

Analogy-Making as a Core Primitive in the Software Engineering Toolbox

Matthew Sotoudeh
masotoudeh@ucdavis.edu
University of California, Davis
USA

Aditya V. Thakur
avthakur@ucdavis.edu
University of California, Davis
USA

Abstract

An *analogy* is an identification of structural similarities and correspondences between two objects. Computational models of analogy making have been studied extensively in the field of cognitive science to better understand high-level human cognition. For instance, Melanie Mitchell and Douglas Hofstadter sought to better understand high-level perception by developing the Copycat algorithm for completing analogies between letter sequences. In this paper, we argue that analogy making should be seen as a core primitive in software engineering. We motivate this argument by showing how complex software engineering problems such as program understanding and source-code transformation learning can be reduced to an instance of the analogy-making problem. We demonstrate this idea using SIFTER, a new analogy-making algorithm suitable for software engineering applications that adapts and extends ideas from Copycat. In particular, SIFTER reduces analogy-making to searching for a sequence of update rule applications. SIFTER uses a novel representation for mathematical structures capable of effectively representing the wide variety of information embedded in software.

CCS Concepts: • Computing methodologies → Reasoning about belief and knowledge; • Software and its engineering → Software maintenance tools.

Keywords: fluid analogies, software maintenance, static analysis, program understanding

ACM Reference Format:

Matthew Sotoudeh and Aditya V. Thakur. 2020. Analogy-Making as a Core Primitive in the Software Engineering Toolbox. In *Proceedings of the 2020 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! '20)*, November 18–20, 2020, Virtual, USA. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3426428.3426918>



This work is licensed under a Creative Commons Attribution International 4.0 License.

Onward! '20, November 18–20, 2020, Virtual, USA
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8178-9/20/11.
<https://doi.org/10.1145/3426428.3426918>

1 Introduction

An *analogy* is defined as “a comparison between two objects, or systems of objects, that highlights respects in which they are thought to be similar” [5]. Humans make complex, fluid analogies in everyday communication. For example, a recent CNN headline [84] states that “the ‘beating hearts’ of these pulsating stars create music to astronomers’ ears,” noting correspondences between pulsation of stars, the beating of a heart, and the rhythm of music. In fact, analogy-making is such a fundamental skill that a major portion of the original SAT exam was dedicated to having test-takers complete analogies such as “Paltry is to significance as X is to Y.”

Analogies have been shown to be a useful instructional tool for improving student learning [23]. For instance, analogies are frequently used to explain concepts in physics [11, 66], geography [56], mathematics [71, 72], and computer science [10, 15, 16, 19, 24, 31, 47]. For example, a professor of an introductory computer science course might explain the concept of a program by forming an *analogy between* the execution of a program by a computer and the following of a recipe by a cook, likening steps in a recipe to statements in a computer program, ingredients in a recipe to resources in a program or user inputs — analogies are not always unambiguous.

The important role of analogy-making in high-level cognition has been argued by many researchers in both cognitive science and computer science [17, 18, 53]. Professor Douglas Hofstadter and his Fluid Analogies Reasoning Group (FARG) [32, 33] argue that humans are constantly making analogies, comparing features of their current situation with previously-encountered scenarios to decide what to do next. To study the analogy-making process in more detail, they have developed *computational models of analogy making*. One such algorithm, *Copycat* [34], can complete analogies over strings (Section 2). Copycat can answer questions such as “*abc* is to *abd* as *efg* is to what?”

This paper explores the role of analogy-making in software engineering (SE). We argue that a number of SE problems can be framed as analogy-making. We introduce a dedicated analogy-making algorithm, SIFTER, which can be used as a core primitive to solve such SE problems. We believe that future research on analogy-making algorithms like SIFTER will solidify analogy-making as a core primitive in the software engineering toolbox.

The first application of analogy-making we consider is the problem of *program understanding* (Section 3.1). One fundamental way that humans understand new source code is by analogy to source code that they already understand. As an example, consider a programmer who has become familiar with the source code of the Bourne-Again Shell (bash) [7] and now wishes to add a new feature to the Friendly Interactive Shell (fish) [6]. A reasonable first step would be to read the source code of fish and attempt to identify functions in its implementation that play a similar role to more familiar functions in the implementation of bash. In other words, the programmer will *form an analogy* between the two source code repositories and use the analogy to determine where to make the desired modification.

The second application of analogy-making we consider is the problem of *generalizing a source code optimization* (Section 3.2). As an example, consider a scenario where an optimized matrix multiplication implementation should be used when the matrix sizes satisfy certain conditions. After seeing a small number of code corrections replacing the sub-optimal multiplication routine with the optimized one, most trained programmers would begin to internalize the pattern. Consequently, when they come across new code that can be optimized they would be able to form an analogy to those corrections and modify the code to use the optimized version. Here, the programmer needs to form an analogy *between pairs of programs*, before and after the transformation, for example stating that “all of the sub-optimal code calls matrix multiplication routine X with inputs like Y , and the corresponding optimized code replaces X with Z .”

Finally, the third application of analogy-making we consider is the problem of *API migration* (Section 3.3). As an example, consider the scenario where a library that has updated its public interface to change its error codes and remove a now-redundant parameter from each relevant function. After updating a small number of functions to use the new API, a programmer might quickly realize that all of the changes they need to make are “fundamentally the same:” lookup the new error code in the documentation, switch the old error code for the new wherever it is checked for, and then remove the redundant parameter. The programmer has, thus, formed an analogy between the different edits, which they can use to quickly migrate similar code elsewhere in the project.

In all of these examples, and many more discussed in Section 3.4, the programmer reasons about the source code and relations between different parts of the source code to form an *analogy* highlighting the fundamental similarities between a set of examples.

Motivated by this, we developed SIFTER, an analogy-making algorithm suitable for analogy-making on source code (Section 4). SIFTER takes as input a description of the source code in a mathematical format using arbitrary relations. This description is expressive enough to represent syntactic as well as semantic information about the source code. SIFTER can

also take as input and reference in its analogies non-code sources like documentation. The design of SIFTER is principled and ultimately reduced the analogy-making problem to that of a search over possible rule applications for rewriting a workspace described as a *triplet structure*.

We believe that SIFTER can form a powerful primitive in the SE toolbox. SIFTER is general enough to handle arbitrary relations, and, hence, can make analogies synthesizing semantic, syntactic, and natural-language information. SIFTER’s output is explainable, as it internally solves analogies by symbolic manipulation and identifying corresponding facts. Finally, by reducing many distinct problems to SIFTER, improvements to it will directly pay dividends across a wide number of applications. The implementation of SIFTER is available at <https://github.com/95616ARG/sifter>.

2 Analogy Making over Strings

This section explains the notion of analogy making. Following Hofstadter, we focus first on analogies between *letter strings*, sequences of letters and symbols. Section 2.1 then describes analogy completion. We end by describing Copycat (Section 2.2), a system from prior work for completing analogies involving such letter strings, which influenced the design of the SIFTER system introduced in this paper.


Informally, analogy making over strings entails determining in what respects two given strings are similar. For instance, consider the strings abc and efg . For the string abc , we have that the second letter is the successor in alphabetical order of the first, and the third is the successor of the second. The exact same property holds for the string efg .

More formally, analogy making entails inferring properties that hold for both strings. Consider the binary relation NextTo , where $\text{NextTo}(x, y)$ implies that the letter x is the left of the letter y in a given string. For the string abc , we have $\text{NextTo}(a, b)$ and $\text{NextTo}(b, c)$. Furthermore, consider the binary relation LetterSuccessor where $\text{LetterSuccessor}(x, y)$ holds if the letter y is the successor of letter x in alphabetical order. In this example, we have $\text{LetterSuccessor}(a, b)$, $\text{LetterSuccessor}(b, c)$, $\text{LetterSuccessor}(e, f)$, and so on. Using these two relations, we can state that the elements $S = \{a, b, c\}$ of the string abc satisfy the following property φ : $\forall x, y \in S. \text{NextTo}(x, y) \implies \text{LetterSuccessor}(x, y)$. We see that the elements of the string efg satisfy this same property. Hence we have formed a meaningful analogy between abc and efg .

However, not all analogies can be succinctly expressed via a first-order logic formula such as φ . Consider making an analogy between the strings abc and gfe . In this case, the elements of abc satisfy φ , but the elements of gfe satisfy the property φ' : $\forall x, y \in S. \text{NextTo}(x, y) \implies \text{LetterSuccessor}(y, x)$. In particular, the order of the arguments to LetterSuccessor in φ and φ' are different. Intuitively, the property φ' reads the string from right to left, instead of

left to right. Consequently, the analogy maker needs to be able to express such “slips:” the strings *abc* and *gfe* satisfy *almost* the same property, *except that* one reads the string from left to right and the other from right to left.

Such properties can become even more difficult to state in a first-order logic notation when *grouping* is involved. Consider making an analogy between the strings *aaabbc* and *ddddccbbba*. There are many reasonable analogies between these strings. For example, we might associate the *group of letters* *aaa* in the first string with *a* in the second string, noting that all of the letters in the first and the one letter in the second satisfy the unary *IsLetterA* predicate. Alternatively, we might associate *aaa* in the first string with *dddd* in the second string, as they are both groups of a single repeated letter occurring at the start of their corresponding string. In particular, the strength of an analogy lies less in the number of features the two strings have in common than in the overlap of relational structure between the two strings [22]. Regardless, it is not clear how one might naturally express the fundamental properties that these two strings share in a standard logic notation. Instead, as humans we might be tempted to communicate the shared property via a drawing like the following:




In this drawing, G_1 through G_n represent the groups of letters in each string. Each group is made up of letters x_1, \dots, x_m , which all satisfy the *LetterSame* relation with each other. The groups themselves then have some shared binary relation r_1 relating them. For example, in *aaabbc*, we might have *aaa* take the place of G_1 and $r_1(x, y) = \text{LetterSuccessor}(x, y) \wedge \text{NextTo}(x, y)$, while in *ddddccbbba* we might have *a* take the place of G_1 with $r_1(x, y) = \text{LetterSuccessor}(x, y) \wedge \text{NextTo}(y, x)$. Note again that we have flipped the order in the latter *NextTo*, intuitively to read the second string from right-to-left.

Such drawings motivate an alternate way of thinking about analogy making. In this interpretation, analogy-making involves defining an abstract string description that can be instantiated to produce the given strings. As we have seen, such an abstract string description needs to be general enough to handle all of the complex grouping and slipping that can occur in such letter string analogies.

2.1 Analogy Completion on Transformations

One particularly interesting use-case for analogy-making is to make analogies between *pairs* of objects, such as a state before and after some transformation. For example, given the pairs *abc*→*abd* and *efg*→*efh*, we can compare the two pairs, forming an analogy that might be represented by the following abstraction.



In this scenario, the analogy-making process involves learning an *abstract representation* of the transformation performed in each pair of strings. We have said that both pairs of strings correspond to each other because they both share the properties shown in this diagram.

With such a representation of a transformation, one can also perform *Analogy Completion*, like in the original SAT exam. For example, given two pairs of strings *abc*→*abd* and *efg*→*efh* and a *prompt* *ijk*→?, we can ask for a *completion of the analogy*, a value which can replace the ? to make all three pairs form a strong analogy. In this case, one completion would be the string *ijl*.

We can find such a completion by first constructing an analogy between the examples *abc*→*abd* and *efg*→*efh* to form the abstract string drawn above, then we can start to form an analogy with the examples and the prompt *ijk*, from which we might infer that *i* is an instance of x_1 , *j* an instance of x_2 , and *k* an instance of x_3 in the abstract string. We can then *infer* that there should be some letters corresponding to y_1, y_2, y_3 , and that they should satisfy the properties in the drawing. From the drawing, then, we have $\text{LetterSame}(i, y_1)$, hence we should have $y_1 = i$; then $\text{LetterSuccessor}(y_1 = i, y_2)$ to get $y_2 = j$; and finally $\text{LetterSuccessor}(k, y_3)$ to get $y_3 = l$, completing the analogy with the desired string *ijl*.

2.2 The Copycat Algorithm

The Copycat algorithm [34] was developed to solve such string analogy-completion problems. Copycat’s architecture is similar to that of a blackboard-based automated theorem prover. It consists of a *workspace*, or blackboard, which initially contains only the example pairs (such as *abc*→*abd*) and a prompt (such as *ijk*→?). This workspace is modified by a set of *codelets*, which are small programs that operate on the workspace. These codelets can make a variety of modifications to the workspace. Some codelets may *group* letters together, like the *aaa* in *aaabbc*. Other codelets identify *bonds*, or relational facts about symbols, for example noting that $\text{LetterSuccessor}(a, b)$ is true. Still further codelets can build *bridges* between symbols, representing the determined correspondences in the analogy. Once a consistent analogy among the examples and prompt is made, a purpose-built solver is used to construct the corresponding completion.

