

Toward Semantic Foundations for Program Editors

Cyrus Omar¹, Ian Voysey¹, Michael Hilton², Joshua Sunshine¹,
Claire Le Goues¹, Jonathan Aldrich¹, and Matthew A. Hammer³

1 Carnegie Mellon University, Pittsburgh, PA, USA

{comar,iev,sunshine,clegoues,aldrich}@cs.cmu.edu

2 Oregon State University, Corvallis, OR, USA

hiltonm@eecs.oregonstate.edu

3 University of Colorado Boulder, Boulder, CO, USA

matthew.hammer@colorado.edu

Abstract

Programming language definitions assign formal meaning to complete programs. Programmers, however, spend a substantial amount of time interacting with *incomplete* programs – programs with holes, type inconsistencies and binding inconsistencies – using tools like program editors and live programming environments (which interleave editing and evaluation.) Semanticists have done comparatively little to formally characterize 1) the static and dynamic semantics of incomplete programs; 2) the actions available to programmers as they edit and inspect incomplete programs; and 3) the behavior of editor services that suggest likely edit actions to the programmer based on semantic information extracted from the incomplete program being edited, and from programs that the system has encountered in the past.

This paper serves as a vision statement for a research program that seeks to develop these “missing” semantic foundations. Our hope is that these contributions, which will take the form of a series of simple formal calculi equipped with a clear mechanized metatheory, will guide the design of a variety of current and future interactive programming tools, much as various lambda calculi have guided modern language designs. Our own research will apply these principles in the design of Hazel, an experimental *live lab notebook* programming environment designed for data science tasks. We plan to co-design the Hazel language with the editor so that we can explore e.g. edit-time semantic conflict resolution mechanisms and mechanisms that allow library providers to install new library-specific edit-time features.

Digital Object Identifier 10.4230/LIPIcs.SNAPL.2017.23

1 Introduction

Language-aware program editors (e.g. Eclipse or Emacs, with the appropriate extensions installed [15]) offer programmers a number of useful editor services. Simple examples include (1) syntax highlighting, (2) type inspection, (3) navigation to variable binding sites, and (4) refactoring. More sophisticated editors provide context-aware suggestions to the programmer (using various code completion, program synthesis and program repair techniques.) Many editors also offer *live programming* [28, 5] services, e.g. by displaying the run-time value of an expression directly within the editor as the program is running.

When these editor services encounter *complete programs* – programs that are well-formed and semantically meaningful according to the definition of the language in use – they can rely on a variety of well-understood reasoning principles and program manipulation techniques. For example, a syntax highlighter for well-formed programs can be generated automatically from a context-free grammar [23, 51] and the remaining editor services enumerated above can follow the language’s type and binding structure as specified by a standard static semantics. Live programming services can additionally refer to the language’s dynamic semantics.



© the authors;

licensed under Creative Commons License CC-BY

Summit oN Advances in Programming Languages (SNAPL 2017).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:12



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The problem, of course, is that many of the edit states encountered by a program editor do not correspond to complete programs. For example, the programmer may be in the midst of a transient edit, or the programmer may have introduced a type error somewhere in the program. Standard language definitions are silent about incomplete programs, so in these situations, simple program editors simply disable various editor services until the program is again in a complete state. In other words, useful editor services become unavailable when the programmer needs them most! More advanced editors attempt to continue to provide editor services during these incomplete states by using various *ad hoc* and poorly understood heuristics based on idiosyncratic internal representations of incomplete programs.

This paper advocates for a research program that seeks to understand both incomplete programs, and the editor services that interact with them, as semantically rich mathematical objects. This will broaden the scope of the “programming language theory” (PLT) tradition, which has made significant advances by treating complete programs, programming languages and logics as semantically rich mathematical objects.

In following the PLT tradition, we intend to start by developing a series of minimal calculi that build upon well-understood typed lambda calculi to capture the essential character of incomplete programs and various editor services of interest. Editor designers can then apply the insights gained from studying these calculi (together with insights gained from the study of human factors and other topics) to design more elaborate program editors. Some of these editors will evolve directly from editors already in use today. In parallel with these efforts, we plan to design a “clean-slate” programming environment, **Hazel**, based directly on these first principles. This will allow researchers to cleanly explore the frontier of what is possible when one considers languages and editors within a common theoretical framework. Such a clean-slate design might also prove useful in certain educational settings, and perhaps even make its way into practice.

Figure 1 shows a mockup of the **Hazel** user interface, which is loosely modeled after the widely adopted IPython / Jupyter lab notebook interface [39]. This figure will serve as our running example throughout the remainder of the paper. Each section below briefly summarizes a fundamental problem that we must confront as we seek to develop a semantic foundation for advanced program editors. For each problem, we discuss existing approaches, including those advanced by our own recent research, and suggest a number of promising future research directions for the community to pursue.

