

Seeking the User Interface

Steven P. Reiss

Department of Computer Science
Brown University
Providence, RI 02912 USA
spr@cs.brown.edu

Yun Miao

Google
1600 Amphitheater Parkway
Mountain View, CA 94043 USA
yunmiao@google.com

Qi Xin

Department of Computer Science
Brown University
Providence, RI 02912 USA
qx5@cs.brown.edu

Abstract

User interface design and coding can be complex and messy. We describe a system that uses code search to simplify and automate the exploration of such code. We start with a simple sketch of the desired interface along with a set of keywords describing the application context. If necessary, we convert the sketch into an SVG diagram. We then use existing code search engines to find results based on the keywords. We look for potential Java-based user interface solutions within those results and apply a series of code transformations to the solutions to generate derivative solutions, aiming to get solutions that constitute only the user interface and that will compile and run. We run the resultant solutions and compare the generated interfaces to the user's sketches. Finally, we let programmers interact with the matched solutions and return the running code for the solutions they choose. The system can be used for exploring alternative interfaces to the initial, for looking at user interfaces in a code repository, and for generating a skeleton user interface for an application.

1. Introduction

User interfaces are always a challenge to design and create. The coding involves understanding complex widget sets, building multiple prototypes to try achieving the best user experience, convoluted, inverted-control-based code, and a variety of layout strategies. The resultant code is often complex, bug-ridden, difficult to maintain, and not particularly transparent. Comparing alterna-

tive interfaces or exploring the design space for a task can be problematic. Testing user interfaces, especially during development, is difficult and time consuming; testing interfaces aesthetics and usability even more so. Yet user interfaces are a critical part of today's applications.

The goal of our research is to simplify and eventually automate the process of exploring and eventually building user interfaces by letting the programmer rely on the growing repositories of already developed and tested open source applications. Essentially we want to let the programmer view, interact with, and explore the code of user interfaces in these repositories. Programmers should simply sketch the user interface they want and then our tool will search the various repositories of open source applications, extract user interfaces from these applications, and return working code that is similar to the programmer's design.

Open source code repositories and systems are growing exponentially. OpenHub now claims over 680,000 repositories with over 35 billion lines of code. (This has doubled in the past 3 years.) For many applications, programmers have developed and tested viable user interfaces and the results are in the repositories. A programmer might want to build on one of these existing designs.

Our work lets the programmer start with a sketch of the user interface along with some context information. We use the context information to search open source repositories for appropriate Java applications using existing search engines. We extract the user interface code from these applications, get the code to compile and run, check whether the generated interface matches the given diagrams, and let the programmer view and check the result interfaces by interacting with them. The actual source code for the generated interfaces can then be returned to the programmer if desired.

This approach can return interfaces that are fully developed, that include interaction code, code to handle different window sizes, and callback hooks. Such interfaces are more substantial than those generated by the user interface builders common to today's programming environments. The approach can also be used to explore the space of interfaces for an application domain, looking at different alternatives and filling in the gap between a preliminary sketch and a usable interface. Finally, the approach has been used to explore user interfaces as an aid to browsing code repositories.

The contributions of this work, in addition to showing the feasibility of using code search for user interface design, are:

- A service for translating user sketches into structured SVG diagrams that can be more easily understood for user interface search.
- A means for translating SVG-based user interface sketches into a form that can be used to check if a given user interface is valid.
- Methods for gathering the appropriate code for a user interface from the simple results returned by code search engines.
- Transformations that take the raw code returned from code search, extract the user interfaces, and then make the code runnable outside of the original context.
- Techniques for matching a user interface sketch with an actual user interface.
- Tools that let the programmer see and interact with candidate interfaces to choose which they want the code for.

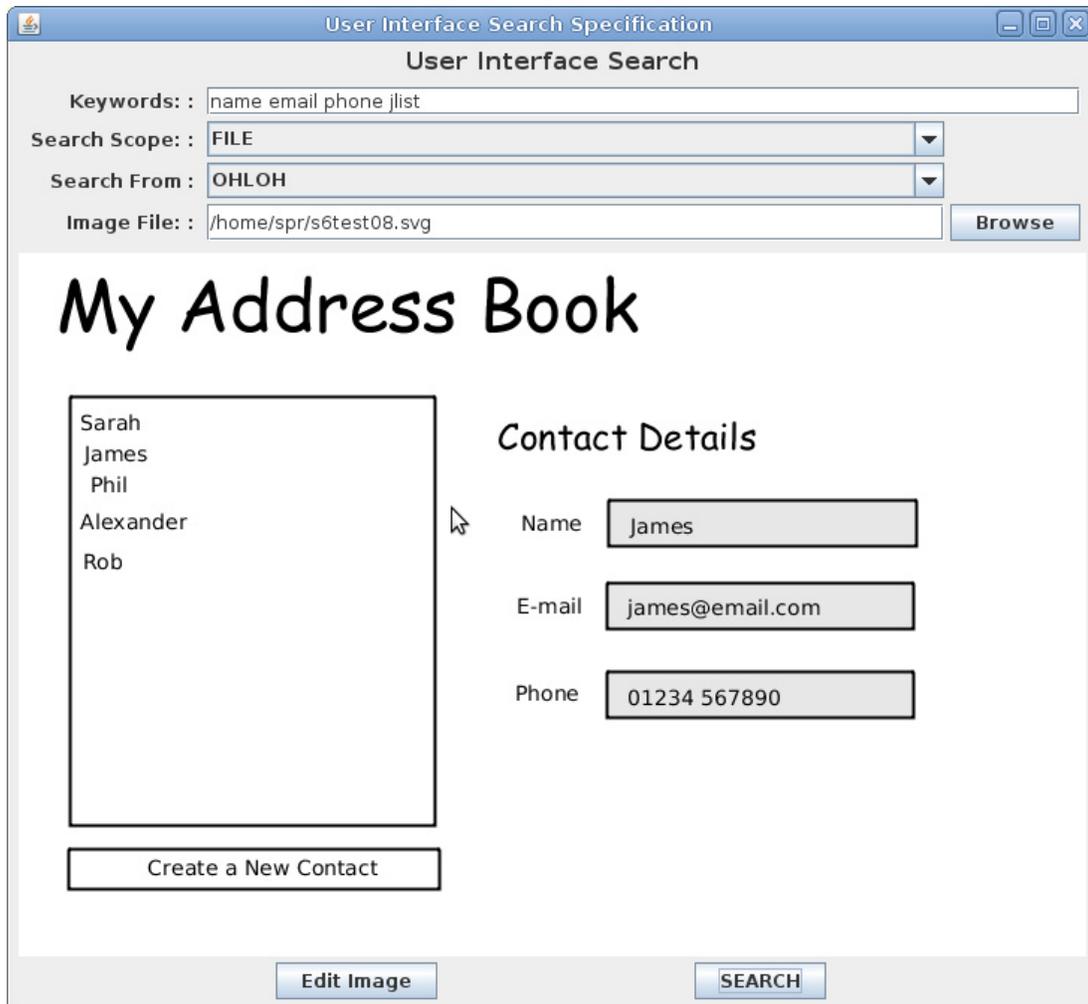


FIGURE 1. The user interface for specifying what to search for. The specification includes keywords and an SVG-based sketch. Search options include which code search engine to use and the scope of the search.

- Integration of the user interface search facility into a code search front end.

2. Overview

Searching for user interfaces can be broken down into three stages: specifying what to search for, generating candidate solution, and validating those solutions.

To specify a user interface, the user provides a sketch of what is desired along with a set of keywords describing the application context of the desired interface. Our tool, SUISE, is shown in Figure 1.

The user interface sketch is provided as an SVG file. If the user has a free-hand sketch of the user interface, we provide a facility to convert that diagram into an appropriate SVG file. This is described in Section 5. An SVG diagram is more structured and easier to interpret than an arbitrary sketch. Moreover, SVG is a common standard, works well with the web, there are many

```

<COMPONENT HEIGHT=416.7938537597656' ID=U_70' TYPES=java.awt.Container' WIDTH=574 X=27' Y=1'>
  <COMPONENT DATA='My Address Book' HEIGHT=51' ID=U_51' LEFT=U_70' TOP=U_70' TYPES=javax.swing.JLabel' WIDTH=386' X=27'
    Y=15' />
    <COMPONENT DATA='E-mail' HEIGHT=10' ID=U_64' TYPES=javax.swing.JLabel' WIDTH=41' X=334' Y=233' />
    <COMPONENT DATA='Contact Details' HEIGHT=20' ID=U_60' TYPES=javax.swing.JLabel' WIDTH=171' X=321' Y=114' />
    <COMPONENT HEIGHT=32 ID=U_62' RIGHT=U_70' TYPES=javax.swing.JTextField' WIDTH=207' X=393' Y=166' />
    <COMPONENT HEIGHT=291' ID=U_52' LEFT=U_70' TYPES=javax.swing.JList, javax.swing.JTextArea, javax.swing.JEditorPane' WIDTH=245
      X=34' Y=97' />
    <COMPONENT DATA='Name' HEIGHT=10' ID=U_61' TYPES=javax.swing.JLabel' WIDTH=39' X=337' Y=178' />
    <COMPONENT HEIGHT=33' ID=U_67' RIGHT=U_70' TYPES=javax.swing.JTextField' WIDTH=207' X=392' Y=283' />
    <COMPONENT BOTTOM=U_70' DATA='Create a New Contact' HEIGHT=29' ID=U_58' LEFT=U_70'
      TYPES=javax.swing.JButton, javax.swing.JMenuItem' WIDTH=250' X=32' Y=403' />
    <COMPONENT DATA='Phone' HEIGHT=11' ID=U_69' TYPES=javax.swing.JLabel' WIDTH=41' X=333' Y=292' />
    <COMPONENT HEIGHT=33 ID=U_65' RIGHT=U_70' TYPES=javax.swing.JTextField' WIDTH=207' X=392' Y=222' />
  </COMPONENT>

```

FIGURE 2. Hierarchical component specification generated from the diagram shown in Figure 1. Each component includes a position and size.

available tools for creating and editing such diagrams, and Apache provides a suite of Java-based tools for SVG.

When the user completes the specification and hits the search button, we build a Java code search request for a modified version of our S^6 search engine [55]. To do this we transform the user's sketch into a *hierarchical component description* as described in Section 4. This description includes the components that should be in the user interface and the relationships among those components. Components can be nested. For each component, the description includes a set of Java Swing/AWT widget types that can be used to implement this particular component. The component description from the example shown in Figure 1 is shown in Figure 2.

Next, S^6 uses the keywords to find candidate solutions from an existing code search engine such as OpenHub, Hunter, CodeExplorer, or Github. S^6 next looks for candidate user interfaces. A candidate solution can be a class that implements a Swing/AWT component or a non-private method that returns such a component. Next S^6 applies transformations to each solution in an attempt to create code that is compilable, runnable, and only contains the user interface. The result of each transformation is a new solution that can also be transformed. The end result of these transformations is a set of candidate user interfaces that might meet the user's criteria. This process is described in Section 6.

We validate these solutions in several ways. First, we ensure the code compiles and runs. Second, the user interface generated by the code needs to match the hierarchical component specification. Third, the interface needs to look and act correctly. For the first two of these, S^6 compiles and runs the code, and then matches the user interface generated in the run against the component specification. The various constraints and values included in the specification are used to generate a matching score which is used to rank the solutions. This is described in the first part of Section 7.

