

# Rewriting the Web with Chickenfoot

Robert C. Miller, Michael Bolin, Lydia Chilton, Greg Little,  
Matthew Webber, and Chen-Hsiang Yu

MIT Computer Science and Artificial Intelligence Laboratory

*(This chapter is a revision of a paper that originally appeared as M. Bolin, M. Webber, P. Rha, T. Wilson, and R.C. Miller, "Automation and Customization of Rendered Web Pages," UIST 2005.)*

The World Wide Web has become a preferred platform for many kinds of application development. Over the past decade, applications which formerly would have been designed for the desktop — calendars, travel reservation systems, purchasing systems, library card catalogs, map viewers, crossword puzzles, and even Tetris — have made the transition to the Web, largely successfully.

The migration of applications to the Web opens up new opportunities for user interface customization. Applications that would have been uncustomizable on the desktop sprout numerous hooks for customization when implemented in a web browser, without any effort on the application developer's part. Displays are represented primarily by machine-readable HTML or XML, navigation and commands are invoked by generic HTTP requests, and page designs can be reinterpreted by the browser and tailored by stylesheets. Unlike desktop applications, web applications are much more exposed and open to modification. Here are some of the customization possibilities that arise when an application is moved to the Web:

**Automating repetitive operations.** Web automation may include navigating pages, filling in forms, and clicking on links. 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 repetitive web browsing to download each paper, print it, and (later) upload a review for it. Tedious repetition is a good 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.

**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.

**Integrating multiple web sites.** The simplest kind of integration is just adding links from one site to another<sup>1</sup>, but much richer integration is possible. For example, many retailers' web sites incorporate maps and directions provided by a mapping service directly into their web pages, to display their store locations

---

1. Microsoft. "Smart Tags and Smart Documents." [msdn.microsoft.com/office/understanding/smarttags/default.aspx](http://msdn.microsoft.com/office/understanding/smarttags/default.aspx)

and provide driving directions. But end users have no control over this kind of integration. 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 is an example of a *mashup*, a combination of data or user interface from more than one site.

These examples involve not only *automating* web user interfaces (clicking links, filling in forms, and extracting data) but also *customizing* them (changing appearance, rearranging components, and inserting or removing user interface widgets or data). The openness and flexibility of the web platform enables customizations that would not have been possible on the desktop.

Previous approaches to web automation used a scripting language that dwells outside the web browser, such as Perl, Python, or WebL<sup>2</sup>. 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 to an agent running outside the web browser. Highly-interactive web applications like Google Mail and Google Maps — sometimes called AJAX applications<sup>3</sup> — because they use asynchronous JavaScript and XML — have made this situation worse.

But perhaps the most telling difference, and the most intimidating one to a user, is the simple fact that outside a web browser, a web page is just raw HTML. Even the most familiar web portal can look frighteningly complicated when viewed as HTML source.

Chickenfoot is a programming system we have developed that provides a platform for automating and customizing web applications through a familiar interface — as web pages rendered in a web browser. The challenge for Chickenfoot is simply stated: a user should not have to view the HTML source of a web page in order to customize or automate it.

Chickenfoot addresses this challenge in three ways. First, it runs inside the web browser, so that the rendered view of a web page is always visible alongside the Chickenfoot development environment. Second, its language primitives are concerned with the web page's user interface, rather than its internal details. For example, Chickenfoot uses commands like `click`, `enter`, and `pick` to interact with forms. Third, it uses novel pattern-matching techniques to allow users to describe components of a web page (targets for interaction, extraction, insertion, or customization) in terms that make sense for the rendered view. For example, `click` identifies the button to be clicked using keywords from its text label, rather than the name it was given by the web page designer. Figure 1 shows Chickenfoot in action, with a script for Google Image Search that uses some of these commands.

---

2. Kistler, T. and Marais, H. "WebL – a programming language for the Web." *WWW7*, 1998, pp. 259–270.

3. Garrett, J. "Ajax: A New Approach to Web Applications." [www.adaptivepath.com/publications/essays/archives/000385.php](http://www.adaptivepath.com/publications/essays/archives/000385.php)

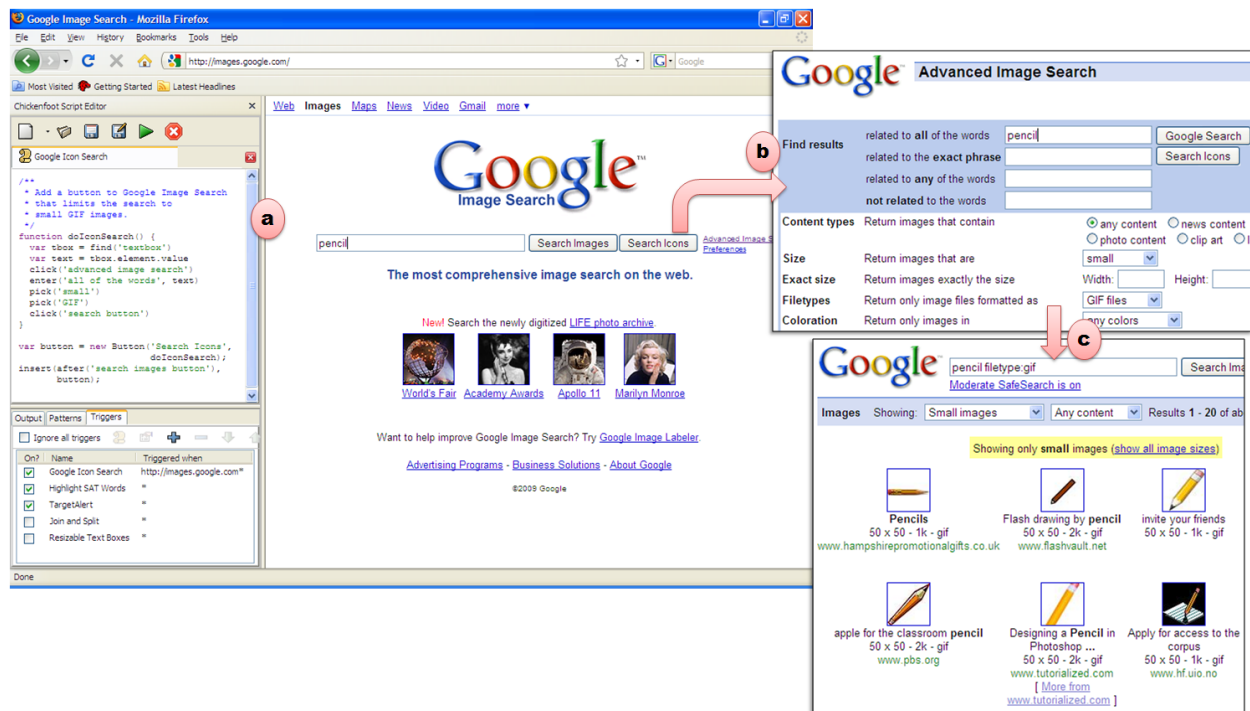


Figure 1. Chickenfoot in action. (a) Chickenfoot sidebar showing a complete script that customizes Google Image Search with a new "Search Icons" button, which (b) automatically fills out the Advanced Search form in order to (c) display only small GIF images matching the query.

Chickenfoot is implemented as an extension for the Mozilla Firefox web browser, written in Java, JavaScript, and XUL. It consists of a development environment, which appears as a sidebar of Firefox, and a library built on top of JavaScript. Chickenfoot customizations are essentially JavaScript programs, so Chickenfoot currently does not support nonprogramming users. We assume that a Chickenfoot user has *some* knowledge of JavaScript and HTML — not an unreasonable assumption, since many power users showed the ability and willingness to learn these during the explosive growth of the Web. The problem Chickenfoot is addressing is not learning JavaScript and HTML syntax, but rather reading and understanding the complex HTML used by today's web applications.

