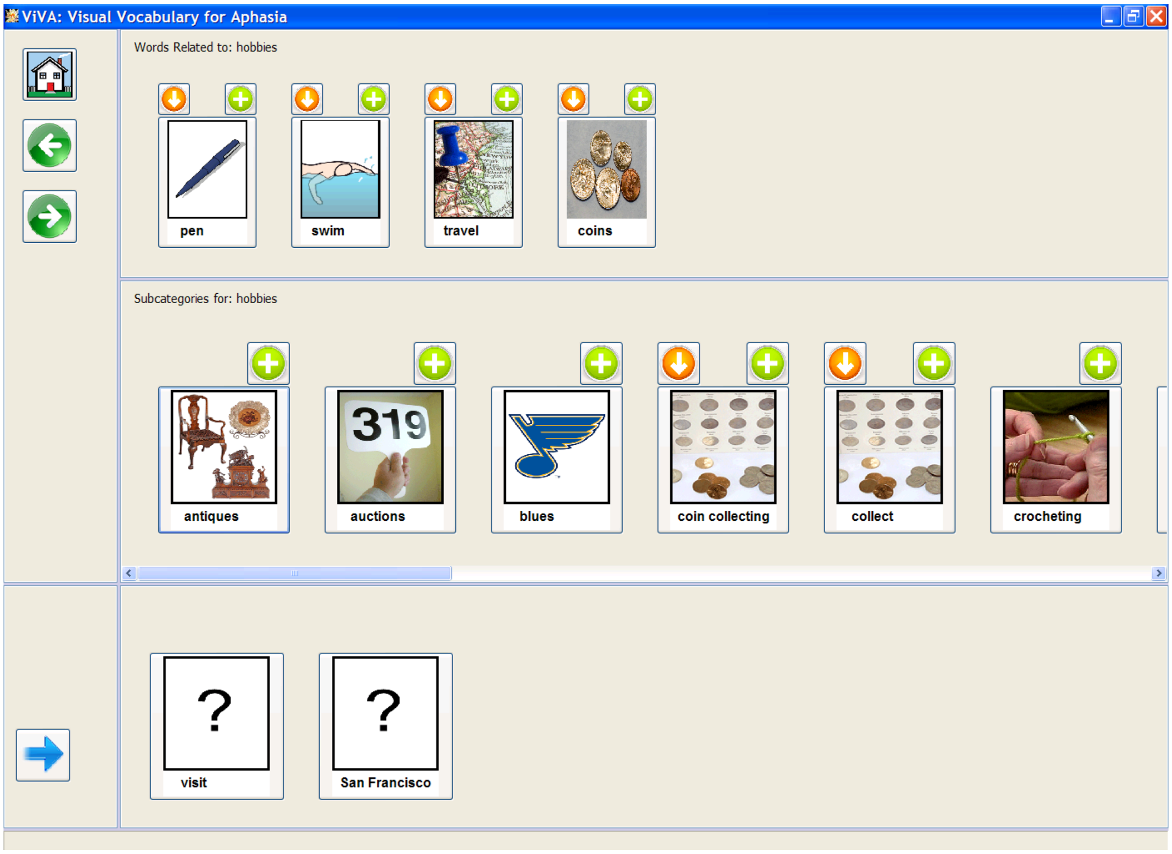

Figure 7.5. ViVA with integrated personal associations displayed in the related-word panel in the upper part of the interface. Because the participant enjoys swimming and collects pens and coins, the words *pen*, *coins*, and *swim* are associated with *hobbies*.

7.4.2 Word-finding with ViVA

To accustom Sam to using ViVA, we conducted an experiment similar to the one described in Chapter 6. We asked Sam to complete, using ViVA, a number of phrases with missing pairs of words. The task consisted of six scenarios comprised of an image followed by a sentence related to the image (refer to Figure 7.6 for an example scenario). We first showed the image to the participant, followed by the sentence which was read out loud. Sam was asked to identify, from a choice of three words, the word that was missing in the sentence. Once he had chosen the correct word, he had to find the missing word and the word it is associated with (*cuts* and *cake* in Figure 7.6). The participant took advantage of the provided associations when navigating to the target words, but we do not present formal results, because this exercise was only meant to familiarize Sam with ViVA.

Figure 7.6. Stimulus used for word-finding task.

Our participant was high-functioning and very comfortable with using a computer. Thus, the next task we gave Sam was adapted from the phrases we used in the study with able participants. We gave Sam 33 sets of words to find in ViVA. Table 7.2 lists the words and the paths he took to find them. We did not tailor the task to encourage use of any of his personal associations in particular. The exercise was designed to make Sam more comfortable with ViVA and to illustrate the role of word associations in ViVA. We observed what paths Sam took to connect the target words and whether he utilized the provided associations. Table 7.2 summarizes the results.

Table 7.2. Target word pairs, ViVA's optimal path, the actual path taken by the participant and utilization of available related words.

Word pair	Optimal ViVA path *(# of clicks)*	Actual ViVA path *(# of clicks)*	Utilization of related words along search path
eat – drink	2	8	attempted, but unsuccessful
football – game	1	1	football→game
sleep – health	1	1	health→sleep
university–class	1	1	university→class
cereal–breakfast	1	1	breakfast→cereal
van–engine	2	2	engine→car→van
meat–steak	2	2	meat→beef→steak
room–kitchen	1	1	kitchen→room
coffee–dessert	1	1	coffee→dessert
dessert–tea	2	6	breakfast→coffee→tea
bathroom–sink	1	1	bathroom→sink
tea–dessert	1	1	tea→dessert
dessert–pudding	1	1	utilized direct link in hierarchy
student–explain	2	7	none
aunt–daughter –uncle	1–1	1–1	aunt→daughter→uncle
breakfast–meat	1	1	breakfast→meat
chair–cushion	2	3	cushion→armchair→office→chair
food–smell	2	14	attempted, but unsuccessful
tasty–breakfast	2	2	tasty→food→breakfast

travel–drive	1	1	travel→drive
hungry–eat	2	10	food→house→dining room→eat
buses–trains	1	1	buses→trains
tea–sugar–honey	1–1	1–1	tea→sugar→honey
bacon–tasty	1	1	bacon→tasty
cook–bake	1	1	cook→bake
call–write	1	1	call→write
brown–rice–meat	1–2	1–2	brown→rice→ham→meat
plants–animals	1	1	plants→animals
pill–swallow	2	13	attempted, but unsuccessful
bread–tasty–butter	1–2	1–2	bread→tasty→sugar→butter
motorcycle–drive	2	4	motorcycle→car→speed→fast→drive
food–garlic– smell	1–1	1–12	food→garlic→…→senses→smell (partially shown path)
tea–coffee–milk	1–2	1–2	tea→coffee→cream→milk

As seen in Table 7.2, Sam used many of the provided association to navigate to the target words. He emphasized that the associations are useful for him, because he cannot spell and thus he had difficulty searching by typing (in Lingraphica, for example). Finding words through the categories is often confusing for him, because they are broadly defined (e.g., finding *eat* in the *actions* category). In addition, Sam revealed that it is especially difficult for him to retrieve verbs. Thus, he found the cross-part-of speech associations, e.g. *food–eat* and *drive–car*, particularly useful. As he was navigating through the vocabulary, he recognized some of the personal associations that we had incorporated and enjoyed pointing them out (e.g., his *car* is *new*, his *wife* is a *teacher*, and *travelling* is *fun*). As in the controlled study with aphasic participants, ViVA was able to guide Sam in connecting words once he had found the first word in a pair, but it was still challenging to navigate the upper layers of the hierarchy where the categories appear more abstract or broadly-defined. Confirming these difficulties with Sam brought about the idea of introducing personalized access points to the vocabulary to complement the main access (through the dictionary icon seen on Figure 7.7). We hypothesized that introducing personalized access to the vocabulary based on the topics Sam likes to discuss, will enable him to find words faster as well as encourage him to browse the vocabulary and be more expressive as opposed to using single words to communicate.

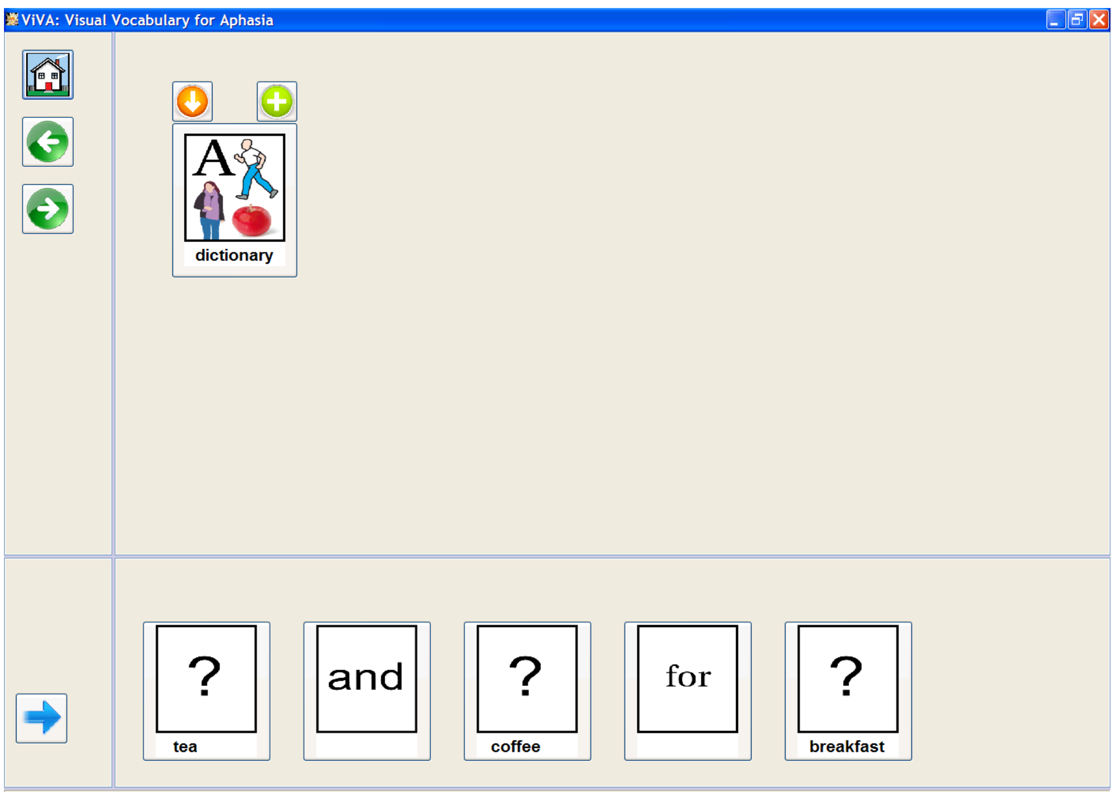

Figure 7.7. The main access point for the vocabulary is the dictionary icon which enables the user to descend the core hierarchy.

