

# Type-Directed Code Reuse using Integer Linear Programming

Yuepeng Wang  
University of Texas at Austin  
ypwang@cs.utexas.edu

Arati Kaushik  
University of Texas at Austin  
arati@cs.utexas.edu

Yu Feng  
University of Texas at Austin  
yufeng@cs.utexas.edu

Isil Dillig  
University of Texas at Austin  
isil@cs.utexas.edu

Ruben Martins  
University of Texas at Austin  
rmartins@cs.utexas.edu

Steven P. Reiss  
Brown University  
spr@cs.brown.edu

## ABSTRACT

In many common scenarios, programmers need to implement functionality that is already provided by some third party library. This paper presents a tool called HUNTER that facilitates code reuse by finding relevant methods in large code bases and *automatically synthesizing* any necessary wrapper code. The key technical idea underlying our approach is to use *types* to both improve search results and guide synthesis. Specifically, our method computes *similarity metrics* between types and uses this information to solve an *integer linear programming (ILP)* problem in which the objective is to minimize the cost of synthesis. We have implemented HUNTER as an Eclipse plug-in and evaluate it by (a) comparing it against S<sup>6</sup>, a state-of-the-art code reuse tool, and (b) performing a user study. Our evaluation shows that HUNTER compares favorably with S<sup>6</sup> and significantly increases programmer productivity.

## 1. INTRODUCTION

In many common scenarios, programmers need to implement functionality that is already provided by some third party library. In such cases, accepted wisdom dictates that it is preferable to reuse existing code rather than reimplementing the desired functionality from scratch. Indeed, there are at least two benefits to software reuse: First, it increases programmer productivity, allowing developers to focus on more creative tasks. Second, implementations provided by existing APIs are typically well-tested; hence, code reuse decreases the likelihood that the implementation is buggy.

Unfortunately, there is often a wide gap between the adage “software reuse is good” and the reality of software development: In many cases, programmers end up reimplementing functionality that already exists in a library. This duplicated effort is sometimes inadvertent because programmers may not know whether the desired functionality exists or which third-party library provides it. In other cases, the exposed interface may not quite fit the developer’s needs, tempting programmers to reimplement the desired functionality.

This paper seeks to address both of these problems by presenting a tool called HUNTER for automatically reusing existing code. Given a few *test cases* and a *search query* provided by the user, HUNTER searches existing code bases for candidate adaptee methods and *automatically synthesizes* any necessary wrapper code. The key technical idea underlying our approach is to use *types* to both improve search results and guide synthesis. Specifically, our method computes *sim-*

*ilarity metrics* between types and uses this information to solve an *integer linear programming (ILP)* problem in which the objective is to minimize the cost of synthesis.

The high-level architecture of the HUNTER tool is shown in Figure 1. Our approach consists of three components, namely *code search*, *interface alignment*, and *synthesis*. Given a candidate adaptee identified through code search, *interface alignment* refers to the problem of finding a suitable mapping from parameters in the adaptee to those in the adapter. Given a solution to the interface alignment problem, HUNTER uses existing synthesis tools to generate the desired wrapper code. If the synthesized program does not pass the provided test cases, HUNTER backtracks by finding another candidate adaptee or a different solution to the interface alignment problem.

The key technical contribution of this paper is a type-directed approach for solving the interface alignment problem. Specifically, HUNTER computes *distance metrics* between types in order to estimate how likely it is that an object of type  $\tau_1$  can be used in the place of another object of type  $\tau_2$ . Specifically, the larger the distance  $d$ , the less likely it is that  $\tau_1$  can be “converted” to  $\tau_2$ . Given such a *type distance matrix*, HUNTER solves an optimization problem in which the goal is to minimize the cost of converting between the type signature of the adaptee and that of the desired method. Specifically, we formulate interface alignment as an integer linear programming problem that is parametrized by the type distance matrix.

We have implemented HUNTER as an Eclipse plug-in and evaluate it by performing three sets of experiments. First, we compare HUNTER against S<sup>6</sup>, a state-of-the-art code reuse tool, and show that HUNTER can reuse more code compared to S<sup>6</sup>. Second, we perform a user study and show that HUNTER enables users to finish programming tasks quicker and with fewer bugs. Third, we show that the type distances computed by HUNTER can be gainfully used to improve the results of existing code search tools. To summarize, this paper makes the following contributions:

- We describe a novel code reuse tool, HUNTER, that combines code search with program synthesis. HUNTER is available as an Eclipse plug-in and can automatically generate adapter methods that reuse existing Java methods in massive open source repositories.
- We propose a technique for measuring similarity between arbitrary Java types and leverage it to re-rank search results obtained from existing code search tools.

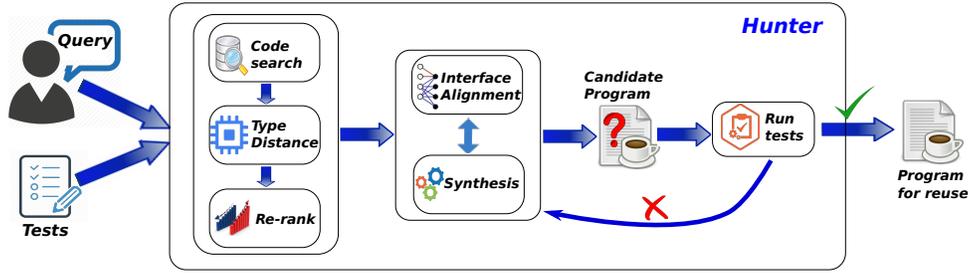


Figure 1: Workflow of the HUNTER tool

```

@Test public void test() {
    Vector<MyPoint> v1 = new Vector<MyPoint>();
    v1.add(new MyPoint(0, 0)); v1.add(new MyPoint(1, 1));
    v1.add(new MyPoint(2, 1)); v1.add(new MyPoint(3, 2));
    v1.add(new MyPoint(4, 2)); v1.add(new MyPoint(5, 3));
    Bresenham.drawLine(new MyPoint(5, 3), res);
    assertEquals(v1, res);
}

```

Figure 2: Test cases for the desired `drawLine` method

- We define the *interface alignment problem* for matching a pair of method signatures and show how it can be solved using integer linear programming (ILP).
- We perform an extensive evaluation of HUNTER by comparing it against a state-of-the-art code reuse tool (S<sup>6</sup>), and conducting a user study. Our evaluation shows that HUNTER compares favorably with S<sup>6</sup> and significantly increases programmer productivity.

## 2. OVERVIEW

In this section, we give a high-level overview of the HUNTER system through a simple motivating example.

### 2.1 Motivating Example

Consider a user, Alice, who needs to draw a straight line from the origin to a point. Specifically, Alice would like to implement in Java the following `drawLine` method:

```
void drawLine(MyPoint pt, Vector<MyPoint> res)
```

Here, `pt` is the point specified by the user, and `res` is a vector of points on the raster that should be selected to approximate a straight line between the origin and `pt`. Alice knows about Bresenham’s line drawing algorithm that could be used to implement this functionality, but she does not know exactly how it works. While she would like to reuse an existing implementation of Bresenham’s algorithm, she cannot find an implementation that quite fits her needs.

The HUNTER tool could help a user like Alice by finding an existing implementation of Bresenham’s algorithm and automatically synthesizing all the wrapper code needed to integrate it into the desired `drawLine` interface. In order to use HUNTER, Alice needs to provide the method signature shown above as well as a brief natural language description, such as “Draw a line between the origin and specified point using Bresenham’s algorithm”. Alice also needs to provide a small test suite (e.g., the one shown in Figure 2) that HUNTER can use to validate the synthesized code.

HUNTER first queries a code search engine using Alice’s description and retrieves the top  $k$  results. To simplify the

example, suppose that the search engine yields a single function with the following signature:

```
Point[] bresenham(int x0, int y0, int x1, int y1)
```

Specifically, this function returns an array of `Point`s to approximate a line starting from  $(x_0, y_0)$  to  $(x_1, y_1)$ . Note that there are several differences between Alice’s desired `drawLine` interface and the existing `bresenham` function:

- Alice’s interface uses `MyPoint` to represent the user-specified point, while the existing function uses two integers (namely, `x1` and `y1`) to represent the end point.
- Alice’s interface assumes the origin as a starting point, whereas `bresenham` takes `x0` and `y0` as input.
- The existing `bresenham` function returns an array of `Point`’s, whereas Alice’s interface “returns” the line by storing the result in `res`. Furthermore, Alice represents points using a custom type called `MyPoint`, whereas `bresenham` uses a different type called `Point`.