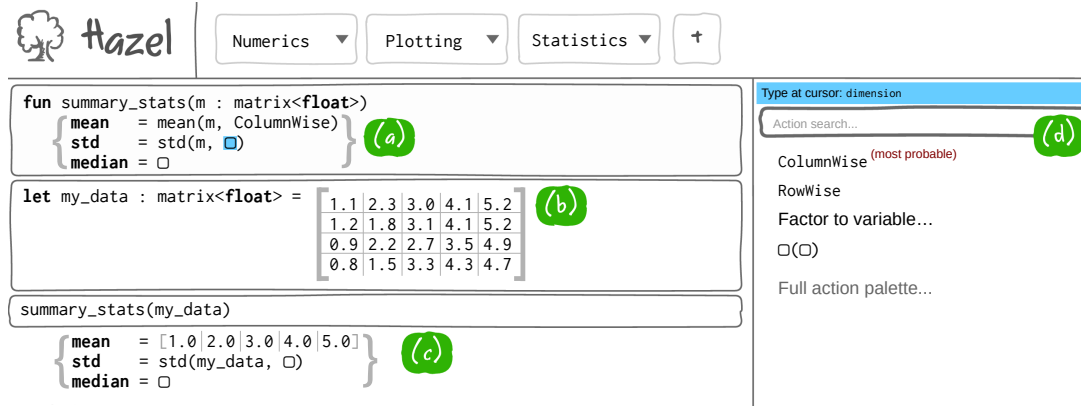
2 Problem 1: Syntactically Malformed Edit States

Textual program editors frequently encounter edit states that are not well-formed with respect to the textual syntax of complete programs. For example, consider a programmer constructing a call to a function `std`:

```
std(m,
```

There is a syntax error, so editor services that require a syntactically complete program must be disabled. This is unsatisfying.

Sophisticated editors like Eclipse, and editor generators like Spoofox [21], use *error recovery* heuristics that silently insert tokens so that the editor-internal representation is well-formed [1, 8, 16, 20]. These heuristics are typically provided manually by the grammar designer, though certain heuristics can be generated semi-automatically by tools that are given a description of the scoping conventions of the language or the secondary notational conventions (e.g. whitespace) [20, 14]. Error recovery heuristics require guessing at the programmer’s intent, so they are fundamentally *ad hoc* and can confuse the programmer [20].



■ **Figure 1** A mockup of Hazel.

A cleaner alternative approach, and the approach that we plan to take for Hazel, is to build a *structure editor* – a program editor where every edit state maps onto a syntax tree, with *holes* representing leaves of the tree that have not yet been constructed. This representation choice sidesteps the problem of syntactically malformed edit states. Notice that in Figure 1, the program fragment in cell (a) contains holes, appearing as squares. This design also permits non-textual *projections* of expressions, e.g. the 2D projection of a matrix value in cell (b). We will return to the topic of structure editors and projections below.

Structure editors have a long history. For example, the Cornell Program Synthesizer was developed in the early 1980s [49]. Although text-based syntax continues to predominate, there remains significant interest in structure editors today, particularly in practice. For example, Scratch is a structure editor that has achieved success as a tool for teaching children how to program [44]. `mbeddr` is an editor for a C-like language [55], built using the commercially supported MPS structure editor workbench [54]. TouchDevelop is an editor for an object-oriented language [50]. Lamdu [26] and Unison [9] are open source structure editors for functional languages similar to Haskell. Most work on structure editors has focused on the user interfaces that they present. This is important work – presenting a fluid user interface involving higher-level edit actions is a non-trivial problem, and some aspects of this problem remain open even after many years of research. There is reason to be optimistic, however, with recent studies suggesting that programmers experienced with a modern keyboard-driven structure editor (e.g. `mbeddr`) can be highly productive [3, 56].

Researchers have also explored various “hybrid” approaches, combining the concept of holes with text based editing. For example, recent work on *syntactic placeholders* envisions a textual program editor where edit actions cause textual placeholders (a.k.a. holes) to appear, rather than leaving the program transiently malformed [2]. This “approximates” the experience of a structure editor in common usage, while allowing the programmer to perform arbitrary text edits when necessary. Some programming systems, e.g. recent iterations of the Glasgow Haskell Compiler (GHC) [46] and the Agda proof assistant [35], support a workflow where the programmer writes down holes manually at locations in the program that remain under construction. Another approach would be to perform error recovery by attempting to insert holes into the internal representation used by the program editor, without including them in the surface syntax exposed to programmers. If “pure” structure editing proves too rigid as we design Hazel, we will explore hybrid approaches like these.

$$\begin{aligned} \text{HType} : \hat{\tau} &::= \hat{\tau} \rightarrow \hat{\tau} \mid \text{num} \mid \emptyset \\ \text{HExp} : \hat{e} &::= x \mid \lambda x. \hat{e} \mid \hat{e}(\hat{e}) \mid \underline{n} \mid \hat{e} + \hat{e} \mid (\hat{e} : \hat{\tau}) \mid \emptyset \mid \langle \hat{e} \rangle \end{aligned}$$

■ **Figure 2** Syntax of H-types and H-expressions in the Hazelnut calculus [37].

