

Naming Page Elements in End-User Web Automation

Robert C. Miller and Michael Bolin

MIT CSAIL

32 Vassar St

Cambridge, MA 02139 USA

1-617-324-6028

{rcm,mbolin}@mit.edu

ABSTRACT

Names of commands and objects are an important part of the usability of a programming system. We are developing a web automation system in which users need to identify web page elements, such as hyperlinks and form fields, in pages written by other designers. Using a survey of 40 users asking them to provide names for page elements, we were able to develop name resolution techniques that automatically find the desired element from the user's name for it.

Categories and Subject Descriptors

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, Experimentation, Human Factors, Languages.

Keywords

End-user web automation, web browsers.

1. INTRODUCTION

The names given to software components – such as variables, functions, classes, and commands – are an important part of the user interface of an end-user programming system. Choice of names, whether made by the system's designers or by end-user programmers themselves, can affect learnability, recall, readability, and maintainability of programs. Professional programmers recognize the importance of names, and naming conventions are the result (e.g., [3],[11]). A classic study of naming by Furnas *et al.* [2] was somewhat discouraging, showing that command names chosen by different people had a very low likelihood being consistent with each other. The solution proposed by Furnas *et al.* was *unlimited aliasing*, allowing "many, many alternate verbal routes" to the same functionality.

In this paper, we discuss how we have faced the naming problem in the design of an end-user programming system for automating and customizing interaction with the Web. The main question we consider is how a user should refer to elements on a web page (such as hyperlinks and form fields) in customization or scripting, particularly when the web page was authored by another designer, not by the user. In the next section (section 2), we discuss a number of design principles that interact in the choice of a name.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2005 ACM 1-58113-000-0/00/0004...\$5.00.

In section 3, we describe Chickenfoot, the end-user web automation system we are developing. In section 4, we present a pilot study we conducted to learn how users might name web page elements. Finally, in section 5, we outline a name resolution algorithm that implements a form of limited aliasing that performs well on the kinds of names we discovered in the study.

2. DESIGN PRINCIPLES

The goal of a name, whether used in programming or in natural language, is to identify a thing, so that both the writer and the reader agree about which thing is under discussion. Unlike natural languages, however, programming languages have very different kinds of readers with very different needs: software and humans.

To software tools, such as compilers or interpreters, the most important property of a name is its *precision*: i.e., that it uniquely identifies the thing intended by the user. A precise name identifies exactly one thing. Naming systems in software are generally designed to minimize ambiguity, rejecting attempts to introduce names that would be imprecise. For example, the Windows file system refuses to allow two files of the same name in the same directory. In Java, two variables in the same scope may not share the same name, and two classes with the same name may not be imported simultaneously. For a software tool, *name collisions* are the worst kind of failure that can occur, since they leave the software unable to resolve references to the name.

Since software engineering is also concerned with the correctness of a program over its entire lifecycle (maintainability) and in other contexts (reuse and extensibility), a well-chosen name in a well-designed naming system should *remain* precise as a program is modified and combined with code written by other programmers, not necessarily talking to each other. The need for precision over time and space, a property that we'll call *robustness*, is what drives naming systems to introduce scoping and package mechanisms, in order to isolate one module's names from another's. When names chosen by different programmers may have to mingle together, naming conventions are developed that reduce the chance of a collision, usually by referring to an external source of unique identifiers. For example, Java programmers are encouraged to prefix their package names with their organization's domain name, such as com.sun.*, since the uniqueness of domain names is guaranteed by domain name registrars [10]. An extreme form of this is the use of universally unique identifiers (UUIDs), constructed from a network card's MAC address and the current clock time. UUIDs are used in Microsoft COM to name classes and interfaces, and in RDF to name objects and properties.

If every name were a UUID, precision and robustness would be satisfied, and complicated scoping and namespace rules would be unnecessary. (Indeed, many source code analysis tools internally

rename all the user's messy names with fresh unique identifiers to simplify managing these rules.) Human programmers, on the other hand, would find this intolerable, since they depend on other properties of a name. The most important of these properties is *suggestiveness*, the extent to which a name describes the content, use, and type of the thing it identifies. A variable named *radius* is more suggestive than one named *r*. Suggestiveness depends strongly on shared experience between the writer of the code and its reader, and also on the context of the code. In code dealing with polar coordinates, *r* may be just as suggestive as *radius*.