The resulting analogy is determined by the order and type of codelets used, along with where each one “focuses.” The behavior of the codelets is controlled by what is essentially a sophisticated set of heuristics wrapped into a structure known as a *Slipnet*.

3 Applications of Analogy Makers in SE

In this section, we demonstrate the use of our analogy maker SIFTER by applying it to three SE problems and list specific challenges addressed by SIFTER. A large number of further applications are then discussed in Section 3.4. We defer the description of the design of SIFTER to Section 4.

3.1 Comparative Program Understanding

Suppose you are a programmer who is quite familiar with the implementation (source code) of bash [7], and you would like to add a new feature to a *different* shell, such as fish [6]. Because you have been adding features to the bash shell for many years, you may know exactly which function(s) to modify in the source code of bash to add the desired feature. However, being new to the source code of fish, you face a significant challenge understanding the implementation of fish before you can even begin writing your new feature.

You might start by reading the source code of fish, looking for functions and objects that play similar roles to ones you are more familiar with in bash. This is fundamentally an analogy-making problem, where we are attempting to form an analogy between the source code of bash (which we are familiar with) and that of fish (which we are not).

This problem is similar to that of forming letter analogies in Section 2, where, given the strings `abc` and `efg`, we found that `b` and `f` corresponded to each other. Here, `abc` is instead the source code of bash, while `efg` is the source code of fish. `b` and `f` are likewise functions in bash and that which play similar roles.

SIFTER is designed to make such analogies on programs. We first load both source repositories into SIFTER, then ask it to identify an analogy between the two. This analogy will effectively be a *mapping between the repositories*, identifying functions, classes, and statements in the source code of bash with those in the source code of fish. Such an analogy can form an invaluable guide, allowing you to look up, for example, the function in fish which SIFTER thinks plays the most similar role to the one you would have modified in bash.

3.1.1 Challenge 1: Compositional Structure. Program source-code relies heavily on compositional structure. For example, the meaning and interpretation of any function depends not only on its immediate body, but also on the body of all functions that it calls, and the functions they call, and so forth. On the other hand, the role a function plays in a large piece of software usually depends on the functions that call it.

SIFTER can make use of this compositional structure, by building on top of analogies it has already made. For example, in Figure 1 we have provided snippets of the source code of the bash and fish shells. SIFTER begins by associating the functions `cd_builtin` (B1) and `builtin_cd` (F1) because of their similar names and signatures. Then, once it has decided that these two functions correspond, SIFTER can start

to make inferences that the places where they are used are likely to correspond as well. For example, it might note that both `shell_builtins` (B3) and `builtin_datas` (F2) contain a struct with a field of `cd_builtin` or `builtin_cd` respectively, and mark those two objects as corresponding in the analogy. It can similarly infer that functions using those objects, such as `builtin_address_internal` (B4) in bash and `builtin_lookup` (F3) in fish, correspond.

In this way, SIFTER can build up analogies made about parts of the program to begin to make stronger and stronger inferences about how the rest of the source code corresponds. Although we have not demonstrated it in this example, SIFTER can also make analogies in a top-down fashion, e.g., by starting at the main function in both programs, or by alternating between such top-down and bottom-up strategies. Notably, such compositional structure was not needed to find analogies between the letter groups in Section 2, demonstrating how analogy-making on programs can be richer and more challenging than on letter strings.

3.1.2 Challenge 2: Multiple Representations. Another challenge with analogy-making on programs is that semantically equivalent programs can have multiple syntactic representations. For example, functions can be inlined or `if/else` conditionals can be inverted. SIFTER can handle such scenarios by applying *transformation rules*, such as function inlining, to transform either source repository it is given. SIFTER searches through different representations of each source repository until it finds ones that are amenable to forming strong analogies.

Because letter strings do not have an assumed semantics, there is no equivalent notion of semantics-preserving transformation rules for the examples in Section 2. However, operationally, the process of grouping letters, e.g., in `aaabbc`, can be seen as such a transformation, where the internal representation of the individual letters `a`, `a`, and `a` are transformed into a single *group* of letters `aaa`.

3.2 Generalizing Program Transformations from Examples

Suppose we have a linear algebra library with multiple General Matrix Multiply (GEMM) routines for computing matrix multiplications. Some routines, such as `gemm_large`, are optimized for the case where the input matrices are relatively large, say with over 1,000 rows each, while others like `gemm_skinny` are optimized for “skinny” inputs, e.g., where the inner dimension is half the size of either of the outer dimensions.

For a particular team working on a particular codebase, it may be the case that most matrices are usually quite large and so `gemm_large` might become the de-facto routine that developers use in new code without thinking too deeply about matrix sizes, or simply used due to copy and paste from existing code. While this might be a reasonable default

```

1 int cd_builtin (list) WORD_LIST *list; { ... } //B1
2
3 struct builtin static_shell_builtins[] = { ...
4   { "cd", cd_builtin, ... }, //B2
5   ... }
6 struct builtin *shell_builtins = static_shell_builtins
7   ; //B3
8
9 struct builtin * builtin_address_internal
10   (name, disabled_okay)
11   char *name; int disabled_okay; { ... //B4
12   j = shell_builtins[mid].name[0] - name[0];
13   ... }

```

(a) Bash source

```

1 int builtin_cd(parser_t &parser, io_streams_t &streams
2   , wchar_t **argv) { ... } //F1
3
4 static const builtin_data_t builtin_datas[] = {
5   ...
6   {L"cd", &builtin_cd, ...}, //F2
7   ... }
8
9 static const builtin_data_t *builtin_lookup(const
10   wchar_t *name) { //F3
11   const builtin_data_t *array_end = builtin_datas +
12     BUILTIN_COUNT;
13   ... }

```

(b) Fish source

Figure 1. Comparative program understanding between bash and fish source code. SIFTER can form an analogy between these implementations, helping explain the code for fish to a programmer used to the codebase of bash by noting where objects and functions (like `builtin_address_internal` in bash and `builtin_lookup` in fish) play similar roles in each.

for this team, it is likely the case that *some* matrices in a program are better suited for `gemm_skinny` — in that case, the instinctive default would be sub-optimal, and another programmer might notice during code review that `gemm_skinny` would be a better choice.

The question we would like to consider is: given two example code pairs where `gemm_large` has been transformed into `gemm_skinny`, can we automatically optimize new code in the same manner? This is fundamentally an analogy problem, where we would like to compare the pairs of pre- and post-replacement code to learn the *core transformation* that explains all of them. We can then use this analogy to infer, for some new sub-optimal code, the corresponding optimized code. This is similar to completing the analogy $abc \rightarrow abd$, $efg \rightarrow efh$, $ijk \rightarrow ?$ in Section 2.1.

This scenario is shown in Figure 2. The first two rows in that figure show pairs of examples of the desired source code transformation provided to SIFTER, which play the same role as $abc \rightarrow abd$ and $efg \rightarrow efh$ in Section 2.1. The code in the left-hand column is sub-optimal because it calls `gemm_large` on matrices with dimensions that would be better suited for `gemm_skinny`. The code in the right-hand column has been optimized by replacing the call to `gemm_large` with a call to `gemm_skinny`. In the third row of Figure 2, we have provided SIFTER with a new piece of code on the left and ask it to *complete the analogy*, i.e., produce the corresponding piece of code on the right that makes all three rows the most similar. This plays the same role as the $efg \rightarrow ?$ input in the letter analogy example. The code produced by SIFTER is shown on the bottom right of Figure 2 in green, where we see it has correctly replaced the call to `gemm_large` with a call to `gemm_skinny`.

3.2.1 Challenge 3: Avoiding False Positives. SIFTER is forming an analogy between the rows in Figure 2, including between the before code on the left-hand side. In Figure 2, for example, SIFTER has noted as part of its analogy that all of the left-hand code snippets call the function `gemm_large`

with the last argument at least twice that of the second-to-last argument. This behavior can be thought of as learning to *recognize code that can be optimized*, and can be used to avoid false positives. For example, suppose instead of the sub-optimal prompt code given in the bottom-left of Figure 2, we gave SIFTER the code:

```

1 assert(k > 0);
2 int outer = k * 10;
3 int inner = k * 10;
4 read_mat(outer, inner, &A);
5 read_mat(inner, outer, &B);
6 gemm_large(A, B, &C, outer, inner, outer);

```

In that scenario, SIFTER would attempt to form an analogy between this code and the example before-transformation code on the left-hand side of the first two rows of Figure 2. While it may succeed in proposing an analogy, SIFTER will note as part of its output that the analogy is not particularly strong. This is because this new code does not share the property that the last argument to `gemm_large` is at least twice that of the second-to-last one.

If examples of already-optimized code are available, given a new instance SIFTER can also try to form an analogy using these *negative* examples. A threshold can be set based on a comparison with the negative vs. positive examples to determine whether to apply the transformation.

3.2.2 Challenge 4: Using Semantic Information. Recognizing sub-optimal code relies on *semantic information* about the possible values a variable can take on. In particular, we only want to apply the transformation when the inner matrix dimension is at most half the size of the outer dimensions. This would cause difficulty for syntax-based tools like GetAFix [4]. However, SIFTER takes as input an *arbitrary structure* consisting of symbols and relations between the symbols. This means that, in addition to providing the source code, we can annotate the structure representing the source code with the results of a program analyzer, which allows us to include information about semantic properties of the code. In this example, we can annotate the structure

<pre> 1 assert(k > 0); 2 int outer = k * 100; 3 int inner = k * 10; 4 read_mat(outer, inner, &A); 5 read_mat(inner, outer, &B); 6 gemm_large(A, B, &C, outer, inner, outer); </pre>	<pre> 1 assert(k > 0); 2 int outer = k * 100; 3 int inner = k * 10; 4 read_mat(outer, inner, &A); 5 read_mat(inner, outer, &B); 6 gemm_skinny(A, B, &C, outer, inner, outer); </pre>
<pre> 1 assert(k > 1); 2 int outer = k, A_cols = k / 2; 3 read_mat(outer, A_cols, &A); 4 read_mat(A_cols, outer, &B); 5 while (!done(A, B)) { 6 read_row(&A); 7 read_col(&B); 8 outer++; } 9 gemm_large(A, B, &C, outer, A_cols, outer); </pre>	<pre> 1 assert(k > 1); 2 int outer = k, A_cols = k / 2; 3 read_mat(outer, A_cols, &A); 4 read_mat(A_cols, outer, &B); 5 while (!done(A, B)) { 6 read_row(&A); 7 read_col(&B); 8 outer++; } 9 gemm_skinny(A, B, &C, outer, A_cols, outer); </pre>
<pre> 1 assert(k > 5); 2 int AB_rowcol = k * k; 3 int inner = k; 4 read_mat(AB_rowcol, inner, &A); 5 read_mat(inner, AB_rowcol, &B); 6 gemm_large(A,B,&C, AB_rowcol, inner, AB_rowcol); </pre>	<pre> 1 assert(k > 5); 2 int AB_rowcol = k * k; 3 int inner = k; 4 read_mat(AB_rowcol, inner, &A); 5 read_mat(inner, AB_rowcol, &B); 6 gemm_skinny(A,B,&C, AB_rowcol, inner, AB_rowcol); </pre>

Figure 2. Optimizing program source code with SIFTER. The left column shows before the optimization and the right column after the optimization. The first two lines are the examples given to SIFTER, while the green code is generated by SIFTER to complete the analogy for the last line. SIFTER includes in its analogy semantic information like the fact that all of the outer dimensions in the examples are at least twice that of the inner dimensions, helping it to avoid false positives.

with relations that some variable is always at least twice that of another. The analogy-making algorithm that forms the core of SIFTER is entirely indifferent to the underlying relations being used, and will use such semantic relations to form analogies just like it would more syntactic relationships describing the source code.

3.3 API Migration

For the final task, consider the two versions of a camera library documented in Figure 3, which has been updated to automatically determine the resolution to use as well as changed the error codes. Suppose you have a program that uses version 1 of this camera library. Snippets of this program are shown on the left column of Figure 4. Your colleague has partially migrated your code to use version 2 of the camera library, as shown in the first two boxes on the right column of Figure 4. You would now like to migrate the bottom-left code in Figure 4 to use version 2.