3 Problem 2: Statically Meaningless Edit States

No matter how an editor confronts syntactically malformed edit states, it must also confront edit states that are syntactically well-formed but statically meaningless. For example, the following value binding has a type inconsistency:

```
val x : float = std(m, ColumnWise)
```

because `std` has type `matrix<float> * dimension -> vec<float>`, but the type annotation on `x` is `float`, rather than `vec<float>`. This leaves the entire surrounding program formally meaningless according to a standard static semantics.

In the presence of syntactic holes, the problem of reasoning statically about incomplete programs becomes even more interesting. Consider the incomplete expression `std(m, □)` from cell (a) in Figure 1. Although it is intuitively apparent that the type of this expression, after hole instantiation, could only be `vec<float>` (the return type of `std`), and that the hole must be instantiated with values of type `dimension`, the static semantics of complete expressions is again silent about these matters.

Various heuristic approaches are implemented in Eclipse and other sophisticated tools, but the formal character of these heuristics are obscure, buried deep within their implementations. What is needed is a clear static semantics for incomplete programs, i.e. programs that contain holes (in both expressions and types), type inconsistencies, binding inconsistencies (i.e. unbound variables), and other static problems. Using such a static semantics, Hazel could communicate to the programmer (see right column of Figure 1) that the expression at the cursor, highlighted in blue in cell (a), must be of type `dimension`. The incomplete function `summary_stats` as a whole could be assigned the following incomplete function type: `matrix<float> -> { mean : vec<float>, std : vec<float>, median : □ }`.

We have investigated a subset of this problem, by defining a static semantics for a simply typed lambda calculus (with a single base type, `num`, for simplicity) extended with holes and type inconsistencies, but no binding inconsistencies. The results are detailed in a paper to appear at POPL 2017 [37]. Figure 2 defines the syntactic objects of this calculus – *H-types*, $\hat{\tau}$, are types with holes (\emptyset), and *H-expressions*, \hat{e} , are expressions with holes (\emptyset), and marked type inconsistencies, $\langle \hat{e} \rangle$. We call marked type inconsistencies *non-empty holes*, because they mark portions of the syntax tree that remain incomplete and behave semantically much like empty holes. Types and expressions that contain no holes are *complete types* and *complete expressions*, respectively.

We will not reproduce the details in this submission. Instead, let us simply note some interesting connections with other work.

First, type holes behave much like unknown types, `?`, from Siek and Taha’s pioneering work on gradual typing [45]. This discovery is quite encouraging, given that gradual typing is also motivated by a desire to make sense of one class of “intermediate edit states” – programs that have not yet been fully annotated with types.

Empty expression holes have also been studied formally, e.g. as the *metavariables* of contextual modal type theory (CMTT) [33]. In particular, expression holes can have types and are surrounded by contexts, just as metavariables in CMTT are associated with types and contexts. This begins to clarify the logical meaning of a typing derivation in Hazelnut –

it conveys well-typedness relative to an (implicit) modal context that extracts each expression hole’s type and context. The modal context must be emptied – i.e. the expression holes must be instantiated with expressions of the proper type in the proper context – before the expression can be considered complete. This relates to the notion of modal necessity in contextual modal logic.

For interactive proof assistants that support a tactic model based directly on hole filling, the connection to CMTT and similar systems is quite salient. For example, Beluga [41] is based on dependent CMTT and aspects of Idris’ editor support [4] are based on a similar system – McBride’s OLEG [27]. As we will discuss below, our notion of a program editor supports actions beyond hole filling.

There are a number of future research directions that are worth exploring.

Binding inconsistencies. In the simple calculus developed so far, all variables must be bound before they are used, including those in holes. We must extend this system to support reasoning when a variable is used before it is bound (as is a common workflow.) Dagenais and Hendren also studied how to reason statically about programs with binding errors using a constraint system, focusing on Java programs whose imports are not completely known [13]. They neither considered programs with holes or other type inconsistencies, nor formally mathematically specified their technique. However, they provide a useful starting point.

Expressiveness. The simple calculus discussed above is only as expressive as the typed lambda calculus with numbers. We must scale up the semantics to handle other modern language features. Our plan is to focus initially on functional language constructs (so that Hazel can be used to teach courses that are today taught using Standard ML, OCaml or Haskell.) This will include recursive and polymorphic functions, recursive types, and labeled product (record) and sum types. We also propose to investigate ML-style structural pattern matching. All of these will require defining new sorts of holes and static inconsistencies, including: (1) non-empty holes at the type level, to handle kind inconsistencies; (2) holes in label position; and (3) holes and type inconsistencies in patterns.