Despite these significant differences, HUNTER is able to automatically generate the following implementation of Alice’s `drawLine` interface:

```

void drawLine (MyPoint pt, Vector<MyPoint> res) {
    int v1 = pt.getX();
    int v2 = pt.getY();
    Point[] v3 = external.bresenham(0, 0, v1, v2);
    for (Point v4 : v3) {
        int v5 = v4.getX();
        int v6 = v4.getY();
        MyPoint v7 = new MyPoint(v5, v6);
        res.add(v7);
    }
}

```

Observe that the code generated by HUNTER first deconstructs the `MyPoint` object `pt` into a pair of integers by invoking the appropriate getter methods. Also note that HUNTER can supply default values for `x0` and `y0` even though there are no corresponding parameters in the `drawLine` interface. Finally, after invoking the existing `bresenham` method, HUNTER can synthesize code to convert the array of points into a vector of `MyPoints`.

### 2.2 System Overview

The high-level structure of our code reuse algorithm is presented in Algorithm 1. The CODEREUSE procedure takes as input a desired signature  $\mathcal{S}$ , a natural language description  $\mathcal{D}$ , and a test suite  $\mathcal{T}$ . The return value of CODEREUSE is an implementation  $\mathcal{R}$  of  $\mathcal{S}$  that passes all test cases if one exists, and  $\perp$  otherwise.

---

**Algorithm 1** Code Reuse Algorithm

---

```
1: procedure CODEREUSE( $\mathcal{S}, \mathcal{D}, \mathcal{T}$ )
2:   Input: Signature  $\mathcal{S}$  of desired method, natural
3:     language description  $\mathcal{D}$ , and tests  $\mathcal{T}$ 
4:   Output: adapter  $\mathcal{R}$  or failure  $\perp$ 
5:    $[\mathcal{E}_0] := \text{CODESEARCH}(\mathcal{S}, \mathcal{D})$ 
6:    $\Lambda := \text{COMPUTETYPEDISTANCE}(\mathcal{S}, [\mathcal{E}_0])$ 
7:    $[\mathcal{E}] := \text{RERANK}([\mathcal{E}_0], \mathcal{S}, \Lambda)$ 
8:   for all  $\mathcal{E} \in [\mathcal{E}]$  do
9:     do
10:       $\mathcal{M} := \text{GETBESTALIGN}(\mathcal{S}, \mathcal{E}, \Lambda)$ 
11:       $\mathcal{R} := \text{ADAPTERGEN}(\mathcal{M}, \mathcal{E}, \mathcal{S})$ 
12:      if  $\text{RUNTESTS}(\mathcal{R}, \mathcal{T})$  then
13:        return  $(\mathcal{R}, \mathcal{E})$ 
14:      while  $\mathcal{M} \neq \emptyset$ 
15:   return  $\perp$ 
```

---

As shown in Algorithm 1, our technique first uses a code search engine to retrieve a ranked list  $[\mathcal{E}_0]$  of relevant methods that could serve as possible adaptees. Next, we invoke the `COMPUTETYPEDISTANCE` procedure to compute a distance between each pair of types  $(\tau, \tau')$ , where  $\tau$  is a type used in signature  $\mathcal{S}$  and  $\tau'$  is a type used in the signature of some  $\mathcal{E}_0 \in [\mathcal{E}_0]$ . The return value  $\Lambda$  of `COMPUTETYPEDISTANCE` is a matrix that maps each pair of types  $(\tau, \tau')$  to a distance. The larger the distance, the less similar  $\tau$  is to  $\tau'$  and the less likely it is that an argument of type  $\tau$  will be mapped to an argument of type  $\tau'$ .

Once we compute the *type distance matrix*  $\Lambda$ , HUNTER re-ranks the original search results and tries to synthesize wrapper code for each candidate adaptee  $\mathcal{E} \in [\mathcal{E}]$ . Towards this goal, we first solve the interface alignment problem using integer linear programming. Specifically, the call to the `GETBESTALIGN` procedure at line 10 returns a lowest cost mapping between the parameters in the adapter and those in the adaptee. For example, the alignment between `drawLine` and `bresenham` from Section 2.1 is given by the mapping:

```
x1: int    -> pt: MyPoint
y1: int    -> pt: MyPoint
v3: Point[] -> res: Vector<MyPoint>
```

Intuitively, this alignment minimizes the objective function in our ILP problem according to the costs given by  $\Lambda$ . Note that every parameter in the adapter must be mapped to some parameter in the adaptee because an implementation of  $\mathcal{R}$  that does not use one of  $\mathcal{R}$ 's arguments is extremely unlikely to be correct. On the other hand, we do *not* require that every parameter in the adaptee to be mapped to a parameter in the adapter, since the implementation of the adaptee may be more generic.

Given a candidate alignment  $\mathcal{M}$ , our algorithm invokes a procedure called `ADAPTERGEN` to synthesize wrapper code based on this alignment. For instance, for our running example, `ADAPTERGEN` synthesizes code to “convert” the `pt` object of type `MyPoint` to an integer `x1` of type `int`. Since there are several existing tools [24, 18, 14, 12] for type-directed code synthesis, our `ADAPTERGEN` procedure uses existing synthesis techniques to generate a code snippet for “converting” an object of type  $\tau$  to an object of type  $\tau'$ .

Once HUNTER generates wrapper code for a given candidate alignment, we run the provided test cases to check if any test fails. If so, we backtrack and ask the ILP solver for the next best alignment, if one exists. If we exhaust all possible alignments for a given candidate  $\mathcal{E}$ , we also backtrack and try the next best candidate adaptee. The `CODEREUSE` algorithm terminates when we find an implementation that passes all test cases or we run out of possible adaptees.

### 3. TYPE DISTANCE

In this section, we describe our technique for computing distances between a pair of types. As mentioned earlier, our approach uses type distances to both re-rank search results and compute a best alignment between the signature of the desired method and that of the candidate adaptee.

#### 3.1 Multiset Representation of Types

In order to compute a distance metric between a pair of types, our approach represents each Java type using a *multiset (bag) representation*. Given a type  $\tau$ , we use the notation  $\psi(\tau)$  to denote  $\tau$ 's multi-set representation. Intuitively,  $\psi(\tau)$  represents  $\tau$  as a bag of *features*. In this context, a feature is either a Java type or a boolean attribute, such as *numeric*, *collection* etc. As we will see in the next subsection, our approach computes a distance between two types  $\tau$  and  $\tau'$  using the multi-set representations  $\psi(\tau)$  and  $\psi(\tau')$ .

We define the multiset representation  $\psi(\tau)$  of each type  $\tau$  using the inference rules shown in Figure 3. The first rule, *Primitive I*, states that the multi-set representation of a numeric primitive type  $\tau$ , such as `int` or `double`, includes both the type name as well as the attribute `numeric`.<sup>1</sup> On the other hand, if a primitive type is not numeric, then *Primitive II* states that its multiset representation consists of only the type name. For instance, we have  $\psi(\text{int}) = \{\text{int}, \text{numeric}\}$ , and  $\psi(\text{String}) = \{\text{String}\}$ .

To capture the different collection types, we define `category`( $\tau$ ) as a function that returns the collection type of  $\tau$ . Specifically, we classify Java built-in collections into four different categories, namely `Vector`, `List`, `Set`, and `Map`. For instance, we have `category(ArrayList) = List` and `category(TreeMap) = Map`. Hence, according to the *Collection* rule, the multi-set representation for a collection  $\tau$  is given by  $\{\text{category}(\tau), \text{collection}\}$ .

The next rule describes the multi-set representation for arrays. First, since arrays and collections are often used interchangeably, we also represent arrays using the attribute `collection`. Second, if  $\tau$  is an array with element type  $\tau'$ , observe that  $\psi(\tau)$  also includes the multiset representation of  $\tau'$ . For example,  $\psi(\text{int}[])$  is given by:

$$\{\text{int}, \text{numeric}, \text{Array}, \text{collection}, \text{Array}, \text{collection}\}$$

The next two rules labeled *Wildcard* and *TypeParam* allow us to handle generic types. Specifically, we define  $\psi(\tau) = \emptyset$  for both the wild card type `?` as well as a type parameter, such as `E`. For parametrized types  $\tau$  (e.g., `List<int>`), the *Parametrized* rule defines  $\psi(\tau)$  to be the union of the multiset representation of the declaring raw type as well as its type arguments.

**Example 1.** *The multiset representation of a generic list, `List(E)`, is  $\psi(\text{List}) \cup \psi(\text{E}) = \{\text{List}, \text{collection}\}$ . The*

<sup>1</sup>Boxed types of primitives are considered as their corresponding primitive types.