This is an analogy completion problem, where we want to form an analogy between the examples that effectively describes how to migrate code to use the new API. We then want to find a *completion*, or source code to fill in to the bottom-right of Figure 4 that makes the last row analogous to the first. Because this involves transforming source code, the input/output to SIFTER will be similar to that of Section 3.2. However, as we will see, we face different challenges here.

Just like in Section 3.2, we can provide the first two rows of Figure 4 to SIFTER as examples of the desired migration, then ask it to *complete the analogy* by producing the migrated code corresponding to the prompt (unmigrated) code in the

bottom left of Figure 4. The code generated by SIFTER is shown in green on the bottom left of Figure 4.

3.3.1 Challenge 5: Using Documentation. There is one particularly pressing challenge we wish to highlight here: given *only* these code pairs, there is no reasonable way to complete this analogy, because it depends on knowing the new error code for `record_frame`, which did not appear in any of the examples. To address this, we can *give SIFTER the before/after documentation* in addition to just the source code. Again, because SIFTER takes any structure as input, we can encode documentation just as easily as we can encode source code. Then, in making analogies, SIFTER can *refer to the documentation* and include looking up in the documentation as part of the analogy.

3.3.2 Challenge 6: Using DNN Models. Once we begin involving arbitrary text (e.g., in documentation), we need to start being able to handle fuzziness inherent in human languages. For example, in the documentation for the first two functions, the corresponding error message could be found with a relatively simple search, because it was prefaced with “On error ... returns.” However, for the `record_frame` function, the corresponding sentence uses “failure” instead of “error,” which could cause SIFTER to lose confidence in its analogy. To help increase SIFTER’s confidence in its analogy and guide it towards the right answer, we can use existing natural-language tools such as DNN-based sentiment analysis models. These models take a paragraph, sentence, or word and estimate how positive or negative it is. Just like with the results of a program analyzer, we can annotate this

<pre> 1 # CameraLib v1.0 2 ## `record_video(buffer, buffer_size, resolution)` 3 Records video from the main camera into `buffer` until `buffer_size` bytes are reached. On error returns -1. 4 ## `record_audio(buffer, buffer_size, resolution)` 5 Uses the main camera's microphone to record audio into `buffer` until `buffer_size` bytes have been recorded. On error returns -5. 6 ## `record_frame(buffer, buffer_size, resolution)` 7 Uses the main camera to record a single image to ` buffer`. On failure returns -3. </pre>	<pre> 1 # CameraLib v2.0 2 ## `record_video(buffer, buffer_size)` 3 Records video from the main camera into `buffer` until `buffer_size` bytes are reached. On error returns -4. 4 ## `record_audio(buffer, buffer_size)` 5 Uses the main camera's microphone to record audio into `buffer` until `buffer_size` bytes have been recorded. On error returns -2. 6 ## `record_frame(buffer, buffer_size)` 7 Uses the main camera to record a single image to ` buffer`. Automatically sets the resolution to fit in `buffer_size`. On failure returns -6. </pre>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 3. API documentation before after migration for the camera API. This documentation is provided to SIFTER to complete the analogy shown in Figure 4.

<pre> 1 ... 2 void try_record_video() { 3 int result = record_video(buffer, BUFFER_SIZE, 4 RES_AUTO); 5 if (result == -1) { 6 printf("Could not record video.\n"); } 7 ... } </pre>	<pre> 1 ... 2 void try_record_video() { 3 int result = record_video(buffer, BUFFER_SIZE); 4 if (result == -4) { 5 printf("Could not record video.\n"); } 6 ... } </pre>
<pre> 1 ... 2 void try_record_audio() { 3 int result = record_audio(buffer, BUFFER_SIZE, 4 RES_AUTO); 5 if (result == -5) { 6 printf("Could not record audio.\n"); } 7 ... } </pre>	<pre> 1 ... 2 void try_record_audio() { 3 int result = record_audio(buffer, BUFFER_SIZE); 4 if (result == -2) { 5 printf("Could not record audio.\n"); } 6 ... } </pre>
<pre> 1 ... 2 void try_record_still() { 3 int result = record_frame(buffer, BUFFER_SIZE, 4 RES_AUTO); 5 if (result == -3) { 6 printf("Could not record still.\n"); } 7 ... } </pre>	<pre> 1 ... 2 void try_record_still() { 3 int result = record_frame(buffer, BUFFER_SIZE); 4 if (result == -6) { 5 printf("Could not record still.\n"); } 6 ... } </pre>

Figure 4. API migration with SIFTER. The left (right) column shows before (after) the migration. The first two rows are the input to SIFTER, while the green code is generated by SIFTER to complete the analogy for the last row. SIFTER was also given as input the API documentation pair from Figure 3, which is where it looks to find the new error code used in the generated code.

information on top of the structure, e.g., by marking all symbols representing words which are determined to be highly positive or highly negative by the sentiment analyzer with a unary relation like `IsNegativeSentiment`. Even though they are not exactly the same word, the fact that they both satisfy the `IsNegativeSentiment` relation will increase SIFTER's confidence in the analogy.

3.4 Future Software Engineering Applications

Future work can apply analogy completion to more varied input/output domains and languages. For example, analogy completion can *generate documentation* based on existing code, similar to the problem of automated comment generation, which has been addressed by code-clone detection [90]. We can also treat translation between programming languages as analogy completion using examples of semantically equivalent programs in either language. This will likely rely on a number of smaller analogies, matching common

patterns in one programming language and mapping them to idiomatic code for that pattern in another, similar to existing work in this field [48, 51].

A smart editor may make analogies between the user's current editor state and a corpus of code samples to suggest structural code completions, similar to those accomplished using large code corpuses like Aroma [46] or more local history such as Blue Pencil [52]. Because our approach does not rely on AST parsing (see Section 5.1), we can use analogies to automatically improve tooling (error messages, linting, suggestions, bug finding, etc.) for nascent and domain-specific languages (DSLs) that may not yet have a formal grammar [9, 62].

One might also apply analogy-makers to recover higher-level structure from low-level binaries or compiler intermediate representations. Such information would be useful for decompilation of binary programs [2], in-place binary analysis [82], binary rewriting [87], de-obfuscation of code [67],

and identification of code replaceable with highly-optimized libraries [55] or coprocessors [83].

Analogies between existing code with correctness proofs and new code may allow for *proof transfer* to more quickly prove correctness of the latter. A related technique has been proposed for the Coq theorem prover [64].

More varied information sources can be used for analogy making. For example, the use of information from a profiler may be helpful in a code optimization setting, and compiler error-messages may be useful when using analogies to provide edit suggestions for syntactically-invalid code. This information could be used to rank different proposed analogies or to help find them in the first place.

In a classroom setting, analogy-makers can cluster student assignments, an important problem which is currently addressed via a variety of different techniques [29, 36].

SIFTER can also detect where a strong analogy *almost holds*, under a small modification its inputs, e.g., in *abc* and *xyf*. If applied to common coding patterns, this might suggest the existence of a bug in the program’s implementation of this pattern. For example, one binary-search implementation might be *almost* analogous to a reference one, except that it computes the midpoint as $(1+h)/2$ instead of $1+(h-1)/2$, introducing a subtle integer overflow bug that a future version of SIFTER might flag as anomalous in the analogy.

4 Design of SIFTER

In this section, we describe the design of our analogy-making algorithm SIFTER and illustrate how it addresses the challenges discussed in Section 3. Its design was influenced by that of Copycat (Section 2.2). However, as we will discuss in Section 6, Copycat’s implementation was specially designed for the letter-analogy domain, whereas we would like to support *arbitrary* relations and input structures.

At a high level, the behavior of SIFTER is formulated as a number of *update rules* operating on a *workspace*. The workspace initially contains a representation of the source code and other inputs which it is supposed to make analogies about. Update rules gradually modify the workspace, both identifying facts, such as when some letter is a successor of another, and making new analogies. Analogies made during this process are explicitly represented within the workspace, as discussed in Section 4.4, and new analogies can build iteratively on existing analogies in the workspace. When a sufficient analogy is found by the system, it can be read off directly from the workspace and returned to the user.

In Section 4.1 we will describe *triplet structures*, a novel data structure used to represent SIFTER’s workspace. Triplet structures can represent arbitrary relational facts in a standardized way, making them a particularly flexible tool for representing SIFTER’s workspace. In Section 4.2 we will describe how we initialize the triplet structure representing the

SIFTER workspace for an example analogy problem. In Section 4.3, we introduce a domain-specific language for expressing *update rules* that modify triplet structures, and can be used to infer new facts about the objects in question. Section 4.4 describes how analogies are represented in the workspace, while Section 4.5 describes update rules that can be used to automatically find such analogies.

4.1 Triplet Structures

A *triplet structure* is a novel data structure used to represent the state of SIFTER’s workspace. A triplet structure:

1. Represents objects and facts in a standardized form, so that code for operating on the workspace does not have to worry about details like arity of relations.
2. Is able to naturally represent *partial facts*, e.g., we can represent the state “I know letter O is the predecessor of *something*, but I’m not sure exactly what yet.”
3. Supports efficient lookups and queries, so that operations on the structure can be performed quickly.

Definition 4.1. A *triplet structure* is a pair of sets (S, F) where $F \subseteq S \times S \times S$. We call each member of S a *node*, each member of $S \times S \times S$ a *triplet fact*, and F the set of triplet facts in the structure.

We can encode any finite mathematical structure as a triplet structure with polynomial increase in size. First, for every n -ary relation R , we add n nodes to the triplet structure representing slots in the relation. Generally, for an n -ary relation R we can always add nodes $R:1$ through $R:n$, although we will usually use more descriptive names in our examples. Second, each fact in the original structure gets a *fact node* in the triplet structure, which is a node in S that *represents the original fact itself* in the triplet structure. For a fact in the original structure of the form $R(x_1, x_2, \dots, x_n)$ corresponding to a fact node f , we then add triplet facts of the form $(f, x_i, R:i)$ for every $i \in \{1, 2, \dots, n\}$.

Fact nodes in triplet structures can be thought of as C-style **structs**, where each fact (f, v, k) asserts that the field k in struct f takes the value v . Alternatively, each fact node f can be thought of as expressing an interpretation of part of the structure, with a fact (f, v, k) asserting that, in the interpretation f , v is of type k .

Example 4.2. Consider a mathematical structure with objects $O = \{x, y, z\}$, a single binary relation R , and two relational facts $R(x, y), R(y, z)$.

To encode this mathematical structure as a triplet structure, we break the binary relation R into two nodes $R:1$ and $R:2$ representing each of its slots. We then create the fact node f_1 for $R(x, y)$ and the fact node f_2 for $R(y, z)$. We also add nodes for each of the original objects in O to get the set of nodes:

$$S = \{x, y, z, R:1, R:2, f_1, f_2\}.$$

Finally, we add triplet facts relating each slot of each fact to arrive at the set of triplet facts in the structure:

$$F = \{(f_1, x, R:1), (f_1, y, R:2), (f_2, y, R:1), (f_2, z, R:2)\}.$$

Triplet structures also have an intuitive graph representation. Nodes in the structure correspond to nodes in the graph. For each triplet fact (f, v, k) , we add an edge $v \rightarrow k$ with label f . The graph for the triplet structure considered in this example is the Structure T1 (Figure 5) shown below.




Figure 5. Triplet Structure T1

Note that, while we have drawn T1 using descriptive names, shapes, and colors, no intrinsic meaning is assigned to any symbol. In the rest of this paper, we will usually only show this visual representation of a triplet structure instead of explicitly listing the nodes and facts. Hence, the reader is encouraged to ensure the connection between the two is well-understood before proceeding.

In addition to directly encoding relational facts, some structures can be more naturally expressed directly as a triplet structure. This is highlighted in the next example.

Example 4.3. Consider encoding the scenario “Homer and Marge are the parents of Bart and Lisa.” We may encode this as four facts of the form $Parent(Homer, Bart)$, $Parent(Homer, Lisa)$, $Parent(Marge, Bart)$, $Parent(Marge, Lisa)$. With triplet structures, we can express this by saying “Homer, Marge, Bart, and Lisa form a family, where Homer and Marge are the parents, and Bart and Lisa are the children.” This scenario is represented by Structure T2 (Figure 6) below.




Figure 6. Triplet Structure T2

Another feature of triplet structures is that they can represent *partial facts*, as demonstrated by the next example.

