

Notes on
Domain-Specific Languages:
Introduction to DSLs

H. Conrad Cunningham

16 April 2022

Contents

1	Introduction to DSLs	2
1.1	Chapter Introduction	2
1.2	What are DSLs?	2
1.2.1	Motivation	2
1.2.2	Examples	3
1.2.3	Definition	3
1.2.4	DSL/GPL Boundaries	4
1.3	External and Internal DSL Syles	4
1.3.1	External	5
1.3.2	Internal	5
1.4	Shallow and Deep Embeddings of Internal DSLs	6
1.4.1	Shallow embedding	7
1.4.2	Deep embedding	7
1.4.3	ELI Calculator Language	8
1.5	DSLs Used to Produce This Document	8
1.6	Possible Advantages of Using DSLs	9
1.7	Possible Disadvantages of Using DSLs	11
1.8	Designing DSLs	12
1.8.1	SCV analysis	12
1.8.2	DSL design guidelines	13
1.8.3	More guidelines for internal DSL design	19
1.9	What Next?	20
1.10	Exercises	21
1.11	Acknowledgements	21
1.12	Concepts and Terms	22
1.13	References	22

Copyright (C) 2014, 2016, 2017, 2018, 2022, H. Conrad Cunningham

Professor of Computer and Information Science
University of Mississippi
214 Weir Hall
P.O. Box 1848
University, MS 38677
(662) 915-7396 (dept. office)

Browser Advisory: The HTML version of this textbook requires a browser that supports the display of MathML. A good choice as of April 2022 is a recent version of Firefox from Mozilla.

1 Introduction to DSLs

1.1 Chapter Introduction

Note: This is mostly a document I used in my Spring 2018 CSci 658, Software Language Engineering, course. In 2022, I integrated the citations and hyperlinks with my other notes, but this Although it has had some modifications in 2022, but the document is still rough and incomplete.

TODO:

- Structure the document better. Check section titles.
- Finish missing sections and add better citations in places.
- Make flow more smoothly.
- Consider whether it should be broken into multiple chapters.
- Consider what are appropriate examples. Add or remove as needed. Expand on explanations of those I do include.
- Better integrate the various lists of guidelines for construction of DSLs. Better tie these to concrete examples.
- Should this consider monadic DSLs in Haskell, macro systems, Template Haskell, staged metaprogramming, or other techniques?

1.2 What are DSLs?

1.2.1 Motivation

Few computer science graduates will design and implement a general-purpose programming language during their careers. However, many graduates will design and implement—and all likely will use—special-purpose languages in their work. These special-purpose languages are often called domain-specific languages [24,87].

Hudak describes a *domain-specific language* (or *DSL*) as “a programming language tailored to a particular application domain” [36], that is, to a particular kind of problem.

General-purpose languages (GPLs) [86], such as Java, Python, C, C++, and Haskell, seek to be broadly applicable across many domains. They can, in theory, compute any function that is computable by a finite procedure; they are said to be *Turing complete* [45,85].

Ideally, a DSL should enable experts in an application area to program without programming—that is, to express the problems they want the computer to solve using familiar concepts and notations, without having to master the intricacies of programming in a general-purpose language [36,74].

DSLs might be Turing complete, but often they are not. DSLs are little languages [3] that “trade generality for expressiveness” [52].

1.2.2 Examples

As discussed in Bentley’s classic column on “Little Languages” [3], the DSL PIC [42,59,88] enables writers to produce line drawings in typeset documents; they can focus on the layout of the drawings without being required to develop programs in C, the primary general-purpose language used on Unix.

The PIC language and tool were built according to the Unix philosophy [84]. This approach focuses on a minimalist and modular approach to software development. The Unix tools have limited functionality but are usually built to read and write streams of bytes so that tools can be readily composed using Unix pipes.