Automation. Although we plan to explore some of these language extensions “manually,” extending our existing mechanization, we ultimately plan to *automatically* generate a statics for incomplete terms from a standard statics for complete terms, annotated perhaps with additional information. There is some precedent for this in recent work on the Gradualizer, which is capable of producing a gradual type system from a standard type system with lightweight annotations that communicate the intended polarities of certain constructs [12]. However, although it provides a good starting point, gradual type systems only consider the problem of holes in types. Our plan is to build upon existing proof automation techniques, e.g. Agda’s reflection [52] (in part because our present mechanization effort is in Agda.)

4 Problem 3: Dynamically Meaningless Edit States

Modern programming tools are increasingly moving beyond simple “batch” programming models by incorporating *live programming* features that interleave editing and evaluation [48, 47, 29]. These tools provide programmers with rapid feedback about the dynamic behavior of the program they are editing, or selected portions thereof [30]. Examples include *lab notebooks*, e.g. the popular IPython/Jupyter [40], which allow the programmer to interactively edit and evaluate program fragments organized into a sequence of cells (an extension of the read-eval-print loop (REPL)); spreadsheets; live graphics programming environments, e.g. SuperGlue [29], Sketch-n-Sketch [11, 10] and the tools demonstrated by Bret Victor in his lectures [53]; the TouchDevelop live UI framework [6]; and live visual

and auditory dataflow languages [7]. In the words of Burckhardt et al., live programming environments “capture the imagination of today’s programmers and promise to narrow the temporal and perceptive gap between program development and code execution” [6].

Our proposed design for Hazel combines aspects of several of these designs to form a *live lab notebook interface*. It will use the edit state of each cell to continuously update the output value displayed for that cell and subsequent cells that depend on it. Uniquely, rather than providing meaningful feedback about the dynamic behavior only once a cell becomes complete, Hazel will provide meaningful feedback also about the dynamic behavior of incomplete cells (and thereby further tighten the “perceptive gap”).

For example, in cell (c) of Figure 1, the programmer applies the incomplete function `summary_stats` to the matrix `my_data`, and the editor is still able to display a result. The value of the column-wise mean is fully determined, because evaluation does not encounter any holes, whereas the standard deviation and median computations cannot be fully evaluated. Notice, however, that the standard deviation computation does communicate the substitution of the applied argument, `my_data`, for the variable `m`.¹

To realize this functionality, we need a dynamic semantics for incomplete programs that builds upon our proposed static semantics. There is some precedent for this: research in gradual typing considers the dynamic semantics of programs with holes in types, and our proposed static semantics for incomplete programs borrows technical machinery from theoretical work on gradual typing [45]. We seek a dynamic semantics for the incomplete programs that also have expression holes (and in the future, other sorts of holes.) Research on CMTT has not yet considered the problem of evaluating expressions under a non-empty metavariable context. Accomplishing this goal will require great formal care, because the presence of type and binding errors would violate the usual notions of type safety – evaluation can “get stuck” because a hole appears in a critical (elimination) position. Normally, this would violate the classical notion of Progress – evaluation can neither proceed, nor has it produced a value. We conjecture that this is resolved by (1) positively characterizing *indeterminate* evaluation states, those where a hole blocks progress at all locations within the expression, and (2) defining a notion of Indeterminate Progress that allows for evaluation to stop at an indeterminate evaluation state.

There are several more practical applications that we aim to explore after developing the initial foundations just described. For example, it would be useful for the programmer to be able to select a hole that appears in an indeterminate state and 1) be taken to its original location; 2) be able to inspect the *value* of a subexpression under the cursor in the environment of the selected hole (rather than just its type.) Again, CMTT’s closures provide a theoretical starting point for this editor service.

It would also be useful to be able to continue evaluation where it left off after making an edit to the program that corresponds to hole instantiation. This would require proving a commutativity property regarding hole instantiation. Fortunately, initial research on commutativity properties for holes has been conducted for CMTT, which will serve as a starting point for this work [34]. There are likely to be interesting new theoretical questions that arise when one adds non-termination and memory effects (e.g. regarding the shape of the heap before and after the edit.) Relatedly, IPython/Jupyter [40] support a feature whereby numeric variable(s) in cells can be marked as being “interactive”, which causes

¹ To avoid exposing the internals of imported library functions, evaluation does not step into functions, like `std`, that have been imported from external libraries indicated by the row at the top of Figure 1 (unless specifically requested, not shown.)

the user interface to display a slider. As the slider value changes, the new value of the cell is recomputed. It would be useful to be able to use the mechanisms just proposed to incrementalize parts of this recomputation.