Naturally, many users would benefit from a web automation system that avoids the need to learn programming language syntax. We regard Chickenfoot as a step towards this goal, but a crucial one, since it provides a level of expressiveness and completeness unavailable in special-purpose web automation systems.

One system similar to Chickenfoot in implementation is Greasemonkey,<sup>4</sup> a Firefox extension that can run user-written JavaScript on web pages just after they are loaded in the browser. Though Greasemonkey enables mutation and automation of web pages in the browser, it still requires users to understand the HTML underlying the page. Platypus<sup>5</sup> is another Firefox extension, designed to work with

4. Greasemonkey. [greasemonkey.mozdev.org](http://greasemonkey.mozdev.org)

5. Platypus. [platypus.mozdev.org](http://platypus.mozdev.org)

Greasemonkey, that allows some customization of rendered web pages, but not automation or integration of multiple web sites.

The rest of this chapter is organized as follows. First we give an overview of the Chickenfoot language and development environment, and describe a range of applications we have built using Chickenfoot. Then we delve deeper into a novel aspect of Chickenfoot: the pattern matching used to identify web page elements for automation or customization. We describe a survey of web users that motivated the design of the pattern matching, and present the algorithm we developed as a result. Finally we discuss some experience and lessons learned from the Chickenfoot project, review related work, and conclude.

## Chickenfoot

Chickenfoot is an extension to the Mozilla Firefox web browser, consisting of a library that extends the browser's built-in JavaScript language with new commands for web automation, and a development environment that allows Chickenfoot programs to be entered and tested inside the web browser. This section describes the essential Chickenfoot commands, including pattern matching, form manipulation, page navigation, and page modification. The section concludes by describing the development environment (Figure 1a).

## Language

Chickenfoot programs are written in JavaScript, using the JavaScript interpreter built into Mozilla Firefox. As a result, Chickenfoot users have access to the full expressiveness of a high-level scripting language, with a prototype-instance object system, lexically scoped procedures, dynamic typing, and a rich class library.

Because Chickenfoot uses JavaScript, web developers can easily transfer their knowledge of JavaScript from web page development over to Chickenfoot. Chickenfoot predefines the same variables available to JavaScript in web pages – e.g., `window`, `document`, `location`, `frames`, `history` – so that JavaScript code written for inclusion *inside* a web page can also be used as a Chickenfoot script *outside* the web page. JavaScript has its own ways of visiting new pages (`location`), manipulating form controls (`document.forms`), and modifying page content (using the Document Object Model, or DOM<sup>6</sup>). These mechanisms can be used by Chickenfoot scripts in addition to the Chickenfoot commands described in the next few sections. We have found that using the native JavaScript mechanisms generally require reading and understanding a web page's HTML source. But by providing access to them, Chickenfoot provides a smooth escape mechanism for script developers that need to do something more low-level.

All current web browsers, including Firefox, implement a security model for JavaScript to protect web users from malicious downloaded scripts. A major part of this security model is the *same-origin policy*, which prevents JavaScript code downloaded from one web server from manipulating a web page downloaded from a different server. This restriction is clearly too severe for Chickenfoot, since its primary purpose is integrating and customizing multiple web sites. As a result, Chickenfoot scripts run at a privileged level, where they have access to the entire web browser, all pages it visits, and the user's

---

6. W3C. "Document Object Model (DOM)." [www.w3.org/DOM/](http://www.w3.org/DOM/)

filesystem and network. Users must trust Chickenfoot code as much as they trust any other desktop application. As a result, Chickenfoot scripts cannot be embedded in downloadable web pages like other JavaScript. But Chickenfoot code *can* inject new behavior into downloaded pages, which is explained in more detail later.

## Pattern Matching

Pattern matching is a fundamental operation in Chickenfoot. To operate on a web page component, most commands take a pattern describing that page component.

Chickenfoot supports three kinds of patterns: keywords, XPath, and regular expressions. A keyword pattern consists of a string of keywords that are searched in the page to locate a page component, followed by the type of the component to be found. For example, "Search form" matches a form containing the keyword Search, and "Go button" matches a button with the word "Go" in its label. The component type is one of a small set of primitive names, including `link`, `button`, `textbox`, `checkbox`, `radiobutton`, `listbox`, `image`, `table`, `row`, `cell`, and `text` (for text nodes). When a keyword pattern is used by a form manipulation command, the type of page component is often implicit and can be omitted. For example, `click("Go")` searches for a hyperlink or button with the keyword "Go" in its label. Case is not significant, so `click("go")` has the same effect.

An XPath pattern uses the XPath pattern matching syntax<sup>7</sup> to identify nodes in the HTML document. For example, the XPath pattern `"/b"` finds all bold elements in the page; `"/b/a"` finds hyperlinks inside bold elements; and `"/div[@class='box']"` finds `<div>` elements that have the class attribute "box". The XPath syntax is far more powerful and precise than keyword patterns, but harder to learn and use correctly, and often requires close examination of the page's HTML. XPath patterns are generally used to identify parts of a page for modification or extraction; form manipulation is often easier to do with keyword commands instead.

An early version of Chickenfoot also supported *text constraint* patterns, which combine a library of primitive patterns (such as `link`, `textbox`, or `paragraph`), literal strings (such as "Go"), and relational operators (e.g., `in`, `contains`, `just before`, `just after`, `starts`, `ends`). Text constraint patterns were used previously in LAPIS.<sup>8</sup> Text constraint patterns were removed from Chickenfoot because the implementation was very slow and unpredictable on the complex pages typical of today's Web. XPath can handle most of the tasks that text constraints do, at the cost of a steeper learning curve and greater complexity.

The third kind of pattern, regular expressions,<sup>9</sup> search the visible text content of the page, using Javascript's regular expression syntax and semantics. For example, `/[A-Z]\w+/"` finds capitalized words in the page. For the purpose of matching, the page content is divided into *blocks* by hard line breaks (`<br>` elements) and HTML-block-level element boundaries. The pattern is applied separately to each paragraph. Thus `/^\w+/"` matches words at the start of a block, and `/ISBN: .*/` will match only part of a block, not the entire page after the first occurrence of "ISBN".

---

7. W3C. "XML Path language (XPath) Version 1.0," 1999.

8. Miller, R.C. and Myers, B.A. "Integrating a Command Shell into a Web Browser." *USENIX*, 2000, pp. 171–182.

9. Friedl, J. *Mastering Regular Expressions*. O'Reilly, 2002.

Chickenfoot's `find` command takes one of these patterns and searches for it in the current page. A keyword pattern is represented simply as a string, while XPaths and regular expressions are represented by XPath and RegExp objects:

```
find("Search form") // finds a form with "Search" somewhere in it
find(new XPath("//b//a")) // finds hyperlinks in bold
find(/\w+/) // finds words in the page
```

`Find` returns a `Match` object which represents the first match to the pattern and provides access to the rest of the matches as a Javascript iterator. Here are some common idioms using `find`:

```
// test whether a pattern matches
if (find(pattern).hasMatch) { ... }

// count number of matches
find(pattern).count

// iterate through all matches
for each (m in find(pattern)) { ... }
```

A `Match` object represents a contiguous region of a web page, so it also provides properties for extracting that region. If `m` is a `Match` object, then `m.text` returns the plain text it represents, i.e., the text that would be obtained by copying that region of the rendered web page and pasting it to a text editor that ignored formatting. Similarly, `m.html` returns the source HTML of the region, which is useful for extracting the region with formatting intact. Finally, `m.element` returns the DOM `Element` object represented by the region, if the region contains a single outermost element. This `element` can be used to get and set element attributes, e.g.:

```
find("link").element.href
```

The `find` command is not only a global procedure, but also a method of `Match` (returned by a previous `find`) and `Tab` (which identifies a Firefox tab or window). Invoking `find` on one of these objects constrains it to return only matches within the page or part of a page represented by the object. Here are some common idioms:

```
// nested finds
for each (t in find("table")) {
  for each (r in t.find("row")) {
    ...
  }
}

// find in a different tab
otherTab.find(pattern)
```

## Clicking and Form Manipulation

The next few commands allow Chickenfoot scripts to click on hyperlinks and fill in and submit forms.