Suggestiveness lies behind recommendations to use long identifiers, including whole words and multiple words, and avoiding unnecessary abbreviations. Suggestiveness drives the naming conventions used in many languages and APIs. In Java, for example, case distinctions are conventionally used to suggest whether a name refers to a variable (*string*), a class (*String*) or a constant (*STRING*). Hungarian notation, first articulated by Charles Simonyi and popularized in the Microsoft Windows API, uses short prefixes to encode the type of a variable in the name, such as *lpzFirstName*, where *lpz* means *long pointer to string terminated by zero*. Hungarian notation can make finer distinctions in type and usage than the C/C++ type system is capable of expressing. For example, *ichFirstName* and *cchFirstName* are both integer variables, but the former would be an *index* into a character array, while the latter would be the *count* of the characters in the array.

Precision is not as important to people, since humans are more tolerant of ambiguity. One way to resolve ambiguity is by appealing to context; for example, in a discussion of Java collections, *List* obviously means the collection class *java.util.List*, not the user interface widget *java.awt.List*. Another way to resolve ambiguity is to engage in a dialogue ("Which *List* do you mean?"), but this is only feasible when the communication is interactive.

Names have other properties that are important for human readers. Some of these properties can be derived from well-known usability design heuristics [7]. *Consistency* means using similar names for similar things, and dissimilar names for dissimilar things. For example, a human reader can more readily recognize an idiom like *for (int i=0; i<n; ++i)* when the loop control variable is consistently named *i*. Conversely, using *s1* to name a string and *s2* for a stream in the same function is liable to cause confusion. Naming conventions help improve consistency, especially between names chosen by different programmers.

Efficiency means that (all other properties equal) a shorter name is better than a longer name. Shorter names are faster to use, whether the user is typing them, reading them, or speaking them aloud (as Tikki Tikki Tembo's parents learned to their regret[7]).

Error prevention is also desirable; an ideal name should not be prone to misspelling or misreading. For example, *weird* may be easily mistyped as *wierd* or misread as *wired*. We noticed this effect in developing a text pattern language which used *containing* as a pattern operator. So many users mistyped the operator as *containg* – even we did, the system's developers – that we eventually added this synonym to the grammar.

Although names in computer systems are primarily used in written form (typed on a keyboard or read on a screen or on paper),

pronunciation also matters, since people often talk about the names. In software development, this may happen in design discussions, code reviews, classes, or in pair programming. Unpronounceable names like *m_lprgchName* seriously inhibit this kind of communication. URLs were not designed with pronunciation in mind: *http://www* is so hard to say that most speakers simply omit it, and web browsers wisely tolerate the omission. (Tim Berners-Lee reads *www* as "wuh-wuh-wuh," but that hasn't caught on.)

Thus far our examples have all been names that can be written and spoken – *textual* names. Another important way to identify a thing is to *point* at it. This approach is the basis for direct manipulation interfaces; in direct manipulation, things are identified by pointing at a visual representation of the thing on the screen. A human user typically points at the thing with a mouse; when software needs to point, it typically uses highlighting. Pointing is easier for humans to learn than textual names, and imposes far less memory burden, which helps to account for the substantially better learnability of direct manipulation interfaces compared to command languages. Infants, for example, use pointing to identify things long before they are able to speak the names of the things.

Pointing is also less error prone, more efficient, and more precise than a textual name, as long as the thing is already in view. But pointing doesn't *scale* well to discriminating among a large number of things. To be pointed at, a thing must be found and brought into view. In large namespaces this can be a very costly operation. On the Web, for example, it is infeasible to try to identify an arbitrary web page purely by pointing (i.e., clicking a sequence of hyperlinks).

3. CHICKENFOOT

We have encountered some of these naming issues in the design of Chickenfoot, an end-user programming system integrated into a web browser. This section describes the goals and design of Chickenfoot.