The task of exploring user interfaces and deciding which are the most appropriate is left to the programmer. Our tool presents the candidate solutions to the user first by showing images of the interface as seen in Figure 3. The user can accept or reject the solutions directly, based on their image. Alternatively, if a solution is clicked on, then the system will run the user interface along with a viewer that lets the user investigate the widget hierarchy and the various events generated by interaction. Once the user has selected a set of acceptable solutions, they can hit the "Show the

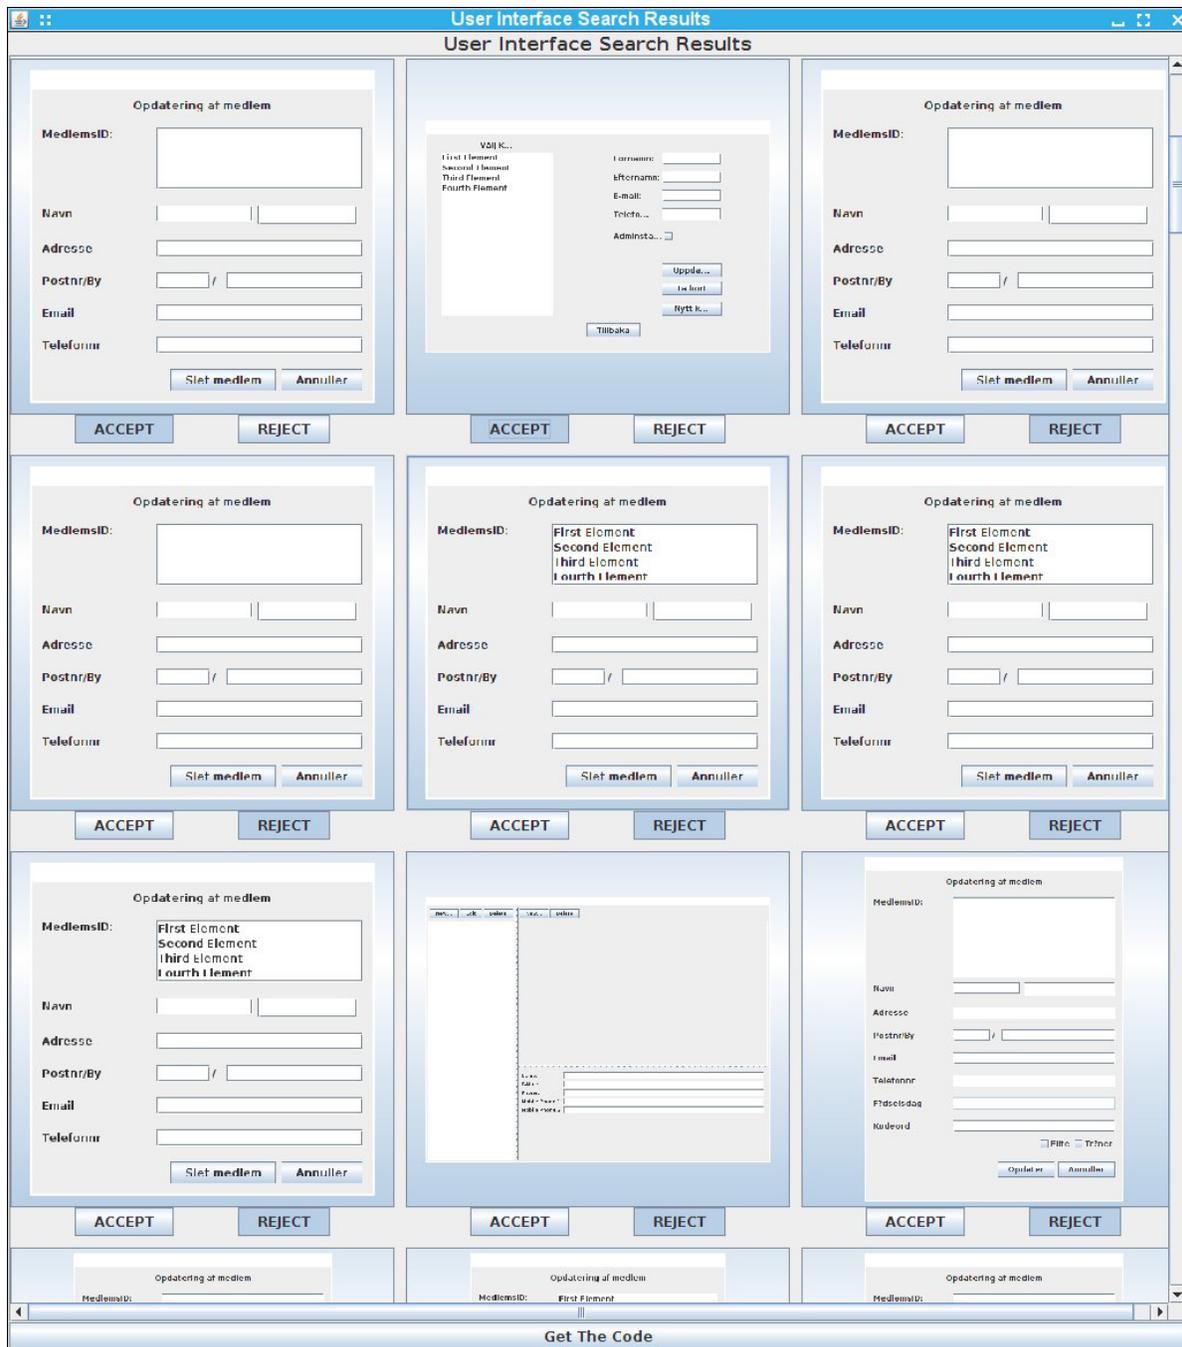
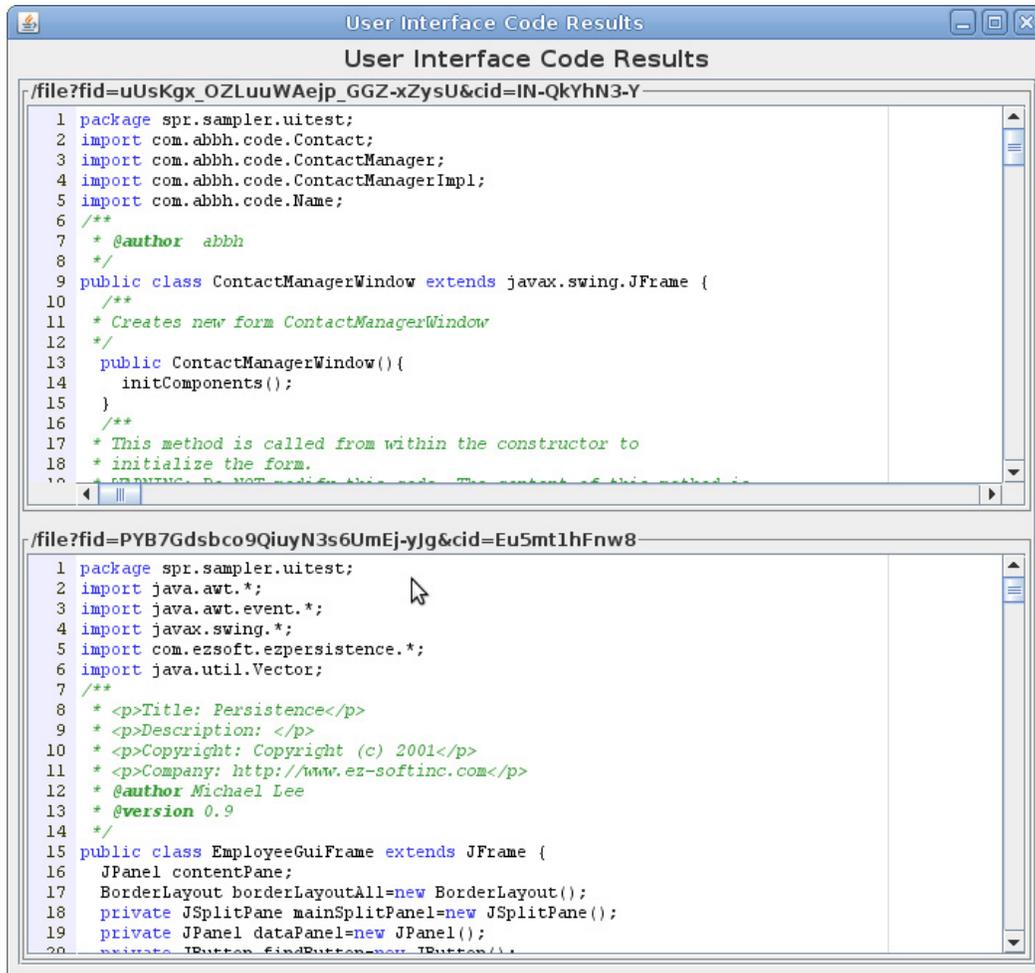


FIGURE 3. The resultant display showing potential solutions for the address book sketch of Figure 1. Each solution can be accepted or rejected by the user. Moreover, the user can experiment with the solution by clicking on it.

Code” button to get a display of the resultant code. This is shown in Figure 4 and explained in Section 7.



```
/file?fid=uUsKgx_OZLuuWAejp_GGZ-xZysU&cid=IN-QkYhN3-Y
1 package spr.sampler.uitest;
2 import com.abbh.code.Contact;
3 import com.abbh.code.ContactManager;
4 import com.abbh.code.ContactManagerImpl;
5 import com.abbh.code.Name;
6 /**
7  * @author abbh
8  */
9 public class ContactManagerWindow extends javax.swing.JFrame {
10     /**
11     * Creates new form ContactManagerWindow
12     */
13     public ContactManagerWindow(){
14         initComponents();
15     }
16     /**
17     * This method is called from within the constructor to
18     * initialize the form.
19     * @param args the command line arguments
20     */
21 }
```

```
/file?fid=PYB7Gdsbco9QiuyN3s6UmEj-yJg&cid=Eu5mt1hFw8
1 package spr.sampler.uitest;
2 import java.awt.*;
3 import java.awt.event.*;
4 import javax.swing.*;
5 import com.ezsoft.ezpersistence.*;
6 import java.util.Vector;
7 /**
8  * <p>Title: Persistence</p>
9  * <p>Description: </p>
10 * <p>Copyright: Copyright (c) 2001</p>
11 * <p>Company: http://www.ez-softinc.com</p>
12 * @author Michael Lee
13 * @version 0.9
14 */
15 public class EmployeeGuiFrame extends JFrame {
16     JPanel contentPane;
17     BorderLayout borderLayoutAll=new BorderLayout();
18     private JSplitPane mainSplitPanel=new JSplitPane();
19     private JPanel dataPanel=new JPanel();
20     private JButton findButton=new JButton();
21 }
```

FIGURE 4. Final display showing the code for the user interfaces the user accepted.

In the next section we describe S⁶ and other related work. Section 8. describes our experiences to date and offers an evaluation of the work. Section 9. then concludes by describing our on-going work.

3. Related Work

Creating graphical user interfaces has been a difficult problem since the 1980's when such interfaces starting to become common. While most interfaces then and now are still hand-coded, there have been a wide variety of tools developed to assist and even attempt to automate the process. A good summary of the state of the art in 2000 is provided by [44].

3.1 User Interface Generation

Modern development environments such as NetBeans, Visual Studio, and Eclipse support user interface generation. They let the developer drag and drop widgets into containers and to set the various properties of the widgets. Once a user interface is designed, the basic code for the interface can be generated. The programmer can modify this code to interact correctly with other por-

tions of the application. These tools provide some simplification of the process, but are not ideal in that a) they don't handle interaction, data validation, or other interface dynamics; b) they often use absolute positions and it is complex to generate easily resizable interfaces; c) they don't handle dynamically generated interfaces where the interface depends on external files or the state of the application; d) the code that is generated may not be in a style or form the programmer desires; and e) once the code is modified it becomes difficult to use the support to update or change the user interface. The latter is a problem because user interface design often involves the rapid iterative design, exploration and comparison of different interface implementations [23].

There are some tools that attempt to generate user interfaces without actually writing code. Some of these involve using non-procedural specifications such as Mozilla's XUL [17]. Others involve developing various models representing the underlying data and the presentation and then generating the interface from these models [41,52,61]. The model driven tools have been more successful when applied to specialized environments [19,47]. There has also been work on automatically adapting user interfaces based on device or user constraints [16,48].

There are also tools that let users take sketches as input for generating user interfaces. JavaSketchIt allows creating user interfaces through hand-drawn geometric shape identified through gesture recognition [8]. SILK lets designers quickly sketch a graphical user interface using an electronic pad and stylus and then recognizes widgets and other interfaces elements as the designer draws them [30]. [59] lets users take a picture of a sketch and maps the sketch elements into mobile widgets. More sophisticated sketch-based tools can generate GUI code. MobiDev provides users with a set of predefined elements to draw sketches of applications and then generates the application based on those elements [60]. REMAUI infers mobile interface code from screen shots of conceptual drawings using OCR and computer vision techniques but only supports the top three Android widgets [46].

