

Translating Keyword Commands into Executable Code

Greg Little & Robert C. Miller
MIT CSAIL
32 Vassar St
Cambridge, MA 02139 USA
{glittle,rcm}@mit.edu

ABSTRACT

Modern applications provide interfaces for scripting, but many users do not know how to write script commands. However, many users are familiar with the idea of entering keywords into a web search engine. Hence, if a user is familiar with the vocabulary of an application domain, we anticipate that they could write a set of keywords expressing a command in that domain. For instance, in the web browsing domain, a user might enter *click search button*. We call expressions of this form *keyword commands*, and we present a novel approach for translating keyword commands directly into executable code. Our prototype of this system in the web browsing domain translates *click search button* into the Chickenfoot code *click(findButton("search"))*. This code is then executed in the context of a web browser to carry out the effect. We also present an implementation of this system in the domain of Microsoft Word. A user study revealed that end-users could use keyword commands to successfully complete 90% of the web browsing tasks in our study without instructions or training. Conversely, we would expect end-users to complete close to 0% of the tasks if they had to guess the underlying Javascript commands with no instructions or training.

ACM Classification H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces; D.3.3 [Programming Languages]: Language Constructs and Features; D.2.6 [Programming Environments]: Interactive environments; H.5.2 [User Interfaces]: User-centered design.

General Terms Algorithms, Design, Experimentation, Human Factors, Standardization, Languages

KEYWORDS: End-user programming, Command languages, Natural Language Processing, Web automation

INTRODUCTION

Many modern applications have scripting interfaces. These interfaces are powerful tools, both for automating tasks within applications and for coordinating tasks between applications. Unfortunately, the learning curve for writing scripts is prohibitive for many users. Three factors contribute to this learn-

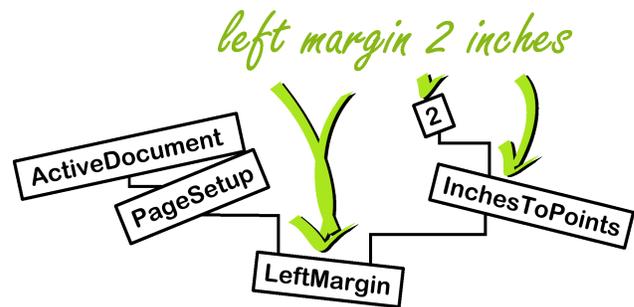


Figure 1: Illustration of keyword command translation.

ing curve. First, users must learn to cope with rigid and seemingly arbitrary syntax rules for formulating expressions. The canonical example is the semi-colon required at the end of each expression in C, but even modern scripting languages like Python and Javascript have similar arbitrary syntax requirements. Second, users must often learn several different scripting languages and be able to switch between them. This is necessary because of the wide variety of languages in use by applications, and it is difficult because the scripting languages can be very similar, but different enough to cause problems. Finally, users must learn the Application Programmer Interface (API) for the application they want to script, but API's can be quite large and it can be difficult to isolate the portion of the API relevant to the current task.

Much end-user programming research has focused on this problem. Two notable approaches include programming-by-demonstration (PBD) [6, 8] and structured editors [7, 11]. The basic idea behind PBD is to infer a program from manual actions taken by the user, like a macro recorder. This approach can be a good first step toward creating a script, but it is often easier to correct problems and generalize the script using a script editor. Structured editors can be a good approach for end users to write and edit scripts, since they force the script to retain a syntactically correct form. They also often provide menu systems so the user can recognize, rather than recall, various commands. However, structured editors tend to require the user to think in a certain order, which can be counter-intuitive. Also, structured editors sacrifice some of the benefits inherent to purely textual environments.

The benefits of purely textual environments are too great to dismiss outright, even for end-user programming. Consider that text is ubiquitous in computer interfaces. Facilities for easily viewing, editing, copying, pasting, and exchange-

ing text are available in virtually every user interface toolkit and application. Plain text is very amenable to editing—it is less brittle than structured solutions. Also, text can be easily shared with other people through a variety of communication media, including web pages, paper documents, instant messages, and e-mail. It can even be spoken over the phone. Tools for managing text are very mature. The benefits of textual languages for programming are well-understood by professional programmers, which is one reason why professional programming languages continue to use primarily textual representations.

The primary hurdle for end-users has been creating textual expressions that the computer will understand. Previous efforts in this regard have focused on natural language programming. Sammet advocated the use of English (or any natural language) as a programming language as early as 1966 [15]. However, interpreting arbitrary English expressions as executable code has proven to be a real challenge. A study by Miller [10] outlines some of these challenges, including the variety of styles humans use to express ideas.

Because of this, many natural language programming systems are built around grammars or templates which impose some constraints on input expressions. Grammar based systems like NLC [1] use formal languages that read like English. The problem with these systems is that they do not permit grammatical errors or extraneous words. Template based systems like NaturalJava [14] try to overcome these problems by searching for recognizable language constructs within an expression. However, these systems still require the user to incorporate recognizable constructs into their expressions.