3.1 Goals

The primary goal of Chickenfoot is to give the user a platform for automating and customizing their interaction with the Web. Although web browsers have a long history of built-in scripting languages, these languages are not designed for the *end user* of a web site. Instead, languages like JavaScript, Java, and Curl [8] are aimed at *designers* of web sites. Granted, many web designers lack a traditional programming background, so they may be considered end-user programmers in that respect. But the needs of a designer, building a web application from whole cloth, differ significantly from the needs of a user looking to tailor or script an existing web site. Current web scripting languages do not serve the needs of web automation.

A second goal of Chickenfoot is to allow the end user to automate and customize web sites using a familiar interface, namely the web site's user interface. Existing approaches to web automation use a scripting language that dwells outside the web browser, such as Perl, Python, screen-scraper [1], and WebL [4]. For an end-user, the distinction is significant. Cookies, authentication, session identifiers, plugins, user agents, client-side scripting, and proxies can all conspire to make the Web look significantly different outside the web browser than inside. But perhaps the most telling

difference, and the most intimidating one for an end user, is the simple fact that outside a web browser, a web page is just raw HTML. Even the most familiar web portal looks frighteningly complicated when viewed as HTML source. The primary challenge for Chickenfoot can be simply stated: the user should never have to view the HTML source of a web site in order to customize or automate it.

Chickenfoot is targeted at the three kinds of automation:

Automating repetitive operations. For example, many conferences now use a web site to receive papers, distribute them to reviewers, and collect the reviews. A reviewer assigned 10 papers to read and review faces a lot of repetitive web browsing to download each paper, print it, and (later) upload a review for it. Tedious repetition is a strong argument for automation. Other examples include submitting multiple queries and comparing the results, and collecting multiple pages of search results into a single page for easy printing or additional sorting and filtering.

Integrating multiple web sites. Some web sites already provide some level of integration with other sites; for example, many retailers use MapQuest to display their store locations and provide driving directions. But end-users have no control over this integration, to their detriment. For example, before buying a book from an online bookstore, a user may want to know whether it is available in the local library—a question that can be answered by submitting a query to the library's online catalog interface. Yet the online bookstore is unlikely to provide this kind of integration, not only because it may lose sales, but because the choice of library is inherently local and personalized to the user. This kind of integration can only be sensibly done by an end-user.

Transforming a web site's appearance. Examples of this kind of customization include changing defaults for form fields, filtering or rearranging web page content, and changing fonts, colors, or element sizes. Web sites that use Cascading Style Sheets (CSS) have the potential to give the end user substantial control over how the site is displayed, since the user can override the presentation with personal stylesheet rules. With the exception of font preferences, however, current web browsers do not expose this capability in any usable way.

3.2 Design

Chickenfoot is being developed as an extension to the Mozilla Firefox web browser. Chickenfoot's design has two parts: (1) a development environment that allows users to enter and test Chickenfoot programs, and (2) a scripting language that extends the browser's built-in Javascript language with new commands for web automation.

Figure 1 shows a screenshot of the development environment presented by the current Chickenfoot prototype, which appears as a sidebar in Firefox. At the top of the sidebar is a text editor which accepts a Javascript program, which may be merely a single expression or command to execute, or a program with function and class definitions. The bottom of the sidebar is a console output window, which displays error messages, printed output, and the result of evaluating the Javascript code (i.e., the value of the last expression). This interface, though minimal, goes a long way toward making the Javascript interpreter embedded in every web browser actually *accessible* to the end-user. Previously, there

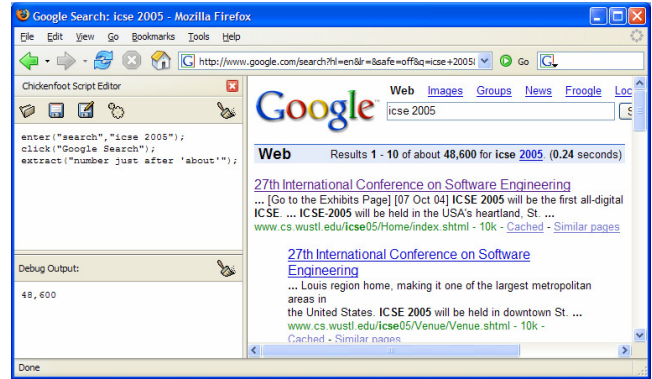


Figure 1. Chickenfoot development environment running inside the Firefox web browser.

were only two ways to run Javascript in a web browser: by embedding it in a web page (generally impossible if the page is fetched from a remote web site, since the user can't edit it), or by using a *javascript:* URL, which requires the entire program to be written on a single line.