3.2 Basic Code Search

There has been significant work done on code search. Early work in this area demonstrated that keywords from comments and variable names were often sufficient for finding reusable routines [15,36]. Later work did query refinement either directly [65], by looking at what the programmer was doing [76,77], using class signatures [26], using an appropriate ontology [75], using the surrounding context [24,73], using learning techniques [12], using natural language [10], using an execution trace [35], using topic graphs [72], using associations [66], using tpestates [43], or using collaborative feedback [70]. Recent approaches, such as Assieme [25], Sorcerer [2,3], Codifier [4], Exemplar [18] and Portfolio [40] expand basic keyword search to consider program structure and semantics. Other recent work has looked at more sophisticated IR techniques [67] and on automatic query reformulation [20,21,62]. More sophisticated search techniques use theorem proving techniques [63,64]. These are in addition to commercial tools such as Krugel (krugel.org), Github (www.github.com) and the currently unavailable OpenHub and Google code searches.

While code search shows much promise, it has not caught on extensively. To some extent, this is because the various code search engines are not particularly effective. However, even with an effective search, the programmer still has to do a significant amount of work in order to use the result. This includes checking whether the code actually does what is desired, adapting the code to

their particular project, possibly debugging the code, and converting the code to their style and formatting standards.

3.3 Semantic Search with S⁶

Our prior work on code search is the semantic code search engine S⁶. S⁶ attempts to address several of the problems with current code search technology by effectively automating the multiple tasks the programmer has to do manually in order to use the output of a code search tool [55,56].

S⁶ can be used to search for either Java classes or methods. It provides a web-based interface that asks the user to first provide a description of what is wanted in terms of keywords and the semantics of the target code. The latter includes the signature for the target class or method, one or more test cases, and optionally contracts (preconditions and postconditions) and security specifications (e.g. the returned code should not do any file I/O).

Once this data is entered, S⁶ processes the request. It first uses the keywords with an existing code search engine (Ohloh, Krugle, Github, GrepCode, or Sourcerer [2]) to get a starting set. It generally takes the first 100-200 files from the search results to build an initial set of solutions. The next step is to apply transformations to each solution to generate new solutions. This is done repeatedly until no more transformations are applicable and no new solutions are generated. These transformations include relatively simple ones such as change the name of the method to match the name in the specified signature or reordering the parameters; moderately complex ones such as replacing a parameter with an appropriate assignment; and complex ones such as extracting functionality from a method by finding a top-level statement computing a value of the return type, doing a backward slice of the code until the only free variables are of the parameter types, and then extracting the resultant code into its own function.

The system next takes all the resultant candidate solutions and does a dependency check. This check adds other code fragments such as field declarations and auxiliary methods from the initial file that might be needed to make the candidate compile. It removes candidate solutions with unmet dependencies that will not compile. For each passing candidate the system generates a test program that tests that candidate against the user's original test cases, contracts, and security constraints. This test program is compiled and run using Apache Ant and JUnit. The system does an additional pass looking at the output from the tests, and will try additional transformations as appropriate, for example, transformations that handle off-by-one or uppercase/lowercase errors.

Finally, the system takes the candidate solutions that pass all the test cases and passes the resultant code back to the user. It gives the user the option of different formatting styles [54] and different orderings for the results (e.g. fastest to slowest, smallest to largest, least to most complex). It also provides license information for each of the fragments. The user can then take the result, cut and paste it into their program and use it with the confidence it actually compiles and passes their test cases.

S⁶ provides a general framework for using code search for different purposes. It starts with keywords to identify a set of initial candidate files. Next, it uses a set of transformations that convert these candidate files into initial candidate solutions. Next it transforms the candidate solutions so that they are likely to compile and run. These transformations are limited by applying an interme-

diate check as to whether the solution is feasible or not. Finally, it needs to compile and validate the resultant solutions. Our search tool implements and specializes this framework for user interfaces.

3.4 Other Search Tools

Another semantic approach involved defining the behavior to searched for. This was originally given as input-output pairs [51], and then generalized to allow slightly more flexible specifications [9,22]. More recent work in this area includes PARSEWeb that does static analysis on code fragments found by a text-based code search engine and then looks at input-output types [68]. Other techniques such as program patterns [49,71] and keyword programming [34] are designed to work at the level of a code fragment.

Several search-based systems use test cases as input. CodeGenie [31,32] lets the user define tests as part of the development process in Eclipse and then uses the method names and signatures from the test to build a query. It uses an internal search engine that understands program structure to find code to test and then presents the result to the user. Other recent code search work on test cases includes [1,27,33] and our S⁶. Test cases and semantics have also been used in a similar fashion for finding web services [14,53], but have the problem that the user must know exactly what is being searched for [28].

4. Specifying User Interfaces

Our goals in specifying a user interface for code search were three-fold. First, we needed to provide an appropriate starting point for the search process. Second, we wanted to use a natural metaphor, starting with sketches as designers typically do. Third, we wanted to be able to check the result against the specification so we could test if a generated search solution was appropriate.

There are two aspects to identifying a starting point for the search. The first is the set of keywords that will be used in conjunction with an existing code search engine to find initial files. The second specifies if the solution should be within a single file, within a single package, or spanning multiple packages.

User interfaces can be implemented in a variety of ways. Simple interfaces and interfaces developed using user interface builders are often implemented within a single file. More complex interfaces, where the user creates custom components, uses custom models for tables or lists, or implements complex internal functionality, are often implemented in multiple classes within a single package. Applications that have multiple user interfaces may use a common user interface package for support code while implementing the actual interface in a separate package within the system.

In order to accommodate these different user interface implementation styles, we support initial solutions that are either file-based, package-based, or system-based. For package based solutions we start with the initial file returned from the code search engine, search for other classes in the same package, and merge the results into a single virtual file for further processing. This merger yields a single Java file containing multiple classes that would typically not compile directly. However, we retain enough information to separate this into multiple files when we need to

compile it. The merger also takes into account the different imports for the different files, yielding a common set of imports by replacing simple names with qualified names where necessary.

In the case where the user interface might span packages, we start with the initial file, add the other files for the package as above, and then use the import clauses and qualified names in the result to identify packages that share a common prefix with the original one. All the classes in these packages are merged with the original file as well and the process is repeated until no new packages are identified. The merging here moves all the files to a single package, updating names and import statements accordingly.

The remainder of the specification is an SVG sketch of the desired interface. We assume that the user either creates the sketch using an SVG editor or makes use of our SVG conversion facility described in the next section. SVG is web-friendly, matching the current S^6 web interface. Moreover, there are several tools available for creating and editing SVG files such as Inkscape and the web-based `svg-edit`. The edit image button on the bottom of the interface will bring up an appropriate editor, either Inkscape if it is installed on the system or GLIPS Graffiti. Finally, the search button at the bottom right of the interface starts the whole user interface search process.

The SVG-based user interface sketch addresses our second criteria, letting the user start with a sketch. To make this usable by S^6 we analyze the sketch and translate it into a hierarchical component description in stages.

The first stage finds potential components. We use Apache's Batik package to map the SVG diagram into drawable components, either shapes (rectangles, rounded rectangles, ellipses or general paths) or text. We further characterize shapes as either boxes, input regions, lines, symbols, icons, rounded regions, or text. Input regions are boxes that are either lightly filled or that have a thicker than normal border. Lines are either lines or are boxes that have essentially one dimension. Symbols are shapes that are small and can represent either a simple button (e.g. a radio button), an icon, or an arrow (for a scroll bar). Icons are larger two dimensional symbols and can represent larger icons or general drawings. Squiggly paths that are long and narrow are taken to represent potential text. Text regions are further characterized as containing single or multiple lines.

The next stage creates a hierarchy of the candidate components. This is done by looking at the bounding boxes of each component and seeing what other components are nested inside. Here we use an approximation to actual nesting to accommodate minor errors in the sketch. For example, if a string happens to lie mostly inside a rectangle, but extends outside slightly, we consider the string to be nested. Once we determine all such nestings, we build a hierarchy by finding the innermost nesting for each component. Finally, if there is no unique top-level component, we create one.

The next stage attempts to merge logical groups of components and to characterize the components so that we can assign potential widget types for checking. This is done with a series of hand-coded checks that assign properties to the components and clean up the hierarchy. The actual checks done here include:

- Looking for components containing only text subcomponents. This characterizes components as buttons if the text is a simple string and is generally centered or if the enclosing region is circular; or as single line or multiple line text (input) regions otherwise. If text is present we check

for asterisks to indicate a password or hidden field and for only numbers to indicate a numeric field. Where components are further characterized, the text subcomponents are removed from the result hierarchy.

- Looking for combo boxes (buttons with a choice of options). These are text boxes with a symbol on the right. If one is found, the symbol and text are removed.
- Looking for toggle buttons such as radio buttons and check boxes. These are buttons or text components with an adjacent symbol. Where these are found, a new component is created spanning both the original components which are removed.
- Looking for menu and tool bars. These are long, narrow regions containing a one or more symbols, buttons or text strings.
- Looking for tables, trees and lists. Tables are characterized as boxes containing both vertical and horizontal lines and possible text elements. Lists can contain horizontal lines or multiple text items. Trees can contain vertical lines and have text areas that are properly offset. Any internal subcomponents are removed.
- Looking for scroll bars. These are either long or narrow regions that contain symbols at the top and bottom and possibly a box or symbols in the middle. Any internal symbols and boxes are removed if a scroll bar is identified.
- Looking for spinners. These are numeric fields with one or two symbols immediately to the right.
- Looking for sliders. We look for a long narrow component with additional symbols on top of it and with potential text immediately above or below. If a slider is identified, all the internal components are replaced with a single slider component.
- Looking for drawing areas. These are characterized as a component containing multiple symbols and shapes but no buttons.

The checks here are designed to be forgiving in order to accommodate minor errors in the original sketch. For example, the checks involving aspect ratios accept an overly broad range of values; checks for horizontal and vertical lines allow ease; and extra marks or boxes that are small or seem irrelevant are ignored.

The fourth stage of component processing takes the resultant set of components and computes a set of relative positional constraints that can be used for checking. Each component can identify another component that is immediately above it, one that is immediately below, one to the left, and one to the right. Nested components can also be assigned a position relative to their parent, for example, a component that is at the top of its parent has the parent identified as the component immediately above it.

The final stage assigns potential widget types to each component. This uses the properties set by the above drawing analysis to create a list of candidate AWT/Swing widgets for each component. This is the only part of the specification stage that is dependent on the user interface being generated for Java AWT/Swing.

The resultant component hierarchy for the input shown in Figure 1 is shown in Figure 2. The box on the left is identified as either a *JList*, *JTextArea* or *JEditorPane*; the three boxes on the right are

identified as *JTextField*s, and the box at the bottom as a *JButton* or *JMenuItem*. The remaining elements are either *JLabels* or the outermost *Container*.

Our user interface for specifying what to search for can be seen in Figure 1. The top three boxes define the starting point for search. The top box contains the keywords; the second box identifies the type of search; the third box selects the search engine to be used. The sketch, selected from a file, is displayed below.

5. Generating SVG Diagrams from User Sketches

For the purpose of UI search, we developed a basic image to SVG converter by building a shape detector and converter. This lets the user start with a free-form sketch of the interface they are interested in, convert it into an appropriate SVG diagram, and then use the converted diagram as input into the search process. The converter can start with an jpeg, or png image.

In order to convert from a sketch into a structured diagram, the tool needs to identify the various user interface components in the original sketch. To do this, the converter uses techniques from machine learning and computer vision. The components that are needed for the search process are text and shapes such as circles, ellipses, rectangles, triangles, and tables. We use a gradient based method to find points of interest to shapes and to infer shape information. A variation of the Hough Transform [13] is used to detect tables. A number of simplifying assumptions are made: each hand-drawn shape or table is closed; the boundaries of the shapes and tables have the thickness of a ball point pen trace; and there is at most one table in the image.

5.1 Text Detection

We use a simple contour based text localization method to detect text in the user images. We use four weak features to train four different linear support vector machines (SVMs), and use adaptive boosting (AdaBoost) to combine the classifiers to detect contours that are likely to be text. The four features are: compactness, defined as the ratio of bounding rectangle perimeter over contour perimeter; solidity, defined as the ratio of contour area over convex hull area; contour area; and horizontal crossings [45]. We first convert text into contours, then we get a list of text regions by merging adjacent contours based on the classification results. Text regions are then excluded from later shape detection. This is similar to other approaches to text recognition in hand-drawn diagrams [7,11,50,74]. Figure 5 shows sample detection results before and after merging positive contours.