5 Problem 4: A Calculus of Edit Actions

Our first two specific aims allow us to assign meaning to intermediate edit states. However, to understand the act of *editing* itself, we need a calculus of *edit actions* that governs transitions between edit states.

Assuming that Hazel will be a structure editor, the goal will be for every possible edit state to be both statically and dynamically meaningful according to the semantics proposed in the previous two sections. This corresponds formally to proving a metatheorem about the action semantics: when the initial edit state is semantically meaningful, the edit state that results from performing an action is as well. In addition to this crucial metatheorem, which we call *sensibility*, there are a number of other metatheorems of interest that establish the interactive expressive power of the action semantics, e.g. that every well-typed term can be constructed by some sequence of edit actions. (In a textual setting, these structured edit actions would need to be supplemented by lower-level text edit actions that do not maintain such invariants.)

In our recent work on Hazelnut, we have developed an action calculus for the minimal calculus of H-types and H-expressions described in Section 3 [37]. We have mechanically proven the sensibility invariant, as well as a number of other metatheorems. What remains is to scale this calculus up, and to investigate *action composition principles*. For example, it would be worthwhile to investigate the notion of an *action macro*, whereby functional programs could themselves be lifted to the level of actions to compute non-trivial compound actions. Such compound actions would give a uniform description of transformations ranging from the simple—like “move the cursor to the next hole to the right”—to quite complex whole program refactorings, while remaining subject to the core semantics. Using proof automation, it should be possible to prove that an action macro implements derived action logic that is admissible with respect to the core semantics. This would eliminate the possibility of “edit-time” errors. This is closely related to work on tactic languages in proof assistants, e.g. the Mtac typed macro language for Coq [57], differing again in that the action language involves notions other than hole filling.

6 Problem 5: Meaningful Suggestion Generation and Ranking

The simplest edit actions will be bound to keyboard shortcuts. However, Hazel will also provide suggestions to help the programmer edit incomplete programs by providing a *suggestion palette*, marked (d) in Figure 1. This palette will suggest semantically relevant code snippets when the cursor is on an empty hole. It will also suggest other relevant edit actions, including high-level edit actions implemented by imported action macros (e.g. the refactoring action in Figure 1.) When the cursor is on a non-empty hole, indicating a static error, it will suggest bug fixes. We plan to also consider bugs that do not correspond to static errors, including those identified explicitly by the programmer, and those related to assertion failures or exceptions encountered when using the live programming features of Hazel. In these situations, we plan to defer to existing automated fault localization techniques [19, 42, 43].

Note that features like these are not themselves novel. Many editors provide contextually relevant suggestions. Indeed, suggestion generation is closely related to several major research

areas: code completion [32, 18], program synthesis [17], and program repair [24, 31, 25, 22].

The problems that existing systems encounter is exactly the problem we have been discussing throughout this proposal: when attempting to integrate these features into an editor, it is difficult to reason about malformed or meaningless edit states. As such, many of these systems use tokenized representations of programs [18]. Because **Hazel** will maintain the invariant that every edit state is a syntactically and semantically meaningful formal structure, we can develop a more principled solution to the problem of generating meaningful suggestions. In particular, we will *prove* that every action suggestion generated for a particular edit state is meaningful for that edit state.

In addition to investigating the problem of populating the suggestion palette with semantically valid actions, we will consider the problem of evaluating the statistical likelihood of the suggestions. This requires developing a statistical model over actions. We will prove that this statistical model is a proper statistical distribution (e.g. that it “integrates” to 1), and that it assigns zero probability to semantically invalid actions. Collectively, we refer to these contributions as a *statistical action suggestion semantics*.

7 Language-Editor Co-Design

With **Hazel**, we are intentionally blurring the line between the programming language and the program editor. This opens up a number of interesting research directions in language-editor co-design. For example, it may be possible to recast “tricky” language mechanisms, like function overloading, type classes, implicit values and unqualified imports, as editor mechanisms. Because we will be treating programming as a structured conversation between the programmer and the programming environment, the editor can simply ask the programmer to resolve ambiguities when they arise. The programmer’s choice is then stored unambiguously in the syntax tree.

Another important research direction lies in exploring how types can be used to control the presentation of expressions in the editor. Following an approach that we developed in a textual setting of *type-specific languages* (TSLs), it should be possible to have the type that an expression is being analyzed against define a *type-specific projection* defined by elaboration to the underlying static semantics and action semantics of **Hazel** [36]. For example, the matrix projection shown in Figure 1 need not be built in to **Hazel**. Instead, the **Numerics** library provider could introduce this logic. This line of research is also related to our work on *active code completion*, which investigated type-specific code generation interfaces [38].