Other DSLs on the Unix (or Linux) platform include:

- AWK [1,60,89] for extracting data and generating reports from streams of text-based data using regular expressions to specify the processing
- Lex [43,44,90] (and GNU Flex) for generating lexical analyzers
- Yacc [44,65,91] (and GNU Bison) for generating parsers
- Make [14,51,92] (and GNU Make) for building software from its various source files
- C preprocessor (cpp) [64,93] for preprocessing C programs to include header files, expand macros, and support conditional compilation
- DOT [66,94] little language for specifying graph structures within the Graphviz [67] set of visualization tools

Markup languages are also DSLs:

- HTML [78,95] is a DSL for formatting documents on the Web.
- LaTeX [68,96] is a powerful markup language used for books and articles (especially in STEM disciplines).
- Markdown [31,97] is a simple markup language for documents that may be needed in various formats, especially HTML. It is often used by wikis and other websites with user-contributed content (e.g., Wikipedia, GitHub). This document is written using the variant of Markdown supported by the Pandoc [46] tool.
- reStructuredText [30,98] is a simple markup language used mostly in the Python community for technical documentation.

The designers of a DSL must select relevant concepts, notations, and processes from the application domain and incorporate them into the DSL design [36].

1.2.3 Definition

Fowler defines a *domain-specific language* as a “computer programming language of limited expressiveness focused on a particular domain” [24].

He explains these terms as follows:

- As a **language**, a DSL has *fluency*. The expressiveness of the language comes not just from the simple expressions but also from how those expressions can be easily composed to form larger units.
- As a **computer programming language**, a DSL is structured so that humans can effectively read and write “programs” in the language and computers can accurately read and “execute” the instructions. It has a well-defined syntax and semantics.
- Having **limited expressiveness** means the DSL includes features that are needed for its purpose and excludes features that are not needed for that purpose.
- Being **focused on a particular domain** means that the DSL has a purpose that is both clearly and narrowly defined.

1.2.4 DSL/GPL Boundaries

What is and is not a DSL is a fuzzy concept. Not all writers agree on the definition. In these notes, we use Martin Fowler’s definition above to consider whether or not a notation is a DSL.

Fowler suggests that one key characteristic of a DSL is its *language* nature [24].

- The language (DSL) is intended to be read and written by humans as well as read by computers.
- The DSL must be *fluent*. That is, the meaning of the language comes not just from the individual expressions (i.e., the words) but also from how the expressions are composed together (i.e., into sentences and paragraphs). It has both a vocabulary and a grammar.

Fowler also suggests two other key characteristics of a DSL—its *limited expressiveness* and its *domain focus* [24].

Fowler argues that a DSL should have only those features needed to support its target domain. A DSL should not attempt to solve all problems for all users for all time. It should not seek to define an entire software system. It should instead focus on providing an effective, uncluttered solution to a specific aspect of the overall system.

Which of these boundaries is important in a particular situation depends upon the style of the DSL.

1.3 External and Internal DSL Styles

Fowler classifies DSLs into two styles [23,24]:

- external

- internal

Although the terminology is relatively new, the ideas are not.

1.3.1 External

An *external DSL* is a language that is different from the main programming language for an application, but that is interpreted by or translated into a program in the main language. The external DSL is a *standalone* language with its own syntax and semantics.

The Unix little languages PIC, Lex, Yacc, and Make exhibit this style. They are separate textual languages with their own syntax and semantics, but they are processed by C programs (and may also generate C programs).

External DSLs may choose from a variety of language processing approaches such as:

- ad hoc techniques (e.g., hand-coded recursive descent parsers)
- parser-generation tools (e.g., Lex and Yacc in the C/Unix environment, Happy [48] and Alex [47] on the Haskell platform, and ANTLR [55,56] on various platforms),
- parsing libraries (e.g., the Haskell library Parsec [33], Scala parser combinator library [54,62], Python parser combinator library Parsita [32], and Lua parsing library LPeg [38,39])

Example external DSLs associated with these course notes include:

- most of the Scala-based State Machine (Secret Panel) DSLs (adapted from Fowler [24])
- one of the Lua- and Python-based Lair Configuration DSLs (adapted from Fowler [22])
- three of the Ruby-based Reader DSLs (adapted from Fowler [18,19])

To distinguish between external DSLs and GPLs, Fowler cites the need for a DSL to limit its features to the minimal set needed for the specific domain [24]. (See the discussion of DSL/GPL boundaries in Section 1.2.4.)

For example, Fowler considers the programming language R [69,79,80] a GPL not a DSL. Although R has special-purpose features to support statistical programming, it has a full set of general -urpose features for a wide range of programming tasks.

1.3.2 Internal

An *internal DSL* transforms the main programming language itself into the DSL—the DSL is *embedded* in the main language [36].

The techniques for constructing internal DSLs vary from language to language.