Example 4.4. Suppose in the previous example that we know Abe is the parent of *someone*, but we are not sure who

yet. We represent this uncertainty in Structure T3 (Figure 7) below by adding a new fact node f_2 , which only states that Abe is a parent, without noting a corresponding child.




Figure 7. Triplet Structure T3

If we later learn that Homer is Abe’s child, we can extend f_2 to include this information as shown in Structure T4 (Figure 8) below.




Figure 8. Triplet Structure T4

4.2 Initializing Workspaces

The SIFTER workspace initially contains only symbols representing input objects (such as a in abc) and information about their relative position. For example, when comparing the strings ab and ef , the SIFTER workspace is initialized as shown in Structure T5 (Figure 9), where x_1 represents a , x_2 represents b , y_1 represents e , and y_2 represents f .

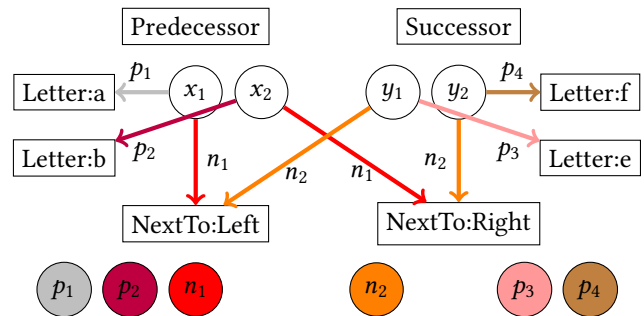


Figure 9. Triplet Structure T5

Notably, we also have nodes like Letter: a representing the Platonic concept of a particular letter; the fact that x_1 is mapped to Letter: a corresponds to asserting the unary

IsLetterA(x_1). We have included nodes for some predicates (specifically Predecessor and Successor representing the slots of the binary LetterSuccessor predicate) that have no incoming or outgoing edges. This indicates that, while SIFTER knows about the concept of predecessor and successor, it has not yet explicitly recognized any letter-successor pairs in the structure. In the next section, we will describe how such facts may be inferred from this initial encoding of the problem via the use of update rules.

4.3 Modifying the Workspace with Update Rules

SIFTER proceeds to modify the workspace in two ways: (1) refining the representation of its inputs, e.g., to infer the facts LetterSuccessor(x_1, x_2) and LetterSuccessor(y_1, y_2) in the above example, and (2) building an analogy between its inputs by comparing such inferred facts. Both types of modifications are implemented using the same framework of update rules. This section introduces our language for expressing update rules using a simple example rule of the first kind, deferring discussion of the second type of modification to Section 4.4. Note that the language described here for expressing update rules works to define update rules for any triplet structure. However, we focus our examples on their use for expressing inference rules for the SIFTER workspace.

Note that, in this section, we will discuss a method of “hard-coding” certain rules to express things like letter-successorship, both because this is how our current implementation operates, as well as to introduce the notion of update rules. Section 4.6 describes more general mechanisms for making such changes without explicitly enumerating all such rules ahead of time.

Recall the initial state of the SIFTER workspace for the example of **ab** and **ef**, shown above as Structure T5 (Figure 9). Consider now the problem of defining a rule that modifies the workspace by identifying when some letter instance is an alphabetical successor of another. For example, we may wish to create a rule that marks instances of the letter ‘a’ and the letter ‘b’ as LetterSuccessor pairs. In first-order logic, we might write the desired rule as $\text{IsLetterA}(v_1) \wedge \text{IsLetterB}(v_2) \implies \text{LetterSuccessor}(v_1, v_2)$. We have developed a visual domain specific language (DSL) for expressing such rules operating on the SIFTER workspace. Our full DSL is capable of expressing rules containing alternating quantifiers and other pattern-matching features. We will describe here only a simplified subset of the language that suffices for the uses in this paper.

Rules in this DSL look like triplet structures themselves (and in fact can be stored as such), although they are annotated with extra information about which nodes represent variables to search for and how the structure should be modified if such variables are found. For example, Rule R1 (Figure 10) below shows a rule which notates **a, b** letter-successor pairs.




Figure 10. Triplet Structure Rule R1

In such rule diagrams, one first looks at the parts not shaded green. In these parts, dashed nodes are variables that should be looked for in the structure, while solid nodes are constants assumed already to exist in the structure. When this pattern is found in the structure, this is called a *rule match* and the green nodes and facts can be added. In this case, the rule expresses that whenever two nodes v_1 and v_2 are found such that v_1 is the letter ‘a’ and v_2 the letter ‘b’, then we can add a new fact node and corresponding triplet facts which express that v_2 is an alphabetical successor of v_1 .

For example, we may apply Rule R1 (Figure 10) to Structure T5 (Figure 9) by taking the rule assignment with $v_1 = x_1$, $v_2 = x_2$, $v_{f1} = p_1$, and $v_{f2} = p_2$. This produces a new fact node n_{f3} which asserts that x_1 is a predecessor of the successor x_2 . Letting s_1 be the generated node corresponding to n_{f3} , Rule R1 (Figure 10) transforms Structure T5 (Figure 9) into Structure T6 (Figure 11) below.

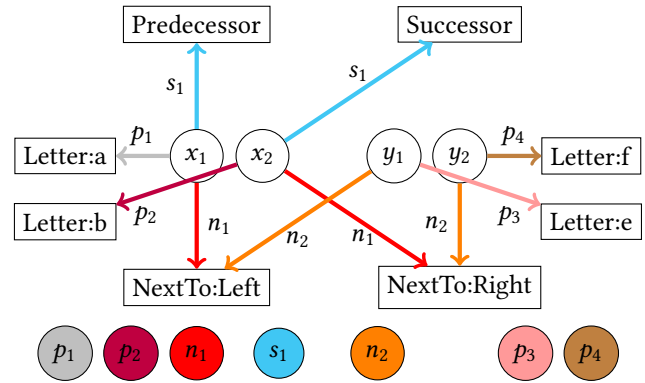


Figure 11. Triplet Structure T6

Similarly, Rule R2 (Figure 12) below identifies **e, f** pairs as successor pairs.




Figure 12. Triplet Structure Rule R2

Applying Rule R2 (Figure 12) to Structure T6 (Figure 11) marks y_1 and y_2 as predecessor and successor respectively, producing Structure T7 (Figure 13) below.

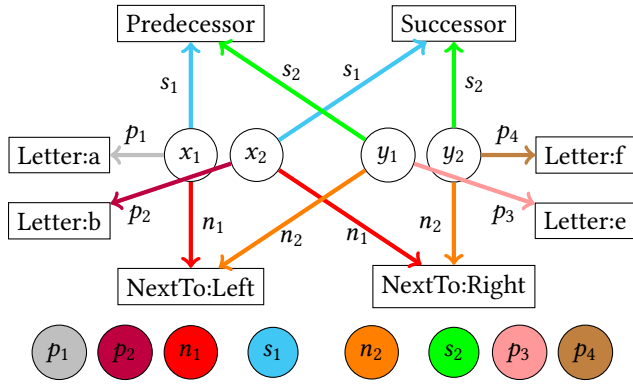


Figure 13. Triplet Structure T7

4.4 Representing Analogies in Triplet Structures

This section discusses how SIFTER represents and makes analogies in its triplet-structure workspace. Analogies are represented as *abstractions*, similar to the abstract letter strings in Section 2. Intuitively, instructed to form an analogy between *abc* and *efg*, SIFTER forms a shared abstract representation of them both, roughly of the form $(?1)(?2)(?3)$ with additional information such as $(?2)$ is a successor of $(?1)$. It then adds facts stating that, for example, both the original *a* and *e* are *instances* of this more abstract $(?1)$ object. Two input objects *correspond* if they are instances of the same abstract node.

Consider the example from Section 4.3, where we are forming an analogy between two letter strings *ab* and *ef*. Suppose the current SIFTER workspace is represented by Structure T8 (Figure 14) below, which is identical to Structure T7 (Figure 13) except with a few of the nodes/relations removed for ease of exposition.

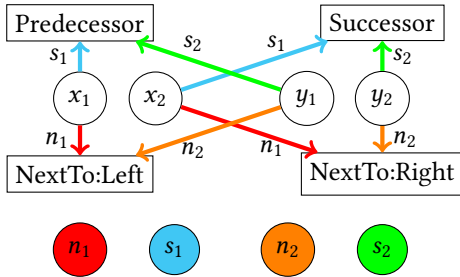


Figure 14. Triplet Structure T8

An *analogy* between the two strings *ab* and *ef* might determine that x_1 corresponds to y_1 and x_2 to y_2 , because they both form instances of a more abstract type of “two

successive letters next to each other.” This is represented in the SIFTER workspace as Structure T9 (Figure 15) below. For clarity, we have only shown the newly-added facts (i.e. those that make up the abstraction), although the facts from Structure T8 (Figure 14) would still be present.

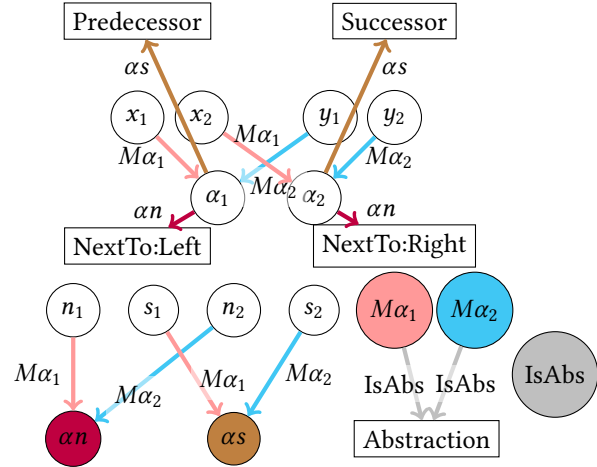


Figure 15. Triplet Structure T9

In Structure T9 (Figure 15), we have added new nodes α_1 and α_2 to represent the abstract type of which x_1, y_1 and x_2, y_2 respectively are instances of. We have also abstracted the fact nodes that correspond to each other into nodes α_n and α_s . These abstract fact nodes each express the same fact about the abstract α_1 and α_2 as the original, or *concrete*, fact nodes expressed about, e.g., x_1 and x_2 . Finally, we have added fact nodes Ma_1 and Ma_2 that map the concrete nodes in each instance to their abstract counterparts. For bookkeeping reasons in the structure, we label each of these as *Abstractions* so we can keep track of which nodes in the workspace are abstract vs. provided in the input.

From Structure T9 (Figure 15) above, we can extract the analogy that x_1 corresponds to y_1 because both are instances of the abstract α_1 node, and similarly for x_2, y_2 , and α_2 .

4.5 Rules for Making Analogies in a Triplet Structure

We now turn our attention to designing rules for forming such analogies. All such rules will be of the form discussed in Section 4.3. Each rule application makes a small change to the structure; for example, abstracting two concrete nodes together, or lifting a single concrete fact to the abstraction. These rules create a search space that can be explored using heuristics.

We have found that all of the rules necessary for abstraction-forming can be formed as variations on the following *Begin Analogy rule*, Rule R3 (Figure 16), which starts a new abstraction.




Figure 16. Triplet Structure Rule R3

Here, the variable nodes A and B represent the concrete nodes that should correspond to each other in the analogy, like x_1 and y_1 in the previous example. The variable node C is the field that they both share. For example, in the previous example C might be `NextTo:Left`, because both x_1 and y_1 are mapped to `NextTo:Left` by fact nodes n_1 and n_2 , which in turn map to M_A and M_B , respectively, in the rule above.

Applying Rule R3 (Figure 16) to Structure T9 (Figure 15) may produce the start of an analogy shown below in Structure T10 (Figure 17). For clarity, we have left out the facts between concrete nodes.

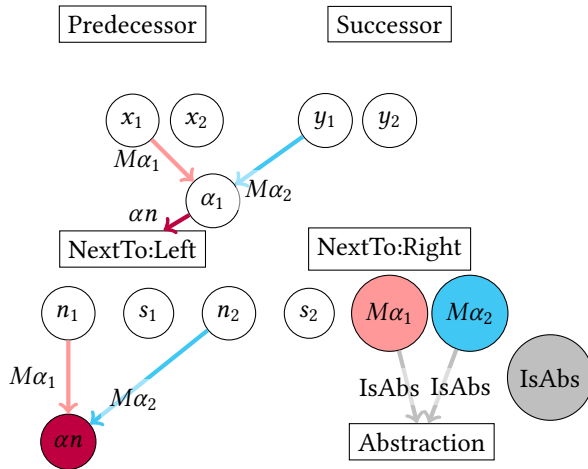


Figure 17. Triplet Structure T10

By shading different subsets of the nodes green, we can modify Rule R3 (Figure 16) into a variety of rules for *extending analogies*. For example, Rule R4 (Figure 18) below “follows” a fact node from an existing analogy to map two new concrete nodes to each other.