$\frac{isPrimitive(\tau), isNumeric(\tau)}{\psi(\tau) = \{\tau, \text{numeric}\}}$	(Primitive I)
$\frac{isPrimitive(\tau), \neg isNumeric(\tau)}{\psi(\tau) = \{\tau\}}$	(Primitive II)
$\frac{isCollection(\tau)}{\psi(\tau) = \{\text{category}(\tau), \text{collection}\}}$	(Collection)
$\frac{isArray(\tau), \tau = \tau' []}{\psi(\tau) = \psi(\tau') \cup \{\text{Array}, \text{collection}\}}$	(Array)
$\frac{isWildcard(\tau)}{\psi(\tau) = \emptyset}$	(Wildcard)
$\frac{isTypeParam(\tau)}{\psi(\tau) = \emptyset}$	(TypeParam)
$\frac{isParameterized(\tau)}{\tau = t\langle\tau_1, \dots, \tau_n\rangle}{\psi(\tau) = \psi(t) \cup \psi(\tau_1) \cup \dots \cup \psi(\tau_n)}$	(Parameterized)
$\frac{isRef(\tau), \tau = t\{L\}}{\psi(\tau) = \{t\} \cup \psi(L)}$	(RefType)
$\frac{L = \emptyset}{\psi(L) = \emptyset}$	(FieldList I)
$\frac{L = \tau_1; L_1, isRef(\tau_1)}{\psi(L) = \{\tau_1\} \cup \psi(L_1)}$	(FieldList II)
$\frac{L = \tau_1; L_1, \neg isRef(\tau_1)}{\psi(L) = \psi(\tau_1) \cup \psi(L_1)}$	(FieldList III)

Figure 3: Multiset representation of Java types

*multiset representation of a parameterized vector, such as `Vector(Vector<Integer>)`, is  $\psi(\text{Vector}(\text{Vector}(\text{Integer}))) = \psi(\text{Vector}) \cup \psi(\text{Vector}) \cup \psi(\text{int}) = \{\text{Vector}, \text{collection}, \text{Vector}, \text{collection}, \text{int}, \text{numeric}\}$ .*

For reference types (i.e., user-defined classes), we perform one-level unrolling for the fields using the rules *RefType* and *FieldList I - III*. For example, given a type `Point` with two fields of type `int`, we have  $\psi(\text{Point}) = \text{Point} \cup \psi(\text{int}) \cup \psi(\text{int})$ . Similarly, for a recursive type defined as:

```
class ListNode { int key; ListNode next; },
```

we have  $\psi(\text{ListNode}) = \{\text{ListNode}, \text{int}, \text{numeric}, \text{ListNode}\}$ .

### 3.2 Computing Distances Between Types

We now consider how to compute the *distance* between a pair of types from their multiset representation. Given two types  $\tau$  and  $\tau'$ , we write  $\delta(\tau, \tau')$  to denote the distance between  $\tau$  and  $\tau'$  and define it to be the normalized cardinality of the *symmetric difference* between  $\psi(\tau)$  and  $\psi(\tau')$ . More formally,  $\delta(\tau, \tau')$  is defined as follows:

**Definition 3.1** (Type distance).

$$\delta(\tau, \tau') = \frac{|(\psi(\tau) - \psi(\tau')) \cup (\psi(\tau') - \psi(\tau))|}{|\psi(\tau) \cup \psi(\tau')|}$$

Intuitively,  $\delta(\tau, \tau') \in [0, 1]$  represents the fraction of elements in  $\psi(\tau) \cup \psi(\tau')$  that are not shared between  $\psi(\tau)$  and  $\psi(\tau')$ . Hence,  $\delta(\tau, \tau')$  is the normalized cardinality of the *symmetric difference* between  $\psi(\tau)$  and  $\psi(\tau')$ . Observe that, if  $\tau$  and  $\tau'$  are the same type, then  $\delta(\tau, \tau')$  is always zero. Also observe that  $\delta(\tau, \tau')$  is always the same as  $\delta(\tau', \tau)$ . Furthermore, if  $\delta(\tau, \tau_1) < \delta(\tau, \tau_2)$ , this indicates that  $\tau_1$  is more similar to  $\tau$  than  $\tau_2$  is to  $\tau$ .

**Example 2.** Consider the types  $\tau_1: \text{ArrayList}\langle\text{Integer}\rangle$ ,  $\tau_2: \text{LinkedList}\langle\text{Double}\rangle$ ,  $\tau_3: \text{HashSet}\langle\text{Double}\rangle$  and their corresponding multiset representations:

$$\begin{aligned} \psi(\tau_1) &: \{\text{int}, \text{numeric}, \text{List}, \text{collection}\} \\ \psi(\tau_2) &: \{\text{double}, \text{numeric}, \text{List}, \text{collection}\} \\ \psi(\tau_3) &: \{\text{double}, \text{numeric}, \text{Set}, \text{collection}\} \end{aligned}$$

We have  $\delta(\tau_1, \tau_2) = 0.25$  and  $\delta(\tau_1, \tau_3) = 0.5$ . Hence,  $\tau_2$  is more similar to  $\tau_1$  compared to  $\tau_3$ .

In the rest of this paper, we will slightly abuse this notation and write  $\delta(x, y)$  to denote the distance between the *types* of variables  $x$  and  $y$ .

## 4. INTERFACE ALIGNMENT

We now define the *interface alignment* problem and show how to find an optimal alignment using integer linear programming.

### 4.1 Problem Definition

Consider an adapter method with name  $\mathcal{R}$  and type signature  $\tau_1 \times \dots \times \tau_n \rightarrow \tau$ . Also, suppose that we have the source code of an adaptee method with name  $\mathcal{E}$  and signature  $\tau'_1 \times \dots \times \tau'_m \rightarrow \tau'$ . In order to synthesize an implementation of  $\mathcal{R}$  using  $\mathcal{E}$  as an adaptee, we must first identify a candidate mapping between the parameters of  $\mathcal{R}$  and those of  $\mathcal{E}$ . For example, the first argument of  $\mathcal{E}$  may correspond to the second parameter of  $\mathcal{R}$ . Furthermore,  $\mathcal{R}$  and  $\mathcal{E}$  need not necessarily have the same number of parameters. For example, if  $\mathcal{R}$  has signature `Point`  $\rightarrow$  `void` and  $\mathcal{E}$  has signature `int`  $\times$  `int`  $\rightarrow$  `void`, we may need to use both parameters of  $\mathcal{E}$  to construct the first argument of  $\mathcal{R}$ . Hence, a precursor to synthesizing adapter code is to identify a suitable mapping between the parameters of  $\mathcal{R}$  and  $\mathcal{E}$ . We refer to this problem as *interface alignment*.

**Definition 4.1** (Interface alignment). Consider adapter  $\mathcal{R}$  and adaptee  $\mathcal{E}$  with the following signatures:

$$\begin{aligned} \mathcal{R} &: r_1 : \tau_1 \times r_2 : \tau_2 \times \dots \times r_n : \tau_n \rightarrow r_0 : \tau \\ \mathcal{E} &: e_1 : \tau'_1 \times e_2 : \tau'_2 \times \dots \times e_m : \tau'_m \rightarrow e_0 : \tau' \end{aligned}$$

where  $x : \tau$  indicates  $x$  has type  $\tau$ . Let  $P(\mathcal{R})$  denote the set  $\{r_1, r_2, \dots, r_n, r_0\}$  if the return value of  $\mathcal{R}$  is not `void` and  $\{r_1, r_2, \dots, r_n\}$  otherwise. The *interface alignment problem* is to find a mapping  $\mathcal{M}$  from  $P(\mathcal{E})$  to  $P(\mathcal{R})$  such that (i)  $\mathcal{M}$  is surjective, and (ii)  $\mathcal{M}$  is not many-to-many.

Intuitively, we require  $\mathcal{M}$  to be *surjective* because all parameters of  $\mathcal{R}$  should be used in  $\mathcal{R}$ 's implementation. Furthermore, we do not require  $\mathcal{M}$  to be injective because different arguments of the adaptee can be mapped to the same argument in the adapter. For instance, for the motivating example from Section 2.1, both `x1` and `y1` from the adaptee are mapped to argument `pt` of the adapter.

It is important to emphasize that a solution  $\mathcal{M}$  to the interface alignment problem does not need to be a function. In particular, we allow many-to-one mappings. To see why this is necessary, consider an adaptee that has a single argument `pt` of type `Point` and an adapter that takes two arguments `x`, `y` of type `int`. In this case, since we can construct a `Point` from two `int`'s, a suitable alignment maps `pt` to both `x` and `y`, meaning that we construct `pt` from `x` and `y`. Finally, since  $\mathcal{M}$  does not have to be a function, some parameters of the adaptee may also be unassigned. For instance, for the motivating example from Section 2.1, parameters `x0` and `y0` of `bresenham` are not assigned to any parameters of `drawLine`. If a parameter `x` of the adaptee is unassigned in  $\mathcal{M}$ , we synthesize a constant value of `x` using a pre-defined set of constants for each type (e.g.,  $\{0, 1\}$  for `int` and  $\{true, false\}$  for `boolean`).