5.2 Shape Detection

We develop a novel gradient-based method to detect interest points in an image, and then use the interest points to infer shape types. By using gradients, our method is scale invariant and can detect shapes of all sizes.

We first find inflection points using gradients. We compute a gradient vector for each point on a contour. We use this to compute the gradient angle change at each point. We then choose the local maximum gradient angle change in a two point window. This effectively filters the contour points by suppressing points with non-maximal local gradient change. This strategy works well because rectangles and the triangles have straight edges. The gradient change between two points on a

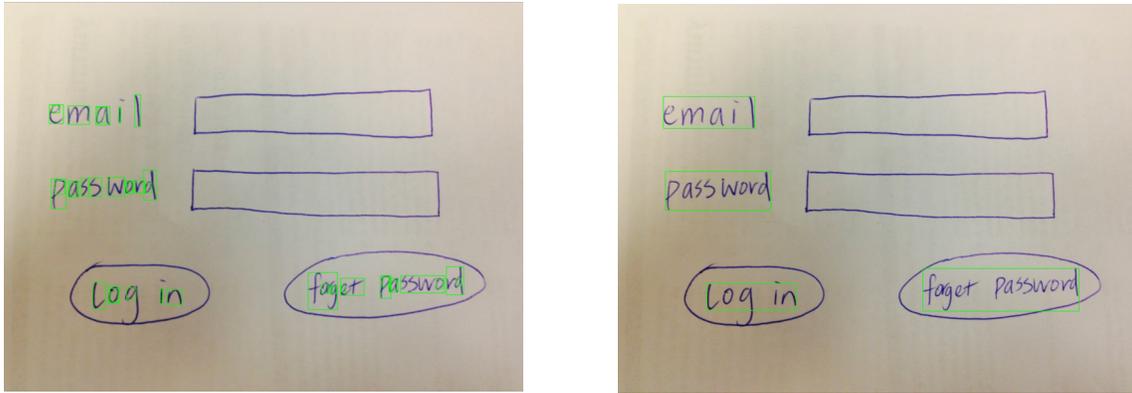


FIGURE 5. Text detection results for a sample diagram. The boxes on the left indicate the individual text contours identified using machine learning. The boxes on the right show the result after merging positive contours.

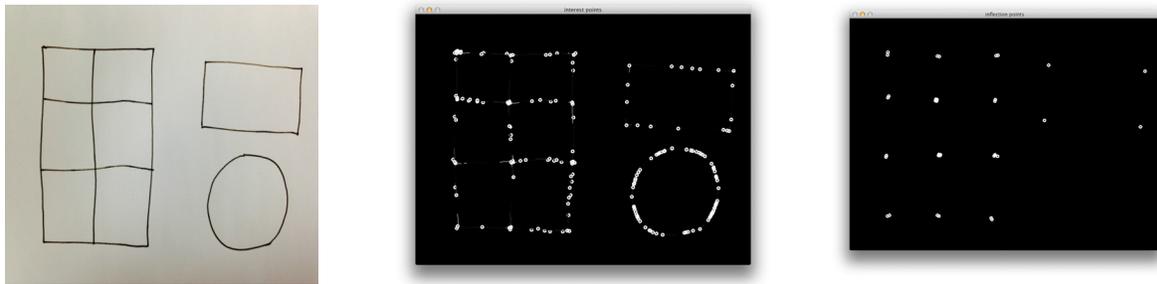


FIGURE 6. A shape detection example. The image on the left shows the original image. The center image shows the points of interest detected using gradients. The image on the right shows the corner points that are detected based on the points of interest.

straight line is close to zero. The gradient only changes significantly at the corner. This observation helps us eliminate a many unhelpful contour points. The above three steps give us a list of interest points which are potential corner points on a contour, see Figure 6.

We next obtain contour points by calculating the gradient angle change again among the identified points. This gives us the corner points on the contour. These corner points are then used to determine the shape types and table parameters [69].

5.3 Table Detection

Identifying tables is slightly more complex. A table usually build from a set of overlapping rectangles. After the shape information is inferred, we use the list of rectangles to find if there exists a table in the sketch. This is done by finding if there are two or more rectangles that are adjacent either horizontally or vertically. When a table exists, we collect the corner points from all rectangles and use the corner points to determine the position of rows and columns in the table. We use a simplified version of the Hough Transform to find the rows and columns in a table by assuming that our table contains only horizontal and vertical lines. An example is shown in Figure 7.

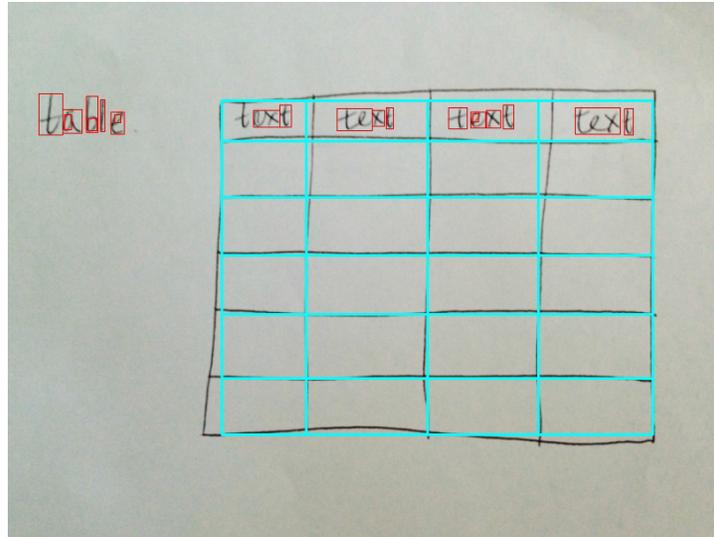


FIGURE 7. Sample figure showing what is detected from a figure showing a table with some text.

5.4 Converting the Image into SVG

During shape and text detection, we obtain a list of shapes and text regions. Each shape contains necessary parameters for SVG generation. For example, a line shape contains its two end points while an ellipse shape contains its center, width, height and its X-axis rotation angle. We generate SVG diagrams using the python library *svgwrite*. When generating text regions in SVG files, we made the decision to use the general word `text` to represent any text content in the user input, since we currently do not implement character recognition, although this could be done in the future. Instead, we leave it to the user to replace the placeholder text in the generated SVG diagram with the proper text. . Figure 8 shows two sample initial images (on the left) with the identified shapes marked. The result, with generic text, is shown in the center. The right-hand images show the resultant SVG diagram with the proper text.

6. Generating Solutions

Once we have an SVG diagram, either directly or indirectly, we are ready to search for similar interfaces. To do the actual work of searching for a user interface, we substantially modified the S⁶ search engine. The modifications generally fall into three categories. The first is handling packages and systems rather than individual functions or classes. These were described in the previous section. The second involves restricting the code to that relevant to the user interface by eliminating unnecessary elements. The third involves getting the resultant code to compile and run, effectively duplicating what a programmer might do when extracting the interface from the code.

S⁶ for user interface search starts with the code generated from either files, packages, or multiple packages based on the initial code search. Each of these code files (with the latter ones being considered single files after all the code has been merged), is considered a candidate solution.

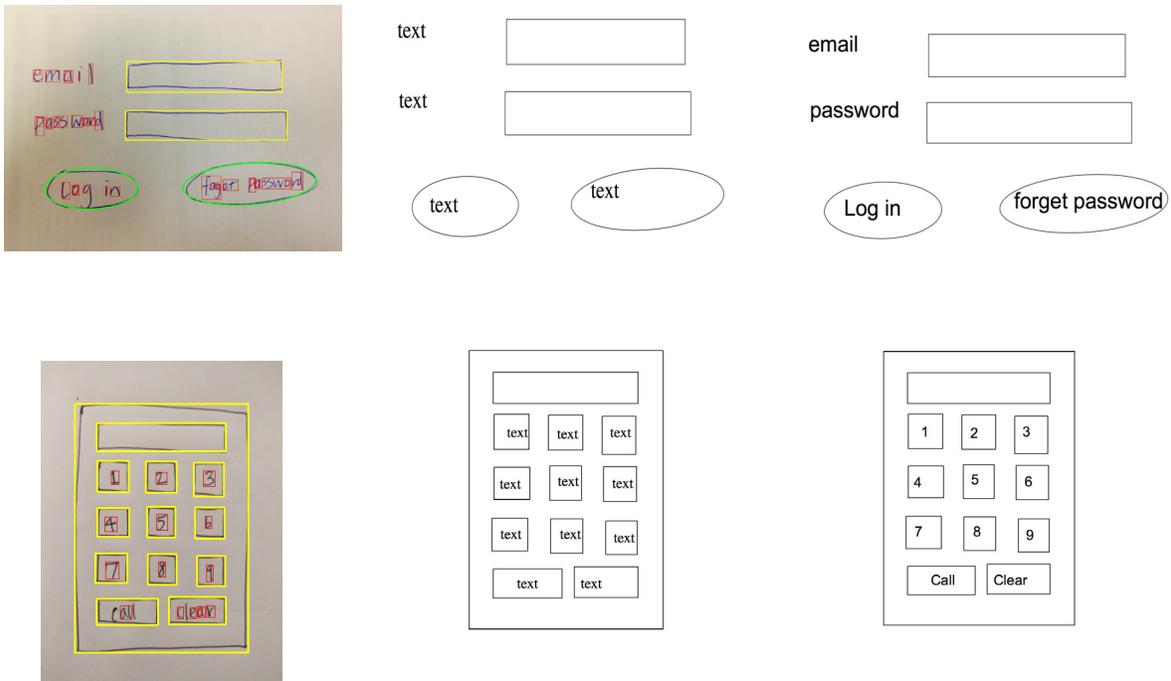


FIGURE 8. Examples of the overall process converting a hand-drawn sketch into an appropriate SVG diagram.

The next step is to identify potential user interfaces in each solution and generate separate solutions for each. This is done using S^6 code transforms. We first convert the package name to a standard one for the user interface. Next we find all candidate interfaces. These are non-private constructors for any class that extends *java.awt.Container* and any non-private methods of a class that return an object that extends *Container*. For each such candidate, we create a new solution by creating a new class with a standard name that either calls the appropriate constructor or first builds the class and then calls the identified method. Where there are multiple possible constructors, we generate separate solutions for each, using logical default values for any parameters on the constructor. If the identified methods take parameters, we generate separate solutions using different default values for those parameters.

Each potential user interface solution is restricted by a transformation that eliminates any code that cannot be reached from the class added for the solution. The result of this is a set of candidate solutions that implement a potentially relevant user interface and that are restricted to the code needed for that interface.

The next step uses existing S^6 transforms along with Swing-specific transforms to take these solutions and build new solutions that have a greater possibility of compiling and meeting the user's needs.

The transformations that we have added for handling multiple classes, user interfaces and Swing include:

6.1 Swing-Specific Transforms

The first set of transforms try to modify the code to compile while preserving the resultant user interface. These attempt to replace uncompileable code that involve Swing or AWT calls with equivalent generic code that will compile.

The first transform looks at AWT or Swing method calls and class instance creations (new expressions) for AWT or Swing classes that involve invalid expressions. Invalid expressions can arise from field, method or type references to classes that are not included in the search solution, e.g. are outside of the file, package or system or that are in non-standard libraries the project uses. The transform replaces each invalid parameter the original statements with an appropriate value. Any undefined string parameter is replaced with a unique string; integer and Boolean parameters are replaced with a 0 or false respectively; color parameters are replaced with a valid color; image and icon parameters are replaced with a known image or icon; parameters that are subtypes of *java.awt.Component* are replaced with a simple label; and other parameters are replaced with null. An example of such a call would be a method that sets the text associated with a widget based on the result of an external call. In this case, the transform would effectively change the string shown in the user interface to some S⁶-specific string but would otherwise preserve the interface.

A second transform deals with anonymous classes that inherit from or implement a Swing or AWT interface that would not compile because they contain invalid or undefined expressions. This transform either removes the offending internal methods or replaces them with methods that will compile by returning default values.