Figure 18. Triplet Structure Rule R4

Applying Rule R4 (Figure 18) to our running structure with $A = x_2$, $B = y_2$, and $C = \text{NextTo:Right}$ would extend the analogy to include x_2 and y_2 by “following” the `NextTo:Right` relation, producing Structure T11 (Figure 19) below.

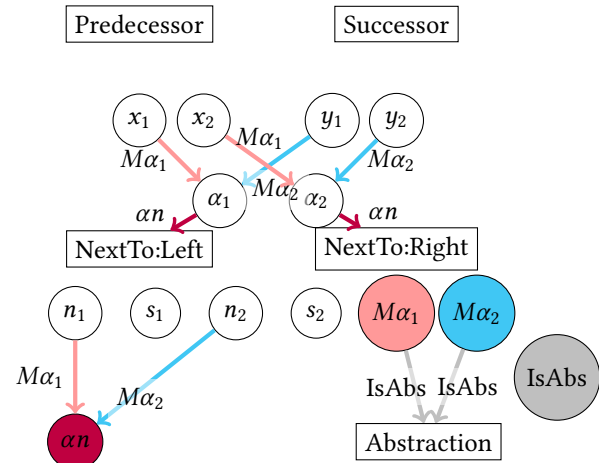


Figure 19. Triplet Structure T11

Similarly, by shading just the αM_{AB} node, we get Rule R5 (Figure 20) that adds a new fact node to the abstraction.




Figure 20. Triplet Structure Rule R5

Applying Rule R5 (Figure 20) to the previous abstraction with $A = x_1$, $B = y_1$, $\alpha_{AB} = \alpha_1$, $C = \text{Predecessor}$, $M_A = s_1$, and $M_B = s_2$ allows us to associate s_1 and s_2 with each other, producing Structure T12 (Figure 21) as shown below.

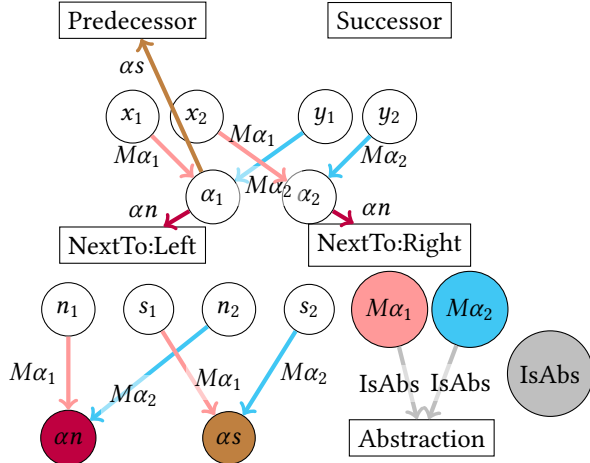


Figure 21. Triplet Structure T12

Finally, by shading just the blue fact edge from αAB to C , we get Rule R6 (Figure 22) that lifts a single fact into the abstraction.




Figure 22. Triplet Structure Rule R6

Applying this rule with $\alpha AB = \alpha_2$, $C = Successor$, and $\alpha M_{AB} = \alpha s$ completes the abstraction, giving the final abstraction we saw earlier in Structure T9, reproduced below.

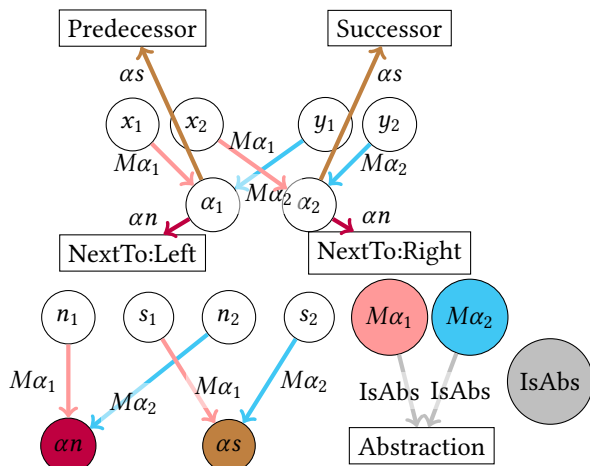


Figure 23. Triplet Structure T9

We encourage the motivated reader to consider interpretations of different shadings. Of particular interest is if a node like B is shaded green, which corresponds to *completing* an analogy: constructing a concrete node that plays a particular role in an existing analogy. This shading of the rule prototype is how SIFTER generated the code in the bottom right of Figure 2 and Figure 4.

4.6 Higher-Order Analogies and Slips

We described in Section 4.3 that rules could be used to infer new facts, such as $LetterSuccessor(x_1, x_2)$. Such facts could later be used in analogies, e.g., to compare ab and ef as “two-letter strings where the letters satisfy the $LetterSuccessor$ relation”. To do this, we had to first explicitly add the facts $LetterSuccessor(x_1, x_2)$ and $LetterSuccessor(y_1, y_2)$. In general, this approach requires us to explicitly enumerate rules for all such relations used in our analogies. This section considers a more general approach based on forming analogies between types in the structure.

At first glance, it is tempting to resolve the issue using a general transitivity rule such that, for example, if x_1 is an instance of T_1 , x_2 is an instance of T_2 , and there is some fact $R(T_1, T_2)$, then we can add $R(x_1, x_2)$ as well. For example, if $LetterSuccessor(Letter:a, Letter:b)$ and x_1 was an instance of $Letter:a$, x_2 an instance of $Letter:b$, then the rule would infer $LetterSuccessor(x_1, x_2)$ as desired.

However, facts about *types* may not be valid or well-defined when applied to *instances of those types*. For example, when forming an analogy involving both numerical value and color we may have two nodes be instances of opposite numbers, e.g., -1 and 1 , or opposite colors, e.g., black and white. If we were to directly use transitivity to say that the two nodes were simply “opposites,” we would lose important information because we would not know whether they were opposite *numbers* or *colors*. Similarly, consider forming an analogy between the pairs $abc \rightarrow cba$ and $efg \rightarrow gfe$. Fundamentally, what we want to express is that the letters in the first string satisfy *either* $LetterSuccessor$ *or* $LetterPredecessor$, and that those in the second string satisfy *the opposite*.

To express such scenarios naturally, we need a way to *include the types in the analogy*, i.e., make a *type slip*. In the mapping rules shown so far, we require the both concrete nodes A and B to be of the same type C . However, we can define new mapping rules, using the template of Rule R7 (Figure 24), that allow the type itself to be abstracted as part of the analogy. In the example discussed, we could have $LetterSuccessor$ be C_1 and $LetterPredecessor$ be C_2 .




Figure 24. Triplet Structure Rule R7

Although allowing for creative analogies, such an approach significantly increases the search space. To control this, one can make *compound analogies*, using type slips only for small sub-analogies with smaller search spaces. Then those analogies are used to define types with which to build up larger ones. For example, we could use a type slip to learn the abstract type of “pairs of nodes which are instances of types that have *Successor* relation,” i.e., effectively re-learn the LetterSuccessor relation on its own. Analogy-making rules could then be used to note that both x_1, x_2 and y_1, y_2 form an instance of this abstract type, and then the abstract type can be used exactly like Predecessor and Successor in future analogies.

4.7 Prototype Implementation and Optimizations

We have implemented a proof-of-concept version of SIFTER in Python, with hotspots written in C++. We have run the demonstrations discussed in Section 3 on our prototype to verify that such analogies can be found, represented, and completed by SIFTER.

The project is divided into distinct components, including: TSLib, a library for declaring triplet structures and rules operating on them; Abstracter, a collection of such rules which can be applied to build up analogies; and TSRuntime, an interface for efficiently applying rules to a triplet-structure workspace and can optimize pattern matching, e.g., by only checking parts of the structure that have changed since the pattern was last checked against.

In its full generality, our update rule DSL is Turing-complete, and capable of expressing rules matching complex patterns. One useful feature of our system is its ability to express *consistency constraints* or desired invariants on the structure. For example, we may want to ensure that any symbol can *either* be an instance of Letter:a or Letter:b, but never both at the same time. SIFTER supports *consistency rules*, which are just like normal update rules except (1) they are checked for matches *every time* the structure is modified, and (2) they force SIFTER to backtrack when they match, i.e. undo the last modification.

We optimized our implementation for efficient search through possible applications of update rules. The most impactful optimization has been the use of *differential matching*, where the design of our update rules allows us to restrict our search to only those assignments that make use of the facts added since we last checked for assignments, significantly reducing the amount of redundant time spent searching. More optimizations are possible in the future, both in speeding up such a tree-search approach to analogy-making as well as by investigating other architectures for making analogies within SIFTER (see Section 5.3).

5 Efficient Software Analogies with SIFTER

This section discusses practical considerations with the application of SIFTER to make analogies of the form shown in Section 3. We focus on three particular factors, (i) the representation of source code as a triplet structure, (ii) the use of other sources of reasoning and information, and (iii) heuristics for finding analogies. In each section, we begin with a description of our current solution, then discuss our vision of what a future implementation may be able to accomplish.

5.1 Representing Source Code as Triplet Structures

Currently, given a source file we perform a lightweight lexical analysis before encoding it in the structure. If more information about the meaning of some of the resulting lexemes is known, we can include that as well. For example, given a source file consisting only of the statement `name = user.name`, we might encode it as shown in Structure T13 (Figure 25).




Figure 25. Triplet Structure T13

For brevity, 11 fact nodes are not explicitly shown. However, their existence is implied by the colors and labels on the edges. We first create a node in the structure representing the file. For each lexeme in the file we add a corresponding node. Each lexeme node is marked as a member of the corresponding file, and their relative positions are specified using NextTo:Left and NextTo:Right. There are four nodes representing Platonic strings (or ‘tokens’), which play the same role as the Letter:a nodes in Section 4.3 or a unary IsLetterA(x) predicate. In this example, we assume additional

information about the language, namely that `user.name` represents an access of the name field of the user object.

Notably, such a lexical analysis can usually be developed quite quickly even for new programming languages. At its simplest, it can be implemented as just splitting the source file based on whitespace and special characters such as `*`. This allows for SIFTER to be applicable to nascent DSLs and other languages where a full compiler and AST generator has not yet been developed, or to work with syntactically-invalid programs. As the language tooling grows in maturity or parts of the programs become syntactically valid, more detailed information can be produced from the lexing pass and included in the structure, such as the object-field access notated in the above example.

Future Work: Full ASTs. While the progressive-lexing style of encoding strikes a nice balance between flexibility and richness, if a full AST is readily available for the code in question, then this can be used to produce a richer encoding of the structure. Structure T14 (Figure 26) shows how an AST for the `name=user.name` example might be encoded as a triplet structure.

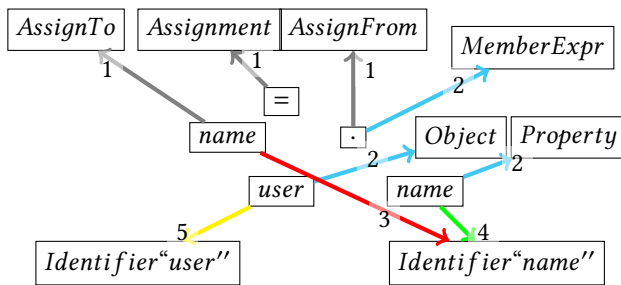


Figure 26. Triplet Structure T14

Future Work: Multiple Granularities. One particularly exciting area of future work for our encoding is to allow the granularity of the encoding to change dynamically, as the analogy process is proceeding. In this model, the workspace would initially begin with only a listing of the names of files and folders in the root directory of the project(s). As analogy-making proceeds, the contents of files may be added to the structure either randomly or according to activity from the analogy-making process itself. For example, if two files are named the same, they may be mapped together indicating their importance to the analogy and hinting to the system that the file contents might be important as well. If more semantic information about the programming languages is known, this type of multi-granularity encoding can be used at that level as well. For example, a file could be loaded first as just a list of functions contained in it. If two functions seem similar based on their signatures, then we expand them and include the full associated code in the workspace.