7.5 Phase 3: Introducing Personalized Access Points

Most assistive communication tools enable the creation of folders where users can store items for quick access. These items can be words, sentences and personal photographs. Sam, for example, has customized folders in his Lingraphica (called pages in Lingraphica terminology) about his family and certain events. As we mentioned earlier, one folder holds sentences related to a car accident he was in; another folder contains pictures of the members of his family and a third one reveals where his family likes to go on vacations. Enabling this type of customization is beneficial, because it provides quick access to frequently used words and sentences. However, to add to the folder and expand his/her relevant vocabulary, the user still needs to browse the main hierarchy in search of appropriate words.

To provide quick access to frequently-used items and improve overall word-finding using these items, we introduced personalized access points to the vocabulary. The access points are words or concepts that already exist in the vocabulary hierarchy and are brought up to its highest level. Once the user clicks on a personalized access point, s/he sees its subcategories and related words. If the access point represents a concept or a topic that the user refers to often, clicking on it can help him navigate to related words faster. We concentrated on exploring the use of such personalized access points in navigating the vocabulary during our last three meetings with Sam.

7.5.1 Word-finding with Personalized Access Points

The access points we created for Sam were aligned with the topics defined in phase one of the study – *car*, *family*, *hobbies*, *travel* and *work*. In the vocabulary interface, they were represented by five additional icons on the home screen (see Figure 7.8). The user could navigate the vocabulary through one of these access points or fall back on traditional access through the dictionary icon.

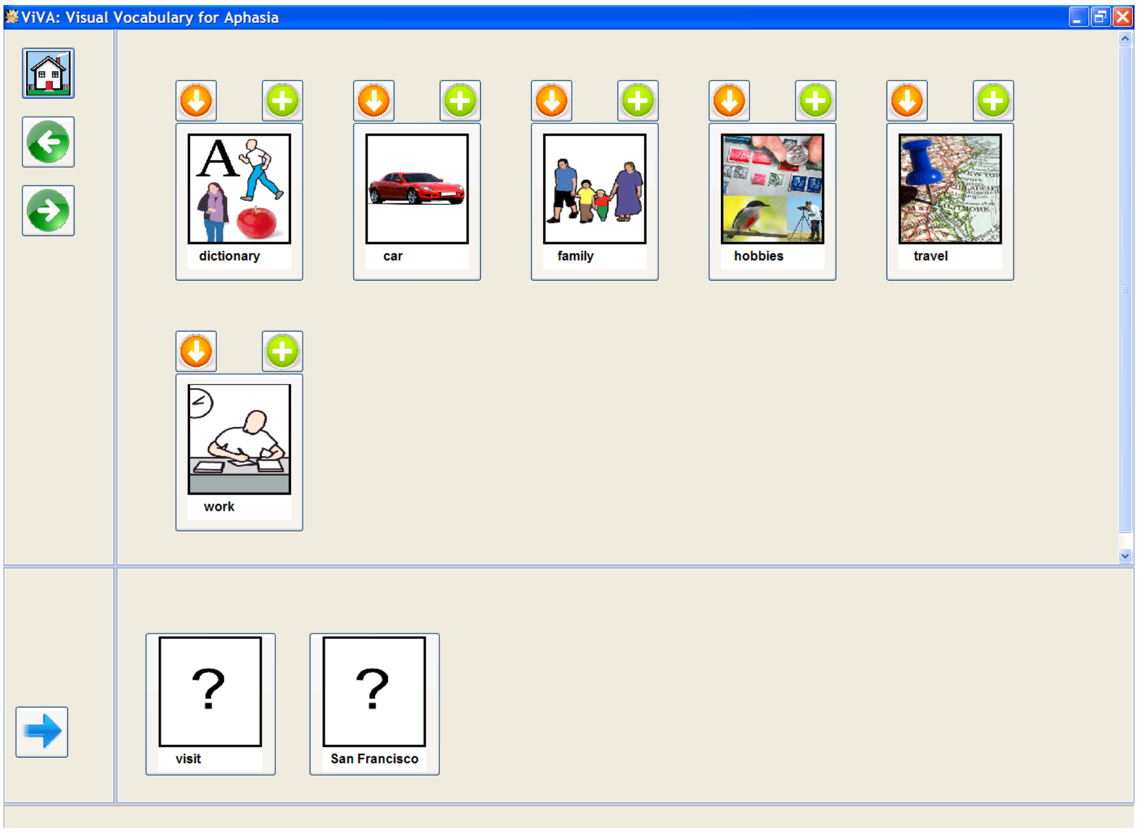

Figure 7.8. ViVA with personalized access points to the vocabulary (compare to Figure 7.7).

7.5.1.1 Task 1

To test whether the personalized access points will improve word-finding for Sam, we created two sets of word-finding tasks. The first set was a brief scenario followed by a sentence with two of the words in the sentence highlighted. We asked Sam to find the two highlighted words. The scenarios and target words were as follows:

1. Someone asks what your son does and you want to answer: My son is a doctor.
    a. Find **son** and **doctor**;
2. You want to ask Julie to call the doctor and set an appointment for you.
    a. Find **call** and **doctor**;
3. Someone asks whose dog this is. You explain it is your daughter's.
    a. Find **daughter** and **dog**;
4. You see a small dog running on the street.
    a. Find **small** and **dog**;
5. Someone asks you how you like to exercise and you tell them you go often to the pool to swim.
    a. Find **pool** and **swim**;
6. You are in a restaurant and you want cheesecake for dessert.
    a. Find **cheesecake** and **dessert** ;
7. Swimming is good for your health.
    a. Find s**wim** and **health**;
8. You are telling someone you just bought a new car.
    a. Find **new** and **car**;
9. Someone with a fast car is racing past you on the highway.
    a. Find **car** and **fast**;
10. You like to travel with your family.
    a. Find **travel** and **family**;
11. When I go to Bulgaria I travel by plane.
    a. Find **travel** and **plane**;
12. Someone asks you how many children you have. You explain you have two sons and a daughter.
    a. Find **son** and **daughter**;
13. Your son's children are young.
    a. Find **children** and **young**;
14. You were telling me you had a Porsche and now have a BMW.
    a. Find **Porsche** and **BMW**;
15. A woman baked a cake.

    a. Find **bake** and **cake**;

16. You told me where you have travelled. Let’s say you want to tell someone else you visited Greece and Turkey.
    a. Find **visit** and **Greece**;
17. You were telling me you like to drive fast.
    a. Find **drive** and **fast**;
18. Children often play in the park.
    a. Find **play** and **park**;
19. Someone asks you about your hobbies and you tell them you collect coins.
    a. Find **collect** and **coins**;
20. The other night I went out for dinner with a friend and we needed coins to park.
    a. Find **coins** and **parking meter**;
21. I am from Bulgaria.
    a. Find **Bulgaria**.

All target word pairs were retrieved from the word list complied during phase one of the study. Some of the scenarios were relevant to Sam’s life, such as scenario 5; others were hypothetical situations (e.g., scenario 9). The scenarios relevant to the participant implied utilizing the personalized access points; the hypothetical scenarios let us observe how Sam would search for the targeted words (through the personalized access points or through the main dictionary icon). Table 7.3 presents the paths taken to reach the target words and Sam’s utilization of the personalized access points for each set of target words.

Table 7.3. The participant utilized the personalized access points to access the first word in a pair, and the related words to connect the two words.