A third transformation looks for calls that use resources by checking for reference to the *java.util.ResourceBundle* interface. Since the retrieved code does not include resource files or the local resources, calls to standard resource routines will compile correctly but would eventually produce run time errors. Moreover, these are frequently used in user interfaces to internationalize the code. We replace calls to *ResourceBundle* with calls to an mplementation of a resource bundle in our test framework that will return valid values for all requests at run time. We also remove calls to *java.util.Properties.load* which would fail because of the non-existent resource files.

A fourth transformation looks for calls that would cause the user interface to hang or become untestable. For example, if the application attempts to run a modal dialog, the test code, the method returning the component that needs to be checked would never return. Our test code would not be able to access the resultant widget hierarchy. Moreover, the test case will eventually time out and fail, even though a valid interface might be built. Here we remove statements that contain calls to the various *javax.swing.JOptionPane* methods. We also remove calls to *java.awt.Dialog.setVisible* and *System.exit* where possible.

A fifth transformation replaces list, tree, and table models offered by the user with simple internal models provided by our test framework. Many of the compilation and run time problems that are encountered when trying to run the retrieved code arose because the code attempted to use incomplete or unavailable models. This occurs frequently because such models are themselves integrated with the application and not the user interface. They may be built from external files or from an external database. They might involve classes outside the package or subsystem that was retrieved. Note that this transformation adds additional potential solutions to be considered. It does not remove the original solutions. Thus if the original models would work in the returned

code, that solution would still be accepted. The new solutions handle the cases where the original models would not work.

6.2 General Transformations

In addition to the Swing-specific transformations, there are several generic transforms provided by S⁶ that are useful in generating a working user interface from the retrieved code. These generally are designed to clean up the code and ensure it will compile and be testable.

The first relevant transform handles undefined types or variables. It removes any code that references such variables by removing the corresponding statement when possible. There are cases, however, where this is not possible. For example, if one were to remove a return statement at the end of a method, the result would not compile. In these special cases, the undefined value is replaced with a default value, with both 0 and 1 (for numerics), true and false (for Booleans), and null and non-null (for objects) being tried. The transformation works repeatedly, so that if a declaration is removed a later pass will remove all references to the now-undefined variable. The transformation also removes empty statements and private methods that have no remaining statements. Note that the AWT/Swing repair transforms cited above are run before this transformation so that this does not remove any direct user interface code.

Removing code to get it to compile can have various side effects. Additional transformations attempt to handle these. For example, if a variable assignment was removed, the compiler might detect an error in that the variable might be used before it is assigned to. Similarly, removing a statement that might have thrown an exception could make a try-catch block be no longer needed. Additional transformations check for these and similar cases and clean up the code. For example, they will add an initial assignment in the declaration of variables that might be used before assigned to, and remove unneeded try-catch blocks.

Other transforms attempt to clean up the class structure. A simpler class hierarchy will sometimes allow code to compile that otherwise would fail. One transformation here converts inner classes to stand-alone classes, adding additional fields (for the implicit this) and changing the constructor as appropriate. Another transformation removes unneeded implements clauses. A third class-structure related transform merges a subclass with its superclass to form a single standalone class. In addition to handling potential future compilation problems, these transformations tend to yield simpler resultant code.

6.3 Filtering Solutions

To keep the number of solutions reasonable, the system applies a filter that eliminates solutions that cannot be transformed to match. In the case of user interfaces, it checks to see if the code has a reference to one of each of the sets of component types needed by each of the user-specified components. This reference can either be direct or indirect (i.e. might be to a subclass rather than the class itself). If there is some user component that cannot possibly be implemented by the solution, the solution is discarded. For example, if the user diagram contains a list, the solution would have to contain either component of type *javax.swing.JList*, *java.awt.List*, or *javax.swing.JTable*. The check ignores labels since these are not critical to the resultant interface.

The number of candidate solutions can vary considerably, but generally doesn't become excessive. For example, the search involved with Figure 1 considered 116 files derived from the OpenHub search engine, generated 236 initial user interface solutions, and found 569 solutions to test out of a total of 4,122 that were generated by the various transformations, and tested the first 500 of those to produce the results shown in Figure 3. The ordering of solutions to determine which to test is a part of S^6 and is a function of the initial rank returned by the search engine, the number of transforms done, and a random value to encourage breadth.

7. Validating Solutions

The next step involves testing whether the code that was extracted as a potential user interface solution actually matches the user's sketch and meets their needs.

We take a two-step approach here. First, we match the generated user interface against the user's sketch. This match first checks that all the components of the user's sketch appear in the generated interface. If they do, then the match computes a score describing the quality of the match. The second step is to present the interfaces to the programmer, first by showing a screen shot of the interface, and second by actually running the interface and letting the programmer interact with it, explore its widget hierarchy and callbacks, and do some simple editing.

To match a generated user interface against the user's sketch, we run the generated solution and investigate the widget hierarchy that results. The code generated for each potential user interface solution returns a user interface object (instance of *java.awt.Container*) from which we extract the hierarchy using the basic methods of *Container*.

In addition to looking at the hierarchy, we ensure that the display is runnable and supports interaction. This includes putting non-window widgets inside a frame and ensuring that dialogs are non-modal. It also involves determining and setting a reasonable size for the resultant window, checking if the window can be resized and making sure the top level user component is visible.

Both the generated widget hierarchy and the user's hierarchy are trees. We use a modified form of tree matching to compare the two. The comparison is loose in that the generated hierarchy is likely to have many additional components and hierarchy levels. For example, it might be organized as multiple panels to effect a better layout; a widget might be contained in a scrolled region (which adds the scroll pane, the viewport, the scroll bars); or the top level might be a root pane with all its associated components. In addition, the actual implementation might include additional widgets that the user's sketch didn't account for. For example, in the address book example, there might be additional fields (e.g. telephone or office address) that other implementations included but the user hadn't thought of (and might want). We also allow a little leeway in the match by permitting a small set of original components (one or two, depending on the total number of components), to not be matched explicitly.

The tree matching we do effectively considers all logical assignments of the user specified components to actual widgets in the implementation. Matches need to satisfy four criteria:

- Each component has to be matched with a widget. This constraint can be relaxed to allow a small number of non-matched components.
- The top-level component needs to match the top-level widget.

- The widget matched with the component must be an object of a class that is either one of the types associated with that component (in the last stage of the user interface specification), or must be a subclass of that type. Java reflection is used to check subtypes.
- The hierarchy specified by the user's components must be reflected in the widgets. If component A is a child of component B in the specification, then the widget associated with A must be a child, either directly or indirectly, of the widget associated with B.

Because there can be an exponential number of matches (consider 20 user labels that can match 20 actual labels), the search is designed to find a reasonable match fast and will stop once a maximum number of solutions (currently 1000) have been found.

Once a match is found, we compute a heuristic score for that match. For each specified component this score takes into account

- Whether the component matched. The score is increased by 200 if so.
- How close the width and height of the component matches that of the widget. For both the width and height, if the actual value is within 100 of the user sketch value, the score is increased by 100 minus the delta.
- If text is associated with the component, the editing distance of that text versus any text associated with the wizard. Here we use reflection to call the *getText* method of the widget. The score is increased by a value between 0 and 100 depending on the quality of the match and the length of the text.
- If left, right, top, or bottom positional constraints are specified for the component, the distance in the implementation between this widget and the widget associated with the constraint. For each specified relationship, if the actual widgets are within 10 pixels, the score is increased by 50.
- Actual components that are not matched by a widget are penalized. We subtract 20 from the score for each extra label, 40 for each extra button or combo box, and 60 for each text field, list, table or tree.

The scoring tries to take into account the relative importance of each factor in assessing the match. It is designed so that obvious matches will have the highest score and be shown to the user first. Because the number of matches to date has not been excessive, the particular values chosen for scoring are not that important.

The next step is to get the user's opinion and validation for each of the matched interfaces. For each solution that matches the specification, S⁶ creates two results for further matching. The first is an image of the widget as a PNG file and the second is a runnable JAR file that can be used to explore the widget. These are passed back to the front end along with a unique identifier and the score for each solution.

The interface for asking the programmer about the interfaces is shown in Figure 3. The programmer is shown the static images of each of the candidate solutions along with an *Accept* and a *Reject* button for each. The solutions are ordered according to their score. In addition, by clicking on the solution itself, the programmer will bring up two windows, one containing the solution that the user can interact with, and a second one that displays a tree showing the widget hierarchy of the solution at the top and a display of all the events that occur when the user interacts with the

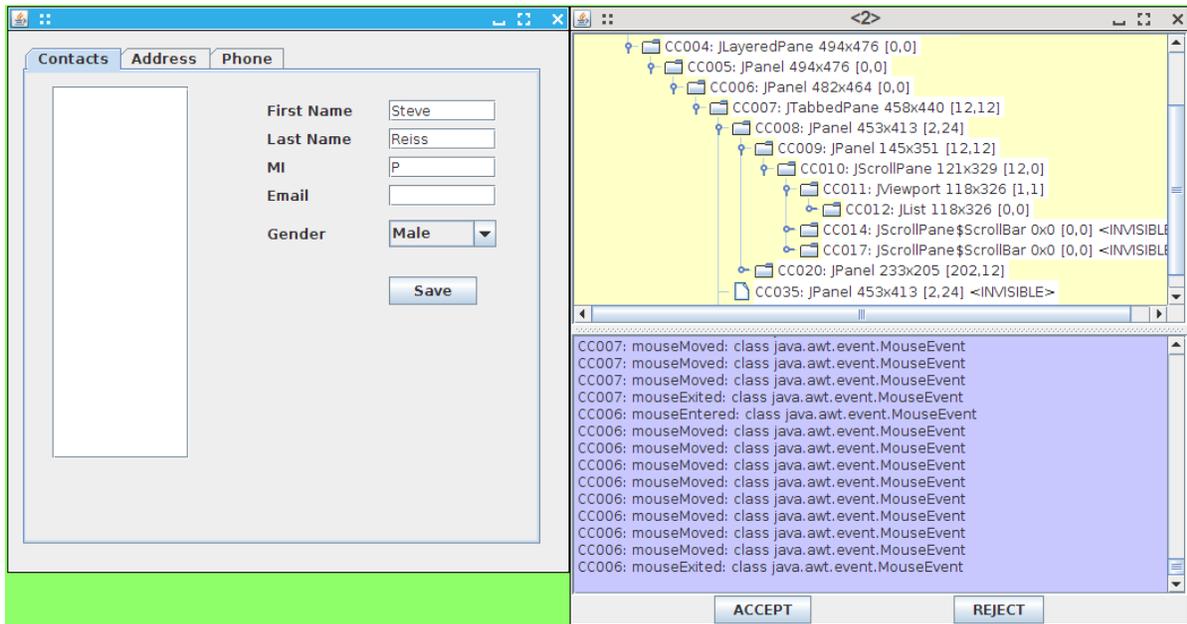


FIGURE 9. A sample solution in an interactive window on the left and the exploration window showing the widget hierarchy at the top and the events from interaction.

window at the bottom. An example can be seen in Figure 9. The interaction window provides the user with the option of accepting or rejecting the given solution.

Once the user is done perusing the returned solutions, they can hit the button at the bottom of the panel in Figure 3 to get corresponding code for any accepted solutions. If no solutions are accepted, then the back end will attempt to continue the search to find additional solutions. If solutions are accepted, the code will be returned in a browsable window such as that shown in Figure 4. The user can cut and paste the code from here into their application. Along with the actual code, information is available about the license under which the code is released.

8. Experience

To test and evaluate our approach for generating user interfaces using code search, we first obtained sketches of user interfaces. We did this by doing a web search for images using “user interface sketch”. We then culled the result for usable sketches that represented potential Java applications (as opposed to web pages or phone applications) and manually converted those sketches into SVG files. In addition to the address book example shown in Figure 1, we used the sketches shown in Figure 10. The test cases then were:

- Login: a sample login screen with a remember me button.
- Pizza: an interface for ordering a pizza with different options.
- Pizza1: similar to Pizza except we only have one list for ingredients rather than two.
- Phone: an interface for making a phone call.
- Mail: a mail reader interface.
- Student: a front end to a student information system.
- Comment: an interface for entering comments in a guest book.