Another interesting problem is that of semantic, interactive documentation. In particular, in **Hazel**, references to program structures that appear in documentation will be treated in the same way as other references and be subject to renaming and other operations. Documentation will also be capable of containing expressions of arbitrary types (e.g. of the **Image** or **Link** type.) Together with the type-specific projection mechanism mentioned above, we hope that this will allow **Hazel** to function not only as a structured programming environment, but also as a structured document authoring environment!

8 Conclusion

To summarize, there are a number of interesting semantic questions that come up in the design of program editors. We advocate a research program that studies these problems using mathematical tools previously used to study programming languages and complete programs. This work will both deobfuscate the design of editors and open up the doors for a number of new advanced editor features.

References

- 1 Alfred V. Aho and Thomas G. Peterson. A minimum distance error-correcting parser for context-free languages. *SIAM J. Comput.*, 1(4):305–312, 1972. URL: <http://dx.doi.org/10.1137/0201022>, doi:10.1137/0201022.
- 2 Luís Eduardo de Souza Amorim, Sebastian Erdweg, Guido Wachsmuth, and Eelco Visser. Principled syntactic code completion using placeholders. In *SLE*, 2016.
- 3 Dimitar Asenov and Peter Müller. Envision: A fast and flexible visual code editor with fluid interactions (overview). In *IEEE Symposium on Visual Languages and Human-Centric Computing, VL/HCC 2014, Melbourne, VIC, Australia, July 28 - August 1, 2014*, pages 9–12, 2014.
- 4 Edwin Brady. Idris, A General-Purpose Dependently Typed Programming Language: Design and Implementation. *Journal of Functional Programming*, 23(05):552–593, 2013.
- 5 Sebastian Burckhardt, Manuel Fähndrich, Peli de Halleux, Sean McDirmid, Michal Moskal, Nikolai Tillmann, and Jun Kato. It’s alive! continuous feedback in ui programming. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI ’13*, pages 95–104, New York, NY, USA, 2013. ACM. URL: <http://doi.acm.org/10.1145/2491956.2462170>, doi:10.1145/2491956.2462170.
- 6 Sebastian Burckhardt, Manuel Fähndrich, Peli de Halleux, Sean McDirmid, Michal Moskal, Nikolai Tillmann, and Jun Kato. It’s alive! continuous feedback in UI programming. In *PLDI*, 2013.
- 7 Margaret M. Burnett, John W. Atwood Jr., and Zachary T. Welch. Implementing level 4 liveness in declarative visual programming languages. In *Proceedings 1998 IEEE Symposium on Visual Languages, Nova Scotia, Canada, September 1-4, 1998*, pages 126–133, 1998.
- 8 Philippe Charles. A practical method for constructing efficient lalr (k) parsers with automatic error recovery, 1991.
- 9 Paul Chiusano. Unison. <http://www.unisonweb.org/>. Accessed: 2016-04-25.
- 10 Ravi Chugh. Prodirect manipulation: bidirectional programming for the masses. In *Proceedings of the 38th International Conference on Software Engineering, ICSE 2016, Austin, TX, USA, May 14-22, 2016 - Companion Volume*, pages 781–784, 2016. URL: <http://doi.acm.org/10.1145/2889160.2889210>, doi:10.1145/2889160.2889210.
- 11 Ravi Chugh, Brian Hempel, Mitchell Spradlin, and Jacob Albers. Programmatic and direct manipulation, together at last. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*, pages 341–354, 2016.
- 12 Matteo Cimini and Jeremy G. Siek. The gradualizer: a methodology and algorithm for generating gradual type systems. In *POPL*, 2016.
- 13 Barthélémy Dagenais and Laurie J. Hendren. Enabling static analysis for partial java programs. In *Proceedings of the 23rd Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2008, October 19-23, 2008, Nashville, TN, USA*, pages 313–328, 2008.
- 14 Maartje de Jonge, Emma Nilsson-Nyman, Lennart C. L. Kats, and Eelco Visser. Natural and flexible error recovery for generated parsers. In *Software Language Engineering, Second International Conference, SLE 2009, Denver, CO, USA, October 5-6, 2009, Revised Selected Papers*, pages 204–223, 2009.
- 15 Erich Gamma and Kent Beck. *Contributing to Eclipse: principles, patterns, and plug-ins*. Addison-Wesley Professional, 2004.
- 16 Susan L Graham, Charles B Haley, and William N Joy. *Practical LR error recovery*, volume 14. ACM, 1979.
- 17 Sumit Gulwani. Dimensions in program synthesis. In *Proceedings of the 12th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming, PPDP*