A Javascript program running in the Chickenfoot sidebar operates on the web page shown in the main part of the window. Unlike most Javascript, Chickenfoot scripts run with no security restrictions, since they are developed and run by the end-user, not downloaded from a potentially malicious remote site. A Chickenfoot script is therefore free to interact with web pages from arbitrary sites and examine any aspect of the web browser's history or user interface.

Chickenfoot extends the standard client-side Javascript with a number of commands to simplify web automation. Some of these commands simulate actions that a web user can perform on the hyperlinks and forms of a web page:

- click (*link-or-button*)
- enter (*textbox, value*)
- pick (*menu-or-list, option*)
- check (*checkbox-or-radiobutton*)
- uncheck (*checkbox*)

These commands raise the question at the heart of this paper: what *name* should we use for the page object (link, button, or other widget) that a command should act on?

For a form widget, like a textbox or a checkbox, one possibility is the name assigned to the widget by the web page designer. This is the name used by Javascript embedded in the web page, and in the HTTP request sent back to the web server when the form is submitted. One key drawback of this name is that it isn't readily available to a web user without examining the HTML source, which contradicts one of the goals of Chickenfoot. The Chickenfoot development environment could solve this problem (e.g., by making form field names visible in the page on command). But these names have a second problem: since they are not chosen by the web user nor intended to be seen by the web user, they are not likely to be *suggestive*. For example, Google forms use names like *as_q*, *as_qdr*, and *as_occt*; MapQuest fields look like *2c* and *2s*. These names are virtually opaque to a user.

Another possibility is to use pointing to identify a page object, rather than a textual name. Indeed, this approach makes a lot of sense for authoring web automation, and our future plans include developing a programming-by-demonstration system on top of Chickenfoot, so that the user's clicks and keystrokes are translated automatically into Chickenfoot statements. But even if the user points at page objects to generate Chickenfoot code, there remains the question of what names to display in the generated code. Although visual representations of the code are possible (e.g. [5],[6]), a compact textual name would be more efficient with screen real estate and more pronounceable.

We chose to explore a third option: using visible labels in the page to identify page objects. For example, hyperlinks and buttons typically contain a visible text label that can be used with the *click* command:

```
click ("Google Search")
```

Other form widgets, such as textboxes and lists, have captions adjacent to the widget that can be used with other commands:

```
enter ("User name", "john_smith_1170@hotmail.com")
enter ("Password", "bri56ght")
click ("Sign In")
```

Visible labels are very likely to be suggestive names, because they are chosen by a web site designer to be read and understood by a user, and also because the user is likely to be familiar with them from manually interacting with the web site. Visible labels are not a complete solution, however. One challenge for this technique is the use of inline images for labeling hyperlinks and buttons. Fortunately, well-designed web sites offer ALT text for these images, intended to help visually-impaired users with screen readers, but which can help Chickenfoot as well. For images with no ALT text, we must fall back to other naming methods, such as internal page names or pointing.

4. NAMING SURVEY

To explore the usability of visible labels as names, we conducted a small pilot survey to find out what kinds of names users would generate for form fields, and whether they could comprehend names based on visible labels.

4.1 Method

Our survey concerned a narrow aspect of the web automation domain: what names users give to textboxes on web pages. Presented over the Web, the survey consisted of three parts, always in the same sequence. Part 1 explored freeform generation of names: given no constraints, what names would users generate? Each task in Part 1 showed a screenshot of a web page with one textbox highlighted in red, and asked the user to supply a name that "uniquely identified" the highlighted textbox. Users were explicitly told that spaces in names were acceptable. Part 2 tested comprehension of names that we generated from visible labels. Each task in Part 2 presented a name and a screenshot of a web page, and asked the user to click on the textbox identified by the given name. Part 3 repeated Part 1 (using fresh web pages), but also required the name to be composed only of "words you see in the picture" or "numbers" (so that ambiguous names could be made unique by counting, e.g. "2nd Month").