The language Lisp (which was defined in the 1950s) [2,49,50,58,81] supports *syntactic macros*, a convenient mechanism for extending the language by adding application-specific features that are expanded at compile time. The Lisp macro approach has been refined and included in languages such as Scheme [15,27,28,57,58,82], Clojure [16,34,83], and Elixir [13,71].

Internal DSLs in the language Ruby [61,72] exploit the language’s *flexible syntax*, runtime *reflexive metaprogramming* facilities, and *blocks* (closures). The Ruby on Rails web framework includes several such internal DSLs.

Haskell’s algebraic data type system has stimulated research on “embedded” DSLs for several domains including reactive animation and music [36,37].

In object-oriented languages, internal DSLs may also exploit object structures and subtyping.

Example internal DSLs associated with these course notes include:

- the Scala-based Computer Configuration DSLs (adapted from Fowler [24])
- the Scala-based Email Message Building DSL (adapted from Fowler [24])
- most of the (Lua- and Python-based) Lair Configuration DSLs (adapted from Fowler [22])
- the Ruby-based Survey DSL [9] (partially motivated by the sidebar in Bentley [3])
- the Haskell-based Sandwich DSL project and similar Scala-based Sandwich DSL project

To distinguish between an internal DSL and a standard *command-query* [20,53,103] *Application Programming Interface* (API) [104], Fowler cites the need for a DSL to be fluent [24]. The command-query API provides the vocabulary; the DSL provides the grammar for composing the vocabulary “words” into “sentences”.

The implementation of a DSL is often supported by an API with a *fluent interface* [21,24,101], a vocabulary of operations designed to be composed smoothly into larger operations.

TODO: Point out which of my DSLs above have a fluent interface.

1.4 Shallow and Deep Embeddings of Internal DSLs

An internal DSL may be implemented with a shallow embedding or a deep embedding. The difference between shallow and deep embeddings concerns the relationship between the implementations of a DSL’s syntax and its semantics.

1.4.1 Shallow embedding

In a *shallow embedding* of an internal DSL, the implementation's types and data structures directly represent a single interpretation of the semantics of the domain but do not represent the syntactic structure of the domain objects [29,73].

For example, the regular expression package from Thompson [73:12.3] is a shallow embedding of the regular expression concept in Haskell. It models the semantics but not the syntax of regular expressions. It uses functions to represent the regular expressions and higher order functions (combinators) to combine the regular expressions in valid ways.

Similarly, the Scala-based Computer Configuration and Email Message Building DSLs and the Lua-based internal Lair DSLs are relatively shallow embeddings of the DSLs.

An advantage of a shallow embedding is that it provides a simple implementation of the semantics of the domain. It is usually straightforward to modify the semantics by adding new operations. If these capabilities are all that one needs, then a shallow embedding is convenient.

A disadvantage is that it is sometimes difficult to relate what happens during execution to the syntactic structure of the program, especially when errors occur.

1.4.2 Deep embedding

In a *deep embedding* of an internal DSL, the implementation's types and data structures model both the syntax and semantics of the domain. That is, it represents the domain objects using *abstract syntax trees (ASTs)* [29,73]. These can be interpreted in multiple ways as needed.

For example, Thompson [73:19.4] redesigns the Haskell regular expression package as a deep embedding. It introduces types that represent the syntactic structure of the regular expressions as well as their semantics.

The advantage of a deep embedding is that, in addition to manipulating the semantics of the domain, one can also manipulate the syntactic representation of the domain objects. The syntactic representation can be analyzed, transformed, and translated in a many ways that are not possible with a shallow embedding.

As an example, consider the deep embedding of the regular expression DSL. It can enable replacement of one regular expression by an equivalent simpler one, such as replacing $(a^*)^*$ by a^* .

Of course, the disadvantages of deep embedding are that they are more complex to develop, understand, and modify than shallow embeddings.

The Haskell-based Sandwich DSL, Scala-based Sandwich DSL, and Haskell-based Exam DSL give other examples of how one can create a deeply embedded DSL in a simple situation.

TODO: Are others of my DSLs deeply embedded?

1.4.3 ELI Calculator Language

What about the ELI Calculator language in the textbook *Exploring Languages with Interpreters and Functional Programming* (ELIFP) [10]?