Target Words	Optimal Path Length *(# of clicks)*	Actual Path Length *(# of clicks)*	Utilization of Personalized Access Points *(personalized access points are in bold)*
son–doctor	1	1	**family**→son→doctor
call–doctor	1	1	**family**→people→hospital **people**→doctor→call
daughter–dog	1	1	**family**→daughter→dog
small–dog	1	1	**family**→daughter→dog→small
pool–swim	1	1	**hobbies**→swim→pool
cheesecake–dessert	1	2	*Found dessert through the main dictionary access: dictionary→things→food→dessert→ cheesecake*
swim–health	2	2	**hobbies**→swim→exercise→health
new–car	1	1	**car**→new
car–fast	2	2	**car**→drive→fast
travel–family	1	1	**travel**→family
travel–plane	2	2	**travel**→transportation→plane
son–daughter	1	2	**family**→daughter→brother→son
children–young	1	1	**family**→children→young
Porsche–BMW	1	1	**car**→BMW→Porsche
bake–cake	1	1	*Found cake through the main dictionary access; bake was a related word: dictionary→things→food→dessert→ cake→bake*
visit–Greece	2	5	**travel**→places→Greece→Turkey→ country→places→travel→visit
drive–fast	2	2	car→drive→fast
play–park	1	1	**family**→children→play→park
collect–coins	1	1	**hobbies**→collect→coins
coins–parking meter	1	1	**hobbies**→coins→parking meter
Bulgaria	n/a	n/a	**family**→places→Europe→Bulgaria

The participant got accustomed to using the personalized access points quickly, acknowledging their benefits. As seen in Table 7.3, he accessed the vocabulary through them for most of the scenarios. Sam commented that it is noticeably more difficult to find words he does not associate with as much (e.g., the pairs *bake–cake* and as *cheesecake–dessert*). The participant also found it natural to use some of his personal association to navigate to more general terms, for example: *hobbies→coins→parking meter*.

7.5.1.2 Task 2

To investigate further whether associations accessible through the personalized access points we introduced can guide word-finding successfully, we designed one last task for Sam. We asked the participant to find a set of word pairs not presented in any specific context. Based on our conversations with the participant, we chose words that were related to but not immediately associated with the topics relevant to his personal life. We also asked Sam to indicate on a 5-point Likert scale the strength of evocation between the two words in a pair after he had found them in the vocabulary.

Table 7.4 shows the optimal path connecting the two terms, the path Sam took to connect them and his utilization of the personal access point.

Table 7.4. Utilizing the personal access points for pairs of words not presented in specific context showed potential for improving word-finding.

Target Words	Optimal Path Length *(in clicks)*	Actual Path Length *(in clicks)*	Strength of Evocation	Utilization of Personalized Access Points
teacher–explain	2	2	4	family→wife→teacher→student→explain
vacation–island	2	3	5	travel→vacation→places→Turkey→island
speed–accelerate	1	2	3	car→speed→fast→accelerate
coins–old	1	1	5	hobbies→coins→old
Cayman Islands–beach	1	1	5	travel→vacation→Cayman Islands→beach
doctor–medication	2	2	4	family→people→hospital people→doctor→prescription→medication
buy–present	1	1	3	hobbies→coins→buy→present
dog–animal	2	2	5	family→daughter→dog→cat→animal
stamp–collect	1	1	4	hobbies→collect→stamp
train–old	4	n/a	3	travel→transportation→train→*…did not find old*
repair–toolbox	2	5	3	car→repair→glue→*back arrow* →car→garage→toolbox
computer–company	2	2	5	work→computer→office→company
visit–San Francisco	3	9	5	travel→places→San Francisco→…→travel→visit
sell–company	1	1	5	work→office→company→sell
lifeguard–swim	2	2	5	travel→visit→places→Cayman Islands→swim→beach→lifeguard

talk–doctor	2	2	4	family→people→hospital people→doctor→call→talk
pen–write	1	1	5	hobbies→pen→write
eat–restaurant	3	6	5	*Found eat through the main dictionary access: dictionary→things→food→eat*
university–learn	1	1	4	family→son→university→learn
visit–friend	1	1	3	family→people→friend→visit
university–fraternity	1	1	5	family→son→university→fraternity
order–Chinese food	1	1	5	family→food→Chinese food→order
New York City–fun	1	1	5	travel→places→New York City→fun
buy–house	3	7	3	family→house … dictionary→actions→need→money→buy
write–list	1	1	4	work→office→pen→write→list
friend–Vermont	3	3	5	family→people→friend→visit→places→Vermont

As seen in Table 7.4, Sam utilized the personalized access points for most of the target words. The paths he took were not always the ones we expected or predicted he would follow. For example, instead of taking the shorter path to *doctor*: *family→son→doctor*, he preferred the path: *family→people→hospital people→doctor*. When he did not utilize the personalized access points, he navigated to the target word through the dictionary icon, but still took advantage of the provided related words. For example, he found *eat* in the pair *eat–restaurant* through: *dictionary→things→food→eat* where *eat* is a member of the related words list for *food*. The results show that the personalized access points make accessing the first word in a pair easier and also provide the start of a more efficient path to other related words that are not necessarily immediately associated with the concept represented by the access point.

7.6 Summary of Results and Discussion

Our observations and the results from the semi-controlled tasks Sam completed, showed that he took advantage of both associations due to the evocation data, representing averaged ratings of human judgments of semantic relatedness, as well as due to his background, experiences and the things he likes to talk about. He found the personalized access points to the vocabulary particularly useful. The results from the last task showed that the personalized access points help him navigate the vocabulary faster (by shortening the browsing path between associated concepts) and the paths he follows reflect associations based on his profile. For example, for the pair *buy–present* which Sam found moderately associated, he navigated along the path: *hobbies→coins→buy→present*, because he often buys old US coins that he collects.

The main drawback of the work we presented in this chapter is that the single-case study makes it difficult to generalize the results. We chose to run the study with one participant, because the exploratory nature of the research requires significant time investment. Even though Sam drives, he preferred meeting at his home which translated into a four-hour drive each week combined with approximately three hours of setting up and conducting the meetings. In addition, the research materials we used often had to be prepared in between meetings based on the progress at the time. Thus, working with more than one participant simultaneously would have been overwhelming if not impossible. However, now that we have established ViVA's benefits, especially the potential of personalized access points, the grounds have been set for further investigation involving more potential users. It would also be beneficial to conduct an evaluation with a high-fidelity prototype that can be left with the participant over an extensive period of time (in the range of a few months).

Working closely with Sam made us realize that we were optimistic to think that by speeding up word-finding, we can ease him into a state where he uses the assistive vocabulary as the primary means for communication during a conversation. Even with ViVA running in front of him, he preferred writing down whatever he could and waiting for his conversation partner to guess what he meant. Part of this preference is probably due to the fact that Sam is accustomed to this compensatory communication strategy. In addition, writing on paper made the conversation more interactive, because it was easier for the conversation partner to observe what was being written than to follow Sam's browsing the vocabulary. This observation does not diminish the value of such a tool, but it redefines our expectations for its purpose and use. ViVA would be most useful for constructing sentences that will be used in future communication, similar to what Sam had done with his Lingraphica and SmallTalk. Our approach to vocabulary organization could also be extended to enhance and assist information organization in utterance-based AAC systems which offer a selection of prestored messages (e.g., the system proposed by McCoy et al. 2010). Improved word-finding will ensure more independence for the user when composing the messages and can encourage him/her to be more expressive, because it is easier to find cross-part-of-speech associations in ViVA and thus link nouns, adjectives and verbs. ViVA can also provide faster retrieval of semantically related messages and suggest words for adapting the message content when the communication is happening.

At the end of the study Sam indicated his satisfaction with ViVA and requested a copy of the application.

Sam: “That stuff, me, that stuff good for me this way. You that stuff, me, you, that stuff very good for me. What is this program?”

Researcher: “You like it?”

Sam: “Yes! That stuff, you, me.”

Researcher: “You want it?”

Sam: “Yes!”

Since ViVA’s current version is not stable enough to be deployed on Sam’s computer, we are sharing our research results with Lingraphicare Inc. who, as a start, is interested in incorporating the evocation data into their vocabulary. Thus, Sam’s Lingraphica (and of all others) could have the semantic associations (personal and general) in the future.

Chapter 8

Conclusion

The work presented in this dissertation addressed the currently cumbersome and ineffective word-finding in Augmentative and Alternative Communication tools. The majority of assistive vocabularies provide extensive word collections that are organized in deep hierarchies or common categories that often cannot meet the needs of individual users. We argue that to enable effective word-finding, an assistive vocabulary needs to adapt to individual user's word usage patterns and to the semantic associations present in a speaker's mental lexicon, where words are stored and organized in ways that allow efficient access and retrieval.

In support of our thesis, we built the Visual Vocabulary for Aphasia (ViVA). ViVA improves word-finding by modeling a user's mental lexicon. ViVA's design relies on psycholinguistic theories that propose a semantic network structure of the lexicon and spreading activation as supported by semantic priming. ViVA provides dynamic networks centred around words in the vocabulary where links between terms reflect rich relationships based on human judgments of semantic relatedness, vocabulary usage statistics and semantic associations based on the large-scale lexical database WordNet (Fellbaum 1998, Miller 1990).

Through two controlled experiments, we showed that ViVA helps able and aphasic users find words faster. We demonstrated its ability to improve word-finding in assistive tools as well as its potential to improve communication tools for broader users groups. Participants found ViVA less confusing to navigate than a commercial assistive vocabulary and agreed that the associations that it provided helped them find words faster.

By running a six-week single-case study with an aphasic participant, we also investigated which approaches to customization can improve word-finding in ViVA even further. The idea of introducing personalized access points to the vocabulary evolved as a result of our work during the case study. Conducting an evaluation of ViVA with the personalized access points showed that word-finding is improved when the user is guided by both familiar on non-familiar context.

We conclude with a summary of the main contributions of this work and outline promising directions for future work.

8.1 Thesis Contributions

The high-level goal of this dissertation is to improve word-finding in assistive communication vocabularies. As means to achieving this goal, we addressed the following research questions:

1. How can we improve word-finding in assistive communication vocabularies?
    a. We showed that we can improve word-finding by incorporating semantic word associations retrieved from existing large-scale lexical databases such as WordNet.

   b. We also showed that we can improve word-finding by incorporating human judgments of semantic relatedness that we collected through a large-scale online experiment.