Our approach goes one step further by eliminating the need for language constructs all together, and focuses on the presence of keywords in a command expression. We call such expressions *keyword commands*. Consider the keyword command *left margin 2 inches* in the context of Microsoft Word. To a human reader, this suggests the command to make the left margin of the current document 2 inches wide. Such a command can be expressed in the formal language of Visual Basic as *ActiveDocument.PageSetup.LeftMargin = InchesToPoints(2)*. A prototype of our system can make this translation automatically from the original keyword command (see Figure 1).

Several key points to note are: First, the user did not need to worry about strict requirements for punctuation and grammar in their expression. For instance, they could have said something more verbose, like *set the left margin to 2 inches*, or they could have expressed themselves in a different order with *2 inches, margin left*. Second, the user did not need to know the syntactic conventions for method invocation and assignment in Visual Basic. The same keyword command would work regardless of the underlying scripting language. Finally, the user did not have to search through the API to find the exact name for the property they wanted to access. They also didn't need to know that *LeftMargin* was a property of the *PageSetup* object, or that this needed to be accessed via the *ActiveDocument*.

One advantage of this approach over PBD is that it allows users to create scripts that access functionality they may not know how to access manually. For instance, assume a user wants to create 3 columns in their Word document as part of a script, but they do not know which dialog affords this change. They would need to discover the proper dialog before they could demonstrate the command to a PBD system. However, our approach would allow them to guess a keyword command like *3 columns*, and the system would do the work of searching through the API to construct an expression likely to achieve this goal.

This example also illustrates a key drawback of structured editors. If the user wants to express the command *3 columns* in a structured editor, they need to begin by filling in the first slot in an expression builder, and this slot is unlikely to accept the number 3. In fact, they may not know whether to begin with an assignment template or a function template, depending on how this particular property is set in the API. The main point here is that Structured Editors require some planning on the part of the user toward building their expression, even if they can pick out the pieces with menu systems. Our approach relaxes this restriction.

Another advantage of keyword commands over programming-by-demonstration and structured editors is that it accommodates *pure text*. This affords all the benefits cited above, and also allows these expressions to serve as meta-URLs for bookmarking application states. One virtue of the URL is that it's a short piece of text—a command—that directs a web browser to a particular place. Because they are text, URLs are easy to share and store. Keyword commands offer the same ability for arbitrary applications—you can store or share a small set of keyword commands that will put an application into a particular state. On the web, this could be used for bookmarking *any* page, even if it requires a sequence of browsing actions. It could be used to give your assistant the specifications for a computer you want to buy, with a set of keyword commands that fill out forms on the vendor's site in the same way you did. In a word processor, it could be used to describe a conference paper template in a way that is independent of the word processor used (e.g. *Arial 10 point font, 2 columns, left margin 0.5 inches*).

To test our idea, we have implemented two prototypes. One prototype operates in the web browsing domain, and translates keyword commands into Chickenfoot [4] commands. These commands allow users to navigate to web pages and interact with web forms. The other prototype operates in the domain of Microsoft Word, and translates keyword commands into Visual Basic commands. These commands can be used to access document properties and issue commands like Save and Print.

We also conducted a formative user study of the web domain prototype to gauge the ability of end users to form keyword commands on their own, without instructions or training. We chose the web domain because end users were likely to be familiar with the capabilities of a web browser—though this is not likely to be the target domain for such a system.

We found that users were able to generate successful key-

word commands for 90% of the tasks, and that their first attempt succeeded 73% of the time.

The next section addresses related work. This is followed by a description of the user interface for keyword commands. After that, we discuss the implementation of the two keyword command prototypes. Next, we present the user study. Finally, we end with future work and concluding remarks.

RELATED WORK

Interest in *natural programming* was renewed recently by the work of Myers, Pane, and Ko [12], who have done a range of studies exploring how both non-programmers and programmers express ideas to computers. These seminal studies drove the design of the HANDS system, a programming environment for children that uses event-driven programming, a novel card-playing metaphor, and rich, built-in query and aggregation operators to better match the way non-programmers describe their problems. Event handling code in HANDS must still be written in a formal syntax, though it resembles natural language.

Bruckman's MooseCrossing [2] is another programming system aimed at children that uses formal syntax resembling natural language. In this case, the goal of the research was to study the ways that children help each other learn to program in a cooperative environment. Bruckman found that almost half of errors made by users were syntax errors, despite the similarity of the formal language to English [3].

More recently, Liu and Lieberman have used the seminal Pane & Myers studies of nonprogrammers to reexamine the possibilities of using natural language for programming, resulting in the Metafor system [9]. This system integrates natural language processing with a common-sense knowledge base, in order to generate "scaffolding" which can be used as a starting point for programming. Keyword commands likewise rely on a knowledge base, but representing just the application domain, rather than global common sense.

USER INTERFACE

The heart of the user interface is keyword command generation. The user needs to conjure some command that the computer is likely to understand. Toward this end, the user must have some idea of the capabilities of the system, along with the conventional vocabulary used to describe these capabilities (e.g. the white-space surrounding a document is called the "margin").

Beyond this, our hope is that the system is natural and intuitive. That is, a user should not have to read the following sections in order to *use* the keyword command language.

Functions

The user can invoke a function by including keywords in their command which are present in the name of the function. For instance, in the example *left margin 2 inches*, the word *inches* appears in the function *InchesToPoints*. A single word is sufficient to identify a function, even if the word appears in other functions, so long as only one function fits well into an interpretation of the entire command. The framework also allows functions to have synonyms, giving the user some slack in remembering the exact name of a function.