The ELI Calculator language’s concrete syntax is defined in Chapter 41, abstract syntax and semantics in Chapter 42, mapping from concrete syntax to abstract syntax (i.e., parsing) in Chapter 44, and compilation to a Stack Virtual Machine in Chapter 46. The programs that implement the language are implemented in Haskell.

The concrete syntax and semantics of the ELI Calculator language is different from its host language Haskell, so at that level it is an external DSL. The parsers recognize valid arithmetic expressions in the input text and create an appropriate abstract syntax tree inside the Haskell program. The abstract syntax tree differs from the textual arithmetic expression and from its parse tree.

However, the abstract syntax tree does capture the essential aspects of the syntax. And the abstract syntax tree itself can be considered a deep embedding of an internal DSL for the abstract syntax. For the remainder of the processing of the expression, the abstract syntax tree preserves the important syntactic (structural) features of the arithmetic expressions.

The ELI Calculator language case study transforms the abstract syntax trees by simplifying them (and by generating their symbolic derivatives) and by generating equivalent instruction sequences for a Stack Virtual Machine.

1.5 DSLs Used to Produce This Document

When we look around, we can see DSLs everywhere!

Consider how I create this document on DSLs. I write the text with the text editor Emacs, which includes a number of extensions, perhaps written in DSL(s). I indicate the document’s structure using a DSL, Pandoc’s dialect [46] of the markup language Markdown [31,97]. If I add mathematical notation to the document, I write these in a subset of the LaTeX [68,96]. I then execute the Pandoc tool [46] on the input file.

Pandoc (which itself is written in Haskell) reads the Markdown input, converts it to an abstract syntax tree (AST) internally, and then writes an appropriate HTML [78,95] (a DSL) output file (that you are likely reading). I also direct Pandoc to write a LaTeX (a DSL) output file, on which I execute the tool `pdflatex` to create a PDF document. I could generate documents in other standard formats such as Microsoft Word’s `.docx` format [12,100] (essentially a DSL) and EPUB [17,99] (a DSL).

For the HTML output, I direct Pandoc to generate MathML [75,95], a standard DSL in the XML family that describes mathematical expressions for display on

the Web. (It is currently supported by the FireFox browser but not all others, which is why I recommend viewing these documents with FireFox.)

If I wish to include a drawing of a graph structure in the document, I may express the graph in the DOT language [66,94] supported by the Graphviz package [67]. Then I can translate it into a Scalable Vector Graphics (SVG) [76] drawing, which is encoded with another dialect of XML.

I wrote the original version of this document directly in HTML before I started using Pandoc. Other documents used in my courses were written originally using LaTeX or Word. In some case, I used Pandoc to convert these documents to Markdown, which gave me the starting point for my recent changes.

To add a new input format to Pandoc requires a new *reader* program that can *parse* the input and generate an appropriate abstract syntax tree.

To add a new output format to Pandoc requires a new *writer* program that can access the abstract syntax tree and generate appropriately formatted output.

Pandoc's abstract syntax tree is made available to writers using either Haskell algebraic data types or JavaScript Object Notation (JSON) structures [8,11].

JSON is an external DSL in that it uses a subset of JavaScript's concrete syntax to express the structure of the data. But, because it is JavaScript code, it also defines an equivalent internal data structure, so it also has aspects of an internal DSL. Pandoc could use a JSON Schema [40] to define the supported format. A JSON Schema is a JSON document with a specific format (a DSL) that defines the format of other JSON documents (other DSLs).

The whole conversion process in Pandoc revolves around the abstract syntax tree. Pandoc enables users to write their own filters that transform one Pandoc AST to another.

Thus my workflow uses many DSLs directly or indirectly.

1.6 Possible Advantages of Using DSLs

Fowler and others give several possible advantages for using DSLs. These include:

1. *DSLs can facilitate communication between domain experts and the programmers* [24].

A DSL may be a small, simple language that can express important aspects of the domain in a manner that nonprogrammers can *read* and understand.

In some cases, users may be able to *write* the DSL programs, allowing them to adapt the application to their specific needs without the intervention of programmers. But designing a DSL that users can write effectively is considerably harder than one they can read effectively.

This is an attack on the effects of the essential complexity of software development [4,5].

2. *DSLs can help encode domain knowledge in concise, concrete forms readable by humans* [63].

This seeks to mitigate the problem of deciding what to build, lessening the effects of the essential complexity of software development [4]. It also enables this knowledge to be reused more readily.