Observe that a given solution  $\mathcal{M}$  to the interface alignment problem describes a particular strategy for generating wrapper code. In particular, if  $\mathcal{M}(x : \tau) = \{y_1 : \tau_1, \dots, y_n : \tau_n\}$ , this means that we should construct parameter `x` of the adaptee using parameters `y1, ..., yn` of the adapter. Our approach uses an existing type-directed synthesis tool [12] to generate code snippets that produce an object of type  $\tau$  using objects of type  $\tau_1, \dots, \tau_n$ .

## 4.2 Optimal Interface Alignment

Given an adapter  $\mathcal{R}$  and a candidate adaptee  $\mathcal{E}$ , there are, in general, *many* possible solutions to the interface alignment problem. One of the key insights underlying our approach is to use type distances to find a *good* alignment. Given a mapping  $\mathcal{M}$ , we define  $Cost(\mathcal{M})$  as follows:

**Definition 4.2** (Cost of mapping). *Let  $\mathbb{I}(x)$  be 1 if  $|M(x)| \neq 1$  and 0 otherwise. Then, the cost of  $\mathcal{M}$  is given by:*<sup>2</sup>

$$Cost(\mathcal{M}) = \left( \frac{\sum_{x \in \text{dom}(\mathcal{M})} \mathbb{I}(x) \cdot \delta(x, \mathcal{M}(x)) + \sum_{y \in \text{dom}(\mathcal{M}^{-1})} (1 - \mathbb{I}(y)) \cdot \delta(y, \mathcal{M}^{-1}(y))}{2} \right)$$

To understand the rationale behind this definition, consider the case where  $\mathcal{M}$  is a many-to-one mapping (e.g.,  $x : \text{int} \mapsto pt : \text{Point}, y : \text{int} \mapsto pt : \text{Point}$ ). In this case, an argument (e.g., `pt`) of  $\mathcal{R}$  corresponds to multiple arguments (e.g., `x, y`) of the adaptee  $\mathcal{E}$ . In our example, this means that we must “convert” a point object to a pair of integers; hence, the cost of this mapping should be  $\delta(\text{Point}, \{\text{int}, \text{int}\})$  rather than  $2 * \delta(\text{int}, \text{Point})$ . Hence, intuitively, the cost of the mapping depends on whether  $M$  is one-to-many or many-to-one. To capture this intuition, we weight each term  $\delta(x, \mathcal{M}(x))$  with  $\mathbb{I}(x)$ , meaning that we should only pay the cost  $\delta(x, \mathcal{M}(x))$  if  $\mathcal{M}$  is one-to-many. Otherwise, we pay the cost  $\delta(y, \mathcal{M}^{-1}(y))$  using the reverse mapping.

**Example 3.** *Consider the following signatures:*

$\mathcal{E} : x : \text{int} \times y : \text{int} \rightarrow \text{void}$   
 $\mathcal{R} : p : \text{Point} \rightarrow \text{void}$

as well as the following mappings:

$\mathcal{M}_1 : [x \mapsto \{p\}, y \mapsto \{p\}]$   
 $\mathcal{M}_2 : [x \mapsto \{p\}, y \mapsto \emptyset]$

Here, we have  $Cost(\mathcal{M}_1) = \delta(p, \{x, y\})$ . On the other hand,  $Cost(\mathcal{M}_2) = \delta(y, \emptyset) + \delta(p, x)$ .

<sup>2</sup>If  $\mathcal{M}(x)$  is not a singleton, we treat  $\mathcal{M}(x)$  as a list and compute  $\delta(x, \mathcal{M}(x))$  using the *FieldList* rules from Figure 3.

Based on this notion of cost, we can now define what it means for an alignment to be optimal:

**Definition 4.3** (Optimal alignment).  *$\mathcal{M}^*$  is an optimal solution to the interface alignment problem if there does not exist another solution  $\mathcal{M}$  such that  $Cost(\mathcal{M}) < Cost(\mathcal{M}^*)$ .*

In other words, an optimal solution minimizes the cost of the mapping according to the type distance matrix. In the next section, we show how to find an optimal solution using 0-1 Integer Linear Programming.

## 4.3 ILP Formulation

A 0-1 Integer Linear Programming (ILP) consists of a set of linear constraints  $\mathcal{C}$  over boolean variables and an objective function  $c$ . The goal is to find an assignment such that all constraints are satisfied and the value of the objective function  $c$  is optimized.

**Definition 4.4** (0-1 Integer Linear Programming). *The 0-1 ILP problem is defined as follows:*

$$\min c : \sum_j c_j x_j \quad \text{s.t.} \quad \mathcal{C} : \bigwedge_i \sum_j a_{i,j} x_j \Delta b_i,$$

with  $\Delta := \{\leq, =, \geq\}$ , and  $x_j \in \{0, 1\}$ .

We formulate the problem of finding an optimal interface alignment using 0-1 ILP. Specifically, constraints  $\mathcal{C}$  encode that  $\mathcal{M}$  is a surjective mapping from  $P(\mathcal{E})$  to  $P(\mathcal{R})$  and that it is not many-to-many. The objective function expresses that we want to minimize the cost of this mapping. In what follows, we describe our encoding in more detail.

**Variables.** The variables in our 0-1 ILP formulation correspond to all possible mappings from  $P(\mathcal{E})$  to  $P(\mathcal{R})$ . In particular, a boolean variable  $x_{e \rightarrow \vec{r}}$  indicates a one-to-many mapping from parameter  $e \in P(\mathcal{E})$  to a set of parameters  $\vec{r} \subseteq P(\mathcal{R})$ .<sup>3</sup> Similarly, boolean variables  $x_{\vec{e} \rightarrow r}$  indicate a many-to-one mapping from parameters  $\vec{e} \subseteq P(\mathcal{E})$  to parameter  $r \in P(\mathcal{R})$ . Note that only parameters of reference type are allowed to have many-to-one or one-to-many mappings: If two parameters  $r_1, r_2 \in P(\mathcal{R})$  are mapped to a parameter  $e_1 \in P(\mathcal{E})$  of reference type, this means that  $e_1$  is constructed using  $r_1$  and  $r_2$ . Similarly, if two parameters  $e_1, e_2 \in P(\mathcal{E})$  are mapped to a single parameter  $r_1 \in P(\mathcal{R})$  of reference type, this means that  $e_1, e_2$  are deconstructed using  $r_1$  (e.g., via getter methods). Hence, while many-to-one and one-to-many mappings make sense for reference types, it does not have a sensible interpretation for primitive types and collections.

We use  $\mathcal{V}$  to denote all variables used in the encoding. Given a parameter  $r \in P(\mathcal{R})$ ,  $\mathcal{V}_r$  denotes the set of variables where  $r$  occurs. Similarly, given a parameter  $e \in P(\mathcal{E})$ ,  $\mathcal{V}_e$  denotes the set of variables where  $e$  occurs.

**Example 4.** *Assume we have an adapter method  $\mathcal{R}$  with signature  $(r_1 : \text{Point} \times r_2 : \text{long}) \rightarrow (r_0 : \text{void})$  and an adaptee method  $\mathcal{E}$  with signature  $(e_1 : \text{int} \times e_2 : \text{int} \times e_3 : \text{long}) \rightarrow (e_0 : \text{void})$ . Then we have  $P(\mathcal{E}) = \{e_1, e_2, e_3\}$  and  $P(\mathcal{R}) = \{r_1, r_2\}$ . Thus, the set of variables are:*

$$\mathcal{V} = \{x_{e_1 \rightarrow r_1}, x_{e_1 \rightarrow r_2}, x_{e_2 \rightarrow r_1}, x_{e_2 \rightarrow r_2}, x_{e_3 \rightarrow r_1}, x_{e_3 \rightarrow r_2}, x_{e_1, e_2 \rightarrow r_1}, x_{e_1, e_3 \rightarrow r_1}, x_{e_2, e_3 \rightarrow r_1}, x_{e_1, e_2, e_3 \rightarrow r_1}\}.$$

<sup>3</sup>We consider a one-to-one mapping a particular case of the one-to-many mapping where  $\vec{r}$  only contains one parameter.