The `click` command takes a pattern describing a hyperlink or button on the current page and causes the same effect as if the user had clicked on it. For example, these commands click on various parts of the Google home page:

```
click("Advanced Search") // a hyperlink
click("I'm Feeling Lucky") // a button
```

Keyword patterns do not need to exactly match the label of the button or hyperlink, but they do need to be unambiguous. Thus, `click("Lucky")` would suffice to match the I'm Feeling Lucky button, but in this case, `click("Search")` would be ambiguous between the Google Search button and the Advanced Search link, and hence would throw an exception. Exact matches take precedence over partial matches, however, so if there were a single button labeled Search, then the `click` command would succeed.

Another way to eliminate ambiguity is by prefixing the pattern with an index, like "first", "second", "3rd", and so on, so

```
click("second Search")
```

will click on the second search button.

When a button or link has no text label, but an image instead, then the search can use keywords mentioned in their ALT text, if any. The keyword matching algorithm is described in more detail later in this chapter.

The `click` command can take a `Match` object instead of a pattern, if the button or hyperlink to be clicked has already been found. For example, to identify the I'm Feeling Lucky button using an XPath pattern, the user might write something like this:

```
click(new XPath("//input[@name='btnI']"))
```

This example illustrates why XPath is generally less readable than keywords.

The `enter` command enters a value into a textbox. Like `click`, it takes a pattern to identify the textbox, but in this case, the keywords are taken from the textbox's caption or other visible labels near the textbox. For example, to interact with the Amazon login page, a script might say:

```
enter("e-mail address", "rcm@mit.edu")
enter("password", password)
```

When the page contains only one textbox, which is often true for search forms, the keyword pattern can be omitted. For example, this sequence does a search on Google:

```
enter("EUP book")
click("Google Search")
```

Checkboxes and radio buttons are controlled by the `check` and `uncheck` commands, which take a keyword pattern that describes the checkbox:

```
check("Yes, I have a password")
uncheck("Remember Me")
```

The `pick` command makes a selection from a listbox or drop-down box (which are both instantiations of the HTML `<select>` element). The simplest form of `pick` merely identifies the choice by a keyword pattern:

```
pick("California")
```

If only one choice in any listbox or drop-down on the page matches the keywords (the common case), then that choice is made. If the choice is not unique, `pick` can take two keyword patterns, the first identifying a listbox or dropdown by keywords from its caption, and the second identifying the choice within the listbox:

```
pick("State", "California")
```

Like `find`, all the clicking and form manipulation commands are also methods of `Match` and `Document`, so that the keyword search can be constrained to a particular part of a page:

```
f = find("Directions form")
f.enter("address", "32 Vassar St")
f.enter("zip", "02139")
f.click("Get Directions")
```

The form manipulation commands described so far permit setting the value of a form widget, which is the most common case for web automation. To read the current value of a widget, a script can use `find` to locate it, and then access the value of its `Element` object in conventional JavaScript fashion, e.g.:

```
find("address textbox").element.value
```

Chickenfoot also provides a `reset` command which resets a form to its default values, though `reset` is rarely needed.

Finally, Chickenfoot can also simulate more low-level mouse and keyboard input. The `click` command can simulate a mouse click on any page component. By default it searches only through buttons and hyperlinks, but this default can be overridden by specifying a component type. For example, instead of using `check` or `uncheck` to set a checkbox, one can toggle it with `click`:

```
click("Remember Me checkbox")
```

Similarly, the `keypress` command sends low-level keystrokes to a particular page component. Unlike `enter`, which takes a string of text to fill a textbox, `keypress` takes a sequence of keywords describing individual keypresses. Valid keywords include characters (a, A, %, etc.), key names (enter, fl, home, tab, backspace, etc.), and modifier keys (shift, ctrl, control, etc.). The key sequence is interpreted from left to right, and modifier keys affect the first character or key name that follows them, e.g. "ctrl x c" produces Control-X followed by "c". Examples include:

```
keypress("password", "enter") // presses Enter on the password textbox
keypress("ctrl s") // sends Ctrl-S to the only textbox on the page
keypress("E U P space b o o k") // types "EUP book" in the only textbox
```



The last `keypress` command above is almost equivalent to `enter("EUP book")`, except that `enter` discards the old contents of the textbox if any, while `keypress` inserts the keystrokes wherever the textbox's selection happens to be.

## Navigation and Page Loading

In addition to loading pages by clicking on links and submitting forms, Chickenfoot scripts can jump directly to a URL using the `go` command:

```
go("http://www.google.com")
```

If the string is not a valid URL, `go` automatically prefixes it with `"http://"`.

Conventional browser navigation operations are also available as Chickenfoot commands: `back`, `forward`, and `reload`.

The `openTab` command opens a page in a new tab, returning a `Tab` object representing the tab:

```
google = openTab("www.google.com")
```

To retrieve a page without displaying it, the `fetch` command can be used. It returns a `Tab` object representing the invisible page:

```
google = fetch("www.google.com")
```

The JavaScript `with` statement is convenient for performing a sequence of operations on a different tab, by implicitly setting the context for Chickenfoot pattern matching and form manipulation:

```
with (fetch("www.google.com")) {  
    enter("syzygy")  
    click("Google Search")  
}
```

A tab created with `openTab` and `fetch` can be brought to the front by calling its `show` method and closed with its `close` method. The `Tab` object representing the current tab is stored in a readonly variable, `tab`.

Pages retrieved by `go`, `openTab`, `fetch`, and `click` are loaded asynchronously by the browser, while the Chickenfoot script continues to run. Thus, a script can fire off several `fetch` requests in parallel, without forcing each request to complete before the next one starts. When a subsequent Chickenfoot command needs to access the content of a page, such as `find`, the command automatically blocks until the page is fully loaded. The `wait` and `ready` commands make this blocking available to programmatic control. Both commands take a `Tab` object or an array of `Tab`s as an argument. With no arguments, the default is the current tab. `Wait` blocks until at least one of the specified tabs is fully loaded, and returns that tab. `Ready` returns a loaded tab only if it has already completed, otherwise it immediately returns `null`.

## Page Modification

Chickenfoot offers three primitive commands for changing the content of web pages: `insert`, `remove`, and `replace`.

The `insert` command takes two arguments: a location on a page and a fragment of web page content that should be inserted at that location. In its simplest form, the location is a keyword pattern, and the web page content is simply a string of HTML. Since insertion must use a single point in the page, not a range of content, the `before` or `after` commands must be used to reduce the pattern match to a point:

```
insert (before ("textbox"), "<b>Search: </b>")
```

The location can also be a `Match` object:

```
t = find("textbox")
insert (after (t), "<b>Search: </b>")
```

The page content to be inserted can also be a `Match` or `Node`, allowing content to be extracted from another page and inserted in this one:

```
map = mapquest.find("image") // where mapquest is a tab showing a map
insert (after ("Directions"), map)
```

The `remove` command removes page content identified by its argument, which can be a pattern or `Match` object. For example:

```
remove ("Google image")
```

The `replace` command replaces one chunk of page content with another. It is often used to wrap page content around an existing element. For example, the following code puts every number in the page in boldface:

```
for each (num in find(/\d+/)) {
    replace (num, "<b>"+num.text+"</b>")
}
```

## Widgets

When a Chickenfoot script needs to present a user interface, it can create links and buttons and insert them directly into a web page. Hyperlinks are created by the `Link` constructor, which takes a chunk of HTML to display inside the hyperlink and an event handler to run when the link is clicked:

```
new Link("<b>Show All</b>", showAll)
```

The event handler should be a Javascript function object. Buttons are created similarly by the `Button` constructor.