- '10, pages 13–24, New York, NY, USA, 2010. ACM. URL: <http://doi.acm.org/10.1145/1836089.1836091>, doi:10.1145/1836089.1836091.
- 18 Abram Hindle, Earl T. Barr, Zhendong Su, Mark Gabel, and Premkumar Devanbu. On the naturalness of software. In *Proceedings of the 34th International Conference on Software Engineering, ICSE '12*, pages 837–847, 2012.
 - 19 James A. Jones, Mary Jean Harrold, and John Stasko. Visualization of test information to assist fault localization. In *International Conference on Software Engineering (ICSE)*, pages 467–477, Orlando, FL, USA, 2002. doi:10.1145/581339.581397.
 - 20 Lennart C. L. Kats, Maartje de Jonge, Emma Nilsson-Nyman, and Eelco Visser. Providing rapid feedback in generated modular language environments: adding error recovery to scannerless generalized-lr parsing. In *Proceedings of the 24th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2009, October 25-29, 2009, Orlando, Florida, USA*, pages 445–464, 2009. URL: <http://doi.acm.org/10.1145/1640089.1640122>, doi:10.1145/1640089.1640122.
 - 21 Lennart C. L. Kats and Eelco Visser. The spoofax language workbench: rules for declarative specification of languages and ides. In *Proceedings of the 25th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2010, October 17-21, 2010, Reno/Tahoe, Nevada, USA*, pages 444–463, 2010.
 - 22 Yalin Ke, Kathryn T. Stolee, Claire Le Goues, and Yuriy Brun. Repairing programs with semantic code search. In *Proceedings of the 30th IEEE/ACM International Conference On Automated Software Engineering (ASE)*, Lincoln, NE, USA, November 2015.
 - 23 Holger Krahn, Bernhard Rumpe, and Steven Völkel. Monticore: Modular development of textual domain specific languages. In *Objects, Components, Models and Patterns, 46th International Conference, TOOLS EUROPE 2008, Zurich, Switzerland, June 30 - July 4, 2008. Proceedings*, pages 297–315, 2008.
 - 24 Claire Le Goues, ThanhVu Nguyen, Stephanie Forrest, and Westley Weimer. GenProg: A generic method for automatic software repair. *IEEE Transactions on Software Engineering (TSE)*, 38:54–72, 2012. doi:10.1109/TSE.2011.104.
 - 25 Fan Long and Martin Rinard. Automatic patch generation by learning correct code. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '16*, pages 298–312, 2016. URL: <http://doi.acm.org/10.1145/2837614.2837617>, doi:10.1145/2837614.2837617.
 - 26 Eyal Lotem and Yair Chuchem. Project Lamdu. <http://www.lamdu.org/>. Accessed: 2016-04-08.
 - 27 Conor McBride. *Independently typed functional programs and their proofs*. PhD thesis, University of Edinburgh. College of Science and Engineering. School of Informatics., 2000.
 - 28 Sean McDermid. Living it up with a live programming language. In *Proceedings of the 22Nd Annual ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications, OOPSLA '07*, pages 623–638, New York, NY, USA, 2007. ACM. URL: <http://doi.acm.org/10.1145/1297027.1297073>, doi:10.1145/1297027.1297073.
 - 29 Sean McDermid. Living it up with a live programming language. *SIGPLAN Not.*, 42(10):623–638, October 2007.
 - 30 Sean McDermid. Usable live programming. In *Proceedings of the 2013 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software, Onward! 2013*, pages 53–62, New York, NY, USA, 2013. ACM.
 - 31 Sergey Mechtaev, Jooyong Yi, and Abhik Roychoudhury. Angelix: Scalable multiline program patch synthesis via symbolic analysis. In *Proceedings of the 38th International Conference on Software Engineering, ICSE '16*, pages 691–701, 2016.
 - 32 Kıvanç Muşlu, Yuriy Brun, Reid Holmes, Michael D. Ernst, and David Notkin. Improving IDE Recommendations by Considering Global Implications of Existing Recommendations.