Given the parameter  $e_1$ ,  $\mathcal{V}_{e_1}$  denotes the set of variables where  $e_1$  occurs:

$$\mathcal{V}_{e_1} = \{x_{e_1 \rightarrow r_1}, x_{e_1 \rightarrow r_2}, x_{e_1, e_2 \rightarrow r_1}, x_{e_1, e_3 \rightarrow r_1}, x_{e_1, e_2, e_3 \rightarrow r_1}\}.$$

A mapping is defined by the variables that are assigned to 1. For instance, consider an assignment  $\sigma$  that assigns  $x_{e_1, e_2 \rightarrow r_1}$  and  $x_{e_3 \rightarrow r_2}$  to 1, and all other variables to 0. Then,  $\sigma$  corresponds to the following mapping:  $(e_1, e_2) \rightarrow (r_1)$ ,  $(e_3) \rightarrow (r_2)$ .

Observe that the number of variables used in the encoding grows quadratically for non-reference types and exponentially for reference types. However, since the number of parameters is usually small, our encoding introduces a manageable number of variables in practice.

**Constraints.** While the variables describe all possible mappings between parameters of an adaptee  $\mathcal{E}$  and a desired method  $\mathcal{R}$ , not all mappings can occur simultaneously. In particular, we must enforce that any satisfying assignment to  $\mathcal{C}$  corresponds to a surjective mapping from  $P(\mathcal{E})$  to  $P(\mathcal{R})$ . Furthermore, types also impose *hard constraints* that limit which variables in  $P(\mathcal{E})$  can be mapped to which ones in  $P(\mathcal{R})$ . We enforce these hard constraints by generating a system of linear constraints  $\mathcal{C}$  as follows:

1. First, we divide all Java types into 3 categories, namely *primitive*, *collection*, and *reference*, and only allow type conversions marked as “Y” in the following table:

	primitive	collection	reference
primitive	Y	N	Y
collection	N	Y	N
reference	Y	N	Y

If two parameters  $r \in P(\mathcal{R})$  and  $e \in P(\mathcal{E})$  are not compatible due to their types, the boolean variables where these parameters occur are always set to 0:

$$x_{e \rightarrow \vec{r}} = 0 \text{ if } x_{e \rightarrow \vec{r}} \in (\mathcal{V}_e \cap \mathcal{V}_r), (\text{type}(r), \text{type}(e)) \text{ is "N"}$$

$$x_{\vec{e} \rightarrow r} = 0 \text{ if } x_{\vec{e} \rightarrow r} \in (\mathcal{V}_e \cap \mathcal{V}_r), (\text{type}(r), \text{type}(e)) \text{ is "N"}$$

We note that the above constraints are somewhat connected with the power of the synthesis tool used for generating wrapper code. In particular, we disallow mappings between references and collections since HUNTER’s underlying synthesis engine does not support converting between references and collections.

2. For each parameter  $r \in P(\mathcal{R})$ , we impose that there is *exactly one* mapping that contains  $r$ :

$$\forall r \in P(\mathcal{R}). \sum_{x_{e \rightarrow \vec{r}} \in \mathcal{V}_r} x_{e \rightarrow \vec{r}} + \sum_{x_{\vec{e} \rightarrow r} \in \mathcal{V}_r} x_{\vec{e} \rightarrow r} = 1$$

Effectively, these constraints enforce that any solution of  $\mathcal{C}$  corresponds to a surjective mapping.

3. For each parameter  $e \in P(\mathcal{E})$ , we impose that there is *at most one* mapping that contains  $e$ :

$$\forall e \in P(\mathcal{E}). \sum_{x_{e \rightarrow \vec{r}} \in \mathcal{V}_e} x_{e \rightarrow \vec{r}} + \sum_{x_{\vec{e} \rightarrow r} \in \mathcal{V}_e} x_{\vec{e} \rightarrow r} \leq 1$$

These constraints enforce that a parameter  $e \in P(\mathcal{E})$  can only be used in at most one mapping. In particular, it is not necessary that every  $e \in P(\mathcal{E})$  be mapped to some  $r \in P(\mathcal{R})$ .

**Example 5.** Consider the same methods  $\mathcal{R}$  and  $\mathcal{E}$  from Example 4. Constraints  $\mathcal{C}$  for this example are given by:

$$p_1 : x_{e_1 \rightarrow r_1} + x_{e_2 \rightarrow r_1} + x_{e_3 \rightarrow r_1} + x_{e_1, e_2 \rightarrow r_1} + x_{e_1, e_3 \rightarrow r_1} + x_{e_2, e_3 \rightarrow r_1} + x_{e_1, e_2, e_3 \rightarrow r_1} = 1$$

$$p_2 : x_{e_1 \rightarrow r_2} + x_{e_2 \rightarrow r_2} + x_{e_3 \rightarrow r_2} = 1$$

$$p_3 : x_{e_1 \rightarrow r_1} + x_{e_1 \rightarrow r_2} + x_{e_1, e_2 \rightarrow r_1} + x_{e_1, e_3 \rightarrow r_1} + x_{e_1, e_2, e_3 \rightarrow r_1} \leq 1$$

$$p_4 : x_{e_2 \rightarrow r_1} + x_{e_2 \rightarrow r_2} + x_{e_1, e_2 \rightarrow r_1} + x_{e_2, e_3 \rightarrow r_1} + x_{e_1, e_2, e_3 \rightarrow r_1} \leq 1$$

$$p_5 : x_{e_3 \rightarrow r_1} + x_{e_3 \rightarrow r_2} + x_{e_1, e_3 \rightarrow r_1} + x_{e_2, e_3 \rightarrow r_1} + x_{e_1, e_2, e_3 \rightarrow r_1} \leq 1$$

Here,  $p_1$  and  $p_2$  enforce that each parameter  $r \in P(\mathcal{R})$  appears in exactly one mapping. Constraints  $p_3$ ,  $p_4$ ,  $p_5$  guarantee that each parameter  $e \in P(\mathcal{E})$  appears in at most one mapping. These constraints enforce that any solution to  $\mathcal{C}$  corresponds to a surjective mapping.

**Objective function.** The goal of the objective function  $c$  is to find an *optimal* alignment with the lowest cost. Specifically, we define the objective function  $c$  as follows:

$$\sum_{e \in P(\mathcal{E})} \sum_{x_{e \rightarrow \vec{r}} \in \mathcal{V}_e} \delta(e, \vec{r}) \cdot x_{e \rightarrow \vec{r}} + \sum_{r \in P(\mathcal{R})} \sum_{x_{\vec{e} \rightarrow r} \in \mathcal{V}_r} \delta(r, \vec{e}) \cdot x_{\vec{e} \rightarrow r}$$

Each mapping has an associated cost using type distances from Section 3. Observe that this objective function directly encodes the cost metric from Definition 4.2.

**Example 6.** Consider the same methods  $\mathcal{R}$  and  $\mathcal{E}$  from Example 4 and the constraints from Example 5. Here, we have a total of 9 valid solutions to the constraint system. Each of these mappings, along with their cost, is listed below:

id	mapping	cost
$m_1$	$(e_1, e_2) \rightarrow (r_1), (e_3) \rightarrow (r_2)$	$1/9 + 0 \simeq 0.11$
$m_2$	$(e_1) \rightarrow (r_1), (e_3) \rightarrow (r_2)$	$3/7 + 0 \simeq 0.43$
$m_3$	$(e_2) \rightarrow (r_1), (e_3) \rightarrow (r_2)$	$3/7 + 0 \simeq 0.43$
$m_4$	$(e_1, e_3) \rightarrow (r_1), (e_2) \rightarrow (r_2)$	$3/9 + 2/4 \simeq 0.83$
$m_5$	$(e_2, e_3) \rightarrow (r_1), (e_1) \rightarrow (r_2)$	$3/9 + 2/4 \simeq 0.83$
$m_6$	$(e_1) \rightarrow (r_1), (e_2) \rightarrow (r_2)$	$3/7 + 2/4 \simeq 0.93$
$m_7$	$(e_2) \rightarrow (r_1), (e_1) \rightarrow (r_2)$	$3/7 + 2/4 \simeq 0.93$
$m_8$	$(e_3) \rightarrow (r_1), (e_1) \rightarrow (r_2)$	$5/7 + 2/4 \simeq 1.21$
$m_9$	$(e_3) \rightarrow (r_1), (e_2) \rightarrow (r_2)$	$5/7 + 2/4 \simeq 1.21$

Since our objective function is to minimize the cost, the first result we get is  $m_1$ , which is exactly what we want.

## 5. DESIGN AND IMPLEMENTATION

In this section, we describe the design and implementation of the HUNTER tool for automated code reuse. HUNTER is publicly available as an Eclipse plug-in at the Eclipse marketplace and consists of approximately 12,000 lines of Java code. Our implementation uses the Sat4J [9] tool for solving 0-1 ILP problems.