2. How can we tailor word-finding to a specific user?
   a. We showed that incorporating word frequency usage statistics in the structure of the vocabulary shortens the browsing paths needed to connect words by a specific user.
   b. We were also able to improve word-finding by predicting semantic associations based on word-usage and semantic-word-associations data.
3. How does a static hierarchical vocabulary organization compare to a dynamic (adaptive) organization?
   a. We presented empirical evidence that able and aphasic users can find words faster in an adaptive organization with provides shorter browsing paths between semantically related concepts.
   b. We also showed that users find the adaptive organization we implemented in ViVA better and more intuitive to navigate in comparison to a static organization comprised of common categories such as *food* and *clothes*.

In summary, this work introduced a novel approach to vocabulary organization in assistive communication tools that enables effective word-finding. We compiled a large dataset of human judgments of semantic relatedness that can be used to model a user's mental lexicon. We also presented empirical evidence that enhancing a basic vocabulary hierarchy with semantic word-associations improves word-finding compared to existing alternatives.

8.2 Directions for Future Work

We outline directions for future work that would build on the contributions of this dissertation as well as for research that may be of broader interest.

8.2.1 Exploring Additional Resources of Semantic Associations

To model the organization of a user's mental lexicon, we built ViVA such that it provides networks of words populated with meaningful relationships between words. Our initial design of ViVA was guided by WordNet, a large-scale lexical database, which models the mental lexicon as a semantic network that connects related words to one another. Incorporating only WordNet associations in ViVA was insufficient because noun–verb and noun–adjective connections are sparse; combining words from different parts of speech is integral to functional communication. To address this problem, we collected, from thousands of English speakers, scores of the strength with which a given word evokes another, (e.g. *eat–hungry*) (Nikolova et al. 2009).

As a next step in modelling a user's mental lexicon, it would be worth exploiting resources that capture the way people reason about everyday life and describe it in words. Commonsense knowledge, essential to communication, is often left unstated. For example, a *book* can be *read* and *fire* can *burn* things. It would be particularly interesting to mine ConceptNet (Havasi et al. 2007), a repository of such world knowledge gathered from human contributors over the Web, for generalizable semantic information. This task is challenging, because culture-specific associations and individual idiosyncrasies need to be taken into consideration, e.g. *tulips–Holland*, *turkey–Thanksgiving*, and *grandmother–cookies*. Nonetheless, extracting semantic associations encoded in commonsense data will enable us to design intelligent assistive tools that support functional communication by guiding the user and helping him/her find concepts effectively to create thousands of new links that go beyond the relations specified in WordNet.

8.2.2 Semantic Associations, Technology and Language Rehabilitation

As discussed in the related work chapter (Chapter 2), semantic feature analysis is a recognized treatment in the field of aphasiology. Both researchers and speech therapists have recognized the value in methodical treatment of impaired semantic networks surrounding a concept to strengthen any damaged links and thus improve word retrieval.

ViVA has the potential to assist therapist in such treatments by providing rich semantic networks that model the mental lexicon. Introducing ViVA in the initial stages of rehabilitation and using it for treatment early on can help the user become comfortable with the tool and also provide an opportunity for gradual adaptation of the vocabulary. That way the vocabulary can start off with a small collection of words and limited semantic networks and grow as the patient improves with treatment.

An early introduction of the assistive vocabulary by incorporating it closely into initial treatment will also very likely enable better adoption. Customizing the vocabulary to fit the user's profile can help easier transition into using it independently for both rehabilitation and communication once therapy is no longer available (usually medical insurance covers only the first six months of therapy after which most patients are left to adapt to their inability to communicate on their own). These ideas should be developed further only with close collaboration with speech therapists and speech pathology researchers.

8.2.3 Context-Aware AAC Tools

Context-based adaptation of content holds much potential due to the pervasiveness of sophisticated mobile devices such as smart phones. Smart phones enable individuals to share information about their location, experiences, and activities as they occur. This phenomenon has improved our understanding of how people interact with their environment and how to use context to improve the interaction.

There are a number of new commercial mobile tools, for example, that study shoppers' behaviour to provide the location of desired products in a store and real-time coupon alerts (Rosenbloom 2010). Context-based adaptation for language and communication support is another promising field. A smart mobile AAC tool could suggest words and phrases needed for ordering food, (e.g. *vegan option*, *peanut allergies*, *medium rare steak*) when the user is in a restaurant, for example. An adaptive vocabulary such as ViVA could tailor the suggestions to the user's profile; the user should also be able to explore associated concepts to help him/her compose relevant communications faster. Smart phones present an additional advantage in that they enable capturing context through images (as well as through video or audio recordings) which can help a person be more expressive and fluent in recounting experiences at a later time (Boyd-Graber et al. 2006, Sellen et al. 2007). The tool could guide the user in storing the captured material within the appropriate context and associating it with concepts that would facilitate efficient retrieval when the story is recounted.

8.2.4 Users' Response to Adaptation

Adapting to usage patterns and user preferences has been shown to lessen the burden on users when interacting with complex user interfaces (for example, in improving menu navigation (Findlater et al. 2009). However, the benefits of adaptive interfaces are still largely unknown. Proponents (e.g., Benyon 1993 and Gajos 2008) argue that adaptation offers the potential to optimize the human-computer interaction. Critics (e.g., Findlater & McGrenere 2004), on the other hand, caution that the unpredictability resulting from adaptation may offset the benefits.

We have shown that adapting the organization of an assistive vocabulary based on usage frequency and semantic associations can greatly facilitate word-finding. More broadly, we believe that adaptive functionality can significantly improve usability, user satisfaction, and adoption of user interfaces, especially in the context of empowering people beyond their capabilities. Some efforts have already been dedicated to adaptive user interfaces for web-accessibility (e.g., Fink et al. 1998) and for assisting users with motor disabilities (e.g., Gajos 2008), bur more basic research is still necessary for the advancement of this field. For example, we find it would be beneficial to explore how different degrees of and approaches to adaptation affect users with a range of cognitive profiles, for example users in different age groups or of different gender. Based on participants' feedback and our observations in the study we described in Chapter 6, we could detect certain profiles of search attitudes emerging. Some participants explored the vocabulary by clicking on icons without any evident plan of action. Others memorized the organization while browsing to help them find words faster in subsequent tasks. A few participants were guided by the associations and based their next move solely on what was displayed on the current screen. Some participants formed associations and expectations in advance and let their intuition about where a word should be found guide them. Studying these different approaches and user profiles further will enable us to understand if they can be generalized as well as how specific to or influenced by cognitive disabilities they are.

Finally, we implemented a very basic adaptive algorithm and showed the potential of adapting the vocabulary structure using predictions of semantic words associations in improving word-finding. Due to the simplicity of the task we gave our participants, it we did not delve into analysis of the effect of accuracy (how well the algorithm predicts the associations) and predictability (how well can participants model the adaptive algorithm's behavior), two terms discussed by (Gajos et al. 2008). These two factors can influence people with different cognitive abilities differently so further investigation is necessary.

8.2.5 Multi-Modal Associations

In our work, we concentrated on improving word-finding in visual vocabularies of AAC tools by abstracting from the words' non-textual representations. This was possible, because we were able to utilize Lingraphica's multi-modal icons which, despite their drawbacks, comprise a well-designed set of visual representations combined with speech audio and text.

Multi-modal representations of concepts that consist of text, pictures, videos, speech audio and other sounds are an integral part of AAC tools, because they can evoke words that help the user communicate. However, it is challenging to illustrate abstract concepts and categories, e.g., *actions* or *feelings*. Users, especially the ones that rely primarily on the visual representation to navigate the vocabulary, can find such icons difficult to interpret and be easily confused. Thus, overlaying ViVA's semantic networks with a multi-modal collection better than Lingraphica's will ensure more powerful assistive communication. There are ongoing efforts to provide such and improved collection, but the area is still under-investigated. Examples include ImageNet (Deng et al. 2009) which maps high-quality evocative images to WordNet synsets and SoundNet (Ma et al. 2010) which is a non-speech audio-to-concept semantic network. Similar to providing an effective vocabulary organization for all users, providing multi-modal representations that meet the needs of each individual will require a flexible solution.

8.2.6 Broader Applications of ViVA

We believe that in addition to the contributions we summarized above, the lessons learned during the design and evaluation of ViVA can have some broader applications. German (German 1986) defines six different groups of students who may experience word-finding difficulties:

- Students who have specific learning disabilities;
- Students who have reading difficulties;
- Students who have specific language difficulties;
- Students who have fluency difficulties;
- Students who have known brain pathology;
- Students who have attention difficulties and/or are hyperactive.

These six groups of potential users of assistive communication tools are among 3.5 million Americans who find it difficult to use speech to communicate (Beukelman & Mirenda 2006). Thus, in addition to the people who have aphasia, many others can benefit from improved AAC technologies. Our approach to improving word-finding by building customized semantic networks can be applied to the design of communication tools for users with other language impairments as well as to the design of educational tools for children and foreign-language learners.

Bibliography

Adams, H., Davis, P., Torner, J., Grimsman, K., & Vande Berg, J. 2006. Cookie-theft: Image for extinction and inattention item. http://www.uihealthcare.com/topics/medicaldepartments/neurology/stroke/cookie.html. Available Online May 2010.

Adler Aphasia Center. http://www.adleraphasiacenter.org/. Available Online June 2010.

Ageless Project. http://jenett.org/ageless/. Available Online May 2010, Copyright 2001–2010.