The user can also invoke a function without naming it, merely by including its arguments. This is a sufficient suggestion in those cases where there is only one function which is likely to take the given arguments. For instance, consider *UIST 2006 into the search textbox*. This suggests entering "UIST 2006" into the textbox labeled "search", even though we didn't explicitly say *enter*. Our web domain prototype understands this expression because it has only one command that accepts a string and a textbox (namely *enter*).

Basic Data Types

Like functions, the user can also create basic data types by including keywords in their command which represent these types. For instance, the user can create the integer 2 with any of the keywords *2*, *two*, *2nd*, *second*, etc... depending on how many variations the interpreter has in its database.

Strings themselves are basic data types, and can be created by just including the words of the string in the command. However, a sequence of words is considered more likely to be a string if the user places quotes around it. We discuss this more below when we talk about resolving ambiguity.

The exact set of basic data types depends on the domain of the interpreter. Our web domain prototype includes types for integers, strings, booleans, URLs, and keyword-lists (which identify objects in the webpage).

Identifying Arguments

If a command takes multiple arguments of the same type, then the user may need to supply words in their expression to identify the arguments.

One method is to name the arguments. Argument names can appear immediately before or after the words used to express the argument. Another method is to include the arguments in the correct order. For instance, if we had a *click(integer x, integer y)* command, then the system would translate both *click 300 y 200 x* and *200 300 click* into *click(200, 300)*.

Prepositions can serve as intuitive names for some arguments. For instance, the command *copy(path toDestination, path fromSource)* allows for expressions like *copy A:\my_paper.doc to C:\my_backup*. In this case, the arguments are out of order, but the system picks up on the word *to*, which is part of the name for the *toDestination* argument. In practice, we include many synonyms for these prepositions to support variations like *copy A:\my_paper.doc into C:\my_backup*.

It is also possible for the system to disambiguate arguments of the same type based on domain specific heuristics. Continuing the example above, we could say *A:\my_paper.doc C:\my_backup copy*, where the arguments are out of order, but we do not provide a preposition. In this case, the system can disambiguate the arguments using the heuristic that if only one argument is a directory, then it is likely to be the destination.

Nested Functions

The system supports nested functions. For instance, *pick the 4GB RAM listitem* translates into *pick(findListitem("4GB RAM"))*. Note the nested invocation of *findListitem*.

The one restriction regarding nested functions is that a nested function can only be formed from a contiguous portion of the keyword command. This example obeys this restriction because *4GB RAM listitem* is a contiguous string within the command. We could have also exchanged the order to get *listitem 4GB RAM*, but we could not split up the subexpression to get *listitem pick 4GB RAM*.

In some cases, the system may make the correct translation, even if the subexpression is split. In the case of *listitem pick 4GB RAM*, the system fails to find a use for *findListitem()* as a nested function with no arguments, and so it favors an interpretation involving just the keywords *pick 4GB RAM*. The interpreter then introduces the unnamed function *find-Listitem*, since the user supplied arguments suggesting this command, namely *4GB RAM*.

Even without this fallback mechanism, this restriction does not appear to be prohibitive. No users in our study formed expressions violating this rule (and users correctly formed 31 expressions which could have violated the rule). There is also an intuitive argument for why users are unlikely to generate expressions violating this rule. That is, context-free grammars have been used to achieve good approximations of English grammar, and a context-free grammar would ensure the sort of contiguous nesting of subexpressions that our system requires.

Resolving Ambiguity

Sometimes an expression needs to use words in an ambiguous way. The most common example is trying to quote text, when the text itself contains words with other likely interpretations. Consider *enter binary search textbox*. Do we want to enter the word “binary” into the “search” textbox, or do we want to enter “binary search” into the only textbox on the page? The system supports a couple of techniques for resolving such ambiguities.

Quotes: The first technique is to include quotes. When the system considers the possibility that a sequence of words represents a string (or keyword-list), it gives this possibility more weight if the user includes quotes on either side of the sequence. We could therefore say *enter “binary” search textbox*, and the system would favor an interpretation where “binary” was treated as a string separate from *search*.

When using quotes, it is not necessary to place escape characters in front of quote symbols which are embedded within a quote. The system will resolve this ambiguity based on how many string arguments it requires to form a valid interpretation for the whole command. For instance, the expression *enter "history of "foo bar"" into the search textbox* resolves to *enter("history of \"foo bar\"", findTextbox("search"))*, despite the nested quotes in the original expression.

Argument Names: Another technique is to include argument names to identify an argument, as discussed earlier. Using this technique we could say *enter binary into search textbox*, and the system would favor an interpretation where the subexpression *into search textbox* was treated as the *into* argument of the *enter* command.

Selecting From List: When both of these techniques fail (or the user fails to employ them), the system is left with an ambiguous expression. In such a case, it can present the most likely candidates to the user for inspection. This is discussed more in the *Feedback* and *Graphical User Interface* sections below.

Extraneous Words

We explained that the system recognizes certain stop words in certain situations (e.g. *to*, *into*), but it doesn’t recognize all stop words. Also, the user might include other extraneous words in their expressions. Consider *please enter “search” into the textbox*. What is the system supposed to do with the word “please”? It ignores it. If no interpretation of the expression has a good explanation for a word, then the system considers interpretations which simply overlook that word.