Other widgets can be created by inserting HTML, e.g.:

```
insert(..., "<input type=checkbox>")
```

If an `onclick` attribute is included in this HTML element, however, the code it contains will execute like conventional, untrusted JavaScript code downloaded with the page. Commands defined by Chickenfoot would be unavailable to it. Instead, the right way to add a Chickenfoot handler to a button is to use the `onClick` command:

```
onClick(button, showAll)
```

Keyboard handlers can also be attached using the `onKeypress` command, which takes an optional page component, a key sequence (specified using words as in the `keypress` command) and a handler function:

```
onKeypress("ctrl s", saveWikiPage)
onKeypress("username textbox", "tab", fillInPassword)
```

## Development Environment

Figure 1a shows a screenshot of the development environment provided by Chickenfoot, which appears as a sidebar in Firefox. At the top of the sidebar is a text editor used to compose Chickenfoot code. The editor supports multiple tabs editing different script files. This simple interface goes a long way toward making the JavaScript interpreter embedded in every web browser more accessible to the user. Previously, there 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. The Chickenfoot sidebar makes it much easier for an end-user to write and run scripts. Recent web debugging tools, such as Firebug,<sup>10</sup> now include multiline script editors for a similar reason.

The bottom of the sidebar has three tabbed panels. The Output panel is an output window that displays output printed by the Chickenfoot `output` command, error messages and stack traces, and the result of evaluating a Chickenfoot script (i.e., the value of the last expression). If the Record Actions checkbox is enabled, then the Output panel also displays Chickenfoot commands corresponding to the manual actions that the user does in the browser, like `click("Search button")` or `enter("Username", "rcm")`. This feature enables some degree of self-disclosure.<sup>11</sup> A manual browsing sequence can be used as the basis for a script by recording it into the Output panel and copying it into a script.

The Patterns panel displays an interface for developing text constraint patterns, which allows the user to enter a pattern and highlight what it matches in the current page. This pane also displays the types of page components that can be used for keyword commands (`link`, `button`, `textbox`, etc.)

The Triggers panel allows a Chickenfoot script to be installed into the browser for regular use. For automatic invocation, a script can be associated with a *trigger*, which is a URL pattern such as `http://www.amazon.com/*`. Whenever a page is loaded, if its URL matches a trigger, then the associated script executes automatically. If a page matches multiple triggers, the associated scripts

---

10. Firebug -- Web Development Evolved. [www.getfirebug.com](http://www.getfirebug.com)

11. DiGiano, C. and Eisenberg, M. "Self-disclosing design tools: a gentle introduction to end-user programming." *DIS '95*, pp. 189–197.

execute in the fixed order given by the Triggers pane. Trigger scripts can also be invoked manually by right-clicking on a page and selecting Run Trigger on This Page. The Triggers panel provides an interface for adding and removing trigger scripts and enabling and disabling triggers.

Finally, the Triggers panel provides a Package Script As Firefox Extension command, which binds one or more trigger scripts and supporting files, together with an embedded version of the Chickenfoot library, and exports the result as a Firefox extension (XPI) file. Any Firefox user can install this extension in their browser in order to use the triggers, without having to install Chickenfoot first. Thus customizations developed in Chickenfoot can be easily exported as full-fledged Firefox extensions.

## Applications

This section describes a few of the example scripts we have created using Chickenfoot.

### Adding File Type Icons to Links

When a hyperlink points at a resource other than a web page (such as a PDF document, a ZIP archive, or a Word document), it is often helpful for the link to be visually distinguished – first, because the user may be actively scanning the page for one of these resources, and second, because they may want to avoid them while casually surfing. Only a few web sites provide a visual cue to the file type of a hyperlink. TargetAlert<sup>12</sup> is a Firefox browser extension, developed by one of us before Chickenfoot was created, that adds file type icons to hyperlinks on any web site.

TargetAlert was originally written in 217 lines of JavaScript and XUL. We rewrote it in 29 lines of Chickenfoot. The essence of the script is the following loop:

```
for each (l in find("link")) {
  href = l.element.href
  if (m = href.match(/\.(\\w+)$/)) {
    extension = m[1]
    src = "moz-icon://." + extension + "?size=16"
    insert(after(l), "<img src='" + src + "'>")
  }
}
```

To simplify retrieving file type icons, the script exploits a feature of Firefox that works only on Windows: URLs of the form `moz-icon://.ext?size=16` return the icon associated with file extension `.ext` in the Windows registry. Firefox normally uses these URLs to display local directories in the browser. The result of running this script on part of a web page is shown in Figure 2.

---

12. TargetAlert. [www.bolinfest.com/targetalert](http://www.bolinfest.com/targetalert)




Figure 2. File type icons added to links in a web page.

## Sorting Tables

Another feature that some web sites have, but many lack, is the ability to sort a table of data by clicking one of its column headers. A Chickenfoot script can add this functionality automatically to most tables, by replacing every table header cell it finds with a link that sorts the table.

Most of the script is concerned with managing the sort, but here is the part that replaces headers with links:

```
column=0
for each (h=table.find(new XPath("//th//@text"))) {
    f = makeRowSorter(table,column++)
    replace(h, new Link(h, f))
}
```

The `makeRowSorter` function returns a function that sorts the specified table by the specified column number.

## Concatenating a Sequence of Pages

Search results and long articles are often split into multiple web pages, mainly for faster downloading. This can inhibit fluid browsing, however, because the entire content isn't accessible to scrolling or to the browser's internal Find command. Some articles offer a link to the complete content, intended for printing, but this page may lack other useful navigation.

We have written a Chickenfoot script that detects a multi-page sequence by searching for its table of contents (generally a set of numbered page links, with Next and Previous). When a table of contents is found, the script automatically adds a "Show All" link to it (Figure 3). Clicking this link causes the script to start retrieving additional pages from the sequence, appending them to the current page. In order to avoid repeating common elements from subsequent pages (such as banners, sidebars, and other decoration), the script uses a conservative heuristic to localize the content, by searching for the smallest HTML element that contains the list of page links (since the content is nearly always adjacent to this list) and spans at least half the height of the rendered page (since the content nearly always occupies the majority of the page). The content element from each subsequent page is inserted after the content element of the current page.




Figure 3. "Show All" link injected for page concatenation.

## Highlighting Vocabulary Words

Students studying for college placement exams, such as the SAT, often work hard to expand their vocabulary. One way to make this learning deeper is to highlight vocabulary words while the student is reading, so that the context of use reinforces the word's meaning. One of our Chickenfoot scripts takes a list of vocabulary words and definitions (posted on the web) and automatically highlights matching words in any page that the user browses. The script uses a `title` attribute to pop up the word's definition as a tooltip if the mouse hovers over it:

```
for each (w=find(/\w+/)) {
  if (w.text in vocab) {
    html = "<span style='background-color: yellow' title='"
          + vocab[w.text] + "'>"
          + w + "</span>"
    replace(w, html)
  }
}
```

Those who are doing most of the complaining, I do  
 they are (not) deliberately striving to sabotage the  
 effort. They are laboring under the delusion that th  
 when we must make prodigious sacrifices -- that t  
 already won and we can be prodigious But  
 folly of that point of view can be measured by the  
 separates our troops from their ultimate objectives  
 Tokyo -- and by the sum of all the perils that lie al

Figure 4. Vocabulary word highlighting.

## Integrating a Bookstore and a Library

This example is a short script that augments book pages found in Amazon with a link that points to the book's location in the MIT library:

```
isbn = find(/ISBN.*: *(\d+)/).range[1]
with (fetch("libraries.mit.edu")) {
  // run on MIT library site
  pick("Keyword")
```

```

    enter(isbn)
    click("Search")
    link = find("Stacks link")
}

// now back to Amazon
if (link.hasMatch) {
    insert(before("Available for in-store pickup"), link.html)
}

```



Figure 5. A link to a local library injected into a bookstore web page.