Allen, M., McGrenere, J., & Purves, B. 2008. The field evaluation of a mobile digital image communication application designed for people with aphasia. In *ACM Trans. Access. Comput., 1: 1*, 1–26.

Amazon.com, Inc. Amazon Mechanical Turk. https://www.mturk.com. Available Online November 2009. Copyright 2005–2010.

AphasiaBank. AphasiaBank Consortium. http://talkbank.org/AphasiaBank/. Available Online May 2010.

Apple Inc. http://www.apple.com/ipodtouch/. Available Online June 2010, Copyright 2010.

AssitiveWare. Proloquo2Go. http://www.assistiveware.com/. Available Online February 2010, Copyright 2010.

Beeson P.M., & Robey R. R. 2006. Evaluating single-subject treatment research: Lessons learned from the aphasia literature. In *Neuropsychol Rev., 16*, 161–169.

Benyon, D. 1993. Adaptive systems: A solution to usability problems. In *User Modeling and User-Adapted Interaction*, *3: 1*, 65–87.

Beukelman, D. R., & Mirenda, P. (Eds.). 2006. Augmentative & alternative communication: Supporting children & adults with complex communication needs, 3rd ed., Baltimore, MD: Paul H Brooks Publishing Co.

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. In *Proc. CHI '06: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 151–160.

Boyd-Graber, J., Fellbaum, C., Osherson, D., & Schapire, R. 2006. Adding dense, weighted connections to WordNet. In *Proc. Global WordNet Conference'06*, 29–35.

Boyle, M., & Coelho, C. A. 1995. Application of semantic feature analysis as a treatment for aphasic dysnomia. In *American Journal of Speech Language Pathology, 4*, 94–98.

Boyle, M. 2004. Semantic feature analysis treatment for anomia in two fluent aphasia syndromes. In *American Journal of Speech-Language Pathology*, *13*, 236–249.

Carmien, S. 2002. MAPS: PDA scaffolding for independence for persons with cognitive impairments. In *Human-computer Interaction Consortium 2002*.

Cherney, L., & Robey, R. 2002. Aphasia treatment: Recovery, prognosis, and clinical effectiveness. In *R. Chapey (Ed.), Language intervention strategies in adult aphasia*, *4th ed.*, Baltimore, MD: Lippincott Williams, and Wilkins, 148–172.

Cohene, T., Baecker, R., & Marziali, E. 2005. Designing interactive life story multimedia for a family affected by Alzheimer's disease: A case study. In *Proc. CHI '05: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1300–1303.

Collins, A. M., & Loftus, E. F. 1975. A spreading-activation theory of semantic processing. In *Psychological Review*, *82: 6*, 407–428.

Daeman, E., Dadlani, P., Du, J., Li, Y., Erik-Paker, P., Martens, J., & Ruyter, B.D. 2007. Designing a free style, indirect, and interactive storytelling application for people with aphasia. In *INTERACT'07*, 221–234.

Davis, G. A. 2000. Aphasiology: Disorders and clinical practice. Boston, MA: Allyn & Bacon.

Davies, R. 2004. The ethnographically informed participatory design of a PDA application to support communication. Master's Thesis, University of British Columbia, Canada.

Davies, R., Marcella, S., McGrenere, J., & Purves, B. 2004. The ethnographically informed participatory design of a PDA application to support communication. In

Proc ASSETS'04: Proceedings of the ACM SIGACCESS Conference on Computers and Accessibility, 153–160.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR09*.

Duffy, J. R. 1995. Motor speech disorders: Substrates, differential diagnosis, and management. In *Defining, Understanding, and Categorizing Motor Speech Disorders, 1st ed.,* St Louis, MO: Mosby, 3–13.

Dynavox Mayer-Johnson. Dynavox. http://www.dynavoxtech.com/. Available Online February 2010, Copyright 2010.

Edmonds, L. A., Nadeau, S. E., & Kiran, S. 2009. Effect of verb network strengthening treatment (VNeST) on lexical retrieval of content words in sentences in persons with aphasia. In *Aphasiology, 23: 3*, 402–424.

Eisenhardt, K. M. 1989. Building theories from case study research. In The Academy of Management Review, 14: 4, 532–550.

Fellbaum, C. 1998. WordNet: An electronic lexical database, chapter: A semantic network of English verbs. MIT Press, Cambridge, MA.

Findlater, L. 2009. Supporting feature awareness and improving performance with personalized graphical user interfaces. Doctoral Thesis, University of British Columbia, Canada.

Findlater, L., & McGrenere, J. 2004. A comparison of static, adaptive, and adaptable menus. In *CHI'04: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 89–96.

Findlater, L., Moffatt, K., McGrenere, J., & Dawson, J. 2009. Ephemeral adaptation: The use of gradual onset to improve menu selection performance. In *Proc. CHI'09: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1655–1664.

Fink, J., Kobsa, A., & Nill, A. 1998. Adaptable and adaptive information provision for all users, including disabled and elderly people. In *New Review of Hypermedia and Multimedia, 4*, 163–188.

Flyvbjerg, B. 2006. Five misunderstandings about case-study research. In *Qualitative Inquiry, 12: 2*, 219–245.

Gajos, K. 2008. Automatically generating personalized user interfaces. Doctoral Thesis. University of Washington, USA.

Gajos, K., Everitt, K., Tan, D., Czerwinski, M., & Weld, D. 2008. Predictability and accuracy in adaptive user interfaces. In *CHI '08: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1271–1274.

Gardner, H., Zurif, B., Berry, T., & Baker, E. 1976. Visual communication in aphasia. *Neuropsychologia, 14: 3*, 275–292.

German, D. 1986. National College of Education Test of Word Finding. Austin TX: PRO-ED.

Goodglass, H. 1993. Understanding aphasia. In *Foundations of Neuropsychology* series. San Diego, CA: Academic Press.

Goodglass H., & Wingfield A. 1997. Anomia: Neuroanatomical and cognitive correlates. San Diego, CA: Academic Press.

Google Inc. Google Images. http://images.google.com/. Available Online May 2010, Copyright 2010.

Greenberg, S., & Witten, I. H. 1985. Adaptive personalized interfaces: A question of viability. In *Behaviour & Information Technology, 4: 1,* 31–45.

Haigh, K., Kiff, L., & Ho, G. 2006. The independent lifestyle assistant (i.l.s.a.): Lessons learned. In *Assistive Technology, 18*, 87–106.

Havasi, C., Speer, R., & Alonso, J. 2007. ConceptNet 3: a flexible, multilingual semantic network for common sense knowledge. In *Proc. Recent Advances in Natural Language Processing*.

Hine, N., Arnott, J., & Smith, D. 2003. Design issues encountered in the development of a mobile multimedia augmentative communication service. In *Universal Access in the Information Society, 2: 3*, 255–264.

Hirsch, T., Forlizzi, J., Hyder, E., Goetz, J., Kurtz, C., & Stroback, J. 2000. The ELDer project: Social, emotional, and environmental factors in the design of eldercare technologies. In *Proceedings of the ACM Conference on Universal Usability*, 72–79.

Horvitz, E. 1999. Principles of mixed-initiative user interfaces. In Proc. CHI'99: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 159–166.

Jiang, J.J., & Conrath, D.W. 1997. Semantic similarity based on corpus statistics and lexical taxonomy. In *Proc. International Conference on Research in Computational Linguistics*.

Kauhanen, M.L., Korpelainen, J. T., Hiltunen, P., Määttä, R., Mononen, H., Brusin, E., et al. 2000. Aphasia, depression, and non-verbal cognitive impairment in ischaemic stroke. In *Cerebrovascular Diseases, 10: 6*, 455–461.

Kautz, H., L. Arnstein, L., Borriello, G., Etzioni, O., & Fox, D. 2002. An overview of the assisted cognition project. In *AAAI-02 Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care*, 60–65.

Lau, C., & O'Leary, S. 1993. Comparison of computer interface devices for persons with severe physical disabilities. In *American Journal of Occupational Therapy, 47*, 1022–1030.

Levinson, R.1997. PEAT: The planning and execution assistant and trainer. In *Journal of Head Trauma Rehabilitation, 12: 2*, 769–775.

Lin, D. 1998. Automatic retrieval and clustering of similar words. In *Proc. of COLING-ACL*, 768–774.

Lingraphicare Inc. Lingraphica. http://www.aphasia.com/. Available Online February 2010, Copyright 2009.

Ma, X., Fellbaum, C., & Cook, P. R. 2010. SoundNet: Investigating a language composed of environmental sounds. In *Proc. CHI'10: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1945–1954.

Mantei, M. M., & Cattell, R. G. 1982. A study of an entity-based database user interface. In *SIGCHI Bull. 14, 1*, 5–16.

Martin, C., Dellatolas, G., Viguier, D., Willadino-Braga, L., & Deloche, G. 2002. Subjective experience after stroke. In *Applied Neuropsychology, 9: 3*, 148–158.

McCoy, K. F., Bedrosian, J. L., & Hoag, L. A. 2010. Implications of pragmatic and cognitive theories on the design of utterance-based AAC systems. In *Proceedings of the NAACL HLT 2010 Workshop on Speech and Language Processing for Assistive Technologies*, 19–27.

Miller, G. A. 1990. Nouns in WordNet: A lexical inheritance system. In *Int. Journal of Lexicography, 3: 4*, 245–264.

Moffatt, K. 2004. Designing technology for and with special populations: An exploration of participatory design with people with aphasia. Master's Thesis, University of British Columbia, Canada.

Moffatt, K., McGrenere, J., Purves, B., & Klawe, M. 2004. The participatory design of a sound and image enhanced daily planner for people with aphasia. In *Proc. CHI'04: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 407–414.