- In *Proceedings of the New Ideas and Emerging Results Track at the 34th International Conference on Software Engineering (ICSE12)*, pages 1349–1352, Zurich, Switzerland, June 2012. doi:10.1109/ICSE.2012.6227082.
- 33 Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. *ACM Trans. Comput. Log.*, 9(3), 2008.
 - 34 Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. *ACM Trans. Comput. Log.*, 9(3), 2008. URL: <http://doi.acm.org/10.1145/1352582.1352591>, doi:10.1145/1352582.1352591.
 - 35 Ulf Norell. Dependently typed programming in agda. In *Advanced Functional Programming*, pages 230–266. Springer, 2009.
 - 36 Cyrus Omar, Darya Kurilova, Ligia Nistor, Benjamin Chung, Alex Potanin, and Jonathan Aldrich. Safely composable type-specific languages. In *ECOOP*, 2014.
 - 37 Cyrus Omar, Ian Voysey, Michael Hilton, Jonathan Aldrich, and Matthew Hammer. Hazelnut: A bidirectionally typed structure editor calculus. In *Principles of Programming Languages*, 2017. URL: <https://arxiv.org/abs/1607.04180>.
 - 38 Cyrus Omar, YoungSeok Yoon, Thomas D. LaToza, and Brad A. Myers. Active code completion. In *International Conference on Software Engineering (ICSE)*, 2012.
 - 39 Fernando Pérez and Brian E. Granger. IPython: a system for interactive scientific computing. *Computing in Science and Engineering*, 9(3):21–29, May 2007. URL: <http://ipython.org>, doi:10.1109/MCSE.2007.53.
 - 40 Fernando Perez and Brian E. Granger. Ipython: A system for interactive scientific computing. *Computing in Science and Engg.*, 9(3):21–29, May 2007. URL: <http://dx.doi.org/10.1109/MCSE.2007.53>, doi:10.1109/MCSE.2007.53.
 - 41 Brigitte Pientka. Beluga: Programming with dependent types, contextual data, and contexts. In *International Symposium on Functional and Logic Programming (FLOPS)*, 2010.
 - 42 Yuhua Qi, Xiaoguang Mao, Yan Lei, and Chengsong Wang. Using automated program repair for evaluating the effectiveness of fault localization techniques. In *International Symposium on Software Testing and Analysis*, 2013. gpfaultloc.
 - 43 Manos Renieris and Steven Reiss. Fault localization with nearest neighbor queries. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 30–39, Montreal, Québec, Canada, October 2003.
 - 44 Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, and Yasmin Kafai. Scratch: Programming for All. *Commun. ACM*, 52(11):60–67, November 2009.
 - 45 Jeremy G. Siek and Walid Taha. Gradual typing for functional languages. In *Scheme and Functional Programming Workshop*, 2006. URL: <http://scheme2006.cs.uchicago.edu/13-siek.pdf>.
 - 46 Simon Peyton Jones, Sean Leather and Thijs Alkemade. Typed holes in ghc. https://wiki.haskell.org/GHC/Typed_holes.
 - 47 Steven L. Tanimoto. VIVA: A visual language for image processing. *J. Vis. Lang. Comput.*, 1(2):127–139, 1990. URL: [http://dx.doi.org/10.1016/S1045-926X\(05\)80012-6](http://dx.doi.org/10.1016/S1045-926X(05)80012-6), doi:10.1016/S1045-926X(05)80012-6.
 - 48 Steven L. Tanimoto. A perspective on the evolution of live programming. In *Proceedings of the 1st International Workshop on Live Programming, LIVE 2013, San Francisco, California, USA, May 19, 2013*, pages 31–34, 2013. URL: <http://dx.doi.org/10.1109/LIVE.2013.6617346>, doi:10.1109/LIVE.2013.6617346.
 - 49 Tim Teitelbaum and Thomas Reps. The Cornell Program Synthesizer: A Syntax-directed Programming Environment. *Commun. ACM*, 24(9):563–573, 1981.

- 50 Nikolai Tillmann, Michal Moskal, Jonathan de Halleux, and Manuel Fahndrich. TouchDevelop: Programming Cloud-connected Mobile Devices via Touchscreen. In *SIGPLAN Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, 2011.
- 51 Mark van den Brand, Arie van Deursen, Jan Heering, H. A. de Jong, Merijn de Jonge, Tobias Kuipers, Paul Klint, Leon Moonen, Pieter A. Olivier, Jeroen Scheerder, Jurgen J. Vinju, Eelco Visser, and Joost Visser. The ASF+SDF meta-environment: A component-based language development environment. In *Compiler Construction, 10th International Conference, CC 2001 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2001 Genova, Italy, April 2-6, 2001, Proceedings*, pages 365–370, 2001.
- 52 Paul Van Der Walt and Wouter Swierstra. Engineering proof by reflection in agda. In *Symposium on Implementation and Application of Functional Languages*, pages 157–173. Springer, 2012.
- 53 B Victor. Inventing on principle, invited talk at the canadian university software engineering conference (cusec), 2012.
- 54 Markus Voelter. Language and IDE Modularization and Composition with MPS. In *International Summer School on Generative and Transformational Techniques in Software Engineering*, pages 383–430. Springer, 2011.
- 55 Markus Voelter, Daniel Ratiu, Bernhard Schaetz, and Bernd Kolb. Mbeddr: An Extensible C-based Programming Language and IDE for Embedded Systems. In *SPLASH*, 2012.
- 56 Markus Voelter, Janet Siegmund, Thorsten Berger, and Bernd Kolb. Towards User-Friendly Projectional Editors. In *International Conference on Software Language Engineering (SLE)*, 2014.
- 57 Beta Ziliani, Derek Dreyer, Neelakantan R Krishnaswami, Aleksandar Nanevski, and Viktor Vafeiadis. Mtac: A monad for typed tactic programming in Coq. *Journal of Functional Programming*, 25:e12, 2015.