The whole survey used 20 web pages: 6 pages in Part 1, 8 in Part 2, and 6 in Part 3. The web pages were taken from popular sites, such as the Wall Street Journal, the Weather Channel, Google, AOL, MapQuest, and Amazon. Pages were selected to reflect a

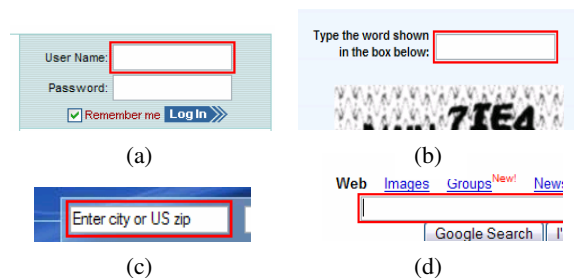


Figure 2. Sample textboxes used in the web survey.

variety of textbox labeling seen across the Web, including explicit captions (Figure 2a), wordy captions (Figure 2b), captions displayed as default values for the textbox (Figure 2c), and missing captions (Figure 2d). Several of the pages also posed ambiguity problems, such as multiple textboxes with identical or similar captions.

Subjects were unpaid volunteers recruited from the university campus by mailing lists. Forty subjects took the pilot survey (18 females, 20 males, 2 unreported), including both programmers and nonprogrammers (24 reported their programming experience as "some" or "lots", 15 as "little" or "none", meaning at most one programming class). All but one subject were experienced web users, reporting web usage at least several times a week.

4.2 Results

We analyzed Part 1 by classifying each name generated by a user into one of four categories: (1) *syntactic* if the name used only words that were visible somewhere on the web page (e.g., "User name" for Figure 2a); (2) *semantic* if at least one word in the name was not found on the page, but was semantically relevant to the domain (e.g., "login name"); (3) *layout* if the name referred to the textbox's position on the page rather than its semantics (e.g., "top box right hand side"); and (4) *example* if the user gave an example of a possible value for the textbox (e.g. "johnsmith056"). About a third of the names included words describing the type of the page object, such as "field", "box", "entry", and "selection"; we ignored these when classifying a name.

Two users consistently used *example* names throughout Part 1; no other users did. (It is possible these users misunderstood the directions; since the survey was conducted anonymously over the Web, it was hard to ask them.) Similarly, one user used *layout* names consistently in Part 1, and no others did. The remaining 37 users generated either *syntactic* or *semantic* names. When the textbox had a concise, visible caption, *syntactic* names dominated strongly (e.g., 31 out of 40 names for Figure 2a were syntactic). When the textbox had a wordy caption, users tended to seek a more concise name (so only 6 out of 40 names for Figure 2b were syntactic). Even when a caption was missing, however, the words on the page exerted some effect on users' naming (so 12 out of 40 names for Figure 2d were syntactic).

Part 2 found that users could flawlessly find the textbox associated with a visible name when the name was unambiguous. When a name was potentially ambiguous, users tended to resolve the ambiguity by choosing the first likely match found in a visual scan of the page. When the ambiguity was caused by both syntactic matching and semantic matching, however, users tended to prefer the syntactic match: given "City" as the target name for Go.com, 36 out of 40 users chose one of the two textboxes explicitly labeled "City"; the remaining 4 users chose the "Zip code" textbox, a semantic match that appears higher on the page.

The user's visual scan also did not always proceed from top to bottom; given "First Search" as the target name for eBay.com, most users picked the search box in the middle of the page, rather than the search box in the upper right corner.

Part 3's names were almost all syntactic (235 names out of 240), since the directions requested only words from the page. Even in syntactic naming, users rarely reproduced a caption exactly; they would change capitalization, transpose words (writing "web search" when the caption read "Search the Web"), and mistype words. Some Part 3 answers also included the type of the page object ("box", "entry", "field"). When asked to name a textbox which had an ambiguous caption (e.g. "Search" on a page with more than one search form), most users noticed the ambiguity and tried to resolve it with one of two approaches: either counting occurrences ("search 2") or referring to other nearby captions, such as section headings ("search products").