Feedback

It is useful to let the user know what action the computer has taken in response to a command. One method is to supply a graphical indication. For instance, when a user clicks a link in a webpage using a textual command, our web prototype animates a green semi-transparent box over the link. This provides the user with assurance that the correct link was chosen.

In some cases, it is not feasible to supply a graphical indicator. In these cases, it is useful to provide textual feedback. This feedback can come in the form of displaying the resulting code generated from the command.

Our framework also provides a mechanism for generating pseudo-natural language representations for commands, which end-users may find easier to read. These interpretations may also act as guides for future expressions since they are themselves interpretable by the system. For instance, if the user enters the command *300GB Hard Drive*, and the system translates this to *select(findRadioButton("300GB Hard Drive"))*, then the pseudo-natural language feedback would look like *select the “300GB Hard Drive” radiobutton*, which is also an expression understood by the system.

Textual feedback is also useful when the user enters an ambiguous command, and the system wants to afford the selection of an interpretation from a list. In these cases, it may be easier to represent the alternatives with text, rather than a graphical indication.

Graphical User Interface

Our web browser prototype consists of a textbox affording input, and an adjacent horizontal bar allocated for textual feedback. The system also generates an animated acknowledgement in the web page around the html object affected by a command (see Figure 2).

We plan to extend this system with a dropdown menu to present a list of likely interpretations for ambiguous commands (the current prototype makes an arbitrary choice in such cases).

We also envision incorporating the system into a script editor as a form of auto-completion. For instance, if a user types a line of text in the Chickenfoot script editor that doesn’t parse



Figure 2: a) command box, b) feedback bar, c) animated acknowledgement

as Javascript, the system could present valid Javascript interpretations of the expression in an auto-completion style popup menu.

IMPLEMENTATION

This section describes how keyword commands are translated into executable code.

Functions

Functions are the building blocks in the system. Basic data types are represented as functions that take no arguments. Each function returns a single value of a certain type, and accepts a fixed number of arguments of certain types. Optional arguments are also permitted.

Functions can have many names. For instance, the *enter* command in the web prototype has names like *type*, *write*, *insert*, *set*, and the = symbol. Functions for basic data types have names corresponding to their textual representations. These names are matched programatically with regular expressions. For instance, integers are matched using the expression “[0-9]+”.

Arguments can also have many names. These usually include prepositions naming the gramatical role of the argument.

Translation Algorithm

The translation algorithm needs to convert an input expression into a likely function tree. One way to do this is to enumerate all possible function trees (to some depth), and then evaluate the likelihood that each tree matches the expression.

Unfortunately this search space is very large for a couple of reasons. First, the number of possible trees is large (growing exponentially with the addition of new functions). Second, the evaluation of the likelihood that a tree matches an expression is itself an exponentially large problem in the number of functions in the tree. Part of the contribution of this work is coming up with an algorithm for quickly finding function trees which correspond to a given expression.

We describe the algorithm in two steps.

Step 1: Tokenize Input

Each sequence of contiguous letters forms a token, as does each each sequence of contiguous digits. All other symbols (excluding white space) form single character tokens.

Letter sequences are further subdivided on word boundaries using several techniques. First, the presence of a lower-case letter followed by an upper-case letter is assumed to mark a word boundary. For instance, *LeftMargin* is divided between the *t* and the *M*. Second, words are passed through a spell checker, and common compound expressions are detected and split. For instance, *login* is split into *log in*.

One potential problem with this technique is that a user might know the full name of a property in an API and choose to represent it with all lower-case letters. For instance, a user could type *leftmargin* to refer to the *LeftMargin* property. In this case, the system would not know to split *leftmargin* into *left margin* to match the tokens generated from *LeftMargin*.

To deal with this problem, the system adds all camel-case sequences that it encounters to the spelling dictionary before splitting them. In this example, the system would add *LeftMargin* to the spelling dictionary. Now when the user enters *leftmargin*, the spell checker corrects it to *LeftMargin*, which is then split into *Left Margin*.

After spelling correction and word subdivision, tokens are converted to all lower-case, and then passed through a common stemming algorithm [13].

Step 2: Recursive Algorithm

The input to the recursive algorithm is a token sequence and a desired return type. The result is a tree of function calls derived from the sequence that returns the desired type.

This algorithm is called initially with the entire input sequence, and the desired return type *void*, since we want the command to *do* something as opposed to returning something.

The algorithm begins by considering every function that returns the desired type. For each function, it tries to find a subsequence of tokens that matches the name of the function. For every such match, it considers how many arguments the function requires. If it requires *n* arguments, then it enumerates all possible ways of dividing the remaining tokens into *n* contiguous subsequences (including subsequences of length 0, to account for optional arguments). Then, for each set of *n* subsequences, it considers every possible matching of the subsequences to the *n* arguments. Now for each matching, it takes each subsequence/argument pair and calls this algorithm recursively, passing in the subsequence as the new sequence, and the argument type as the new desired return type.

The resulting function trees from these recursive calls are grafted as branches to a new tree with the current function as the root. The system then evaluates how well this new tree explains the token sequence (see *Explanatory Power* below). The system keeps track of the best tree it finds throughout this process, and returns it as the result.