## Making a Bug Tracker More Dynamic

For issue tracking in the Chickenfoot project, we use the Flyspray bug tracker<sup>13</sup>, which has the advantage of being free and simple. One downside of Flyspray is that some of its interfaces are inefficient, requiring several clicks and form submissions to use. For example, when we meet to discuss open bugs, we often use Flyspray's list view (figure 6a), which displays bugs in a compact tabular form. But changing a bug's properties from this view -- particularly, its priority, severity, or assigned developer -- requires five clicks and waiting for three different pages to load. What we want instead is direct manipulation of the list view, so that bug properties can be edited immediately without leaving the view.

---

13. Flyspray - The Bug Killer! flyspray.org

ID	Category	Severity	Priority	Summary
495	Bug	High	Immediate	can't click on javascript: links
178	Usability	Low	High	Open File should default to Chickenfoot profile directory.
219	Documentation	Medium	High	Examples on the web site are broken
478	Documentation	Low	High	SAT word example is broken
493	New feature	Low	High	dynamically-scoped withTab
500	Bug	Critical	Immediate	FF2 native libraries break 64-bit Ubuntu
476	Platform bug/Plugin conflict	High	Urgent	Load ChickenSleep manually
411	Usability	Medium	Normal	Recorder should be on all the time
412	Bug	Medium	Normal	Recorder records lots of spurious commands

Figure 6a. Flyspray bug summary display.

495	Bug	High	Immediate	can't click on javascript: links
178	Usability	Low	High	Open File should default to Chickenfoot profile directory.
219	Documentation	Medium	Flash	Examples on the web site are broken
478	Documentation	Low	Immediate	SAT word example is broken
493	New feature	Low	Urgent	dynamically-scoped withTab
500	Bug	Low	High	FF2 native libraries break 64-bit Ubuntu
476	Platform bug/Plugin conflict	Critical	Normal	Load ChickenSleep manually
411	Usability	High	Low	Recorder should be on all the time

Figure 6b. Flyspray augmented by a Chickenfoot script so that bug properties can be edited in place.

We added this behavior to Flyspray using a Chickenfoot script, which replaces values in the table with a dropdown menu (figure 6b). The script automatically discovers the available choices for each dropdown menu by opening the edit page for some bug (invisibly, using `fetch`) and extracting the HTML for each of its dropdown menus:

```

firstBug = find(new XPath("//td[@class='project_summary']/a"))
with (fetch(firstBug.element.href)) {
    click("Edit this task")
    priorityMenu = find("Priority listbox").html
    severityMenu = find("Severity listbox").html
    assignedToMenu = find("Assigned To listbox").html
    ... // for other fields
    close() // closes the fetched edit page
}

```

The script then attaches an `onClick` handler to each cell of the table, so that when a cell is clicked, the value is replaced with the appropriate dropdown menu for that column (figure 6b). If the user changes the value of the property using the dropdown menu, then the script drills into the edit page for that bug (again, invisibly) and changes it:

```

with(fetch(href)) {
    click("Edit this task")
}

```



```

    pick("assigned to", newValue)
    click("Save Details")
    wait() // make sure changes are saved before closing tab
    close()
}

```

The overall effect of this script is to make Flyspray, a traditional web application driven entirely by form submission, feel more like a modern AJAX-driven web application, which responds immediately and updates the server in the background.

It's also worth considering why we used Chickenfoot in this case. Unlike the previous examples, which modified third-party web sites, we control our own installation of Flyspray, which is an open source system written in PHP. Technically, we could modify Flyspray ourselves to add this feature. The main reason we chose not to was the complexity barrier of understanding and using Flyspray's internal interfaces, which we had never seen since we didn't originally develop it. With Chickenfoot, however, the interface we needed to understand and use was the Flyspray *user* interface, which was simple and already very familiar to us as users of Flyspray. Writing the Chickenfoot script took less than an hour; delving into Flyspray and making it AJAX-capable would have been considerably longer.

## Adding Faces to Webmail

Usable security is an area of active research that aims to make computer security easier to use, and use correctly, for end users.<sup>14</sup> One interesting and underexplored problem in this area is *misdirected email*, such as pressing Reply All when you meant to press Reply, or mistyping an email address. Sending email to the wrong recipients is a security error — an access control error in fact, since it inadvertently grants access to the email's contents to a wider group of people than intended. Even secure email suffers from this error, since even digitally-signed and encrypted messages can be sent, securely, to the wrong people.

Using Chickenfoot, we have devised a novel solution to this problem that displays the faces of the recipients directly in the mail composition window of a webmail system. The faces appear automatically as the user is composing the message, acting as a peripheral interface that informs the user without demanding direct attention. Figure 7 shows our interface, called Kangaroo, extending the IMP webmail system. Our studies have found that the face display makes a substantial difference in users' ability to detect whether an email is misdirected.<sup>15</sup>

---

14. L.F. Cranor and S. Garfinkel, eds. *Security and Usability: Designing Secure Systems that People Can Use*. O'Reilly, 2005.

15. E. Lieberman and R. C. Miller. "Facemail: Showing Faces of Recipients to Prevent Misdirected Email." *SOUPS 2007*, pp. 122-131.

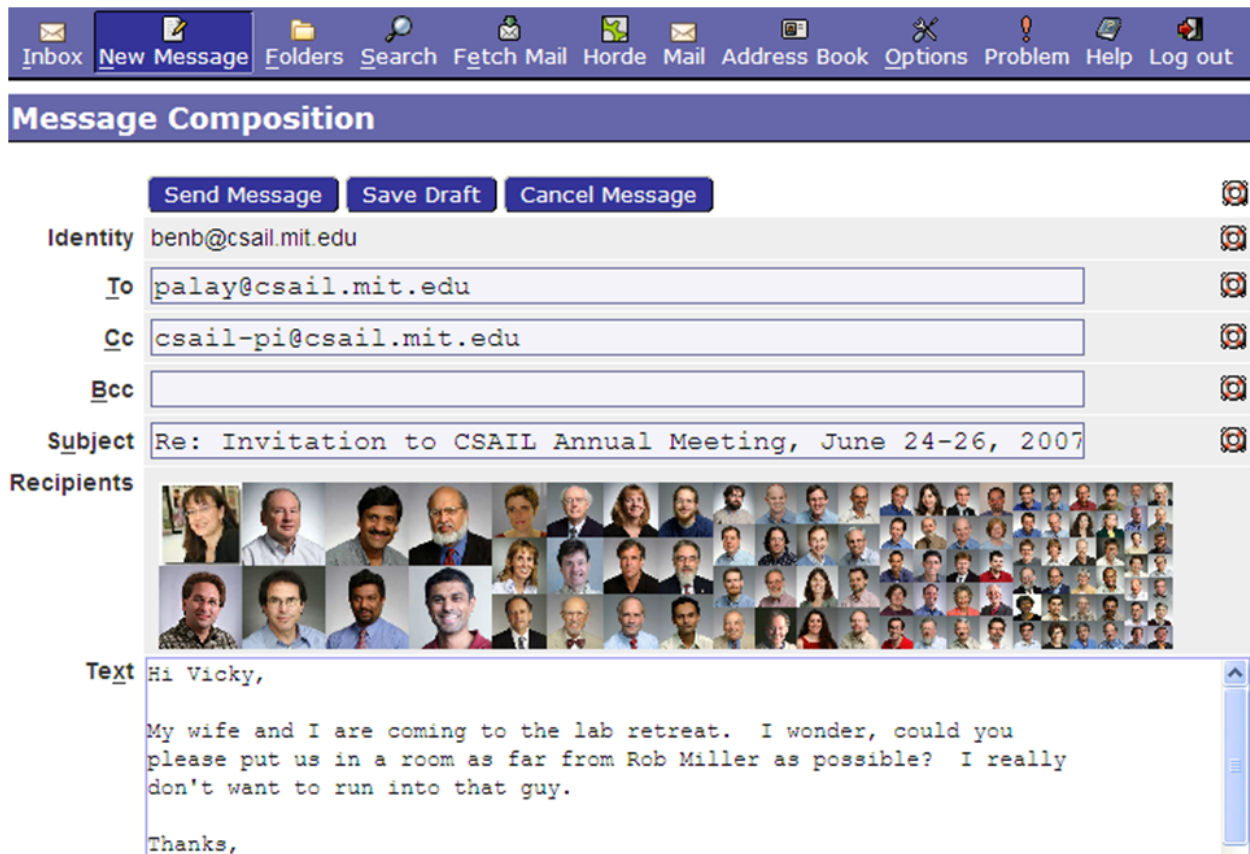


Figure 7. Kangaroo recipient-visualization interface, augmenting the IMP webmail system. In this example, the user pressed Reply-All by mistake, so the crowd of faces is a highly salient reminder about who will read the message.