5. AUTOMATIC NAME RESOLUTION

We have used the names from Part 3 of the survey to develop a heuristic algorithm for resolving names to textboxes in Chickenfoot. Given a name and a web page, the output of the algorithm is a textbox on the page that best matches that name, or no result if no suitable match can be found. For reasons of space, we omit the details of the algorithm, but simply sketch its outline. The name is first normalized by tokenizing it on word boundaries and eliminating capitalization, punctuation, and white space. The page is then searched for text segments containing at least one of the words found in the name. Matching text segments are then ranked according to the number of name words they match (using edit distance for approximate matching, which is tolerant of a small number of typing errors). The highest-ranked text segment is then used to find a corresponding textbox by scanning down and to the right in the two-dimensional layout of the page.

The performance of this algorithm is shown in Figure 3. Each page was tested with the 40 names from Part 3 of the survey, and for each name Chickenfoot either found the right textbox (*Match*), found no matching textbox (*No Match*), or returned the wrong textbox (*Mismatch*). Precision is high for 5 of the 6 pages. An exception is the MIT task; precision is low on this task because it involved an ambiguous caption, and our heuristic algorithm does not yet recognize the disambiguation strategies used for this caption (counting and section headings). This evaluation is only preliminary, since a proper evaluation should use fresh data that wasn't used to drive the development of the algorithm, from a larger selection of web sites. Nevertheless, it suggests that names derived from visible labels can be automatically resolved with high precision.

6. CONCLUSION

We have shown preliminary results that syntactic naming (using words that are visible in the page) is a promising strategy for identifying elements in web pages. Web pages are just one kind of user interface that can be customized and automated; we anticipate that these results will generalize to other kinds of user interfaces as well, as long as the UI contains textual labeling.

Future work includes improving the Chickenfoot development environment so that ambiguous names can be disambiguated




Figure 3. Precision of automatic name resolution.

during code entry, which allows for an ambiguity resolution dialog between the user and the system that wouldn't be sensible at runtime. We are also looking at the robustness of syntactic names against change in web sites. Dealing with web site interfaces that change without warning is a major challenge for web automation, but as yet no one has characterized the kinds of changes that occur.

7. ACKNOWLEDGMENTS

We gratefully acknowledge all the pilot users who took our web survey, as well as Maya Dobuzhskaya, Vineet Sinha, and other members of the LAPIS group who provided valuable feedback on the ideas in this paper. This work was supported in part by the National Science Foundation under award number IIS-0447800. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

8. REFERENCES

- [1] Ekiwi, LLC. screen-scraper: solutions for web data extraction. <http://www.screen-scraper.com/>
- [2] Furnas, G.W., Landauer, T.K., Gomez, L.M., and Dumais, S.T. "The vocabulary problem in human-system communication." *Commun. ACM*, 30, 11 (Nov 1987).
- [3] Green, R. "How To Write Unmaintainable Code." <http://mindprod.com/unmain.html>
- [4] Kistler, T. and Marais, H. "WebL – a programming language for the Web." *Proc. WWW7*, 1998.
- [5] Kurlander, D. "Chimera: Example-based graphical editing." In Cypher, A., ed., *Watch What I Do: Programming By Demonstration*, pp 271–292. MIT Press, 1993.
- [6] Modugno, F., Corbett, A.T., and Myers, B.A. "Graphical Representation of Programs in a Demonstrational Visual SHell - An Empirical Evaluation." *ACM Trans. on Computer Human Interaction*, 4, 3, pp 276-308.
- [7] Mosel, A. *Tikki Tikki Tembo*. Holt & Co, 1968.
- [8] Müffke, F. "The Curl programming environment." *Dr. Dobbs' Journal*, Sept. 2001.
- [9] Nielsen, J. *Usability Engineering*. Academic Press, 1993.
- [10] Simonyi, C. and Heller, M. "The Hungarian Revolution," *BYTE*, 16, 8 (Aug. 1991).
- [11] Sun Microsystems. "Code Conventions for the Java Programming Language." <http://java.sun.com/docs/codeconv/>