HUNTER can be integrated with any code search engine that yields results at the granularity of methods. In our current implementation, we integrated HUNTER with the Pliny code search engine [7] and use a database of over 12 million Java methods collected from open source repositories, such as Github [2] and Bitbucket [1].

		HUNTER	S <sup>6</sup>			HUNTER	S <sup>6</sup>
Success rate:		100.0%	37.5%				
Id	Benchmark	HUNTER	S <sup>6</sup>	Id	Benchmark	HUNTER	S <sup>6</sup>
1	SimpleTokenizer	✓	✓	21	FloodFill	✓	✗
2	QuoteTokenizer	✓	✓	22	FindMedian	✓	✗
3	CheckRobots	✓	✓	23	ListAverage	✓	✗
4	LogBase	✓	✓	24	BresenhamLine	✓	✗
5	RomanNumeral	✓	✓	25	BresenhamCircle	✓	✗
6	RomanToInt	✓	✓	26	Distance	✓	✗
7	PrimeNumber	✓	✓	27	Slope	✓	✗
8	PerfectNumber	✓	✓	28	PrimeSieve	✓	✓
9	DayOfWeek	✓	✓	29	Anagram	✓	✓
10	EasterDate	✓	✓	30	Palindrome	✓	✓
11	MatrixMultiplication	✓	✗	31	PartitionList	✓	✗
12	LcsInteger	✓	✗	32	RotateList	✓	✗
13	RemoveDuplicates	✓	✓	33	ListInsertion	✓	✗
14	TransposeMatrix	✓	✗	34	PalindromeList	✓	✗
15	InvertMatrix	✓	✗	35	SwapNodesPairs	✓	✗
16	MatrixPower	✓	✗	36	InvertBinaryTree	✓	✗
17	DotProduct	✓	✗	37	MinDepthBinaryTree	✓	✗
18	MatrixDeterminant	✓	✓	38	BinaryTreePostorder	✓	✗
19	CountingSort	✓	✗	39	BinaryTreeInorder	✓	✗
20	MatrixAddition	✓	✗	40	SumRootLeafNumbers	✓	✗

Table 1: Comparison between HUNTER and S<sup>6</sup>

**Re-ranking.** Given the initial results provided by the Pliny search engine, HUNTER re-ranks search results using the type similarity matrix described in Section 3. Specifically, for each method  $M$  in the search result, HUNTER computes a *type similarity score* between the signature of  $M$  and that of the desired method  $M^*$ . The type similarity score is the multiplicative inverse of the cost for the optimal alignment between  $M$  and  $M^*$ . Hence, our re-ranking procedure sorts the original search results using the techniques described in Section 4 for finding an optimal alignment. As we demonstrate in Section 6, our re-ranking algorithm based on type similarity significantly improves search results.

**Synthesis of wrapper code.** As mentioned earlier, a solution to the interface alignment problem can be immediately translated into wrapper code using existing synthesis tools [12, 24, 18, 14]. Our current implementation utilizes the SyPET tool for type-directed component-based synthesis [12]. Specifically, let  $\mathcal{M}$  be a solution to the interface alignment problem, and let  $\mathcal{M}(x : \tau)$  be  $\{y_1 : \tau_1, \dots, y_n : \tau_n\}$ . In this case, to synthesize the actual value for  $x$ , we ask SyPET to generate a procedure that takes inputs of type  $\tau_1, \dots, \tau_n$  and produces an output of type  $\tau$ . Since there can be many ways to construct an object of type  $\tau$  from  $y_1, \dots, y_n$ , our implementation tries the top 100 results returned by SyPET before moving on to a different solution to the interface alignment problem (this limit was never reached during our experiments).

Since most type-directed synthesis tools, including SyPET, cannot synthesize loops, we have developed a library of templates for performing conversions between different types of collections. For example, to generate a collection  $x$  of type `Vector<Foo>` from another collection  $y$  of type `Bar[]`, we first retrieve the built-in template for converting an array to a vector and then invoke the procedure synthesized by SyPET on each element.

As mentioned earlier, HUNTER can also synthesize constants in cases where an adaptee parameter is mapped to

the empty set. In our current implementation, each type has a set of “default” values, including `NULL` for references,  $\{true, false\}$  for booleans, and  $\{0, 1\}$  for integers. If an adaptee parameter  $x$  of type  $\tau$  is mapped to  $\emptyset$ , HUNTER tries all default values associated with type  $\tau$ .

**Dependency resolution.** After synthesizing wrapper code for a candidate alignment, HUNTER resolves all dependencies between classes based on the package hierarchy. Specifically, HUNTER starts from the class  $C$  containing the candidate method and transitively identifies all classes  $C_1, \dots, C_n$  that  $C$  depends on. Since this procedure can take a long time for large dependencies, we restrict the maximum number of classes to 10. If the dependencies of a candidate cannot be fully resolved within this limit, HUNTER backtracks and tries the next best candidate.

## 6. EVALUATION

To evaluate the usefulness of HUNTER, we perform three sets of experiments. First, we compare HUNTER against S<sup>6</sup> [29], a state-of-the-art code reuse tool that is publicly available and well-maintained. Second, we use HUNTER to re-rank the search results provided by the Pliny [7] and Greptime [3] search engines and compare the results before and after re-ranking. Finally, we perform a user study and evaluate how long participants take to complete various algorithmic tasks with and without HUNTER. All experiments are conducted on a Lenovo laptop with an Intel i7-5600U CPU and 8G of memory running Ubuntu 14.04.

### 6.1 Comparison with S<sup>6</sup>

S<sup>6</sup> [29] is a state-of-the-art code reuse tool with a web interface that takes the same set of inputs as HUNTER (i.e., method signature, natural language description, and test cases). There are two key differences between S<sup>6</sup> and HUNTER. First, S<sup>6</sup> directly modifies the search result instead of synthesizing wrapper code. Second, HUNTER uses integer linear programming to measure similarity between type signatures



Figure 4: Comparison of running times between HUNTER and S<sup>6</sup>

and uses this information to re-rank search results and facilitate integration with synthesis tools.

We compare HUNTER against S<sup>6</sup> on a variety of benchmarks collected from three different sources: (i) examples used to evaluate S<sup>6</sup> [29], (ii) challenge problems taken from the Mining and Understanding Software Enclaves (MUSE) project [5] and (iii) benchmarks from Leetcode [4], which is an online resource containing programming problems.

Table 1 describes the 40 benchmarks used in our evaluation in more detail. In this table, a ✓ indicates that the tool was able to successfully synthesize the desired code, and an ✗ indicates that synthesis failed. As summarized in Table 1, HUNTER is able to solve *all* benchmarks with an average running time of 14 seconds. In contrast, S<sup>6</sup> can only solve 37.5% of the benchmarks. Since both tools use the same underlying code corpus, these results indicate that HUNTER is more successful at reusing existing code compared to S<sup>6</sup>.

To further investigate running time, Figure 4 compares the performance of HUNTER against that of S<sup>6</sup> on the 15 benchmarks that could be successfully solved by both systems. The numbers on the x-axis refer to the numeric identifier of each benchmark, as presented in Table 1. The y-axis shows the corresponding time that each tool took to solve those benchmarks. As we can see from Figure 4, HUNTER is also much faster compared to S<sup>6</sup>. On average, HUNTER takes 15 seconds to solve these benchmarks, while S<sup>6</sup> takes an average of 83 seconds.

## 6.2 Effect of Re-ranking on Search Results

In this paper, we have argued that type similarity metrics can be used to improve the quality of results returned by code search engines. To substantiate this claim, we perform an experiment to compare the quality of search results before and after using our signature-based re-ranking algorithm. Recall that our re-ranking technique sorts search results based on the cost of the optimal solution to the interface alignment problem. If two methods  $m, m'$  have the same cost according to our signature similarity metric, our re-ranking algorithm preserves the relative ordering between  $m, m'$  based on the original search results.

To evaluate the effectiveness of our re-ranking algorithm, we compare the ranking of the desired solution before and after re-ranking<sup>4</sup>. To show that this improvement is orthogonal to the code search engine, we perform the same

<sup>4</sup>If there are multiple solutions that match a given query, we consider the rank of the highest rank solution. Furthermore, we manually inspected the search results to decide whether a given method implements the desired functionality.

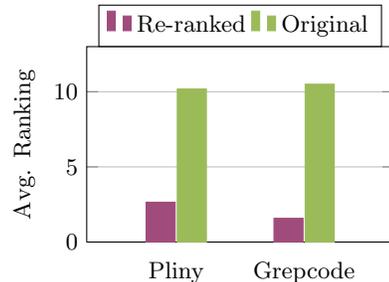


Figure 5: Average ranking of solution

experiment using two different code search engines, namely Pliny [7] and Grepcode [3]. Since Grepcode does not use the same database as Pliny, we were only able to find the correct solution for 24 out of the 40 benchmarks from Table 1. Hence, for Grepcode, we only report the ranking of the solution for the 24 benchmarks. For Pliny, we report the rank of the solution before and after re-ranking for all 40 benchmarks.