Moss, H., & Older, L. 1996. Birkbeck Word Association Norms. In *Psychology Press*.

National Aphasia Association. Aphasia FAQ. http://www.aphasia.org. Available Online February 2010, Copyright 2009.

Natural Language Toolkit. http://www.nltk.org/. Available Online May 2010.

Nikolova, S., Boyd-Graber, J., & Fellbaum, C. (in press). Collecting semantic similarity ratings to connect concepts in assistive communication tools. In A. Storrer et al., Eds. Modeling, Learning and Processing of Text Technological Data Structures. Studies in Computational Intelligence. Heidelberg/New York: Springer.

Nikolova, S., Boyd-Graber, J. Fellbaum, C., & Cook, P. 2009. Better vocabularies for assistive communication aids: Connecting terms using semantic networks and untrained annotators. In *Proc. ASSETS'09: Proceedings of the ACM SIGACCESS Conference on Computers and Accessibility*, 171–178.

Nikolova, S., Ma, X., Tremaine, M., & Cook, P. 2010. Vocabulary navigation made easier. In *Proc. ACM IUI'10: Proceedings of the 14th International Conference on Intelligent User Interfaces*, 361–364.

Noor, K.B.M. 2008. Case study: A strategic research methodology. In *American Journal of Applied Sciences, 5: 11*, 1602–1604.

Norcio, A., & Stanley, J. 1989. Adaptive human-computer interfaces: A literature survey and perspective. In *IEEE Transactions on Systems, Man, and Cybernetics, number, 19: 2*, 399–408.

Patel, R. Pilato, S., & Roy. D. 2004. Beyond linear syntax: An image-orientation communication aid. *Assistive Technology Outcomes and Benefits, 1: 1*, 57–67.

Pollack, M. E., Brown, L., Colbry, D., McCarthy, C. E., Orosz, C., Peintner, B., Ramakrishnan, S., & Tsamardinos, I. 2003. Autominder: An intelligent cognitive orthotic system for people with memory impairment. In *Robotics and Autonomous Systems*, *44*, 273–282.

Rosenbloom, S. 2010. Cellphones let shoppers point, click and purchase. The New York Times. Available February 26, 2010. http://www.nytimes.com/2010/02/27/business/27shop.html?hp.

Schapire, R.E. 2003. The boosting approach to machine learning: An overview. In Denison, D.D., Hansen, M.H., Holmes, C., Mallick, B., Yu, B., eds.: Nonlinear Estimation and Classification. Springer.

Sellen, A., Fogg, A., Hodges, S., & Wood, K. 2007. Do life-logging technologies support memory for the past? An experimental study using SenseCam. In *Proc. CHI '07*: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 81–90.

SmallTalk. Lingraphicare Inc. http://www.aphasia.com/slp/product_smalltalk.aspx. Available Online June 2010, Copyright 2009.

Snow, R., O'Connor, Jurafsky, D., & Ng, A. 2008. Cheap and fast—but is it good? Evaluating non-expert annotations for natural language tasks. In *Proc. EMNLP-08*.

Song, Y. 2009. Designing the ViVA interface: A multi-modal sentence construction tool for aphasia. Undergraduate Independent Work Final Report, Princeton University, USA.

Spradley, J. P. 1979. The ethnographic interview. Wadsworth Group/Thomson Learning.

Steele, R.D., Weinrich, M., Wertz, R. T., Kleczewska, M.K., & Carlson, G.S. 1989. Computer-based visual communication in aphasia. In *Neuropsychologia, 27: 4*, 409–426.

Stephanidis, C., Paramythis, A., Sfyrakis, M., Stergiou, A., Maou, N., Leventis, A., Paparoulis, G., & Karagiannidis, C. 1998. Adaptable and adaptive user interfaces for disabled users in the AVANTI project. In *The 5th International Conference on Intelligence and Services in Networks: Technology For Ubiquitous Telecom Services*, 153–166.

Suchman, L.A. 1987. Plans and situated actions: The problem of human-machine communication. 2nd ed. Cambridge, UK: Cambridge University Press.

Sutcliffe, A., Fickas, S., Sohlberg, M. M., & Ehlhardt, L. A. 2003. Investigating the usability of assistive user interfaces. In *Interacting with Computers*, *15*, 577–602.

Swinney, D. 1979. Lexical access during sentence comprehension: (Re)consideration of context effects. In *Journal of Verbal Learning and Verbal Behavior*, *18*, 645–659.

SWITCHBOARD: A User's Manual, 2007. Accessed January 2010 from http://www.ldc.upenn.edu/Catalog/readme_files/switchboard.readme.html/

Tee, K., Moffatt, K., Findlater, L., MacGregor, E., McGrenere, J., Purves, B., et al. 2005. A visual recipe book for persons with language impairments. In *Proc. CHI ’05: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 501–510.

The British National Corpus, version 3 (BNC XML Edition). 2007. Distributed by Oxford University Computing Services on behalf of the BNC Consortium. URL: http://www.natcorp.ox.ac.uk/

Thorburn, L., Newhoff, M., & Rubin, S. 1995. Ability of subjects with aphasia to visually analyze written language, pantomime, and iconographic symbols. In *American Journal of Speech-Language Pathology, 4*, 174–179.

TouchSpeak & TypeSpeak. http://www.touchspeak.co.uk/. Available Online February 2010, Copyright 2009.

Trnka, K., & McCoy, K. F. 2007. Corpus studies in word prediction. In Proc. ASSETS’07: Proceedings of the ACM SIGACCESS Conference on Computers and Accessibility, 195–202.

Trnka, K., McCaw, J., Yarrington, D., McCoy, K., & Pennington, C. 2009. User interaction with word prediction: The effects of prediction quality. In *ACM Trans. Access. Comput.*, *1: 3*, 1–34.

Video 1. Available May 2010 at http://www.youtube.com/watch?v=zbBu9C_zjaE.

Video 2. Copyright Annenberg Media 1997–2009. Available May 2010 at http://www.learner.org/vod/vod_window.html?pid=758.

Video 3. Available May 2010 at http://www.youtube.com/watch?v=IAWvsfNic6k.

van de Sandt-Koenderman, M. 2004. High-tech AAC and aphasia: Widening horizons? In *Aphasiology, 18: 3*, 245–263.

van de Sandt-Koenderman, M., Wiegers, J., & Hardy, P. 2005. A computerised communication aid for people with aphasia. *Disability and Rehabilitation, 27: 9*, 529–533.

van de Sandt-Koenderman, M., Wiegers, J., Wielaert, M., Duivenvoorden, J., & Ribbers, M. 2007. High-tech AAC and severe aphasia: Candidacy for TouchSpeak (TS). In *Aphasiology*, *21: 5*, 459–474.

Waller, A., Denis, F., Brodie, J., & Cairns, A. Y. 1998. Evaluating the use of TalksBac, a predictive communication device for nonfluent adults with aphasia. In *International Journal of Language and Communication Disorders, 33: 1*, 45–70.

Waller, A. & Newell, A. 1997. Towards a narrative based augmentative communication system. In *European Journal of Disorders of Communication, 32*, 289–306.

Webster J., Morris J., & Franklin S. 2005. Effects of therapy targeted at verb retrieval and the realisation of the predicate argument structure: A case study. In *Aphasiology 19: 8*, 748–764.

WordNet Resrearch @ Princeton. 2010. http://wordnet.cs.princeton.edu/downloads.html. Available Online February 2010.

YouTube. YourTube, LLC. http://www.youtube.com/. Available Online May 2010, Copyright 2010.

Appendix A

Stimulus Phrases Used in Study with Non-Aphasic Participants

The phrases were constructed with frequently used words. The optimal path to a word in the hierarchical organization was balanced across the vocabulary conditions, LG and ViVA (see Chapter 5 for more details).

LG	**ViVA**
eat and ***drink***	***eat*** and ***drink***
the ***elephant*** is a big ***animal***	the ***elephant*** is a big ***animal***
this ***cigarette*** is ***poison*** for you	this ***cigarette*** is ***poison*** for you
I went to a ***football game***	I went to a ***football game***
sleep is ***health***	***sleep*** is ***health***
the ***university team*** won	the ***university class*** is difficult
van is smaller than ***bus***	the ***van*** has a powerful ***engine***
steak is my favourite ***food***	you need good ***meat*** for ***steak***
the hotel ***room*** has yellow ***wall***	this ***room*** will be the ***kitchen***
Milk and honey for tea	***tea*** or ***coffee*** with the ***dessert***
The ***house kitchen*** is small for ***family***	***kitchen*** and ***bathroom*** have broken ***sink***
I like ***pudding***, but when ***hungry***, I prefer other ***food***	***pudding***, not ***tea*** for ***dessert***
put the ***knife*** by the ***plate***	he ***explained*** to the ***student***
my ***aunt*** and ***grandfather*** visited with the ***baby***	my ***aunt*** and ***uncle*** visited their ***daughter***
time for ***coffee*** and ***dessert***	too much ***meat*** for ***breakfast***
chew the ***ham***	the ***chair*** has ***cushion***
you need ***shoes*** and ***umbrella***	I had ***cereal*** for ***breakfast***
no ***sugar*** in the ***jelly***	I ***smell*** the ***food***
increase the ***speed***	this was ***tasty breakfast***