The system also handles a couple of special-case situations:

Extraneous Tokens: If a function takes no arguments, but there are subsequences left over after extracting the function name, then these subsequences are ignored. However, they do subtract some explanatory power from the resulting function tree.

Inferring Functions: If no tokens match any functions that return the proper type, then the system tries all of these functions again. This time, it does not try to find subsequences of tokens matching the function names. Instead, it skips directly to the process of searching for the arguments for each function.

Of course, if a function returns the same type that it accepts as an argument, then this process can result in infinite recursion. Therefore, the system will not infer commands after a certain depth in the recursion.

Explanatory Power

Function trees are evaluated in terms of how well they explain the tokens in the sequence from which they are derived. Tokens are explained in various ways. For instance, a token matched with a function name is explained as invoking that function, and a token matched as part of a string is explained as helping create that string.

Different explanations are given different weights. For instance, a function name explanation for a token is given 1 point, whereas a string explanation is given 0 points (since strings can be created from any token). This is slightly better than tokens which are not explained at all—these subtract a small amount of explanatory power. Also, inferred functions subtract some explanatory power, depending on how common they are. Inferring common functions costs less than inferring uncommon functions.

Web Prototype

The functions in the web prototype map to commands in Chickenfoot [4]. We include 18 functions, with an average of 6.4 names for each function.

Many functions share some of the same names. For instance, *make* is a name for the *pick* function and the *enter* function. These ambiguities are resolved based on argument types.

Word Prototype

The prototype of the system in the domain of Microsoft Word presented some challenges which were not present in the web prototype. First, we were faced with a much larger command set. Chickenfoot has less than 20 commands, whereas the Word API has over two thousand. This meant we could not create all the functions by hand. Instead, we mined Word's type libraries, and converted its properties and methods into functions automatically. However, we were not able to populate the system with function synonyms.

We also had to turn off the ability of the system to infer commands which were not named in the expression if these commands required additional arguments, since this dramatically increased the search space. However, we were able

to achieve some of the same benefits of inferring commands with a few modifications to the algorithm.

First, we wanted to be able to infer commands which returned useful objects in the system like *ActiveDocument.PageSetup*, which holds properties for margin sizes. We noticed that most of these objects were accessible as descendants of either *ActiveDocument* or *ThisApplication*; hence, we implemented an algorithm to enumerate all such descendants, and store them in a list. We then allowed the system to infer commands in this list as if they took no arguments (since we knew which arguments were required to access these objects).

The next modification is best introduced with an example. Consider the expression *columns 2*. The desired interpretation is *ActiveDocument.PageSetup.TextColumns.SetCount(2)*, but this requires inferring the command *SetCount* since no words from this command appear in the original expression. However, we note that *SetCount* has an argument named *NumColumns*, and the word *Columns* does appear in the expression. Hence, in such cases, we mitigate the need to infer a command by including function argument names as synonyms for function names.

The final modification is also motivated by an example. Consider the expression *A4*, which we want to translate to *ActiveDocument.PageSetup.PaperSize = wdPaperA4*. In this case, we need to infer the command *PaperSize*. The way we get around this is to allow the system to search for commands with any return type, instead of just commands returning *void*. Then, if the return is not *void*, we try to fit the return value into some other function that does return *void*.

In our example, the system returns the function *wdPaperA4*, which has the return type *WdPaperSize*. The system then searches for a function which takes something of type *WdPaperSize* as an argument. Remarkably enough, there is only one such function, namely *PaperSize*. This function also requires an argument of type *PageSetup*, but we allow the system to fill in such requirements using the list we built earlier, which supplies the function *ActiveDocument.PageSetup*.

Speed

We implemented the translation algorithm for each prototype in Java. The following running times provide a rough feel for the speed of this algorithm.

In the web prototype, a sample 4 word expression takes 70 milliseconds to translate (on an AMD Athlon 4200+ processor). Increasing this to 8 words takes 270 milliseconds, and increasing it to 12 words takes 800 milliseconds.

The Word prototype translates many short expressions (4 words or less) in less than a second, however the running time increases significantly for longer expressions. This is also very dependant on the words used. Including multiple words which appear in many commands can dramatically increase the translation time. For instance, the expression *left left* takes 4 seconds. This already suggests some room for optimization in the algorithm. For instance, we could make common words like *left* insufficient by themselves to invoke any commands.

Our overall conclusion is that the current algorithm is sufficient for small command sets (under 20 commands) where the user is unlikely to enter large expressions (over 8 words). However, we believe the algorithm already generates some very useful results in larger domains, and we want to explore this potential in future work.

USER STUDY

We believed that our command language was intuitive, and could be used without instructions, provided that the user was familiar with the domain. We therefore chose the web domain (which many end users are familiar with) and implemented a command interface to test how well users could use the system to automate common web browsing tasks.

Participants

Our study involved 9 users, solicited from a public mailing list at a college campus. Seven were between 20 and 30 years old, while the other two were 49 and 56. We had three females and six males. Five were students (4 of these were computer science majors). The other four subjects had a range of occupations. All subjects were compensated for their time.