Our current implementation supports such real-time modifications to the structure, like adding a new file’s contents

halfway during an analogy run. However, the heuristics for knowing when to do such are not yet developed, so we just add the full contents of the files to it at the start (limiting us to small-ish projects).

5.2 Use of Other Engines

Many other reasoning engines for both natural language and software source-code exist, including logic-based techniques (such as Cyc [42]), abstract interpretation [13], and statistical techniques (such as deep learning [25]). We designed SIFTER with the specific goal of easily integrating the knowledge stored in such tools with the analogy-making process. In particular, while we focused in Section 4.3 on the usage of update rules to make inferences about relations such as LetterSuccessor, there is no requirement that modifications to the structure come from such a update rule. Instead, other reasoning engines can provide their own insights into the problem at hand, which can then be translated into triplet facts and added to the structure. Such added facts are used in analogies just like any other facts.

For example, a statistical model might be used to identify when words used are synonyms, which can be encoded into the workspace, e.g., using a Synonym predicate or by stating that both words are instances of some common semantic notion. Similarly, a logical inference engine like Cyc [42] might be used to understand comments in, e.g., function docstrings to identify the semantic meaning behind different functions, or infer the effects of applying multiple different functions in sequence. This information can be encoded back into the triplet structure and used to make analogies.

Future Work: A Unified Workspace. We envision that the triplet structure workspace of SIFTER can serve as a shared workspace among a host of reasoning engines. The core analogy-making rules, which operate structurally, and are in some sense oblivious to the semantics of the relations, can act as glue that can help synthesize inferences across distinct reasoning engines. Reasoning engines can operate independently, inferring new facts and adding them to the structure. Furthermore, they may be able to use the facts inferred by other engines or even predicted to be true via analogy to further their own reasoning, resulting in a virtuous cycle of cooperation between different engines.

Future Work: Probabilistic Triplet Structures. To better interface with statistical models, it may be desirable to associate with each triplet fact a corresponding real-valued probability representing SIFTER’s confidence in that inferred fact. As future inferences are made, their confidence values can be computed as a function of the confidence in that specific step as well as the confidence in the facts it relies on to make that inference. However, as argued by the Cyc authors [42], one should be careful in treating such numbers as a measure of truthfulness of a claim, and instead only as a representation of one’s epistemological uncertainty about immediate observations of the environment.

5.3 Heuristics for and Identification of Strong Software Analogies

To guide the search process, we need a notion of the strength of an analogy. One initial approach to this is to define a stronger analogy as one with more shared facts. This idea can be improved by weighting different types of relations with an importance. For a somewhat extreme example, we may prefer to map two functions together that share a relation of “called by analogous methods” rather than simply “share the same first letter.” This approach is exemplified by the notion of a *Slipnet* in the Copycat architecture, which assign numerical, *a priori* importances to each type of relation.

For scenarios where the goal of the analogy is to generate some *completion* (as in Section 3.2 and 3.3), we have found that an even stronger heuristic measure of analogy depth is to check if a full completion to the original problem can be made from it. In our experience, when we had bugs with our search process, we found that the completions produced in Figure 2 and 4 would either (i) contain very few nodes, or (ii) contain many nodes that the system thought would be there, but was not able to actually infer (consistently) what tokens they should be. This can be used to indicate when an analogy may fit the two examples very well, but does not generalize to the prompt.

Future Work: Cognitive Models. In the long term, we are excited about the possibility of having multiple workers that operate concurrently, either working on separate attempts at making such analogies, finding inconsistencies in analogies made by others, or working together to produce a strong analogy. Such a system may take inspiration from psychology-inspired architectures and theories, like LIDA [20] and Global Workspace Theory [3]. Such a system could be supplemented via the use of something akin to abduction, where rules can request the help of other rules. For instance, an abstraction rule expressing that it would be able to abstract two things together if only one of them were also a letter successor would encourage rules to search for letter successor facts relating to that node. This abductive inference is similar to the execution of the Slipnet and Coderack in the original Copycat algorithm.

Future Work: Meta-Reasoning. SIFTER may be extended to support meta-reasoning similar to that of Metacat [50], where the system learns to recognize common ‘snags’ that it can store and refer back to when it encounters a new problem that is challenging in the same way. It may recognize such “challenging in the same way” instances via a sort of meta-SIFTER, forming analogies between solver states. In such a way, it would be able to effectively introspect on its own solving process via a (copy of) the solver itself. Such a system may be aided by the fact that our SIFTER update rules can be expressed *within triplet structures themselves*.

6 Related Work

Analogy making. Analogies, and the more general class of *metaphors* [40, 41], have been studied extensively in cognitive science. The primary inspiration for this paper was the Copycat algorithm of Melanie Mitchell and Douglas Hofstadter [33, 34], which we generalized into the SIFTER analogy-making algorithm presented here. Mitchell’s original source code [54] is written in a now-defunct dialect of Common Lisp for which we could not obtain an interpreter. Thankfully, there are a number of more-modern ports [8, 30] which we were able to reference and run. We originally intended to adapt these implementations for use on program source code, but quickly found that the Copycat algorithm is highly specific for the letter string domain. Adding support even for upper-case letters, for example, turned out to be a significant project, touching almost every file in the implementation. This is because Copycat implements relations like LetterPredecessor by special-case checks and data structures throughout the code, not as the sort of arbitrary relations that are more familiar in systems based on first-order logic. These issues motivated the construction of our SIFTER analogy-making algorithm.

Beyond Copycat, there are a number of other analogy-making algorithms explored in the cognitive science and philosophy literature, such as SME [18], ACME [35], Winston’s Analogy-Maker [89], and the Geometric Analogy-Maker [17]. Below we will discuss SME, but ACME and Winston’s algorithm have a similar operation. The Geometric Analogy-Maker is interesting and unique, incorporating some amount of grouping and representational manipulation similar to that of Copycat, although it is specific to geometric analogy problems.

The Structure Mapping Engine (SME) [18] is an analogy-making algorithm that gained notoriety for, among other things, “discovering” the Rutherford model of the atom by analogy to a solar system [22]. Its basic operation is similar to that of a sub-graph isomorphism algorithm, in that it looks for an injective mapping between the objects and predicates in one structure into those of another that retains the facts.

The SME algorithm has been criticized by some [33] for its reliance on hand-written input representations, which circumvents the hardest part of analogy-making. For example, when forming an analogy between the strings `aaabbc` and `abbccc`, it might be natural for a human to group the three `aaas` and associate them either with the single `a` or the similar group-of-three `ccc` in the second. However, SME has no conception of modifying the structure, or its representation of the structure, in such a way. The user would have to explicitly group letters before providing it to SME, at which point SME would just look for relations between groups.

The SIFTER workspace represents a its workspace using structures and relations. Similarly, TVLA [43, 74, 75] represents possible program states using a variant of first-order

logic. In TVLA, such structures are abstracted into three-valued structures, somewhat similar to how we abstract concrete instances in the workspace to form analogies.

Comparative program understanding. The problem of comparative program understanding is related to the problem of code detection [21, 37, 38, 49, 73] with recent approaches using deep learning [88]. CP-Miner identifies bugs related to copy-pasted code [44]. The func2vec technique computes function embeddings to learn function synonyms, which are functions that play a similar role in the source code [14]. Such function synonyms are used to identify error-handling bugs in the Linux kernel. The code2vec technique computes an embedding of source code using the AST [1], which can be used to infer names of functions. These techniques represent a research thread that uses “Big Code” [86]. In contrast, SIFTER uses a relatively limited amount of source code to make analogies; however, it can make use of models of source code learned via deep learning and other techniques.

Program transformation-learning. Recent approaches, like GetAFix [4] and Jigsaw [12], have explored the use of *antiunification* of program ASTs to learn program fixes. Related approaches have also applied antiunification to proof transfer [45] and letter analogies [39]. Such antiunification can be seen as a restricted special case of the analogy-as-abstraction process used by SIFTER (see Section 4.4). Existing antiunification-based approaches, however, are generally limited to tree structures (such as ASTs), and cannot make use of additional semantic information. For example, GetAFix would not have been able to accurately learn the transformation in Section 3.2, which relied on semantic properties of the code, or the one in Section 3.3, which relied on referencing the library documentation.

The program transformation-learning problem is related to Programming by Example [26–28, 52, 68, 78–81, 91]. Such techniques typically restrict the program transformation to a limited domain-specific language. Repenning et al. [70] describe how end-users of a programming-by-example system might use analogies to express the desired behavior by comparison to that of an existing program. Perrone et al. [63] propose that implementing code reuse via concrete analogies can be more natural than the use of standard object-oriented programming paradigms and help novice programmers avoid copy-and-pasting code. Recent approaches have explored using natural language and examples as input [69] and using deep learning techniques [85].

API Migration. Many approaches for API migration that use statistical and machine learning techniques have been proposed [57–61, 65].

Cognitive Science in SE. There are numerous works which have highlighted the promise of models from cognitive science in software engineering. Call by Meaning [76] describes a system in which program components (e.g., functions) can

be addressed by their *semantic meaning*, not just their syntactic name. A programmer may describe an existing function using a high-level, semantic description language and the Cyc [42] cognitive model will be used to infer which (composition of) function(s) best matches that description. The Semprola semiotic programming language [77] allows programmers to directly use *signs* instead of the now-dominant focus on textual code symbols. Both Call by Meaning and Semprola are ambitious and exciting projects, requiring a fundamental re-thinking of how and in what languages we write code. While we envision that SIFTER can benefit such approaches in the future, we are excited that, as described in Section 3, SIFTER can wield these cognitive models to more immediately benefit software engineers using the existing programming languages and environments of today.

7 Conclusion

In this paper, we discussed *analogy-making*, a fundamental human ability that involves identifying underlying similarities between two objects. We first described analogy-making through examples in a restricted letter-string domain. We then showed how analogy-making can be used to address a variety of software engineering tasks, namely comparative program understanding, program optimization, and API migration. Finally, we described SIFTER, our proposed algorithm for analogy-making, which is suitable for making analogies about programs. SIFTER relies on a novel triplet-structure representation for its workspace, allowing it to form analogies over arbitrary inputs, such as source code, program analyzer outputs, and documentation. By reducing a variety of problems to analogy-making, improvements to the core analogy-making primitive can pay large dividends across a variety of applications. Software engineering represents a difficult challenge for analogy-making, as it involves a unique balance of unambiguous syntax and semantics of the program, as well as ambiguous information about programmer intent. We hope that this paper serves as a first step towards cementing analogy-making as a core primitive in the software engineering toolbox.

Acknowledgments

We thank the anonymous reviewers as well as Cindy Rubio-González for comments that significantly improved this paper.

References

- [1] Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. 2019. code2vec: learning distributed representations of code. *Proc. ACM Program. Lang.* 3, POPL (2019), 40:1–40:29. <https://doi.org/10.1145/3290353>
- [2] Dennis Andriess, Xi Chen, Victor van der Veen, Asia Slowinska, and Herbert Bos. 2016. An In-Depth Analysis of Disassembly on Full-Scale x86/x64 Binaries. In *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10–12, 2016*.