Figure 5 summarizes the results of this experiment. The green bars (on the right) show the rank of the solution in the original search results, and the purple bars (on the left) indicate the rank of the solution after re-ranking. Without using our re-ranking algorithm, the solution has an average rank of 10 using Pliny and 11 using Grepcode. On the other hand, after re-ranking, the solution has an average rank of 3 for Pliny and 2 for Grepcode. Hence, we believe these results indicate that our signature similarity metric can be used to improve the results of any code search engine that yields results at the granularity of methods.

## 6.3 User Study

To evaluate the impact of HUNTER on programmer productivity, we conducted a user study involving 16 graduate students, 2 professional programmers, and 3 undergraduate students. Among these 21 participants, 13 are experienced Java programmers, and the remaining 8 have limited exposure to Java.

In our user study, we asked the participants to produce a Java implementation for each of the following methods:

- $T_1$ . Double matrix multiplication:

```
void multiply(Vector<Vector<Double>> first,
             Vector<Vector<Double>> second,
             Vector<Vector<Double>> res)
```
- $T_2$ . Integer longest common subsequence:

```
void lcs(Vector<Integer> first,
        Vector<Integer> second, Vector<Integer> out)
```

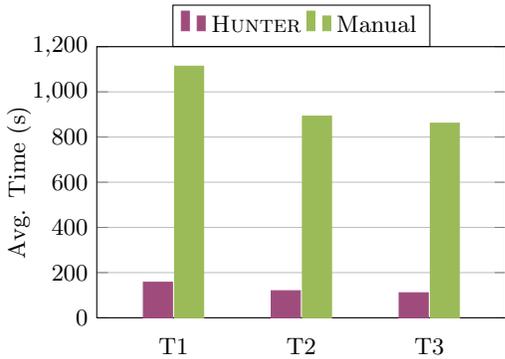


Figure 6: Comparison between HUNTER and manual coding

$T_3$ . Bresenham Line:

```
void drawLine(int x, int y, Vector<MyPoint> p)
```

For each problem, we provided a detailed description of the task to be performed, including a simple input-output example. Since our goal is to compare programmer productivity with and without HUNTER, we instructed participants to complete each task using two different methods:

- **Manual:** In this scenario, participants were asked to complete the programming task manually without using HUNTER. However, participants were explicitly told that they can use any existing code search tool and adapt the search results to their needs. In other words, the participants were aware that they do *not* have to implement the algorithm from scratch.
- **Using Hunter:** In this scenario, participants were asked to use HUNTER to automatically synthesize the desired code by providing a natural language query and implementing appropriate JUnit tests. Since participants did not have any prior familiarity with HUNTER, we gave the participants a brief demo of the tool prior to the user study.

For both scenarios, we asked participants to stop working on a given task after 30 minutes. Hence, any task that the users could not complete within 30 minutes was considered as a “failure”. For the manual implementation case, we did not require participants to write test cases for their code. Hence, there was no overlap between the tasks that the users needed to complete for the two different usage scenarios.

**Programmer productivity.** Figure 6 shows the average time for completing each of the three tasks with and without HUNTER. Overall, participants took an average of 130 seconds using HUNTER and an average of 948 seconds when writing the code manually. Furthermore, while participants were able to successfully complete 100% of the tasks using HUNTER, success rate was only 85% without HUNTER. In particular, participants were not able to manually complete 15% of the tasks (4  $T_1$ , 4  $T_2$ , and 2  $T_3$ ) within the 30 minute time limit. In summary, these results demonstrate that HUNTER allows programmers to be more productive.

To further validate our claim that these results are statistically significant, we also performed a two-tailed paired t-test [11] for each task, ignoring samples that were not completed within the 30 minute time-limit. Since the t-test for each task returned p-values smaller than 0.0001, this evalu-

ation supports our claim that there is a significant difference in average task completion time with and without HUNTER.

**Quality of solutions.** To compare the *quality* of the solutions implemented manually vs. using HUNTER, we also manually inspected the solution and ran the programs on a large test suite. While programs synthesized by HUNTER pass all test cases, 19% of the manually written programs are actually buggy and fail at least one of our test cases. Among those buggy programs, ten of them contain off-by-one bugs, one program implementing task  $T_1$  throws an exception for matrices that are not square, and one program implementing task  $T_3$  produces the wrong result due to a typo in the copy-pasted code. In summary, we believe these results show that programs synthesized by HUNTER are less error-prone compared to their manually written versions.

## 6.4 Threats to Validity

In this section, we discuss possible threats to the validity of the experiments.

**Comparison with  $S^6$ .** One of the threats to validity in our evaluation is that we could only compare against a single test-driven code reuse tool, namely  $S^6$ . In particular, we were not able to compare HUNTER against other test-driven code reuse tools, such as CodeGenie [22] and Code-Conjurer [20], since they are no longer maintained by the original developers.

Another threat to validity is the measurement of the running times when comparing HUNTER and  $S^6$ . Specifically,  $S^6$  is multi-threaded, runs on a remote server, and it terminates when it finds *all* methods that pass all test cases. HUNTER is single-threaded, runs on a common laptop, and it terminates when it finds the *first* method that passes all test cases. We extended HUNTER to terminate after finding all methods, and the average running time increases from 15 seconds to 25 seconds but it is still much faster compared to  $S^6$  (83 seconds).

**Effect of re-ranking.** Another threat to validity is that our re-ranking procedure requires search engines to produce results at the granularity of methods. However, most search engines, such as Github [2], OpenHub [6], and CodeExchange [25], provide solutions either at the granularity of files or require an exact match between method names. We chose to evaluate our type-based re-ranking algorithm using the Pliny and Grepcode search engines since they provide finer-grained support for querying methods.

**User study.** Finally, a threat to the validity of our user study is the choice of programming tasks as well as the background of the participants. Specifically, only 62% of the participants were experienced Java programmers, so it is conceivable that the participants in our study are not fully representative of professional Java developers. However, we believe that the potential users of a tool like HUNTER also include novice programmers.

## 7. RELATED WORK

In this section, we survey related work on code search, code reuse, and program synthesis.

**Code search and reuse.** Existing techniques for code search and reuse can be divided into several categories, in-

cluding textual [25, 33, 10], graph-based [8, 26, 19], test-driven [22, 20, 29, 31], and type-based methods [32, 35, 34].

*Textual approaches* for code search identify relevant methods using keywords, such as comments and variable names [13, 23, 33]. In more recent work, Chatterjee et al. have proposed annotating user code with documentations of API methods in order to improve free-form query search results [10]. Similarly, Exemplar [15] uses help documents to expand queries and achieves multiple levels of granularity. Code-Exchange [25] further refines textual methods by exploiting relationships between successive user queries. Unlike HUNTER, none of these techniques enable *fully automated reuse* of existing code.

*Graph-based approaches* to code search represent relationships between software components as a graph and use a variant of the PageRank algorithm [27] on the resulting graph representation. For example, Sourcerer [8] models software components as a directed graph, where nodes denote classes and edges represent cross-component usage. This representation allows Sourcerer to prioritize widely-used software components in the search results. Similarly, Portfolio [26] generates a graph representation, where nodes correspond to methods and edges denote caller-callee relationships. The key insight underlying Portfolio is that methods that are called more frequently in the corpus are likely to be more relevant.

*Type-based approaches* allow users to formulate code search queries involving types. For example, Prospector [24] and PARSEWeb [32] can synthesize small code snippets that reuse existing functions based on queries of the form  $(\tau_{in}, \tau_{out})$ , where  $\tau_{in}, \tau_{out}$  denote the desired input and output types respectively. Another related line of work is *signature matching*, which aims to retrieve methods that (partially) match a type signature by reducing this problem to first-order unification [35, 34]. HUNTER is closely related to these approaches in that it performs re-ranking and interface alignment based on type signatures. However, our approach differs from these techniques in that (i) we compute a type similarity metric based on multiset representations of classes, and (ii) solve an integer linear programming problem to find the best match. It is worth noting that HUNTER uses an approach similar to Prospector and PARSEWeb for automatically synthesizing adapter code in a type-directed manner.

*Test-driven approaches* use test cases to partially specify the behavior of the desired method. For example, CodeGenie [22] extracts the method signature and search keywords from JUnit tests provided by the user and queries Sourcerer to find relevant methods. It then uses the provided test cases to validate the search result and merges the code into development environment. However, unlike HUNTER, CodeGenie cannot synthesize adapter code.