The subjects were also all experienced web users, and could type reasonably well. Almost every subject claimed to use the web almost every day (except one, who claimed to use the web a few times a week). Every subject had also been to the majority of the web sites involved in the study. Finally, each subject used typing-centered programs (like Word or an Instant Messenger) almost every day (again except for one, but even this user felt reasonably comfortable typing).

Programming experience amongst the users was divided into two groups. Four users had never written a program, or had only written a program for a class. Each of the remaining five users had written multiple programs on their own, and was familiar with a number of different programming languages.

Setup

Each subject sat at a computer loaded with the web domain prototype. We then handed them a set of instructions to read and tasks to complete.

Instructions: The instructions indicated that they should use only the command box (Figure 2a) to do each task, and not click or type directly into the web page. We also indicated that it was up to them to decide what to type into the command box. We did not offer any suggestions about what to type (except for two users, see *Modifications to Study* below).

Tasks: Each of the 36 tasks consisted of a red circle drawn on a screen shot of the web browser, indicating what to do (see Figure 3). For instance, if we wanted the user to navigate to a certain URL, we would circle a screen shot of a location bar loaded with that URL. We did this in order to minimize external hints about how the user should communicate with the system.

Even still, the circled text itself acted as a hint for many tasks, especially for following links and clicking buttons. In these cases, entering the text on the link or button itself was sufficient to click it. However, we did not tell the subjects that this

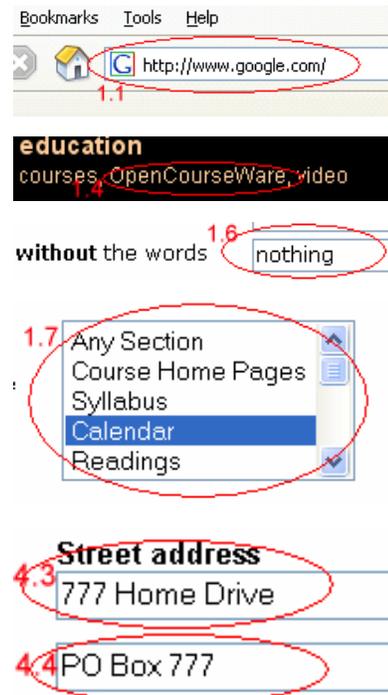


Figure 3: Examples of tasks in the user study. The user had to write a keyword command that would click or set the control circled in red. The red number is the task number from Figure 4.

was the case, and in fact, users often provided unnecessary words for these tasks (like the word “click”).

We also included some more difficult tasks that required words not present in the red circle to complete. However, every task could be accomplished with a single keyword command.

Modifications to Study: After running 6 users through the study, we noticed that people didn’t use many verbs. For instance, subjects tended to enter text into form fields with expressions like *without “nothing”* rather than *enter “nothing” into the without textbox*.

We wondered what effect it would have if we introduced a single initial suggestion of a command involving a verb. We decided to provide a hint for a task that every user had succeeded at (one of the easier tasks). We felt that this wouldn’t overly contaminate the study, while giving us a small hint to satisfy our curiosity.

We therefore suggested to 2 users that they do the first task with *go to google.com*. One user responded with “who needs verbs?”, and proceeded to do this task with just *google.com*. The other user took the suggestion, and used a verb in 10 tasks that no other user used a verb for, which may be cause to investigate this further.

Results

Subjects completed 90% of the tasks successfully, with an average of 1.7 attempts per task. 73% of the tasks were completed on the first try, with only one command. If the system

| tasks | subjects | | | | | | | | | |
|-------|----------|----|---|---|----|----|---|----|---|---|
| 1.1 | 2 | 1 | 2 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| 1.2 | 2 | 1 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 |
| 1.3 | 2 | 1 | 2 | 3 | 1 | 1 | 4 | 1 | 4 | |
| 1.4 | 1 | 2 | 1 | 3 | 5 | 1 | 1 | 1 | 3 | |
| 1.5 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| 1.6 | 3 | 1 | 1 | 3 | 10 | 22 | 1 | 5 | 1 | |
| 1.7 | 2 | 11 | 2 | 0 | 3 | 3 | 1 | 0 | 0 | |
| 1.8 | 1 | 1 | 1 | 1 | 1 | 1 | 2 | 1 | 1 | |
| 1.9 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 2 | |
| 1.10 | 1 | 1 | 1 | 1 | 0 | 1 | 1 | 1 | 1 | |
| 2.1 | 1 | 1 | 1 | 2 | 3 | 1 | 3 | 1 | 1 | |
| 2.2 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 2.3 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 2.4 | 1 | 1 | 1 | 2 | 1 | 1 | 1 | 1 | 1 | |
| 2.5 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 2.6 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 2.7 | 1 | 2 | 1 | 1 | 1 | 2 | 1 | 1 | 1 | |
| 2.8 | 5 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 2 | |
| 3.1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 3.2 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 3.3 | 1 | 1 | 1 | 2 | 1 | 1 | 1 | 1 | 1 | |
| 3.4 | 1 | 7 | 3 | 9 | 4 | 4 | 4 | 10 | 6 | |
| 3.5 | 4 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 3.6 | 1 | 1 | 1 | 1 | 1 | 0 | 0 | 1 | 1 | |
| 4.1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 4.2 | 1 | 1 | 1 | 2 | 1 | 1 | 1 | 1 | 1 | |
| 4.3 | 1 | 1 | 1 | 9 | 1 | 1 | 5 | 6 | 1 | |
| 4.4 | 1 | 13 | 4 | 2 | 6 | 12 | 4 | 1 | 4 | |
| 4.5 | 1 | 2 | 1 | 2 | 1 | 1 | 2 | 2 | 1 | |
| 5.1 | 1 | 2 | 1 | 1 | 1 | 1 | 2 | 1 | 1 | |
| 5.2 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 5.3 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| 5.4 | 1 | 1 | 1 | 3 | 2 | 1 | 1 | 0 | 1 | |
| 5.5 | 1 | 1 | 1 | 4 | 1 | 2 | 2 | 0 | 1 | |
| 5.6 | 1 | 4 | 5 | 1 | 2 | 2 | 5 | 0 | 2 | |
| 5.7 | 1 | 2 | 9 | 2 | 2 | 2 | 3 | 0 | 2 | |