Kangaroo is implemented as a collection of Chickenfoot scripts. One script triggers whenever a mail composition page appears, in order to automatically add the Face section to the page. This script triggers on URLs used by GMail and IMP (the webmail system MIT uses), but it could be readily extended to other webmail systems. The script also attaches event handlers to the To, Cc, and Bcc boxes on the page, so that Kangaroo is notified when the user edits the recipient list. The same simple keyword patterns (e.g. "To textbox") are sufficient to identify these controls across the three webmail systems Kangaroo currently supports, so the porting costs to new webmail systems are likely to be minimal.

A second set of scripts in Kangaroo are concerned with finding faces to display. The database that maps between email addresses and faces is cached locally in the user's web browser, rather than at a server. This database is populated automatically by Chickenfoot scripts that automate searches of various web services, including Google Images and Facebook. Another set of scripts expand mailing list addresses into their subscribers wherever possible, using web interfaces for Mailman and Athena Blanche (a mailing list system used internally at MIT). New scripts can be added to support new sources of faces and new kinds of mailing lists. For example, to use Kangaroo in a corporate email environment, the corporate ID database and mailing list server could readily be connected.

Kangaroo is an example of a substantial and interesting application built on top of Chickenfoot that embodies several of Chickenfoot's key ideas: *changing* existing web applications (in this case GMail), *automating* web services (like Mailman and Facebook), and programming at the level of *rendered user interfaces* ("To textbox") rather than APIs. We also use Chickenfoot's trigger-packaging functionality to make Kangaroo downloadable as a standalone Firefox extension.<sup>16</sup>

In addition to Kangaroo, our research group has used Chickenfoot as a prototyping tool for several other research systems, including minimal-syntax programming,<sup>17</sup> bookmarks for modern AJAX sites,<sup>18</sup> a command line for web applications,<sup>19</sup> and a tool that improves page readability for nonnative English readers<sup>20</sup>. Chickenfoot has proven a flexible and useful tool for exploring new ideas in web user interface.

## User Study of Keyword Matching

One of the novel aspects of Chickenfoot is the use of keyword patterns to identify page elements, such as "Search button" and "address textbox". To explore the usability of this technique, we conducted a small study to learn what kinds of keyword patterns users would generate for one kind of page component (textboxes), and whether users could comprehend a keyword pattern by locating the textbox it was meant to identify.

### Method

The study was administered over the Web. It 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 study 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 the diversity of textbox labeling seen across the Web, including simple captions (Figure 8a), wordy captions (b), captions displayed as default values for the

---

16. Kangaroo. [uid.csail.mit.edu/kangaroo](http://uid.csail.mit.edu/kangaroo)

17. Greg Little and Robert C. Miller. "Translating Keyword Commands into Executable Code." *UIST 2006*, pp. 135-144.

18. Darris Hupp and Robert C. Miller. "Smart Bookmarks: Automatic Retroactive Macro Recording on the Web." *UIST 2007*, pp. 81-90.

19. Robert C. Miller, Victoria Chou, Michael Bernstein, Greg Little, Max Van Kleek, David Karger, and mc schraefel. "Inky: A Sloppy Command Line for the Web with Rich Visual Feedback." *UIST 2008*, pp. 131-140.

20. Froggy. [people.csail.mit.edu/chyu/uid/froggy](http://people.csail.mit.edu/chyu/uid/froggy)

textbox (c), and missing captions (d). Several of the pages also posed ambiguity problems, such as multiple textboxes with similar or identical captions.

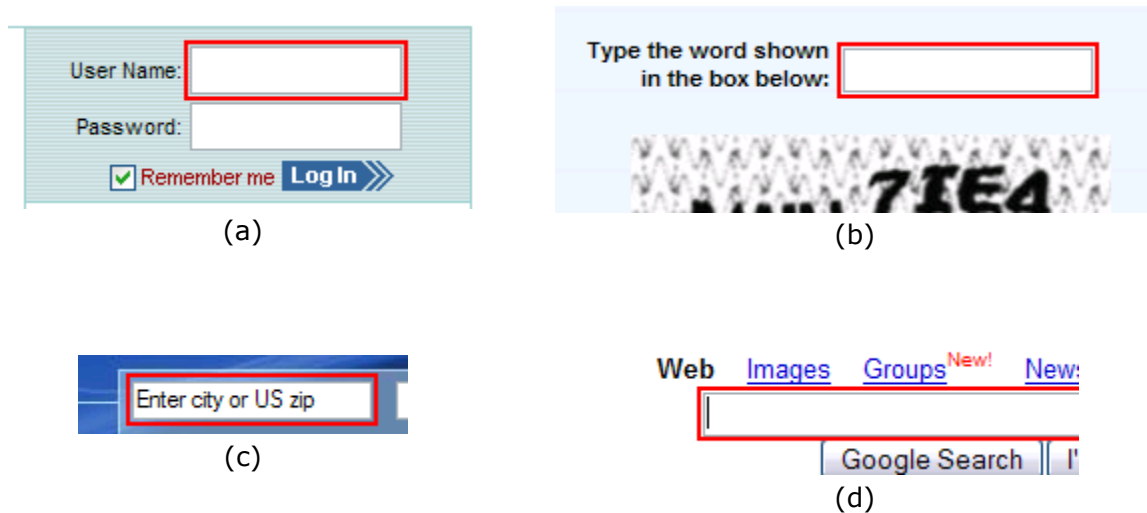


Figure 8. Some of the textboxes used in the web survey.

Subjects were unpaid volunteers recruited from the university campus by mailing lists. Forty subjects participated (20 females, 20 males), 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 that they used the Web at least several times a week.

## Results

We analyzed Part 1 by classifying each name generated by a user into one of four categories: (1) *visible* if the name used only words that were visible somewhere on the web page (e.g., "User name" for Figure 8a); (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 used 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, but since the study 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 *visible* or *semantic* names. When the textbox had an explicit, concise caption, *visible* names dominated strongly (e.g., 31 out of 40 names for Figure 8a were visible). When the textbox had a wordy caption, users tended to seek a more concise name (so only 6 out of 40 names for

Figure 8b were visible). 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 8d were visible).

Part 2 found that users could flawlessly find the textbox associated with a visible name, as long as 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 visible matching and semantic matching, however, users tended to prefer the visible 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 tucked away in the upper right corner.

Part 3's names were almost all visible (235 names out of 240), since the directions requested only words from the page. Even in visible naming, however, 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").

## Keyword Matching Algorithm

We now describe the heuristic algorithm that resolves a keyword pattern to a web page component, such as a textbox. Given a name and a web page, the output of the algorithm is one of the following: (1) a textbox on the page that best matches that name; (2) *ambiguous match* if two or more textboxes are considered equally good matches; or (3) *no match* if no suitable match can be found.

The first step is to identify the text labels in the page that approximately match the provided name, where a *label* is a visible string of content delimited by block-level tags (e.g. <p>, <br>, <td>). Button labels and alt attributes on images are also treated as visible labels. Before comparison, both the name and the visible labels are normalized by eliminating capitalization, punctuation, and white space. Then each label is searched for an approximate occurrence of the name, using a conventional edit distance algorithm to tolerate typos and omitted words. Matching labels are ranked by edit distance, so that closer matches are ranked higher.