I ***drive*** a ***car*** and ***fast motorcycle***	I ***drive*** to ***travel***, I don't like ***buses*** and ***trains***
tea and ***coffee*** for ***breakfast***	***tea*** with ***sugar*** and ***honey***
my ***sister*** and the ***baby sleep***	my ***sister*** thinks ***grandfather*** is ***lonely***
the ***bacon*** is for ***lunch***	the ***bacon*** is ***tasty***
I like most ***vegetables*** but ***broccoli***	I ***cook*** and ***bake***
I will ***eat*** the ***pudding*** later	I will ***eat*** when I am ***hungry***
call him to ***thank*** him	***call*** or ***write*** soon
clean the ***rice***, ***boil*** and ***salt*** it	I like ***brown rice*** with ***meat***
the ***plants*** and ***trees*** are green	I saw many ***plants*** and ***animals***
take a ***pill*** when ***dizzy***	the ***bread*** is ***tasty*** with ***butter***
bread and ***cheese*** for ***breakfast***	***swallow*** the ***pill***
I ***drive*** the ***van***, not the ***bus***	***bread*** and ***sausage*** for ***breakfast***
eat the ***dessert***	I ***drive*** a ***motorcycle***
I ***smell*** the ***cigarette*** and don't ***approve***	I ***smell garlic*** in the ***food***
I drink ***beer*** when ***thirsty***	I drink ***tea*** and ***coffee*** with ***milk***
bake in ***hot oven***	***eat*** when ***hungry***
the ***chocolate*** is ***tasty***	***mow*** the ***grass***
tea comes after ***lunch***	no ***beer*** before ***dinner***
my ***sister*** and ***father travel***	***cook bacon*** for ***lunch***
we had ***rice*** for ***dinner***	She likes ***rice*** and ***eggs***

Appendix B

Stimulus Scenarios Used in Study with Aphasic Individuals

The image scenarios consisted of an image, a phrase related to the image context, a missing word in the phrase and three choices for the missing word. In the experiment, the participant was asked to look at the image, the phrase was read out loud and then s/he had to choose the missing word. The participant was then left with the image shown on a blank page to remind him/her of the context and had to find the chosen word in the provided digital vocabularies, ViVA and LG (see Chapter 6 for more details).

The women ? by  bus

A: travel

B: drive

C: eat

LG Image Scenarios

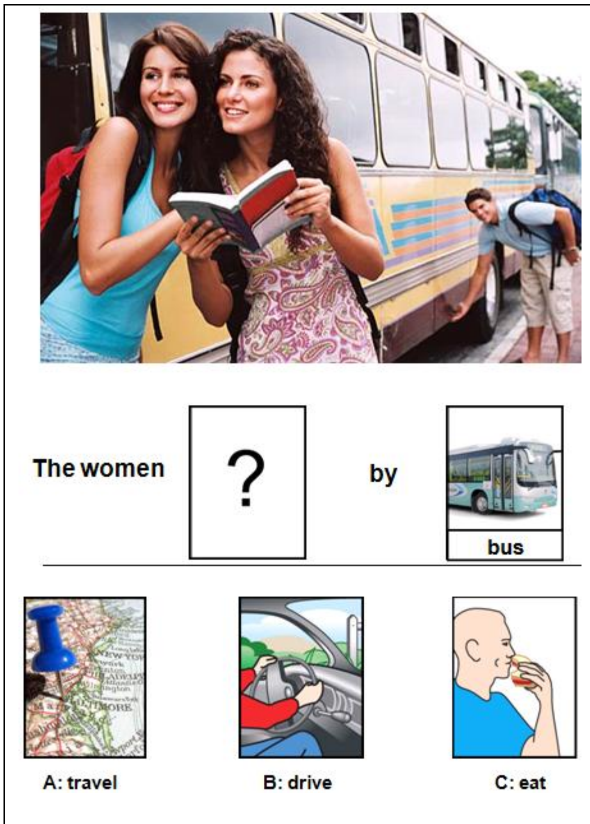

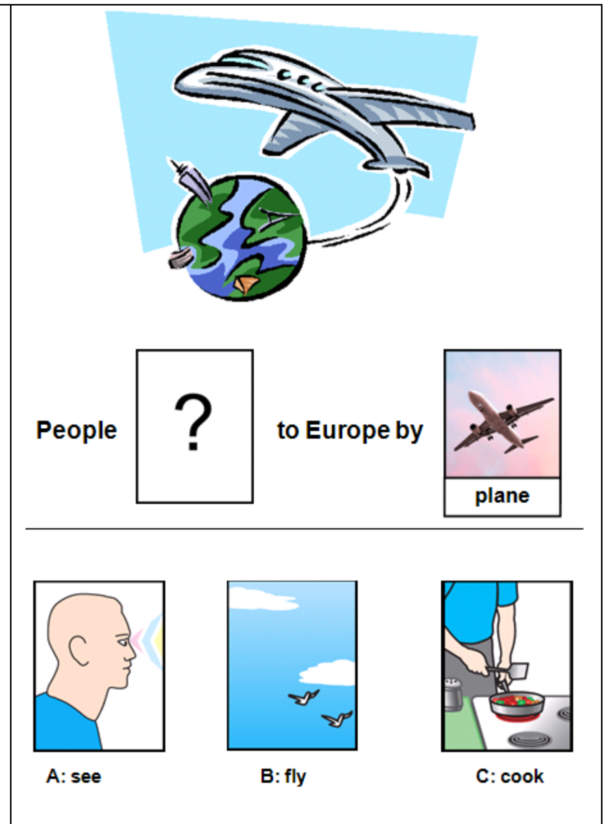

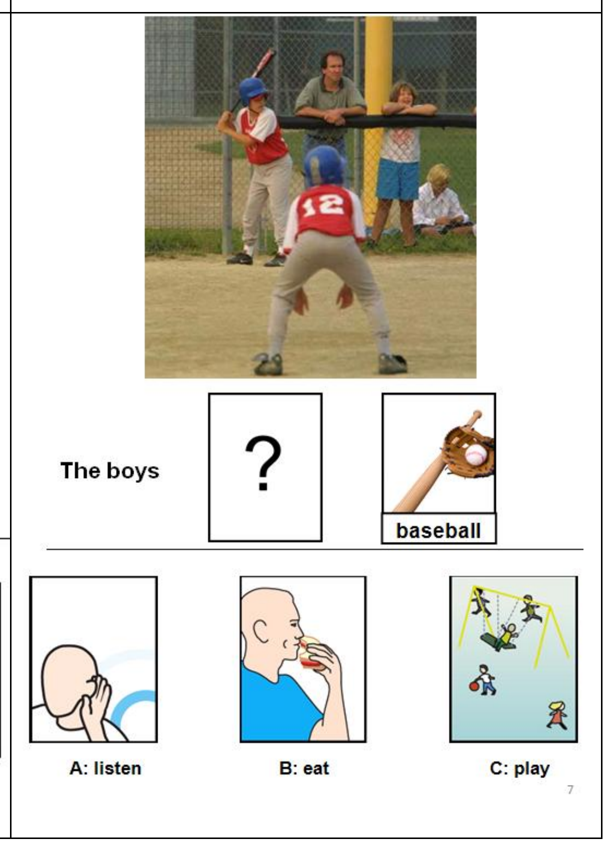

The boy ? to

music

A: listens

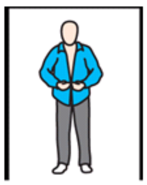

B: dresses

C: sings

People ?

dinner

in a

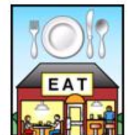

restaurant

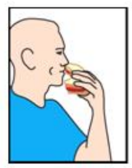

A: eat

B: pay

C: sleep

They ? in the

mall

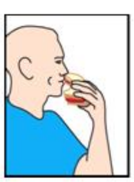

A: eat

B: shop

C: swim

ViVA Image Scenarios

The boy ? the  flower


A: buys


B: smells


C: cooks

The  baby ?

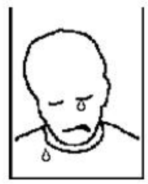
A: cries


B: sleeps

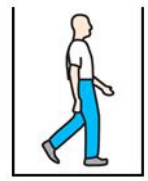
C: walks

The girl ? a  letter


A: opens

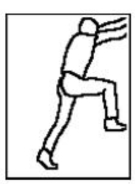
B: jumps


C: writes

Appendix C

Aphasia Friendly Consent Form Used in Word-Finding Study with People with Aphasia

In the experiment described in Chapter 6, participants with aphasia were introduced to the task and the procedure using the aphasia-friendly consent form shown below.