Figure 4: Number of attempts made by each subject for each task. Green cells indicate eventual success.

understood only Javascript, and we had offered no instructions, we would have expected a completion rate around 0%.

Figure 4 shows the number of attempts each user made to accomplish each task. The green cells indicate tasks that were completed successfully, while uncolored cells were not completed.

A zero indicates no attempts. These usually resulted from a subject overlooking a task. One zero resulted from a user entering text directly into a textbox without using the command box. The vertical strip of 0's at the bottom of the chart corresponds to an instance when a webpage was temporarily unavailable.

Discussion

We learned the most from the tasks with the highest failure rates.

Task 1.6: This is the first task with any failed attempts. The task asked the user to enter the word “nothing” into a textbox labeled “without the words.” Only 3 users completed this task on their first attempt.

All of the remaining subjects expected the system to have

a notion of an input focus. They began the task by trying to focus the computer’s attention on the appropriate textbox with commands like *go to without the words*. These attempts would usually result in text being placed elsewhere. Attempts were then made to correct these errors, but these were not counted toward the completion of the task. Through trial and error, 3 of these users eventually succeeded at the task.

It worth saying something about the user with 22 attempts toward this task: most of these attempts were of the form *next* or *up*, presumably in an attempt to move the attention of the computer to the textbox next to or above the most recent visual acknowledgement. This user appears to have mistaken the visual acknowledgments as describing the new focus of the system.

Task 1.7: This was really the second part of a two-part task. The task asked the user to select “Calendar” from a listbox. After doing so, they were meant to notice that the “Any Section” listitem was still selected in the web page, but not selected in the task illustration. At this point, they were meant to deselect the “Any Section” listitem.

Three subjects did not make this attempt, presumably because they thought it was a bug in the instructions or the webpage, or they did not notice it.

Three subjects used keywords which had not been added as synonyms for the proper functions in the system. The lesson learned here is that more synonyms need to be added, but we are encouraged to find that we did not need to add new functions in this case.

The final subject tried to access the listbox by focusing the computer’s attention on it, which, as discussed previously, did not work. However, we believe such a paradigm could be made to work within the bounds of the command language by adding commands like *focusOn*, *focusPrevious*, and *focusNext* (along with appropriate synonyms).

Task 3.4: This task seemed simple: it asked the user to enter a password into a password field. However, it exposed a bug in our choice of weights for the function scoring system. A brief description of this bug should prove instructive. Consider the input *Password bloppy*. 5 users tried this input, and it should have worked. Instead, the system clicked a link with the word “password” in it. Both of these interpretations needed to introduce an unnamed command, but the *enter* command cost 0.2, whereas the *click* command cost only 0.1. Of course, the click interpretation had to treat the word *bloppy* as an extraneous word, which cost another 0.1, whereas the *enter* interpretation treated *bloppy* as a string, costing nothing. Both interpretations explained the word *password* as a keyword-list identifying a textbox or link.

The resulting score for *click(link(“password”))* was 0.8, and the score for *enter(“bloppy”, textbox(“password”))* was also 0.8. However, the *click* interpretation happened to appear first in the list, and the system went ahead with this interpretation, since the prototype afforded no means of disambiguation after the command was entered.

We plan to address this issue by implementing the disam-

biguation dropdown we discussed.

Task 4.4: This was probably the hardest task in the study. It required the user to enter the address “PO Box 777” into the second street address field, shown at the bottom of Figure 3.

People had difficulty identifying this field since it had no label of its own. 5 subjects eventually succeeded, but 4 users did not. It is instructive to examine the potential reasons for these failures.

One user tried the approach of first focusing the computer’s attention, and then typing. This user entered *street address 2* and then *PO Box 777*. It is worth noting that if these commands had been entered as the single expression *street address 2 PO Box 777*, they would have worked. It would also have worked if we had a *focus* command. This command could accept the same arguments as the *enter* command (except for the string argument), and we could add a version of the *enter* command that accepts only a string.

Another user entered *777 Home Drive* into the first textbox successful, and then issued the command *next*. This attempt would also have been helped with the addition of *focus* commands, specifically *focusNext*.

This same user also tried what turned out to be a common approach, which was to try entering both lines with a single command: *777 Home Drive, PO Box 777 Street address*. In fact, 6 users made attempts of this sort.