- Thorsten Holz and Stefan Savage (Eds.). USENIX Association, 583–600. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/andriese>
- [3] Bernard J Baars. 1993. *A cognitive theory of consciousness*. Cambridge University Press.
- [4] Johannes Bader, Andrew Scott, Michael Pradel, and Satish Chandra. 2019. Getafix: learning to fix bugs automatically. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 159:1–159:27. <https://doi.org/10.1145/3360585>
- [5] Paul Bartha. 2019. Analogy and Analogical Reasoning. In *The Stanford Encyclopedia of Philosophy* (spring 2019 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.
- [6] bash 2020. fish - the friendly interactive shell. <https://github.com/fish-shell/fish-shell>. Accessed May, 2020.
- [7] bash 2020. GNU Bash. <https://ftp.gnu.org/gnu/bash/>. Accessed May, 2020.
- [8] J Alan Brogan. 2020. Modern port of Melanie Mitchell’s and Douglas Hofstadter’s Copycat. <https://github.com/fargonauts/copycat>. Accessed May, 2020.
- [9] Fraser Brown, Andres Nötzli, and Dawson R. Engler. 2016. How to Build Static Checking Systems Using Orders of Magnitude Less Code. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '16, Atlanta, GA, USA, April 2-6, 2016*, Tom Conte and Yuanyuan Zhou (Eds.). ACM, 143–157. <https://doi.org/10.1145/2872362.2872364>
- [10] Yam San Chee. 1993. Applying Gentner’s Theory of Analogy to the Teaching of Computer Programming. *Int. J. Man Mach. Stud.* 38, 3 (1993), 347–368. <https://doi.org/10.1006/imms.1993.1016>
- [11] John Clement. 1993. Using bridging analogies and anchoring intuitions to deal with students’ preconceptions in physics. *Journal of research in science teaching* 30, 10 (1993), 1241–1257.
- [12] Rylan Cottrell, Robert J. Walker, and Jörg Denzinger. 2008. Semi-automating small-scale source code reuse via structural correspondence. In *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering, 2008, Atlanta, Georgia, USA, November 9-14, 2008*, Mary Jean Harrold and Gail C. Murphy (Eds.). ACM, 214–225. <https://doi.org/10.1145/1453101.1453130>
- [13] Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages, Los Angeles, California, USA, January 1977*, Robert M. Graham, Michael A. Harrison, and Ravi Sethi (Eds.). ACM, 238–252. <https://doi.org/10.1145/512950.512973>
- [14] Daniel DeFreez, Aditya V. Thakur, and Cindy Rubio-González. 2018. Path-based function embedding and its application to error-handling specification mining. In *Proceedings of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2018, Lake Buena Vista, FL, USA, November 04-09, 2018*, Gary T. Leavens, Alessandro Garcia, and Corina S. Pasareanu (Eds.). ACM, 423–433. <https://doi.org/10.1145/3236024.3236059>
- [15] Sarah A Douglas and Thomas P Moran. 1983. Learning text editor semantics by analogy. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 207–211.
- [16] Enda Dunican. 2002. Making the Analogy: Alternative Delivery Techniques for First Year Programming Courses. In *Proceedings of the 14th Annual Workshop of the Psychology of Programming Interest Group, PPIG 2002, London, UK, June 18-21, 2002*. Psychology of Programming Interest Group, 8. <http://ppig.org/library/paper/making-analogy-alternative-delivery-techniques-first-year-programming-courses>
- [17] Thomas G Evans. 1964. A heuristic program to solve geometric-analogy problems. In *Proceedings of the April 21-23, 1964, spring joint computer conference*. 327–338.
- [18] Brian Falkenhainer, Kenneth D. Forbus, and Dedre Gentner. 1989. The Structure-Mapping Engine: Algorithm and Examples. *Artif. Intell.* 41, 1 (1989), 1–63. [https://doi.org/10.1016/0004-3702\(89\)90077-5](https://doi.org/10.1016/0004-3702(89)90077-5)
- [19] Michal Forisek and Monika Steinová. 2012. Metaphors and analogies for teaching algorithms. In *Proceedings of the 43rd ACM technical symposium on Computer science education, SIGCSE 2012, Raleigh, NC, USA, February 29 - March 3, 2012*, Laurie A. Smith King, David R. Musicant, Tracy Camp, and Paul T. Tymann (Eds.). ACM, 15–20. <https://doi.org/10.1145/2157136.2157147>
- [20] Stan Franklin, Tamas Madl, Sidney K. D’Mello, and Javier Snaider. 2014. LIDA: A Systems-level Architecture for Cognition, Emotion, and Learning. *IEEE Trans. Auton. Ment. Dev.* 6, 1 (2014), 19–41. <https://doi.org/10.1109/TAMD.2013.2277589>
- [21] Mark Gabel, Lingxiao Jiang, and Zhendong Su. 2008. Scalable detection of semantic clones. In *30th International Conference on Software Engineering (ICSE 2008), Leipzig, Germany, May 10-18, 2008*, Wilhelm Schäfer, Matthew B. Dwyer, and Volker Gruhn (Eds.). ACM, 321–330. <https://doi.org/10.1145/1368088.1368132>
- [22] Dedre Gentner. 1983. Structure-Mapping: A Theoretical Framework for Analogy. *Cogn. Sci.* 7, 2 (1983), 155–170. https://doi.org/10.1207/s15516709cog0702_3
- [23] Dedre Gentner, Jeffrey Loewenstein, and Leigh Thompson. 2003. Learning and transfer: A general role for analogical encoding. *Journal of Educational Psychology* 95, 2 (2003), 393.
- [24] Nasser Giacaman. 2012. Teaching by Example: Using Analogies and Live Coding Demonstrations to Teach Parallel Computing Concepts to Undergraduate Students. In *26th IEEE International Parallel and Distributed Processing Symposium Workshops & PhD Forum, IPDPS 2012, Shanghai, China, May 21-25, 2012*. IEEE Computer Society, 1295–1298. <https://doi.org/10.1109/IPDPSW.2012.158>
- [25] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [26] Sumit Gulwani. 2011. Automating string processing in spreadsheets using input-output examples. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, Thomas Ball and Mooly Sagiv (Eds.). ACM, 317–330. <https://doi.org/10.1145/1926385.1926423>
- [27] Sumit Gulwani. 2017. Programming by examples: applications, algorithms, and ambiguity resolution. In *Proceedings of the 19th International Symposium on Principles and Practice of Declarative Programming, Namur, Belgium, October 09 - 11, 2017*, Wim Vanhoof and Brigitte Pientka (Eds.). ACM, 2. <https://doi.org/10.1145/3131851.3131853>
- [28] Sumit Gulwani and Prateek Jain. 2017. Programming by Examples: PL Meets ML. In *Programming Languages and Systems - 15th Asian Symposium, APLAS 2017, Suzhou, China, November 27-29, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10695)*, Bor-Yuh Evan Chang (Ed.). Springer, 3–20. https://doi.org/10.1007/978-3-319-71237-6_1
- [29] Rahul Gupta, Aditya Kanade, and Shirish K. Shevade. 2019. Deep Reinforcement Learning for Syntactic Error Repair in Student Programs. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*. AAAI Press, 930–937. <https://doi.org/10.1609/aaai.v33i01.3301930>
- [30] Joseph Hager. 2017. A translation of Melanie Mitchell’s original Copycat project from Lisp to Python. <https://github.com/ajhager/copycat>. Accessed May, 2020.
- [31] Jeisson Hidalgo-Céspedes, Gabriela Marín Raventós, and Vladimir Lara-Villagrán. 2014. Playing with metaphors: a methodology to design video games for learning abstract programming concepts. In *Innovation and Technology in Computer Science Education Conference 2014, ITiCSE '14, Uppsala, Sweden, June 23-25, 2014*, Åsa Cajander, Mats Daniels, Tony Clear, and Arnold Pears (Eds.). ACM, 348. <https://doi.org/10.1145/2571362.2571368>

- [//doi.org/10.1145/2591708.2602661](https://doi.org/10.1145/2591708.2602661)
- [32] Douglas Hofstadter. 1995. A Review of Mental Leaps: Analogy in Creative Thought. *AI Mag.* 16, 3 (1995), 75–80. <https://doi.org/10.1609/aimag.v16i3.1154>
- [33] Douglas R Hofstadter. 1995. *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought*. Basic books.
- [34] Douglas R Hofstadter and Melanie Mitchell. 1994. The Copycat project: A model of mental fluidity and analogy-making. (1994).
- [35] Keith J. Holyoak and Paul Thagard. 1989. Analogical Mapping by Constraint Satisfaction. *Cogn. Sci.* 13, 3 (1989), 295–355. https://doi.org/10.1207/s15516709cog1303_1
- [36] Shalini Kaleeswaran, Anirudh Santhiar, Aditya Kanade, and Sumit Gulwani. 2016. Semi-supervised verified feedback generation. In *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2016, Seattle, WA, USA, November 13-18, 2016*, Thomas Zimmermann, Jane Cleland-Huang, and Zhendong Su (Eds.). ACM, 739–750. <https://doi.org/10.1145/2950290.2950363>
- [37] Toshihiro Kamiya, Shinji Kusumoto, and Katsuro Inoue. 2002. CCFinder: A Multilinguistic Token-Based Code Clone Detection System for Large Scale Source Code. *IEEE Trans. Software Eng.* 28, 7 (2002), 654–670. <https://doi.org/10.1109/TSE.2002.1019480>
- [38] Raghavan Komondoor and Susan Horwitz. 2001. Using Slicing to Identify Duplication in Source Code. In *Static Analysis, 8th International Symposium, SAS 2001, Paris, France, July 16-18, 2001, Proceedings (Lecture Notes in Computer Science, Vol. 2126)*, Patrick Cousot (Ed.). Springer, 40–56. https://doi.org/10.1007/3-540-47764-0_3
- [39] Ulf Krumnack, Angela Schwing, Helmar Gust, and Kai-Uwe Kühnberger. 2007. Restricted Higher-Order Anti-Unification for Analogy Making. In *AI 2007: Advances in Artificial Intelligence, 20th Australian Joint Conference on Artificial Intelligence, Gold Coast, Australia, December 2-6, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4830)*, Mehmet A. Orgun and John Thornton (Eds.). Springer, 273–282. https://doi.org/10.1007/978-3-540-76928-6_29
- [40] George Lakoff and Mark Johnson. 2008. *Metaphors we live by*. University of Chicago press.
- [41] George Lakoff and Rafael Núñez. 2000. *Where mathematics comes from*. Vol. 6. New York: Basic Books.
- [42] Douglas B. Lenat. 1995. CYC: A Large-Scale Investment in Knowledge Infrastructure. *Commun. ACM* 38, 11 (1995), 32–38. <https://doi.org/10.1145/219717.219745>
- [43] Tal Lev-Ami and Shmuel Sagiv. 2000. TVLA: A System for Implementing Static Analyses. In *Static Analysis, 7th International Symposium, SAS 2000, Santa Barbara, CA, USA, June 29 - July 1, 2000, Proceedings (Lecture Notes in Computer Science, Vol. 1824)*, Jens Palsberg (Ed.). Springer, 280–301. https://doi.org/10.1007/978-3-540-45099-3_15
- [44] Zhenmin Li, Shan Lu, Suvda Myagmar, and Yuanyuan Zhou. 2006. CP-Miner: Finding Copy-Paste and Related Bugs in Large-Scale Software Code. *IEEE Trans. Software Eng.* 32, 3 (2006), 176–192. <https://doi.org/10.1109/TSE.2006.28>
- [45] Jianguo Lu, John Mylopoulos, Masateru Harao, and Masami Hagiya. 2000. Higher order generalization and its application in program verification. *Ann. Math. Artif. Intell.* 28, 1-4 (2000), 107–126. <https://doi.org/10.1023/A:1018952121991>
- [46] Sifei Luan, Di Yang, Celeste Barnaby, Koushik Sen, and Satish Chandra. 2019. Aroma: code recommendation via structural code search. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 152:1–152:28. <https://doi.org/10.1145/3360578>
- [47] Katherine N. Macfarlane and Barbee T. Mynatt. 1988. A study of an advance organizer as a technique for teaching computer programming concepts. In *Proceedings of the 19th SIGCSE Technical Symposium on Computer Science Education, 1988, Atlanta, Georgia, USA, February 25-26, 1988*, Herbert L. Dershem (Ed.). ACM, 240–243. <https://doi.org/10.1145/52964.53024>
- [48] Solomon Maina, Anders Miltner, Kathleen Fisher, Benjamin C. Pierce, David Walker, and Steve Zdancewic. 2018. Synthesizing quotient lenses. *Proc. ACM Program. Lang.* 2, ICFP (2018), 80:1–80:29. <https://doi.org/10.1145/3236775>
- [49] Andrian Marcus and Jonathan I. Maletic. 2001. Identification of High-Level Concept Clones in Source Code. In *16th IEEE International Conference on Automated Software Engineering (ASE 2001), 26-29 November 2001, Coronado Island, San Diego, CA, USA*. IEEE Computer Society, 107–114. <https://doi.org/10.1109/ASE.2001.989796>
- [50] James B Marshall. 2000. Metacat: A self-watching cognitive architecture for analogy-making and high-level perception. (2000).
- [51] Anders Miltner, Kathleen Fisher, Benjamin C. Pierce, David Walker, and Steve Zdancewic. 2018. Synthesizing bijective lenses. *Proc. ACM Program. Lang.* 2, POPL (2018), 1:1–1:30. <https://doi.org/10.1145/3158089>
- [52] Anders Miltner, Sumit Gulwani, Vu Le, Alan Leung, Arjun Radhakrishna, Gustavo Soares, Ashish Tiwari, and Abhishek Udupa. 2019. On the fly synthesis of edit suggestions. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 143:1–143:29. <https://doi.org/10.1145/3360569>
- [53] Marvin Lee Minsky. [n.d.]. Logical vs. Analogical; or Symbolic vs. Connectionist; or Neat vs. Scruffy. ([n. d.]).
- [54] Melanie Mitchell. 2020. Original Copycat Source Code. <https://melaniemitchell.me/ExplorationsContent/how-to-get-copycat.html>. Accessed May, 2020.
- [55] Edoardo Di Napoli, Diego Fabregat-Traver, Gregorio Quintana-Ortí, and Paolo Bientinesi. 2014. Towards an efficient use of the BLAS library for multilinear tensor contractions. *Appl. Math. Comput.* 235 (2014), 454–468. <https://doi.org/10.1016/j.amc.2014.02.051>
- [56] Robert F. Nelson. 1975. Use of Analogy as a Learning-Teaching Tool. *Journal of Geography* 74, 2 (1975), 83–86. <https://doi.org/10.1080/00221347508979866>
- [57] Anh Tuan Nguyen, Hoan Anh Nguyen, Tung Thanh Nguyen, and Tien N. Nguyen. 2014. Statistical learning approach for mining API usage mappings for code migration. In *ACM/IEEE International Conference on Automated Software Engineering, ASE '14, Vasteras, Sweden - September 15 - 19, 2014*, Ivica Crnkovic, Marsha Chechik, and Paul Grünbacher (Eds.). ACM, 457–468. <https://doi.org/10.1145/2642937.2643010>
- [58] Anh Tuan Nguyen, Hoan Anh Nguyen, Tung Thanh Nguyen, and Tien N. Nguyen. 2014. Statistical learning of API mappings for language migration. In *36th International Conference on Software Engineering, ICSE '14, Companion Proceedings, Hyderabad, India, May 31 - June 07, 2014*, Pankaj Jalote, Lionel C. Briand, and André van der Hoek (Eds.). ACM, 618–619. <https://doi.org/10.1145/2591062.2591144>
- [59] Trong Duc Nguyen, Anh Tuan Nguyen, and Tien N. Nguyen. 2016. Mapping API elements for code migration with vector representations. In *Proceedings of the 38th International Conference on Software Engineering, ICSE 2016, Austin, TX, USA, May 14-22, 2016 - Companion Volume*, Laura K. Dillon, Willem Visser, and Laurie A. Williams (Eds.). ACM, 756–758. <https://doi.org/10.1145/2889160.2892661>
- [60] Trong Duc Nguyen, Anh Tuan Nguyen, Hung Dang Phan, and Tien N. Nguyen. 2017. Exploring API embedding for API usages and applications. In *Proceedings of the 39th International Conference on Software Engineering, ICSE 2017, Buenos Aires, Argentina, May 20-28, 2017*, Sebastián Uchitel, Alessandro Orso, and Martin P. Robillard (Eds.). IEEE / ACM, 438–449. <https://doi.org/10.1109/ICSE.2017.47>
- [61] Tien N. Nguyen. 2016. Code migration with statistical machine translation. In *Proceedings of the 5th International Workshop on Software Mining, SoftwareMining@ASE 2016, Singapore, Singapore, September 3, 2016*, Ming Li, Xiaoyin Wang, and Lucia (Eds.). ACM, 2. <https://doi.org/10.1145/2975961.2990477>
- [62] Emma Nilsson-Nyman, Torbjörn Ekman, and Görel Hedin. 2008. Practical Scope Recovery Using Bridge Parsing. In *Software Language Engineering, First International Conference, SLE 2008, Toulouse, France, September 29-30, 2008. Revised Selected Papers (Lecture Notes in Computer*