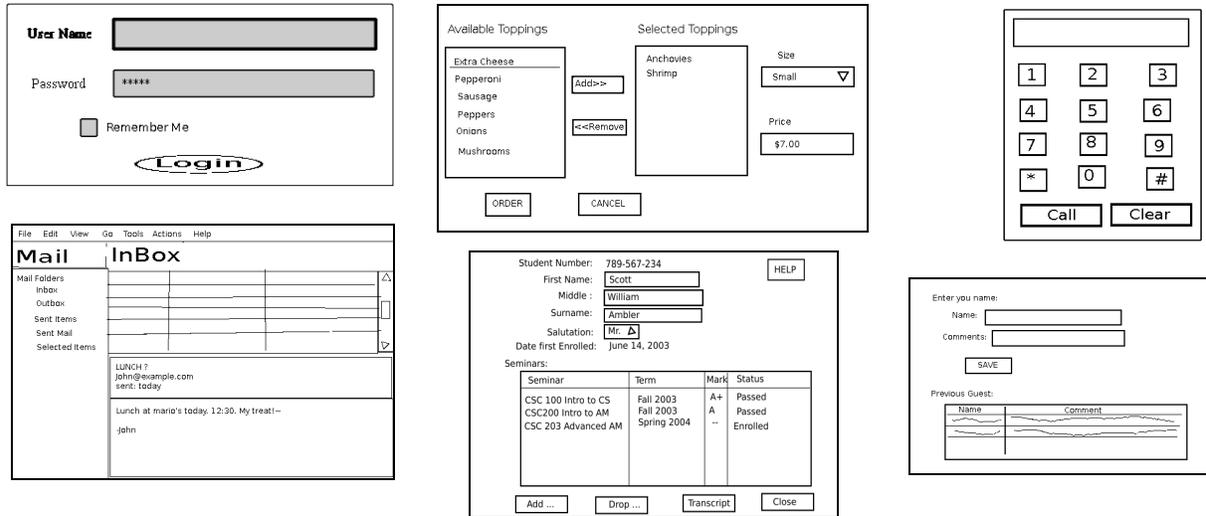


FIGURE 10. User interface sketches used for testing in addition to the address book of Figure 1. From left to right these are Login, Pizza, Phone, Mail, Student, and Comment.

- Address: an interface for maintaining an address book

For each example we tried appropriate keywords. Finding the right set of keywords required some trial and error and we eventually developed a front end for code search that made this easier [58]. We also attempted to find the interface within a file where possible, but did some experiments with larger scopes. For each case we looked at the results that were returned and verified that at least one of the results was a relevant good match for the initial sketch. A summary of the experi-

Table 1: Experimental Results

Test	Keywords	Engine	Scope	Potential Solutions Initial/Total/ Runnable/Tested	Found	Time
Login	login jcheckbox jpassword	OHLOH	FILE	138/2014/453/101	45	3:00
Login	login jcheckbox jpassword	GITHUB	FILE	89/977/231/44	18	1:42
Pizza	jlist jbutton jtextfield jcombobox restaurant	GITHUB+ OHLOH	FILE	15/1284/142/38	17	2:38
Pizza1	jlist jbutton jtextfield jcombobox restaurant	GITHUB	FILE	8/415/20/9	7	0:43
Mail	tree text editor button mail	GITHUB	FILE	88/626/48/15	5	1:02
Mail	tree text editor button mail	GITHUB	PACKAGE	96/46502/351/2	2*	226:33
Mail	tree text editor button mail	GITHUB	SYSTEM	100/47415/340/2	2*	303:27
Phone	phone jbutton	OHLOH	FILE	127/1077/184/9	8	1:15
Phone	phone jbutton	GITHUB	FILE	109/1067/199/14	14	1:23
Student	student jbutton jtext jtable	OHLOH	FILE	91/499/500/47	47*	3:52
Comment	feedback jtextfield jtable	OHLOH	FILE	112/4977/500/178	114*	8:02
Address	name email phone jlist	OHLOH	FILE	133/4122/500/79	131*	4:16
Address	name email phone jlist	GITHUB	FILE	121/3902/500/154	66*	5:48
Address	name email phone jlist	GITHUB	PACKAGE	115/24650/371/25	19*	69:15

ments is shown in Table 1. (Note that OHLOH was renamed OpenHub and is was recently taken off-line.)

The first column of the table indicates the test name. The second column shows the keywords used in the test and the third and fourth indicate the search engine and search scope respectively. The keywords basically describe the application and then include the name of one or more of the expected widgets. This type of search worked well, with the descriptive keywords narrowing the search to appropriate applications and the widget keywords eliminating non-Swing applications. Moreover, where we were doing a file-scoped search, requiring both to occur in the same file gave better results.

The fifth column (*Potential Solutions*) provides some indication of the work done in the search. The first number is the number of starting solutions derived from the returned search results. This is generally the number of unique files found by the search engine. The second number is the total number of solutions that were generated during the search. The third number is the number of solutions that could be compiled and tested while the fourth is the number of solutions that were judged acceptable. The sixth column reports the number of distinct images that were generated and hence the number of distinct passing tests. This is generally not the same as the number of acceptable solutions since there are often solutions that are variants of the same original source and that generate essentially the same user interface. In this case only one is shown to the user initially (although if that is deemed acceptable, the code for all solutions is returned).

The seventh column (*Time*) of the table indicates the wall time (in minutes and seconds) that the search and testing took. The search was run using eight cores of a sixteen core machine. The process is highly parallelizable and these numbers are dependent on the number of threads being used by the search. The time generally does not include wait time in accessing the underlying

search engine since we are caching the initial search results in order to lessen the load on the search engines.

The items which are starred in the sixth (*Found*) column indicate that there might be additional solutions, but that S⁶ stopped looking because an initial set of solutions were found. S⁶ orders solutions based on their initial ranking from the search engines, the number of transforms used, and a random value, and limits the number of solutions it considers at each stage using this ranking. The unchecked solutions are held in abeyance to be checked later if no other solutions are found. The limits include a maximum of 500 solutions to test and a maximum of 2000 active intermediate solutions.

This can be seen most readily when doing searches at the PACKAGE and SYSTEM levels where there are many more potential solutions to consider. This is the reason that these searches returned fewer solutions than the corresponding FILE searches. These searches also take considerable longer. This is due first to the much large set of solutions considered and second to the fact that each of these solutions is substantially larger and hence requires more processing.

While the approach has some problems, it also shows a lot of promise. Using code search it is possible to return working code for a user interface based solely on a sketch and a set of keywords. We see three primary uses for a tool like this.

8.1 Generating User Interface Code

The obvious use is to produce an actual working interface that is similar to the user's original sketch. While this could be done using a user interface builder as is common in many of today's development environments, using existing code as a starting point has several potential advantages. First, user interfaces in the repository have typically been used and tested including user testing in real applications. Second, such interfaces cover conditions that might not have been anticipated by the original sketch. For example, a sample sketch we did for addresses did not have a zip code field while the generated ones did. Third, many interfaces in the repository handle window resizing appropriately, something that many of today's user interface builders have difficulty with. Fourth, interfaces in the repository typically are more consistent with other interfaces and thus with user expectations. Fifth, the generated can often supports interactivity, both between elements of the sketch and with other pieces of the user interface, for example buttons that are only enabled when fields are completed. For example, in a returned phone interface, pushing the buttons entered digits into the display and allowed typing into the display area; one of the returned Pizza1 examples automatically updated the prices as items were added to the order. Sixth, the code that is returned is often more sophisticated and functional that which would be generated by a user interface builder. In particular, the generated code often included proper layout techniques, more sophisticated hierarchies, scrolling where needed, correct adaptation to changing window sizes, etc. Moreover, the returned interfaces often provide validation code that highlighted missing or erroneous fields and that only activated buttons when the inputs needed for them were correct.

While generating user interface code might be a worthwhile goal, actually doing so is problematic. The main issue is the quality and nature of the code that is returned. The system returns compilable, running code that can be copied and pasted into a user project and used directly. However, the quality of user interface code in the repositories varies widely and some of it should probably

not be propagated. Moreover, the code only includes the user interface and hence will need to be modified in order to integrate it into the rest of the system. Many programmers would prefer that code they will have to work on be written in the style and with the conventions they are used to. While some of this can be taken into account by our search tools (they are able to reformat code in standard styles using [54]) or by formatting commands in the user's programming environment, the result will still be code that the users would not want to use in their project.

8.2 Browsing Potential User Interfaces

A more important use for the tool is as a means for exploring user interfaces. User interface sketches, especially those done in the early stages of a project, are usually exploratory rather than definitive. Users often have one or more ideas that they might want to later fill out and complete. The approach taken in SUISE supports this.

First, the user can take a sketch and then view other, existing user interfaces that are similar to the sketch. This lets the user see what might be missing (for example, the zip code field mentioned above), see the range of applications with similar interfaces, and look at alternative layouts and presentations. For example, Figure 3 shows several different layouts for address book information.

Second, the user can interact with the interface representing the sketch and check whether the interaction is appropriate to the application. This can show some of the additional features that are included in the code. It also lets the programmer get a better sense of how the interface might be used. Interaction can also show how the interface deals with different window or screen sizes, how fields and buttons interact, etc.

Third, the tool can be used as a starting point for exploring a broader range of applications for a potential application. The matching algorithm used by SUISE insures that most of the functionality specified by the user is provided, but allows other functionality to be included as well. It also allows considerable leeway on the layout and positioning. The user can sketch a simple form of the interface they are interested in and then use the tool to understand a range of alternatives based on that form.

This is seen in the simplified Pizza interface example. Here the sketch included only the minimum that would be needed for a pizza-ordering interface. The interfaces that match this are typically going to be more sophisticated and include additional functionality. Some of the matching interfaces, for example, are shown in Figure 11. While these don't match the original sketch, they all are functional interfaces for ordering a pizza, include other components beyond what was initially specified (e.g. thick or thin crust) which the user might have overlooked, and illustrate different approaches that the eventual pizza interface could take.

8.3 Browsing User Interfaces

A third use for the tool is as a part of a front end for exploring code repositories. While GitHub and other search engines let the user explore the code in the repository, they do not show the user interfaces that result from that code. The technology included in SUISE can potentially take a user interface code file and transform it into a running example that the user can either simply view or eventually interact with.

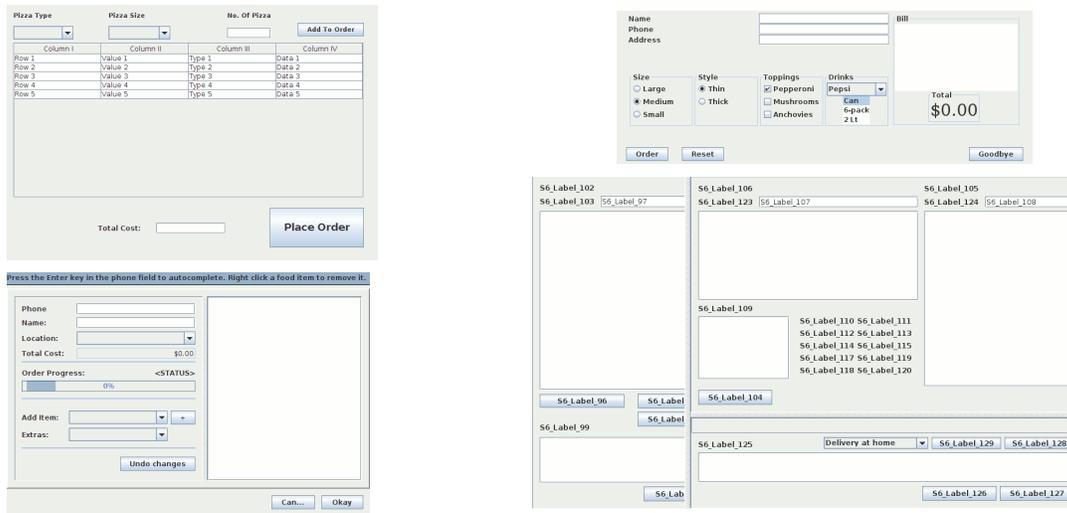


FIGURE 11. Different Pizza ordering interface found by code search for example Pizza1.