This is a harder problem to solve. One solution might consist of an *enter* command with 2 string arguments. However, we should note that all the users who attempted this approach eventually succeeded, except for two. One would have figured out a way if the system obeyed *focus* commands as discussed before. The other would have succeeded if not for a bug in the parser. To understand this bug, it is helpful to know that the most common type of expression that worked was: *Street address 2 “PO Box 777”* (3 people succeeded with commands of this sort). Now the expression of the user in question was: *Street2=PO Box 777*. The problem is that this version of the tokenizer treated *Street2* as a single token. If the user had put a space between *Street* and *2*, this command would have worked. The current prototype now supports this command because it treats alpha and numeric sequences as separate tokens.

Task 1.4: Everyone succeeded at this task, but it exposed another important problem with the tokenizer. The task was to click the “OpenCourseWare” link on MIT’s homepage, which the tokenizer split into 3 tokens. However, this version did not add “OpenCourseWare” to the spelling dictionary. Hence, when users would type *opencourseware*, the system would not split the token into *open course ware*, and it would not match the link.

In the new system, typing *opencourseware* is corrected by the spell checker into *OpenCourseWare*, and the expression is then tokenized correctly.

FUTURE WORK

We plan to continue to develop this technique in several directions. First, we believe there is a lot of room for improvement to the algorithm, now that we’ve convinced ourselves that the algorithm is useful.

Also, we want to add new features to the algorithm. In particular, we would like to support variable declarations which are referenced in the same command. This may allow for for-loop constructs.

A top priority is to integrate the command interpreter into the script editor of Chickenfoot, allowing end users without Javascript experience to create simple scripts. Part of this integration will involve permitting access to variables and user defined functions from the command interpreter.

Another goal is to integrate the Word domain prototype as a plugin for Word. We want to explore the potential of this tool as an alternative to navigating a large menu system. A user study might reveal that users prefer to access some functionality through a command prompt.

We also believe the technique may be useful for programming in Java. The Java API is extensive, and it can take a small time to find a given function. But consider typing *time*, and having the system translate this to *System.currentTimeMillis()*. We believe such translations may be possible, and useful.

Finally, we would like to explore the potential of using the translation system as a generic backend for a speech recognition system. We believe that the general framework may make it easy to expose core sets of functionality of various applications to end users.

CONCLUSION

We have scratched the surface of a domain with great potential: translating keyword commands into executable code. We have described an algorithm for performing such translations in a reasonable time for small applications, and we have demonstrated that end users can form commands that are interpretable by the system, without any training.

ACKNOWLEDGMENTS

We would like to thank all the participants in our user study. We also appreciate all the helpful suggestions we have received from members of the UID group, as well as from other friends and peers. This work was supported in part by the National Science Foundation under award number IIS-0447800. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

1. Ballard, B., and Biemann, A., Programming in Natural Language: NLC as a Prototype, *ACM/CSC-ER Annual Conference*, 228-237. 1979.
2. Bruckman, A., Community Support for Constructionist Learning. Computer Supported Cooperative Work, *Computer Supported Cooperative Work*, 7:47-86, 1998.

3. Bruckman, A., Edwards, E., Should we leverage natural-language knowledge? An analysis of user errors in a natural-language-style programming language, *CHI*, Pittsburgh, PA, 1999.
4. Bolin, M., Webber, M., Rha, P., Wilson, T., Miller, R., Automation and customization of rendered web pages, *User Interface Software and Technology*, Seattle, WA, 2005.
5. Katz, B., Felshin, S., Yuret, D., Ibrahim, A., Lin, J., Marton, G., McFarland, A., and Temelkuran, B., Omnibase: Uniform Access to Heterogeneous Data for Question Answering, *NLDB*, 2002.
6. Cypher, A., Ed. *Watch What I Do: Programming by Demonstration*. MIT Press, Cambridge, MA, 1993.
7. Kelleher, C., Cosgrove, D., Culyba, D., Forlines, C., Pratt, J., and Pausch, R., Alice2: Programming without Syntax Errors, *User Interface Software and Technology*, Paris, France, 2002.
8. Lieberman, H., Ed. *Your Wish is My Command: Programming By Example*. Morgan Kaufmann, San Francisco, CA, 2001.
9. Liu, H., and Lieberman, H., Programmatic Semantics for Natural Language Interfaces. *CHI*, Portland, OR, 2005.
10. Miller, L., Natural Language Programming: Styles, Strategies, and Contrasts *IBM Systems Journal*, 1981.
11. Myers, B., McDaniel, R., and Kosbie, D., Marquise: Creating complete user interfaces by demonstration, *INTERCHI '93: Human Factors in Computing Systems*, Amsterdam, The Netherlands, 1993.
12. Myers, B., Pane, J., and Ko, A., Natural Programming Languages and Environments. *Commun. ACM*, 2004.
13. Porter, M., An algorithm for suffix stripping, *Program*, 14(3) pp 130-137, 1980.
14. Price, D., Riloff E., Zachary J., and Harvey B., Natural-Java: A Natural Language Interface for Programming in Java. *ACM Intelligent User Interfaces Conference*, 2000.
15. Sammet, J., The Use of English as a Programming Language. *Commun. ACM*, 9(3), 228-230. 1966.