- Science*, Vol. 5452), Dragan Gasevic, Ralf Lämmel, and Eric Van Wyk (Eds.). Springer, 95–113. https://doi.org/10.1007/978-3-642-00434-6_7
- [63] Corrina Perrone and Alexander Repenning. 1998. Graphical Rewrite Rule Analogies: Avoiding the Inferior Copy & Paste Reuse Dilemma. In *Proceedings 1998 IEEE Symposium on Visual Languages, Nova Scotia, Canada, September 1-4, 1998*. IEEE Computer Society, 40–46. <https://doi.org/10.1109/VL.1998.706132>
- [64] Frank Pfenning. 1991. Unification and Anti-Unification in the Calculus of Constructions. In *Proceedings of the Sixth Annual Symposium on Logic in Computer Science (LICS '91), Amsterdam, The Netherlands, July 15-18, 1991*. IEEE Computer Society, 74–85. <https://doi.org/10.1109/LICS.1991.151632>
- [65] Hung Dang Phan, Anh Tuan Nguyen, Trong Duc Nguyen, and Tien N. Nguyen. 2017. Statistical migration of API usages. In *Proceedings of the 39th International Conference on Software Engineering, ICSE 2017, Buenos Aires, Argentina, May 20-28, 2017 - Companion Volume*, Sebastián Uchitel, Alessandro Orso, and Martin P. Robillard (Eds.). IEEE Computer Society, 47–50. <https://doi.org/10.1109/ICSE-C.2017.17>
- [66] Noah S Podolefsky and Noah D Finkelstein. 2006. Use of analogy in learning physics: The role of representations. *Physical Review Special Topics-Physics Education Research* 2, 2 (2006), 020101.
- [67] Veselin Raychev, Martin T. Vechev, and Andreas Krause. 2015. Predicting Program Properties from "Big Code". In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, Sriram K. Rajamani and David Walker (Eds.). ACM, 111–124. <https://doi.org/10.1145/2676726.2677009>
- [68] Mohammad Raza and Sumit Gulwani. 2018. Disjunctive Program Synthesis: A Robust Approach to Programming by Example. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, 1403–1412. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17055>
- [69] Mohammad Raza, Sumit Gulwani, and Natasa Milic-Frayling. 2015. Compositional Program Synthesis from Natural Language and Examples. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, Qiang Yang and Michael J. Wooldridge (Eds.). AAAI Press, 792–800. <http://ijcai.org/Abstract/15/117>
- [70] Alexander Repenning and Corrina Perrone. 2000. Programming by example: programming by analogous examples. *Commun. ACM* 43, 3 (2000), 90–97.
- [71] Lindsey E Richland, Keith J Holyoak, and James W Stigler. 2004. Analogy use in eighth-grade mathematics classrooms. *Cognition and instruction* 22, 1 (2004), 37–60.
- [72] Lindsey E Richland, Osnat Zur, and Keith J Holyoak. 2007. Cognitive supports for analogies in the mathematics classroom. *Science* 316, 5828 (2007), 1128–1129.
- [73] Chanchal Kumar Roy, James R. Cordy, and Rainer Koschke. 2009. Comparison and evaluation of code clone detection techniques and tools: A qualitative approach. *Sci. Comput. Program.* 74, 7 (2009), 470–495. <https://doi.org/10.1016/j.scico.2009.02.007>
- [74] Shmuel Sagiv, Thomas W. Reps, and Reinhard Wilhelm. 1999. Parametric Shape Analysis via 3-Valued Logic. In *POPL '99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20-22, 1999*, Andrew W. Appel and Alex Aiken (Eds.). ACM, 105–118. <https://doi.org/10.1145/292540.292552>
- [75] Shmuel Sagiv, Thomas W. Reps, and Reinhard Wilhelm. 2002. Parametric shape analysis via 3-valued logic. *ACM Trans. Program. Lang. Syst.* 24, 3 (2002), 217–298. <https://doi.org/10.1145/514188.514190>
- [76] Hesam Samimi, Chris Deaton, Yoshiaki Ohshima, Alessandro Warth, and Todd D. Millstein. 2014. Call by Meaning. In *Onward! 2014, Proceedings of the 2014 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software, part of SPLASH '14, Portland, OR, USA, October 20-24, 2014*, Andrew P. Black, Shriram Krishnamurthi, Bernd Bruegge, and Joseph N. Ruskiewicz (Eds.). ACM, 11–28. <https://doi.org/10.1145/2661136.2661152>
- [77] Oli Sharpe. 2018. Semprola: a semiotic programming language. In *Conference Companion of the 2nd International Conference on Art, Science, and Engineering of Programming, Nice, France, April 09-12, 2018*, Stefan Marr and Jennifer B. Sartor (Eds.). ACM, 202–213. <https://doi.org/10.1145/3191697.3214330>
- [78] Rishabh Singh. 2016. BlinkFill: Semi-supervised Programming By Example for Syntactic String Transformations. *Proc. VLDB Endow.* 9, 10 (2016), 816–827. <https://doi.org/10.14778/2977797.2977807>
- [79] Rishabh Singh and Sumit Gulwani. 2012. Learning Semantic String Transformations from Examples. *Proc. VLDB Endow.* 5, 8 (2012), 740–751. <https://doi.org/10.14778/2212351.2212356>
- [80] Rishabh Singh and Sumit Gulwani. 2016. Transforming spreadsheet data types using examples. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav Bodik and Rupak Majumdar (Eds.). ACM, 343–356. <https://doi.org/10.1145/2837614.2837668>
- [81] David Canfield Smith. 1975. *Pygmalion: a creative programming environment*. Technical Report. STANFORD UNIV CA DEPT OF COMPUTER SCIENCE.
- [82] Dawn Xiaodong Song, David Brumley, Heng Yin, Juan Caballero, Ivan Jager, Min Gyung Kang, Zhenkai Liang, James Newsome, Pongsin Poosankam, and Prateek Saxena. 2008. BitBlaze: A New Approach to Computer Security via Binary Analysis. In *Information Systems Security, 4th International Conference, ICISS 2008, Hyderabad, India, December 16-20, 2008. Proceedings (Lecture Notes in Computer Science, Vol. 5352)*, R. Sekar and Arun K. Pujari (Eds.). Springer, 1–25. https://doi.org/10.1007/978-3-540-89862-7_1
- [83] Matthew Sotoudeh, Anand Venkat, Michael J. Anderson, Evangelos Georganas, Alexander Heinecke, and Jason Knight. 2019. ISA mapper: a compute and hardware agnostic deep learning compiler. In *Proceedings of the 16th ACM International Conference on Computing Frontiers, CF 2019, Alghero, Italy, April 30 - May 2, 2019*, Francesca Palumbo, Michela Becchi, Martin Schulz, and Kento Sato (Eds.). ACM, 164–173. <https://doi.org/10.1145/3310273.3321559>
- [84] Ashley Strickland. 2020. The 'beating hearts' of these pulsating stars create music to astronomers' ears. <https://www.cnn.com/2020/05/15/world/pulsating-stars-delta-scuti-scn-trnd/index.html>. Accessed May, 2020.
- [85] Marko Vasic, Aditya Kanade, Petros Maniatis, David Bieber, and Rishabh Singh. 2019. Neural Program Repair by Jointly Learning to Localize and Repair. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. <https://openreview.net/forum?id=ByloJ20qtm>
- [86] Martin T. Vechev and Eran Yahav. 2016. Programming with "Big Code". *Found. Trends Program. Lang.* 3, 4 (2016), 231–284. <https://doi.org/10.1561/25000000028>
- [87] Matthias Wenzl, Georg Merzdovnik, Johanna Ullrich, and Edgar R. Weippl. 2019. From Hack to Elaborate Technique - A Survey on Binary Rewriting. *ACM Comput. Surv.* 52, 3 (2019), 49:1–49:37. <https://doi.org/10.1145/3316415>
- [88] Martin White, Michele Tufano, Christopher Vendome, and Denys Poshyvanyk. 2016. Deep learning code fragments for code clone detection. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering, ASE 2016, Singapore, September 3-7, 2016*, David Lo, Sven Apel, and Sarfraz Khurshid (Eds.). ACM, 87–98. <https://doi.org/10.1145/2970276.2970326>

- [89] Patrick H. Winston. 1980. Learning and Reasoning by Analogy. *Commun. ACM* 23, 12 (1980), 689–703. <https://doi.org/10.1145/359038.359042>
- [90] Edmund Wong, Taiyue Liu, and Lin Tan. 2015. CloCom: Mining existing source code for automatic comment generation. In *22nd IEEE International Conference on Software Analysis, Evolution, and Reengineering, SANER 2015, Montreal, QC, Canada, March 2-6, 2015*, Yann-Gaël Guéhéneuc, Bram Adams, and Alexander Serebrenik (Eds.). IEEE Computer Society, 380–389. <https://doi.org/10.1109/SANER.2015.7081848>
- [91] Moshé M. Zloof. 1977. Query-by-Example: A Data Base Language. *IBM Syst. J.* 16, 4 (1977), 324–343. <https://doi.org/10.1147/sj.164.0324>