In a related project, we developed a front end for searching code repositories based on the Code Bubbles programming environment [58]. The goal of this project was to provide a better environment whereby the programmer can explore and work with code from existing repositories. We wanted to provide all the navigation features offered by a programming environment such as Code Bubbles in an environment for browsing code repositories. We also wanted to offer an environment that would provide the feedback and editing necessary for the programmer to extract and reuse code from the repository.

Code Bubbles typically runs as a separate tool on top of Eclipse using a message-based plug-in mechanism [57]. It includes a small Eclipse plug-in which connects to a message bus that the main environment talks to. Integration is achieved using command messages from Code Bubbles to Eclipse and informational messages from Eclipse to Code Bubbles. Code Bubbles uses Eclipse to handle file access and editing, semantic and syntactic searching, compilation, and execution.

Our approach was to replace the Eclipse plug-in with a repository browsing plug-in. This plug-in handles the same basic set of commands that the Eclipse plug-in does except for those related to debugging. The back end also supports a small set of additional search-oriented commands such as initiate a code search against a set of repositories, augment an existing file with other files from same package or project, find a definition of a name in a given project, return the next set of results for a code search, mark a package, file, or section of a file as accepted, and output all accepted definitions to a given directory.

The back end handles going out to various search engines, extracting the source files, and then simulating the corresponding file hierarchy for Code Bubbles. For each source file, it supports editing and both text and semantic-based search. For each project, it supports compilation on demand (as well as after any edits) with feedback of errors. The back end does its own error-tolerant compilation, necessary because individual files separated from their original package and libraries, will generally not compile cleanly.

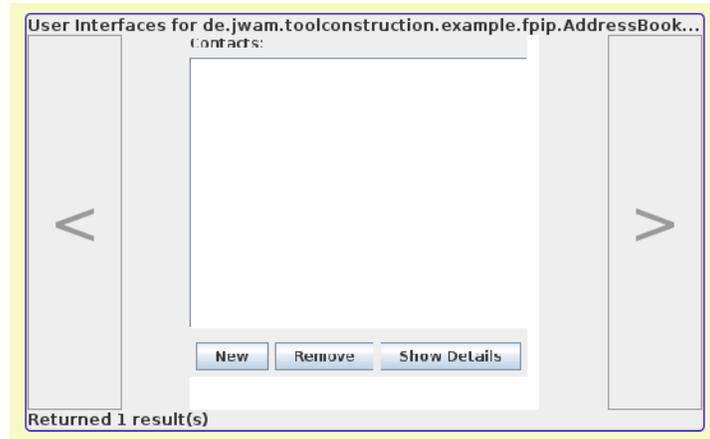


FIGURE 12. User interface bubble showing the user interfaces inside a repository file.

Code Bubbles [5,6] is an attempt to redesign the user interface to programming, making the programming environment conform to the programmer's working model. It does this by displaying and manipulating complete working sets, collections of task-relevant fragments including code, documentation, test cases, notes, bug reports, and other aspects of programming [29,42]. The fragments in a working set may be contained in multiple files, classes, or other modules, so quick and easy viewing of all these at once is complicated in traditional IDEs. Code Bubbles presents fragments in fully manipulatable interface elements in order to provide an intuitive arrangement of working sets. The fragments returned from code search, be they methods, classes or files, are a good match for this type of exploratory environment.

Code Bubbles includes a variety of features to simplify and support navigation in addition to what is typically provided by development environments such as Eclipse or NetBeans. It provides a fast search facility that can be used to look at the code hierarchically or to quickly find classes and methods with given name fragments. It supports popping up bubbles for definition, references, and implementations by pointing at a name and either hitting a function key or selection off a context menu. It provides class overviews that include comment fragments and can be easily refined. It provides typical environmental facilities to show errors and quickly navigate to them. Our search framework was designed to make use of all these facilities, letting users explore repository code just as they would explore their own projects.

We augmented this interface with a facility to show pictures of the user interfaces that might be implemented by a given file. The user selects the class or file in question and right clicks to select Show User Interfaces. The repository version of Code Bubbles knows the original source of the file or class from the repository and sets up a User interface search request similar to what SUISE would do but specific to that particular source. This request includes a very simple hierarchical component specification consisting of a single empty component. The back end then will access the original code from the repository, apply the various transforms to get it to represent a runnable user interface, and then test it. Any user interfaces that match (regardless of the scores) are then returned to the front end as images. A bubble in the front end then can display the different user interfaces in turn. An example of such a bubble is shown in Figure 12.

8.4 Weaknesses

While the approach works in many cases, it still has weaknesses. The first is that the results are sensitive to the initial selection of keywords. This is due to the fact that the search facilities provided by existing code search engines are primitive by modern standards both because of the search techniques used and the difficulty of mapping keywords to programs. Substantial work has been and continues to be done on code search which should address these problems in the future. Code Exchange at UCI, for example, can give very good results initially [37-39]. We expect that these new approaches will find their way into repository search engines over time and that code search will improve significantly. Since our tool is built on top of existing search engines, we can easily piggy back on these improvements. Moreover, the time our tool takes is small enough when doing file-based search so that a programmer can afford to run it multiple times with different keywords.

A second problem is the time taken to do the search, especially if the search is at the PACKAGE or SYSTEM scope. The three cases we considered here took hours to complete. There are several difficulties here. First, they tend to yield a large number of potential solutions that need to be explored. Second, the size of these solutions (500k-5M characters), and the complexity of analyzing and transforming solutions of this size, means that each solution takes significantly longer to be evaluated. Moreover, the large number of candidates from each solution means that solutions returned early by the initial search tend to dominate and solutions that are returned later are not considered fully or sometimes at all. These are problems that can be addressed by doing a better job of restricting the solutions initially and during the search.

A third problem involves the use of user interface libraries. Our efforts to date have concentrated on code that uses Swing and AWT directly. Complex applications often use a third party user interface library, of which there are quite a few. Code that uses such libraries generally can not be made to compile in a useful manner. It would be relatively simple to extend our tool to let the user indicate which if any third party libraries should be allowed and to incorporate those into the search process.

A fourth problem is that more complex interfaces tend to be tightly integrated with the rest of the application and the rest of the application often depends on external packages or external systems such as databases. For example, mail applications would typically be built using a table model that is tied either to a database of mail messages, to a cache front end, or to a sophisticated imap interface. While we have developed a number of transformations to extract user interface code, additional, more sophisticated transformations would let us find and return more running examples.

8.5 Threats To Validity

In addition to the weaknesses cited, there are several things we should note that might affect the utility and efficacy of the approach and the results of the study. These include:

- The set of sketches chosen might not be representative of what programmers are actually interested in. Sketches available on the web tend to be for sample applications, not for real world code.

- The results are dependent on the set of keywords chosen and it is not clear that other users would be able to choose the proper keywords to get appropriate results for a search.
- Our results only look at Java programs with Swing/AWT interfaces. For other languages and user interface libraries the repositories might not have enough samples to let the system find runnable code. Moreover, different and additional transformations would be needed in these cases.
- We have shown that we can return interfaces from code repositories, not that the interfaces that are returned are actually useful, either as starting points or as runnable code. This would require a very different and much more extensive study and is more appropriate after the tool has been further developed.

9. Conclusions

Our work demonstrates that it is possible to generate a complete, working, interactive user interface from a sketch using code search. We have developed the techniques needed to translate a sketch into something checkable, to extract the proper code from existing code search engines, to transform that code into a program that compiles and runs and includes only the user interface, and then to let the user interact with and select the results of interest.

Some of the lessons learned in the process that will be applicable to future work in this area include:

- With existing code search engines, the results returned are very dependent on the selection of search terms.
- The performance of these techniques is acceptable for interfaces that are contained in a single file; where the necessary code is spread across multiple files, improved performance will be needed.
- Additional transforms would yield additional solutions.
- More flexibility in matching the specifications to the generated user interface lets the technique be used for exploration.
- Similarly, it is often better to have the user under specify the interface, both for finding solutions and to facilitate exploration.
- The code returned is interactive and generally does more than the code that would be generated by a simple user interface builder.
- Additional work is required to transform the resultant code into something that programs would feel comfortable including directly into their applications.

The code for our implementation of user interface generation by code search is available as part of the S⁶ code search tool and can be found at <ftp://ftp.cs.brown.edu/u/spr/s6.tar.gz>. The test cases (SVG files) are available upon request.

10. Acknowledgments

This work is supported by the National Science Foundation grant CCF1130822. Additional support has come from Microsoft and Google.

11. References

1. Marat Akhin, Nikolai Tillmann, Manual Fahndrich, Jonathan de Halleux, and Michal Moskal, "Search by example in touch develop: code search made easy," *Proceedings SUITE 2013*, pp. 5-8 (June 2012).
2. Sushil Bajracharya, Trung Ngo, Erik Linstead, Yimeng Dou, Paul Rigor, Pierre Baldi, and Cristina Lopes, "Sourcerer: a search engine for open source code supporting structure-based search," *Proceedings ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications 2006*, pp. 682-682 (October 2006).
3. Sushil Bajracharya, Joel Ossher, and Cristina Lopes, "Sourcerer: an infrastructure for large-scale collection and analysis of open-source code," *Science of Computer Programming* **79** pp. 241-259 (2014).
4. Andrew Begel, "Codifier: a programmer-centric search user interface," *Workshop on Human-Computer Interaction and Information Retrieval*, (October 2007).
5. Andrew Bragdon, Steven P. Reiss, Robert Zeleznik, Suman Karumuri, William Cheung, Joshua Kaplan, Christopher Coleman, Ferdi Adeputra, and Joseph J. LaViola, Jr., "Code bubbles: rethinking the user interface paradigm of integrated development environments," *ACM/IEEE International Conference on Software Engineering 2010*, pp. 455-464 (2010).
6. Andrew Bragdon, Steven P. Reiss, Robert Zeleznik, Suman Karumuri, William Cheung, Joshua Kaplan, Christopher Coleman, Ferdi Adeputra, and Joseph J. LaViola, Jr., "Code bubbles: a working set-based interface for code understanding and maintenance," *Proceedings SIGCHI Conference on Human Factors in Computing Systems*, pp. 2503-2512 (2010).
7. Geoff Bull and Junbin Gao, *Classification of hand-written digits using choriograms* 2011.
8. Anabela Caetano, Neri Goulart, Manuel Fonseca, and Joaquim Jorge, "JavaSketchIt: issues in sketching the look of user interfaces," *AAAI Spring Symposium on Sketch Understanding*, pp. 9-14 (2002).
9. Shih-Chien Chou, Jen-Yen Chen, and Chyan-Goei Chung, "A behavior-based classification and retrieval technique for object-oriented specification reuse," *Software Practice and Experience* **26**(7) pp. 815-832 (July 1996).
10. Shih-Chien Chou and Yuan-Chien Chen, "Retrieving reusable components with variation points from software product lines," *Information Processing Letters* **99** pp. 106-110 (2006).
11. Adam Coates, Blake Carpenter, Carl Case, Sanjeev Satheesh, Bipin Suresh, Tao Wang, David J. Wu, and Andrew Y. Ng, "Text detection and character recognition in scene images and unsupervised feature learning," *IEEE International Conference on Document Analysis and Recognition*, (2011).
12. Christopher G. Drummond, Dan Ionescu, and Robert C. Holte, "A learning agent that assists the browsing of software libraries," *IEEE Transactions on Software Engineering* **26**(12) pp. 1179-1196 (December 2000).
13. R. O. Duda and P. E. Hart, "Use of the Hough transformation to detect lines and curves in pictures," *Communications of the ACM* **15** pp. 11-15 (January 1972).
14. Michael D. Ernst, Raimondas Lencevisius, and Jeff H. Perkins, "Detection of web service substitutability and composability," *WS-MaTe 2006: International Workshop on Web Services -- Modeling and Testing*, pp. 123-135 (June 2006).
15. William B. Frakes and Thomas P. Pole, "An empirical study of representation methods for reusable software components," *IEEE Transactions on Software Engineering* **20**(8) pp. 617-630 (August 1994).
16. Krzysztof Z. Gajos, Daniel S. Weld, and Jacob O. Wobbrock, "Decision-theoretic user interface generation," *In Proceedings of the 22nd AAAI Conf. on Artificial Intelligence (AAAI-08)*, (2008).
17. Ben Goodger, Ian Hickson, David Hyatt, and Chris Waterson, "XML user interface language (XUL) 1.0 Specification," <http://www.mozilla.org/projects/xul/xul.html> (2003).
18. Mark Grechanik, Chen Fu, Qing Xie, Collin McMillan, Denys Poshyvanyk, and Chad Cumby, "A search engine for finding highly relevant applications," *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering*, (May 2010).
19. Saul Greenberg, "Toolkits and interface creativity," *Journal on Multimedia Tools and Applications* **32** pp. 139-159 (2007).
20. Sonia Haiduc, Gabriele Bavota, Rocco Oliveto, Andrian Marcus, and Andrea De Lucia, "Evaluating the specificity of text retrieval queries to support software engineering tasks," *Proceedings of the 2012 International Conference on Software Engineering*, pp. 1273-1276 (2012).

