TECHNOLOGY TRANSFER: PROBLEMS AND PROSPECTS

Jesse W. Fussell

Department of Defense
9800 Savage Road
Fort Meade, MD 20755

1. INTRODUCTION

For at least the last twenty years, DARPA and other government organizations have been sponsoring language processing research. The goal of this research is to develop processes to automate or semi-automate manual speech or text operations and to create new capabilities to make it easier for humans to work with speech and textual data. During that time, a substantial amount of money has been provided to the nation's best organizations which have employed the most innovative and intelligent individuals. However, while there is a large body of objective evidence which demonstrates continuous progress toward the various technical goals of the programs, the fact remains there has been virtually no success to date in transferring any of the research effort into day-to-day operational use by any of the government sponsors. The purpose of this paper is to explore the reasons for the dearth of technology transfer in this technical area in the past, to forecast prospects for technology transfer in the future, and to suggest some ideas for stimulating the process.

2. TECHNOLOGY TRANSFER PROBLEMS

Many reasons can be cited for lack of technology transfer successes in the past. The one which typically is of greatest interest to the research community is that the algorithms developed in the past were not considered adequate for operational use. In the speech recognition area, word recognition error rates were intolerably high unless the speech source was limited vocabulary, isolated word, wide bandwidth and free of background noise. In addition, applications were limited to the

simplest uses because there was little understanding of how to apply language processing results to speech. In the text processing area, problems being addressed were often severely constrained in terms of vocabulary, grammar and application. But it was virtually impossible to find real-world problems that were constrained enough to allow application of the newly developed techniques. The substantial progress that has been made by the research community has changed this situation. The current state-of-the-art, while clearly inadequate for very general problems, has great prospects for certain moderately specialized applications.

An additional problem that has impeded technology transfer has been the high cost and limited capability of the data and signal processing equipment needed to implement the algorithms researchers developed. In the speech processing area, researchers used to be restricted to general-purpose hardware which was tens or hundreds of times slower than real time, or of using special, fixed-point signal processors with very limited memory and dynamic range which were programmed in assembly language. Use of more modern computer languages such as LISP allowed faster design and implementation of experiments, but even with machines specially designed for that environment, processing was very slow. In addition, the graphics needed for a good user interface was limited in display speed and resolution. The net result of these limitations was that, except for the simplest algorithms, it was impractical to transfer research successes into operational use. However, while the thirst of researchers for ever greater computing, storage, networking and display capability still exists, it is greatly slackened. The explosive progress of computing over the past decade, led by the development of modern, high-performance, low-

cost, graphics workstations, now allows real-time or faster operation of fairly computational and memory intensive algorithms on floating point processors programmed in high-level languages.

Another fundamental problem has been that the customer processing environment did not contain the infrastructure needed for the introduction of the techniques being developed. New digital speech processing systems could not be easily introduced into systems in which speech signals were being processed as analog signals and stored on analog tape. Text processing systems could not be made to operate in an environment in which textual information was still being handled on paper and data bases were on file cards. Again, thanks to major advances in data networks, low-cost terminals, optical character recognition, data base software, A/D and D/A conversion, and other technologies, many offices now contain the backbone of the system needed to make use of new text or speech processing techniques. In addition, the adoption of standards for programming languages, operating systems, windowing systems and network protocols often is resulting in the users obtaining computing systems which are compatible or which can easily be made compatible with the systems on which the processes are being developed. Thus, standardization will result in technology transfer being an easier job.

The problems described so far may be classed as technical problems. And, as stated above, these technical problems are rapidly decreasing in importance. However there is one technical problem that still must be overcome before technology transfer can occur on a widespread basis. And unfortunately this remaining major technical problem has received little attention to date, probably because of the way in which the research problem has been structured. That is, the problems reported on by most researchers at this conference are primarily defined by two things: the selected (or created) and marked training and test data, or corpus, and a criterion or criteria for testing the researchers system against that corpus. This standardization of problem domain, goals and test criteria has been a powerful tool of the research managers. It has focussed researchers' attention onto relatively specific goals and objectives and, since the performance of the algorithms produced by different researchers can be directly compared, it has simultaneously created an environment of constructive competition. However this narrow focus has resulted in the neglect of one critical area

of work: the study of the process of converting an algorithm or process from one domain, that is supported by a well defined and documented corpus, to a different, operational domain, which may have little marked training or testing data. The result is there has been little work to even formally define the process necessary to convert from one domain to another and from one objective function to another one, much less to automate or semi-automate that process. In other words, there has been insufficient emphasis on efforts to achieve technology transfer.

There are also a number of political, managerial or psychological problems which need to be addressed if the transfer of language processing techniques into operational use is to be successful. All of these nontechnical problems are associated with the potential customers of the research. And, it is important to recognize that in general, the customers are not the sponsors of the research. The sponsors are usually other researchers or research managers who are supposed to represent the customers. The true customers are the people or organizations that will be the end user of any new product or capability that results from the research effort.

The first problem involving the potential customer is that the customer usually not directly involved in the research efforts of the human language technology program. Or to put the shoe on the other foot, most researchers are not sufficiently familiar with the customers' needs and operating procedures to know whether a test corpus is truly representative of "real" data and whether the research goals will solve any "real" problem.

Government customers also tend to be relatively conservative and sometimes even suspicious of new ideas and new technology. They frequently have developed a well understood routine and procedure for doing their jobs and are reluctant to change. In other words, there is often a lot of inertia which may only be overcome through the use of force. In this context, one such force which comes from demonstrated success. But since the voice and text processing technology is (and probably always will be) imperfect, "success", like "beauty" will be defined in the eye of the beholder. Since customers will be the ultimate judge of our products, we need to ensure that those products are demonstrated in

the best possible manner. Another force for change results from budget reductions. The conflicting pressures of less resources and desire to maintain capability tends to make some customers much more open to accepting new, imperfect technology.

3. RECOMMENDATIONS

First and foremost, the current program of research must be continued. Over the last few years, there has been substantial progress in understanding the fundamental problems of language processing and in developing ever better techniques for addressing those problems. This progress must continue if we are to have any hope of success against the less constrained problems.

However, the goal of this research program should be broadened to explicitly include technology transfer. The task of adapting speech and text understanding processes from one domain to another needs to be specified as one of the program goals. Algorithms need to be developed which can be easily converted from one task to another without requiring years of additional work by highly trained scientists. Techniques need to be developed which will allow supervised adaptation to new situations.

One implication of broadening the program to extend across domains is the need to increase the dimensionality of the research corpora so they also extend across domains. Instead of having data for a single situation which is segmented into portions for training, test and validation, this single set should be considered as the training portion for one domain, with another similar set for a different, but related domain which could be used to test how well the algorithm performs on the new domain, and possibly a third set of data from yet a third domain to validate the results of the domain transfer test.

There also needs to be additional efforts by researchers and research managers to find potential customers in the government, to educate them on the goals and results of the program and to solicit their inputs into those goals. In addition, researchers need to spend more time working with customers to better understand the manual processes currently being used so they can better understand what is needed in order to produce a "successful" language processing capability.

Finally, the overall goal of the program should be broadened, and funding provided to produce pilot or prototype systems which have been designed to be moved into operational situations and used for extended periods of time. Experimental operational prototypes are an absolutely necessary step in any long-term research effort. The trick is to determine the time for that step. The time is now.