For each matching label, we search the web page for textboxes that it might identify. Any textbox that is roughly aligned with the label (so that extending the textbox area horizontally or vertically would intersect the label's bounding box) is paired with the label to produce a candidate (*label, textbox*) pair.

These pairs are further scored by several heuristics that measure the degree of association between the label and the textbox. The first heuristic is pixel distance: if the label is too far from the textbox, the pair is eliminated from consideration. Currently, we use a vertical threshold of 1.5 times the height of the textbox, but no horizontal threshold, since tabular form layouts often create large horizontal gaps between captions and their textboxes. The second heuristic is relative position: if the label appears below or to the right of the textbox, the rank of the pair is decreased, since these are unusual places for a caption. (We

don't completely rule them out, however, because users sometimes use the label of a nearby button, such as "Search", to describe a textbox, and the button may be below or to the right of the textbox.) The final heuristic is distance in the document tree: each *(label, textbox)* pair is scored by the length of the shortest path through the DOM tree from the *label* node to the *textbox* node. Thus labels and textboxes that are siblings in the tree have the highest degree of association.

The result is a ranked list of *(label, textbox)* pairs. The algorithm returns the textbox of the highest-ranked pair, unless the top two pairs have the same score, in which case it returns *ambiguous match*. If the list of pairs is empty, it returns *no match*.

## Discussion

This section reflects on some of the lessons learned and surprises encountered over several years of developing and using Chickenfoot, and mention some of the open areas of future work.

A guiding design principle, mentioned in the introduction of this chapter, was that Chickenfoot users should not have to look at the HTML source of a web page in order to automate and customize it. This is a hard goal to achieve, and Chickenfoot only gets some of the way there. Some success can be seen in the Kangaroo system mentioned earlier, which uses keyword patterns effectively to customize webmail systems in a portable way. But in other cases, getting a good, unambiguous keyword pattern for some web page controls can be very hard, so users fall back on less robust counting patterns ("third textbox") or XPath. It's not uncommon for users to use a web page inspector, like Firebug,<sup>21</sup> to help analyze and understand the web page, and then take what they learned to write a Chickenfoot script.

An aspect of Chickenfoot that turned out to be much more important, however, was the *synchronous, page-independent* nature of Chickenfoot scripts, relative to ordinary Javascript or Greasemonkey. A Chickenfoot script can be written as a simple, straight-line sequence of commands, even if those commands end up affecting different pages. The Chickenfoot interpreter automatically handles waiting for pages to load, which substantially reduces the complexity burden on the programmer. In ordinary Javascript, this simple script might need to be torn apart into many little pieces in event handlers.

Although originally designed for end-user customization of web applications, Chickenfoot was put to a few uses that we didn't anticipate, but that turned out to be very fruitful. One use already mentioned was as a testbed for our own research prototypes. Some of these prototypes customize the *browser itself* with new toolbar buttons or menu items, not just changing the web sites it displays. This is made possible by two incidental facts: first, Firefox's user interface is mostly written in XUL (an XML-based user interface language) and Javascript, which is very similar to the HTML/Javascript model used in web sites; and second, Chickenfoot scripts run with the same privileges as browser-level Javascript, so they have permission to modify XUL and attach event listeners to it. Thus Chickenfoot can be used to prototype full-fledged Firefox extensions. One area of ongoing work is to generalize Chickenfoot's library of commands and patterns so that it works as well on XUL as it does on HTML.

Another unexpected but popular use of Chickenfoot on the web was for functional web testing, i.e., scripts that exercise various functions of a web application to check that they work correctly. Unlike the scenario Chickenfoot was envisioned for, in which end users customize web sites they did not themselves

---

21. Firebug -- Web Development Evolved. [www.getfirebug.com](http://www.getfirebug.com)

create, functional web testing is done by the web site's own developers, who have deep knowledge of the site and its internal interfaces. Chickenfoot turned out to be a useful tool for these users nevertheless, for several reasons. First, it was based on Javascript, so it was familiar to them, and it could directly access and call Javascript interfaces inside the web page. Second, a testing script written at the level of the rendered interface (e.g. `click("search button")`) is decoupled from the internal implementation of that interface, allowing the internal interfaces to change without having to change the testing code. And finally, Chickenfoot's synchronous programming model, which automatically blocks the script while pages load, makes it easy to write test scripts.

Turning now to problems for future work, Chickenfoot's current commands and pattern language seem well designed for *automation*, but are still weak for *data extraction*, also called web scraping. The keyword pattern approach should be applicable here, for identifying captioned data (like "List Price: \$27.95") and tabular data with labeled rows and columns. Currently, most data extraction in Chickenfoot is done with XPath or regular expressions.

Complex or large web pages are hard to automate efficiently and effectively with Chickenfoot. For example, nested `<iframe>` elements (which embed other HTML documents into a web page) cause problems for XPath pattern matching, because Chickenfoot uses the XPath matcher built into Firefox, which expects only one tree with one root, and which cannot dig into an `<iframe>`. For keyword patterns and regular expressions, Chickenfoot searches all embedded documents, but then the problem becomes efficiency. Chickenfoot is written in Javascript, and traversing large DOM trees is very slow using Firefox's current Javascript engine. But browser developers are devoting significant resources to improving the performance of Javascript and DOM manipulation<sup>22</sup>, since they are so critical to modern web applications, so we can anticipate that this situation will improve in the future.

A final question concerns security. Can a malicious web site do harm to a user who is using Chickenfoot scripts? Can a malicious Chickenfoot script be installed, possibly as a trojan horse that pretends to do something useful? The answer to both questions is probably yes, although no cases of either have been observed in the wild so far. Chickenfoot's design includes some decisions aimed at mitigating security risks. For example, Chickenfoot code can't be embedded in web pages like ordinary Javascript, and Chickenfoot puts none of its own objects or interfaces into the namespace seen by ordinary web page Javascript. For comparison, Greasemonkey's architecture had to be rewritten to fix security holes related to the fact that Greasemonkey scripts were injected into the same namespace as the remote, untrusted Javascript.<sup>23</sup>

Chickenfoot also benefits from features of Firefox designed to make general Firefox extension development safer. For example, when a Chickenfoot script obtains a reference to an object from the web page (like the Document object, or Window, or a Node), the reference is hidden by a security wrapper called XPCNativeWrapper, which protects against certain kinds of attacks. But Chickenfoot code runs with full privileges, like other Firefox extensions or indeed like any executable run on the desktop, so it should be subjected to a high level of scrutiny. A thorough and detailed security audit of Chickenfoot remains to be done.

---

22. R. Paul, "New Firefox JavaScript engine is faster than Chrome's V8," September 3, 2008. [arstechnica.com/open-source/news/2008/09/new-firefox-javascript-engine-is-faster-than-chromes-v8.ars](http://arstechnica.com/open-source/news/2008/09/new-firefox-javascript-engine-is-faster-than-chromes-v8.ars)

23. M. Pilgrim. *Greasemonkey Hacks: Tips & Tools for Remixing the Web with Firefox*. O'Reilly, 2005.

## Related Work

Several systems have addressed specific tasks in web automation and customization, including adding links<sup>24</sup>, building custom portals<sup>25</sup>, crawling web sites<sup>26</sup>, and making multiple alternative queries<sup>27,28</sup>. Chickenfoot is a more general toolkit for web automation and customization, which can address these tasks and others as well.

One form of general web automation can be found in scripting language libraries such as Perl's WWW::Mechanize or Python's ClientForm. These kinds of scripts run outside the web browser, where they cannot easily access pages that require session identifiers, secure logins, cookie state, or client-side JavaScript to run.