21. Sonia Haiduc, Gabriele Bavota, Andrian Marcus, Rocco Oliveto, Andrea De Lucia, and Tim Menzies, "Automatic query reformulations for text retrieval in software engineering," *Proceedings of the 2013 International Conference on Software Engineering*, pp. 842-851 (2013).
22. Robert J. Hall, "Generalized behavior-based retrieval," *Proceedings International Conference on Software Engineering*, 93, pp. 371-380 (May 1993).
23. Bjorn Hartmann, Leith Abdulla, Manas Mittal, and Scott R. Klemmer, "Authoring sensor based interactions through direct manipulation and pattern matching," *Proceedings of chi 2007: ACM Conference on Human Factors in Computing Systems*, pp. 145-154 (2007).
24. Emily Hill, Lori Pollock, and K. Vijay-Shanker, "Automatically capturing source code context of NL- queries for software maintenance and reuse," *International Conference on Software Engineering 2009*, (May 2009).
25. Raphael Hoffmann and James Fogarty, "Assieme: finding and leveraging implicit references in a web search interface for programmers," *Proceedings UIST 2007*, pp. 13-22 (October 2007).
26. Oliver Hummel, Werner Janjic, and Colin Atkinson, "Code Conjurer: pulling reusable software out of thin air," *IEEE Software* 25(5) pp. 45-52 ().
27. Werner Janjic, Dietmar Stoll, Philipp Bostan, and Colin Atkinson, "Lowering the barrier to reuse through test-driven search," *SUITE*, 09, pp. 21-24 (May 2009).
28. Werner Janjic and Colin Atkinson, "Leveraging software search and reuse with automated software adaptation," *Proceedings SUITE 2013*, pp. 23-26 (June 2012).
29. Andrew J. Ko, Htet Aung, and Brad A. Myers, "Eliciting design requirements for maintenance-oriented IDEs: a detailed study of corrective and perfective maintenance tasks," *Proceedings of the 27th International Conference on Software Engineering*, pp. 126-135 (2005).
30. J. A. Landay and B. A. Myers, "Sketching interfaces: toward more human interface design," *Computer* 34(3) pp. 56-64 (March 2001).
31. Otavio Lemos, Sushil Bajracharya, Joel Ossher, Ricardo Morla, Paulo Masiero, Pierre Baldi, and Cristina Lopes, "CodeGenie: using test-cases to search and reuse source code," *ASE*, 07, pp. 525-526 (November 2007).
32. Otavio Lemos, Sushil Bajracharya, Joel Ossher, Paulo Masiero, and Cristina Lopes, "Applying test-driven code search to the reuse of auxiliary functionality," *Proceedings ACM Symposium on Applied Computing*, pp. 476-482 (2009).
33. Otavio Augusto Lazzarini Lemos, Sushil Bajracharya, Joel Ossher, Paulo Cesar Masiero, and Cristina Lopes, "A test-driven approach to code search and its application to the reuse of auxiliary functionality," *Information and Software Technology* 53(4) pp. 294-306 (April 2011).
34. Greg Little and Robert C. Miller, "Keyword programming in Java," *Proceedings ASE 2007*, pp. 84-93 (November 2007).
35. Dapeng Liu, Andrian Marcus, Denys Poshyvanyk, and Vaclav Rajlich, "Feature location via information retrieval based filtering of a single scenario execution trace," *Proceedings of the twenty-second IEEE/ACM international conference on Automated software engineering*, pp. 234-243 (2007).
36. Yoelle S. Maarek, Daniel M. Berry, and Gail E. Kaiser, "An information retrieval approach for automatically constructing software libraries," *IEEE Transactions on Software Engineering* 17(8) pp. 800-813 (August 1991).
37. L. Martie, T. D. LaToza, and A. van der Hoek, "CodeExchange: supporting reformulation of code queries in context," *Proceedings of the 30th International Conference on Automated Software Engineering*, (2015).
38. Lee Martie and Andre Van der Hoek, "Toward social- technical code search," *6th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE)*, pp. 101-104 (2013).
39. Lee Martie and Andre Van der Hoek, "Sameness: an experiment in code search," *Kmhsamen*, pp. 76-87 (2015).
40. Collin McMillan, Mark Grechanik, Denys Poshyvanyk, Qing Xie, and Chen Fu, "Portfolio: finding relevant functions and their usage," *Proceeding of the 33rd International Conference on Software engineering*, (May 2011).
41. Gerrit Meixner, Fabio Patern, and Jean Vanderdonckt, "Past, present, and future of model-based user interface development," *i-com* 10(3) pp. 2-11 (2011).
42. B. A. Meyers, A. J. Ko, M. J. Coblenz, and H. H. Aung, "An exploratory study of how developers seek, relate, and collect relevant information during software maintenance tasks," *IEEE Transactions on Software Engineering* 32(12) pp. 971-987 (2006).

43. Alon Mishne, Sharon Shoham, and Eran Yahav, "Typestate-based semantic code search over partial programs," *SIGPLAN Notices* **47**(10) pp. 997-1016 (2012).
44. Brad Myers, Scott E. Hudson, and Randy Pausch, "Past, present and future of user interface software tools," *ACM Transactions on Computer-Human Interaction* **7**(1) pp. 3-28 (March 2000).
45. Lukas Neumann and Jiri Matas, "Real-time scene text localization and recognition," *IEEE Conference on Computer Vision and Pattern Recognition*, (2012).
46. T. A. Nguyen and C. Csallner, "Reverse engineering mobile application user interfaces with REMAUI," *Proceedings of Automated Software Engineering*, pp. 248-259 (November 2015).
47. Jeffrey Nichols and Andrew Faulring, "Automatic interface generation and future user interface tools," *ACM CHI 2005 Workshop on The Future of User Interface Design Tools*, (2005).
48. Stina Nylander, "Semi-automatic generation of device adapted user interfaces," *UIST conference companion*, (October 2005).
49. Santanu Paul and Atul Prakash, "A framework for source code search using program patterns," *IEEE Transactions on Software Engineering* **20**(6) pp. 463-475 (June 1994).
50. Rejean Plamondon and Sargur N. Srihari, "Online and off-line handwriting recognition: a comprehensive survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(1) pp. 63-84 (2000).
51. Andy Podgurski and Lynn Pierce, "Retrieving reusable software by sampling behavior," *ACM Transactions on Software Engineering and Methodology* **2**(3) pp. 286-303 (July 1993).
52. David Raneburger, Roman Popp, and Jean Vanderdonckt, "An automated layout approach for model-driven WIMP-UI generation," *Proceedings of the 4th ACM SIGCHI symposium on Engineering interactive computing systems (EICS '12)*, pp. 91-100 (2012).
53. Steven P. Reiss, "A component model for Internet-scale applications," *Proceedings ASE 2005*, pp. 34-43 (November 2005).
54. Steven P. Reiss, "Automatic code stylizing," *Proceedings ASE '07*, pp. 74-83 (November 2007).
55. Steven P. Reiss, "Semantics-based code search," *International Conference on Software Engineering 2009*, pp. 243-253 (May 2009).
56. Steven P. Reiss, "Specifying what to search for," *Proceedings SUITE 2009*, (May 2009).
57. Steven P. Reiss, "Plugging in and into Code Bubbles," *Proceedings Workshop on Developing Tools as Plug-ins 2012*, pp. 55-60 (June 2012).
58. Steven P. Reiss, "Browsing software repositories," *Unpublished manuscript available at <http://www.cs.brown.edu/people/spr/rebuspaper.pdf>*, (2014).
59. Marco de Sa, Lums Carrigo, Lums Duarte, and Tiago Reis, "A mixed-fidelity prototyping tool for mobile devices," *Proceedings of the Working Conference on Advanced Visual Interfaces*, pp. 225-232 (2008).
60. Julian Seifert, Bastian Pflöging, Elba del Carmen Valderrama Bahamondez, Martin Hermes, Enrico Rukzio, and Albrecht Schmidt, "Mobidev: a tool for creating apps on mobile phones," *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*, pp. 109-112 (2011).
61. Paulo Pinheiro da Silva, "User interface declarative models and development environments: a survey," *Proceeding of the 7th International Conference on Design, Specification, and Verification of Interactive Systems*, pp. 207-226 Springer-Verlag, (2000).
62. Bunyamin Sisman and Avinash C. Kak, "Assisting code search with automatic query reformulation for bug localization," *Proceedings of the 10th Working Conference on Mining Software Repositories*, pp. 309-318 (2013).
63. Kathryn T. Stolee and Sebastian Elbaum, "Toward semantic search via SMT solver," *Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering*, pp. 1-4 (2012).
64. Kathryn T. Stolee, Sebastian Elbaum, and Daniel Dobos, "Solving the Search for Source Code," *ACM Trans. Softw. Eng. Methodol.* **23**(3) pp. 1-45 (2014).
65. Vijayan Sugumaran and Veda C. Storey, "A semantic-based approach to component retrieval," *Advances in Information Systems* **34**(3) pp. 8-24 (2003).
66. Watanabe Takuya and Hidehiko Masuhara, "A spontaneous code recommendation tool based on associative search," *Proceedings of the 3rd International Workshop on Search-Driven Development: Users, Infrastructure, Tools, and Evaluation*, pp. 17-20 (2011).

67. Stephen W. Thomas, "Mining Unstructured Software Repositories using IR Models," *Ph.D. Dissertation from Queen's University, Canada*, (2012).
68. Suresh Thummalapenta and Tao Xie, "PARSEWeb: a programmer assistant for reusing open source code on the web," *Proceedings ASE, 07*, pp. 204-213 (November 2007).
69. Alexander Toshev, Ben Taskar, and Kostas Danilidis, "Shape-based object detection via boundary structure segmentation," *International Journal of Computer Vision* **99**(2) pp. 123-146 (2012).
70. Taciana A. Vanderlei, Frederico A. Duraõ, Alexandre C. Martins, Vinicius C. Garcia, Eduardo S. Almeida, and Silvio R. de L. Meira, "A cooperative classification mechanism for search and retrieval software components," *Proceedings SAC, 07*, pp. 866-871 (March 2007).
71. Jue Wang, Yingnong Dang, Hongyu Zhang, Kai Chen, Tao Xie, and Dongmei Zhang., "Mining succinct and high-coverage API usage patterns from source code," *Proceedings of the 10th Working Conference on Mining Software Repositories (MSR 2013)*, pp. 319-328 (May 2013).
72. Shaowei Wang, David Lo, and Lingxiao Jiang, "Code search via topic-enriched dependence graph matching," *18th Working Conf. on Reverse Engineering*, pp. 119-123 (2011).
73. Doug Wightman, Zi Ye, Joel Brandt, and Roel Vertegaal, "SnipMatch: using source code context to enhance snippet retrieval and parameterization," *Proceedings of the 25th annual ACM symposium on User Interface Software and Technology*, pp. 219-228 (2012).
74. Victor Wu, Raghavan Manmatha, and Edward M. Riseman, "Textfinder: an automatic system to detect and recognize text in images," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22** pp. 1224-1229 (1999).
75. Haining Yao and Letha Eitzkorn, "Towards a semantic- based approach for software reusable component classification and retrieval," *ACMSE, 04*, pp. 110-115 (April 2004).
76. Yunwen Ye and Gerhard Fischer, "Supporting reuse by delivering task relevant and personalized information," *Proceedings International Conference on Software Engineering, 02*, pp. 513-523 (May 2002).
77. Yunwen Ye, "Programming with an intelligent agent," *IEEE Intelligent Systems* **18**(3) pp. 43-47 (May 2003).