Another test-driven tool similar to HUNTER is CodeConjurer [20], which allows users to specify class components using test cases and UML-like interface description. Similar to HUNTER, CodeConjurer can find a suitable mapping between the methods of the candidate class and those of the desired class. However, CodeConjurer tries all possible method-mapping permutations and cannot synthesize adapter code. Since the mapping technique is based on brute-force search, CodeConjurer’s method matching pro-

cedure can take several hours. In contrast, HUNTER solves a different kind of interface alignment problem and finds an optimal solution through integer linear programming.

Among automated code reuse tools, S<sup>6</sup> [29] is the most similar to HUNTER. Specifically, S<sup>6</sup> uses a combination of test cases, method signatures, and natural language description to find relevant methods. Furthermore, similar to HUNTER, S<sup>6</sup> can adapt existing code to fit the desired interface by performing various kinds of transformations (e.g., parameter re-ordering and renaming). As demonstrated in our experiments, HUNTER outperforms S<sup>6</sup>, both in terms of running time as well as the amount of code that it can reuse.

**Program synthesis and code completion.** HUNTER is also related to a long line of work on program synthesis and code completion. In particular, HUNTER uses existing type-directed synthesis techniques to automatically generate any necessary adapter code.

In recent years, there has been a flurry of interest in automated code completion tools [14, 18, 24, 32, 30, 19, 36, 28]. Some of these techniques use type information [18, 24, 32], while others [36, 28] use statistical information about API usage patterns to find relevant completions. We believe that any type-directed code completion tool can be gainfully integrated into HUNTER for automatically synthesizing wrapper code from a candidate alignment.

Another related line of work is component-based synthesis [12, 16, 21, 17], which aims to synthesize programs from a database of underlying components. Component-based synthesis techniques have been used in a variety of application domains, including bit-vector algorithms [16], deobfuscators [21], and geometry constructions [17]. Most recently, Feng et al. have proposed a type-directed approach to component-based synthesis that utilizes Petri net reachability analysis [12]. As mentioned earlier, HUNTER internally uses Feng et al.’s SyPET tool in order to synthesize adapter code in a type-directed manner.

## 8. CONCLUSION

We have presented a tool called HUNTER for finding and automatically reusing existing methods. Given a search query and a set of test cases provided by the user, HUNTER searches massive code repositories for candidate adapter methods and automatically synthesizes any necessary wrapper code. The key technical idea underlying HUNTER is to compute a similarity metric between types and use this information for re-ranking search results and finding a suitable mapping between the parameters of the desired method and that of the adapter. We formulate the problem of finding an optimal (i.e., lowest cost) mapping in terms of integer linear programming (ILP), and automatically synthesize wrapper code from the solution to the ILP problem by leveraging existing type-directed synthesis techniques. Our experimental evaluation shows that HUNTER compares favorably with existing code reuse tools and that our re-ranking algorithm improves the quality of the search results. We have also performed a user study involving 21 participants and 3 different programming tasks. Our study demonstrates that HUNTER increases programmer productivity and helps programmers write less buggy code. The HUNTER tool can be obtained for free from the Eclipse marketplace and used as an Eclipse plug-in to aid software development tasks.

## 9. REFERENCES

- [1] Bitbucket. <https://bitbucket.org>.
- [2] GitHub. <https://github.com>.
- [3] Greppcode. <http://www.greppcode.com>.
- [4] Leetcode Online Judge. <https://leetcode.com>.
- [5] Mining and understanding software enclaves. <http://www.darpa.mil/program/mining-and-understanding-software-enclaves>.
- [6] OpenHub. <https://www.openhub.net>.
- [7] Pliny Project. <http://pliny.rice.edu>.
- [8] S. K. Bajracharya, J. Ossher, and C. V. Lopes. Sourcerer: An infrastructure for large-scale collection and analysis of open-source code. *Sci. Comput. Program.*, 79:241–259, 2014.
- [9] D. L. Berre and A. Parrain. The Sat4j library, release 2.2. *Journal on Satisfiability, Boolean Modeling and Computation*, pages 59–6, 2010.
- [10] S. Chatterjee, S. Juvekar, and K. Sen. SNIFF: A search engine for java using free-form queries. In *FASE*, pages 385–400, 2009.
- [11] W. J. Dixon and F. J. Massey Jr. *Introduction to statistical analysis*. McGraw-Hill, 1957.
- [12] Y. Feng, Y. Wang, R. Martins, A. A. Kaushik, and I. Dillig. Type-directed Component-based Synthesis using Petri Nets. Technical Report TR-16-01, Department of Computer Science, UT-Austin, 2016.
- [13] W. B. Frakes and T. P. Pole. An empirical study of representation methods for reusable software components. *Software Engineering, IEEE Transactions on*, 20(8):617–630, 1994.
- [14] J. Galenson, P. Reames, R. Bodík, B. Hartmann, and K. Sen. Codehint: dynamic and interactive synthesis of code snippets. In *ICSE*, pages 653–663, 2014.
- [15] M. Grechanik, C. Fu, Q. Xie, C. McMillan, D. Poshyvanyk, and C. M. Cumby. A search engine for finding highly relevant applications. In *ICSE*, pages 475–484, 2010.
- [16] S. Gulwani, S. Jha, A. Tiwari, and R. Venkatesan. Synthesis of loop-free programs. In *PLDI*, pages 62–73, 2011.
- [17] S. Gulwani, V. A. Korthikanti, and A. Tiwari. Synthesizing geometry constructions. In *PLDI*, pages 50–61, 2011.
- [18] T. Gvero, V. Kuncak, I. Kuraj, and R. Piskac. Complete completion using types and weights. In *PLDI*, pages 27–38, 2013.
- [19] R. Holmes and G. C. Murphy. Using structural context to recommend source code examples. In *ICSE*, pages 117–125, 2005.
- [20] O. Hummel, W. Janjic, and C. Atkinson. Code conjurer: Pulling reusable software out of thin air. *IEEE Software*, 25(5):45–52, 2008.
- [21] S. Jha, S. Gulwani, S. Seshia, and A. Tiwari. Oracle-guided component-based program synthesis. In *ICSE*, pages 215–224, 2010.
- [22] O. A. L. Lemos, S. K. Bajracharya, J. Ossher, P. C. Masiero, and C. V. Lopes. Applying test-driven code search to the reuse of auxiliary functionality. In *SAC*, pages 476–482, 2009.
- [23] Y. S. Maarek, D. M. Berry, and G. E. Kaiser. An information retrieval approach for automatically constructing software libraries. *Software Engineering, IEEE Transactions on*, 17(8):800–813, 1991.
- [24] D. Mandelin, L. Xu, R. Bodík, and D. Kimelman. Jungloid mining: helping to navigate the API jungle. In *PLDI*, pages 48–61, 2005.
- [25] L. Martie, T. D. LaToza, and A. van der Hoek. Codeexchange: Supporting reformulation of internet-scale code queries in context (T). In *ASE*, pages 24–35, 2015.
- [26] C. McMillan, M. Grechanik, D. Poshyvanyk, Q. Xie, and C. Fu. Portfolio: finding relevant functions and their usage. In *ICSE*, pages 111–120, 2011.
- [27] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: bringing order to the web. 1999.
- [28] V. Raychev, M. T. Vechev, and E. Yahav. Code completion with statistical language models. In *PLDI*, page 44, 2014.
- [29] S. P. Reiss. Semantics-based code search. In *ICSE*, pages 243–253, 2009.
- [30] N. Sahavechaphan and K. T. Claypool. Xsnippet: mining for sample code. In *OOPSLA*, pages 413–430, 2006.
- [31] K. T. Stolee, S. G. Elbaum, and D. Dobos. Solving the search for source code. *ACM Trans. Softw. Eng. Methodol.*, 23(3):26, 2014.
- [32] S. Thummalapenta and T. Xie. Parseweb: a programmer assistant for reusing open source code on the web. In *ASE*, pages 204–213, 2007.
- [33] Y. Ye and G. Fischer. Supporting reuse by delivering task-relevant and personalized information. In *ICSE*, pages 513–523, 2002.
- [34] A. M. Zaremski and J. M. Wing. Signature matching: A tool for using software libraries. *ACM Trans. Softw. Eng. Methodol.*, 4(2):146–170, 1995.
- [35] A. M. Zaremski and J. M. Wing. Specification matching of software components. In *ESEC/FSE*, pages 6–17, 1995.
- [36] H. Zhong, T. Xie, L. Zhang, J. Pei, and H. Mei. MAPO: mining and recommending API usage patterns. In *ECOOP*, pages 318–343, 2009.