In an attempt to access these "hard-to-reach pages",<sup>29</sup> some systems give the user the ability to record macros in the web browser, where the user records the actions taken to require access to a particular page, such as filling out forms and clicking on links. Later, the user can play the macro back to automate access to the same page. LiveAgent<sup>30</sup> takes this approach, recording macros with a proxy that sits between the user's browser and the Web. The proxy augments pages with hidden frames and event handlers to capture the user's input, and uses this information to play back the recording later. Unfortunately, the proxy approach is also limited – for example, pages viewed over a secure connection cannot be seen, or automated, by the proxy. WebVCR<sup>31</sup> is another macro recorder for web navigation, which skirts the proxy problem by using a signed Java applet to detect page loads and instrumenting the page with event-capturing JavaScript after the page loads. Because part of WebVCR runs as an applet inside the browser, it can record all types of navigation. But neither LiveAgent nor WebVCR enable the user to modify the pages being viewed.

CoScripter<sup>32</sup> also uses the macro recording model, but unlike these earlier systems, CoScripter runs directly in the browser as a Firefox extension. In addition, recorded CoScripter scripts are displayed to the user as list of human-readable (and -editable) actions, such as "click on the Google Search button." CoScripter executes these instructions in a manner similar to Chickenfoot keyword pattern matching.

---

24. Microsoft. "Smart Tags and Smart Documents." [msdn.microsoft.com/office/understanding/smarttags/default.aspx](http://msdn.microsoft.com/office/understanding/smarttags/default.aspx)

25. Sugiura, A. and Koseki, Y. "Internet Scrapbook: Automating web browsing tasks by demonstration." *UIST '98*, pp. 9–18.

26. Miller, R.C. and Bharat, K. "SPHINX: a Framework for Creating Personal, Site-Specific Web Crawlers." *WWW7*, 1998, pp. 119–130.

27. Fujima, J., Lunzer, A., Hornbæk, K., Tanaka, Y. "Clip, connect, clone: combining application elements to build custom interfaces for information access." *UIST 2004*, pp. 175–184.

28. J.P. Bigham, A.C. Cavendar, R.S. Kaminsky, C.M. Prince, T.S. Robison. "Transcendence: enabling a personal view of the deep web." *IUI 2008*, pp. 169–178.

29. Anupam, V., Freire, J., Kumar, B., and Lieuwen, D. "Automating web navigation with the WebVCR." *WWW9*, 2000, pp 503–517.

30. Krulwich, B. "Automating the Internet: Agents as User Surrogates." *IEEE Internet Computing*, v1 n4 (July/August 1997), pp. 34–38.

31. Anupam, V., Freire, J., Kumar, B., and Lieuwen, D. "Automating web navigation with the WebVCR." *WWW9*, 2000, pp 503–517.

32. G. Leshed, E. Haber, T. Matthews, T. Lau. CoScripter: automating & sharing how-to knowledge in the enterprise. *CHI 2008*, pp. 1719–1728.



Toolkits such as WBI<sup>33</sup> and Greasemonkey focus on giving the user the ability to modify pages before, or just after, they are loaded in the user's web browser. WBI uses a proxy to intercept page requests, letting user-authored Java code mutate either the request or the resulting page. Giving users the ability to automate pages with Java and all its libraries is a powerful tool; however, WBI is still hampered by the limitations of a proxy.

Though both WBI and Greasemonkey enable the user to mutate pages, neither of them eliminates the need to inspect the HTML of the page to mutate it. For example, the sample scripts on the Greasemonkey site are full of XPath patterns that identify locations in web pages. These scripts are difficult to create because they require the author to plumb through potentially messy HTML to find the XPath, and they are difficult to maintain because they are not resilient to changes in the web site. Chickenfoot avoids this problem by giving users a high-level pattern language based on keyword matching that enables the user to identify pages without knowledge of the page's HTML structure, facilitating development and increasing readability and robustness. In Chickenfoot, users can fall back on XPath expressions to select elements and JavaScript to manipulate a page's DOM, but they are not restricted to these low-level tools.

WebL<sup>34</sup>, a programming language for the Web, also focused on giving users a higher-level language to describe Web page elements. In WebL, the user provides names of HTML elements to create *piece-sets*, where a *piece-set* is a set of *piece* objects, and a *piece* is a contiguous text region in a document. WebL provides various methods to combine *piece-sets* called *operators*, including set operators such as union and intersection, positional operators such as before and after, and hierarchical operators such as in and contain. Although these operators help produce more readable scripts, the language does not eliminate the need to inspect a web page for the names of its HTML elements, as the user must provide those to construct the basic pieces on which the operators work.

Another drawback of WebL, and of most of the aforementioned tools (with the exception of the macro recorders), is that they do not allow scripts to be developed inside the web browser. We consider the ability to experiment with a web site from the script development environment one of the greatest advantages of Chickenfoot – the user does not have to wait to see how it will affect the appearance of the web page, because Chickenfoot gives immediate feedback on the rendered page. LAPIS,<sup>35</sup> a predecessor of Chickenfoot, took a similar approach, giving the user an interactive environment in which to experiment with pattern matching and web automation. Unfortunately, the LAPIS web browser does not support web standards like JavaScript, cookies, and secure connections, so it fails to provide the user with a complete web experience.

In the time since Chickenfoot was originally released, many new systems have appeared for creating *mashups* between web services. Examples include Marmite<sup>36</sup>, Web Summaries<sup>37</sup>, Intel MashMaker<sup>38</sup>, Google Mashup Editor<sup>39</sup>, Microsoft PopFly<sup>40</sup>, and Yahoo Pipes<sup>41</sup>, and Vegemite<sup>42</sup>. Most of these

---

33. Barrett, R., Maglio, P., and Kellem, D.. "How to Personalize the Web." *CHI 1997*, pp. 75–82.

34. Kistler, T. and Marais, H. "WebL – a programming language for the Web." *WWW7*, 1998, pp. 259–270.

35. Miller, R.C. and Myers, B.A. "Integrating a Command Shell into a Web Browser." *USENIX 2000*, pp. 171–182.

36. J. Wong and J. Hong. "Marmite: End-User Programming for the Web." *CHI 2006*.

37. M. Dontcheva, S. M. Drucker, G. Wade, D. Salesin, and M. F. Cohen. "Summarizing personal web browsing sessions." *UIST 2006*, pp. 115–124.

38. R. Ennals and M. Garofalakis. "Mashmaker: Mashups for the Masses." *SIGMOD 2007*.

systems are themselves web services, rather than browser extensions, so they have less access to the user's own data and less ability to transform the user's experience in small or personalized ways, like adding a bit of data or a keyboard shortcut or a new link into an existing web page. These mashup systems tend to have much richer support than Chickenfoot for data extraction and manipulation, however.

## Conclusion

Chickenfoot is a programming system for web automation, integrated into the Firefox web browser. Chickenfoot enables the user to customize and automate web pages without inspecting their source, using keyword pattern matching to name page components. We showed that keyword patterns correspond closely to the names users actually generate for page components, and we presented a heuristic algorithm that implements keyword matching.

As of 2009, Chickenfoot is still being actively developed and used. The latest version is available at <http://uid.csail.mit.edu/chickenfoot/>.

## Acknowledgements

We thank all the students who have worked on Chickenfoot over the years, including Michael Bernstein, Prannay Budhraj, Vikki Chou, Mike Fitzgerald, Roger Hanna, Darris Hupp, Eric Lieberman, Mel Medlock, Brandon Pung, Philip Rha, Jon Stritar, Kevin Su, Tom Wilson, as well as other members of the UID group who provided valuable feedback on the ideas in this chapter, including ChongMeng Chow, Maya Dobuzhskaya, David Huynh, Marcos Ojeda, and Vineet Sinha. This work was supported in part by the National Science Foundation under award number IIS-0447800, and by Quanta Computer as part of the T-Party project. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the sponsors.

---

39. Google Mashup Editor. [code.google.com/gme/](http://code.google.com/gme/)

40. Microsoft Popfly. [popfly.ms](http://popfly.ms)

41. Yahoo Pipes. [pipes.yahoo.com](http://pipes.yahoo.com)

42. J. Lin, J. Wong, J. Nichols, A. Cypher and T. Lau. "End-user Programming of Mashups with Vegemite." *IUI 2009*.