Informed Consent for Research

Project Title: Visual Vocabulary Application:
Word-Finding Study

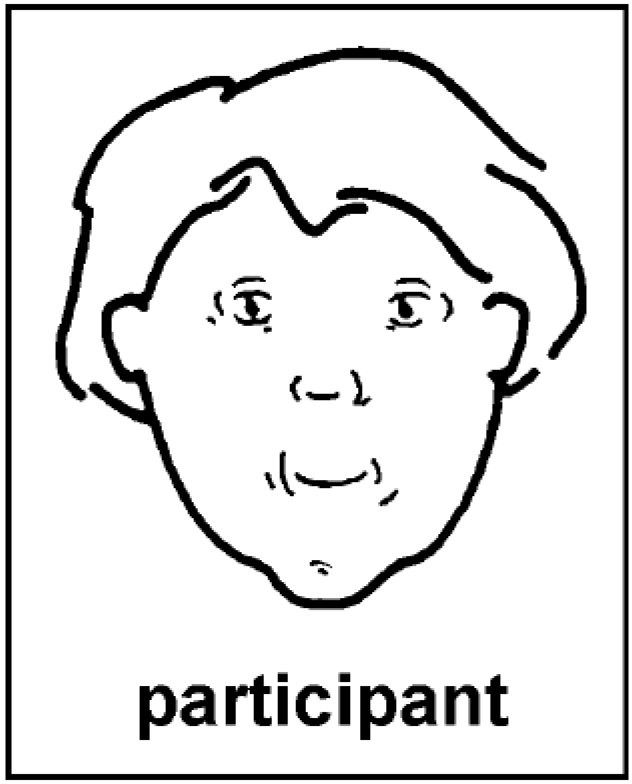

Participant: ___________________________

Researcher: Sonya Nikolova
nikolova@princeton.edu
(609)-933-9084

Advisor: Professor Perry Cook
prc@cs.princeton.edu

Example Task

1. Researcher will show you a picture.

2. Researcher will read a sentence out loud.

They | **?** | **by** |  **car**

3. Researcher will ask you to guess the missing word:

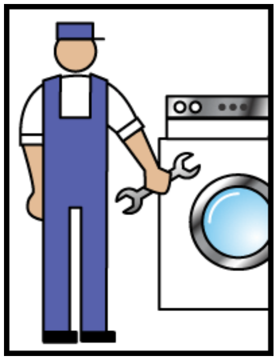

A: repair

B: travel

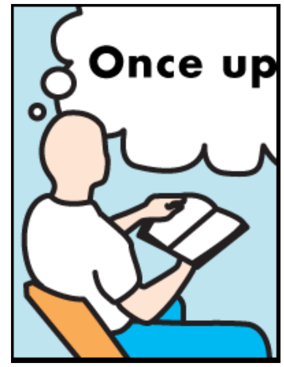

C: read

4. Researcher will explain how to use the computer

5. Researcher will ask you to find the word on the computer:

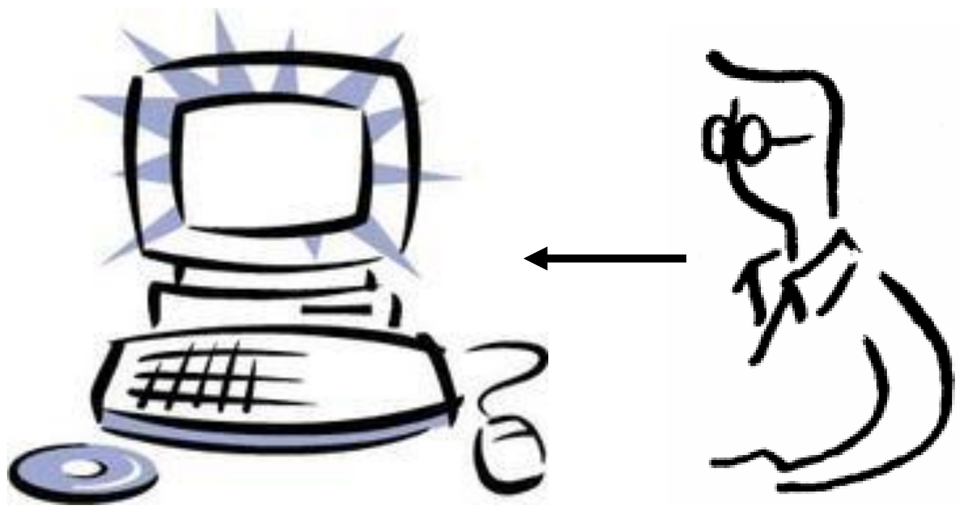

6. Researcher will ask you about your experience during the study.

Your participation and your answers will help us make finding words in assistive communication devices easier.

How long?

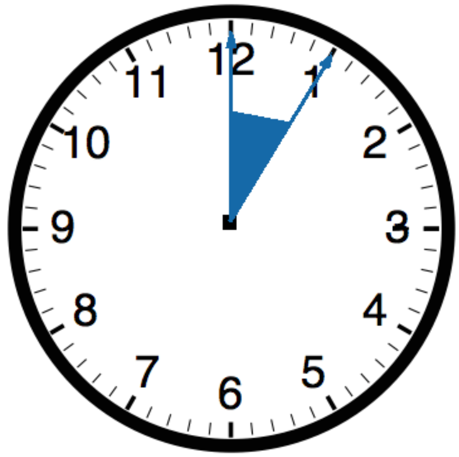

One hour:

- Find a few words;
- Answer a couple of questions;
- Take short break;
- Find more words;
- Answer a few more questions;
- Share what you thought about the study.

Where?

Adler Aphasia Center

Where people touched by Aphasia can comfortably communicate

Right to Withdraw

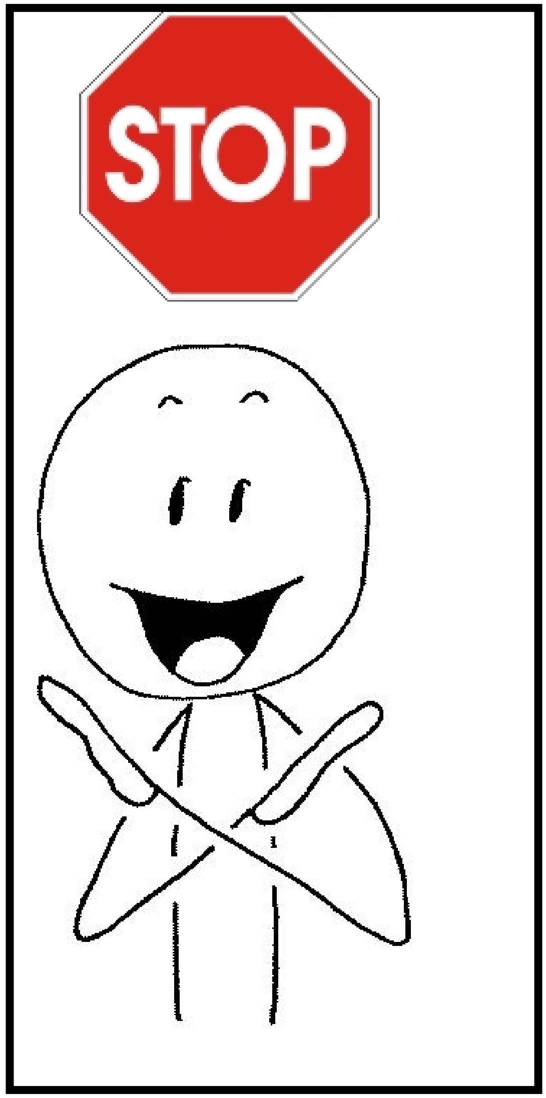

√ You **can stop** at any time

√ It is your **choice**

√ It is **OK** to quit

Potential Risks

✗ There is **NO** danger in participating in this study

✓ Everything is **confidential**

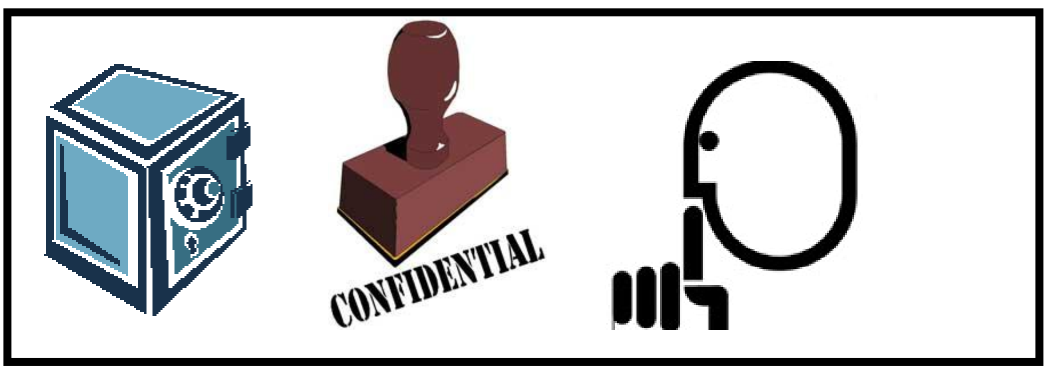

Payment:

Get **$10**

OR

Donate **$10** to the **Adler Aphasia Center**.

This information has been explained to me

YES  NO 

__________ __________

I **agree** to participate in this research project

YES  NO 

__________ __________

I have been given a copy of this form

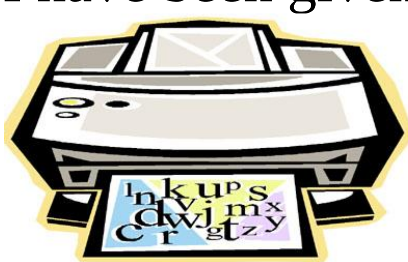

YES  NO 

__________ __________

Participant's signature or initials: ___________

Date:___________________________________

Contact Information

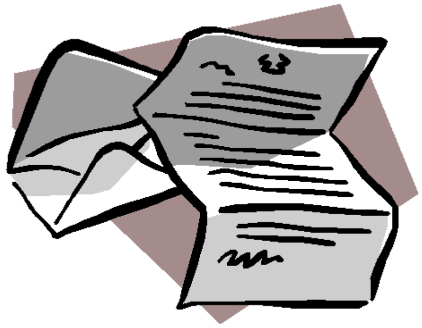

For questions about the research, contact:

Sonya Nikolova, Ph.D. Candidate
35 Olden Street
Computer Science Department
Princeton University
Princeton, NJ 08540
Phone: (609) 933-9084

For questions about your rights as a research participant, contact:

Joseph Broderick
Secretary, Institutional Review Panel for Human Subjects
P.O. Box 36
Princeton University
Princeton, NJ 08544
Phone: (609) 258-3976