3. *DSLs and their implementations can increase opportunities for software reuse* [52].

An implementation of a DSL may generate code that incorporate key data types, software architectures, algorithms, and other domain concepts and processes. These are reused from one use of the DSL to another.

Software reuse is an attack on the effects of the essential complexity of software development [4].

4. *DSLs can improve programmer productivity* [24].

Using a DSL for some narrow, well-defined aspect, programmers can code the computation more quickly and reliably with the DSL than directly in the GPL.

For example, many GPLs have sublanguages for specifying regular expressions. In most cases, programmers can use these regular expression DSLs more productively than programming the pattern recognition directly in the GPL. Also, a mature implementation of the DSL will likely be more reliable than a new implementation directly in the GPL.

5. *DSLs can help shift the execution context between compile time and runtime* [24].

On the one hand, if some aspect of an application needs to be more flexible, we might replace a description in the GPL program (e.g., populating a complex data structure) with a description in a DSL that is interpreted at runtime.

On the other hand, if some aspect of an application needs to be more efficient, we might replace a description of the runtime configuration code with a description in a DSL that is compiled into efficient GPL code.

6. *DSLs can enable convenient use of alternative computational models that are not natively supported by the GPL used in the overall application* [24].

For example, some aspect of the computation might better be expressed as a finite state machine, decision table, dependency network, rule-based system, etc. We can embed such computations within traditional imperative or object-oriented programs by designing appropriate DSLs.

7. *DSLs can promote porting aspects of the application code to a different execution platform* [77].

The front end phases of the DSL need not be changed, just the back end that does code generation or interpretation.

Note: Fowler argues this is a benefit of having a semantic model, not necessarily a DSL.

1.7 Possible Disadvantages of Using DSLs

Fowler also identifies several possible disadvantages of using DSLs [24]. These include:

1. *DSLs can contribute to the **language cacophony**.*

If many, different, difficult-to-learn languages are used in an application or an organization, then this creates considerable work for the software developers to learn them all.

Fortunately, DSLs are usually much easier to learn than GPLs. Also the abstractions represented in the DSLs likely will need to be present in the libraries, APIs, documentation, and tools regardless. Thus using DSLs might not incur as much language learning as it first appears.

2. *DSLs can be costly to build and maintain.*

A DSL is usually built on top of a library, API, or software framework. Designing, implementing, and maintaining the DSL will require some additional work.

Fortunately, DSLs are usually simple and once the developers master the language development tools, designing and implementing a DSL can be done without a huge investment.

As with any tool development or acquisition, the software development organization must decide if the possible benefits are greater than the costs.

3. *DSL can contribute to the **ghetto language problem**.*

If a language uses many languages that are used nowhere else, then it can become difficult to recruit staff.

This can be a significant problem for use of a GPL. (However, some software development organizations choose languages outside the mainstream so they can attract the more aggressive staff willing to take on interesting technical challenges and use leading edge tools.)

But DSLs should be small and limited to a narrow domain, so the problem should not be as significant as for GPLs. The organization should guard against rampant “mission creep” for its DSLs.

4. *DSLs can sometimes lead to narrow thinking.*

The intention of DSL development is to open up the developers to using whatever abstractions are appropriate to the domain, rather than those that are convenient in the GPL.

Organizations that use DSLs should avoid falling back into the same trap with their DSLs. They should develop appropriate new DSLs when needed rather than use an existing DSL with inappropriate abstractions.

1.8 Designing DSLs

In a chapter in the functional programming notes, we examine *families* of related functions to define generic, higher-order functions to capture the computational patterns for each family. We seek to raise the level of abstraction in our programs.

Design of DSLs is similar, except that we seek to design a language to express the family members in some application domain rather than design a higher-order function.

1.8.1 SCV analysis

We should first analyze the domain systematically and then use the results to design an appropriate DSL syntax and semantics. We analyze the domain using *Scope-Commonality-Variability (SCV) [7] and produce four outputs.

1. *scope* – the boundaries of the domain. That is, identify what we must address and what we can ignore.
2. *terminology* – the definitions of the specialized terms, or concepts, relevant to the domain.
3. *commonalities* – the aspects of the domain that do not change from one application to another within the domain. We sometimes call these the *frozen spots*.
4. *variabilities* – the aspects of the domain that may change from one application to another within the domain. We sometimes call these the *hot spots*.

TODO: Expand the explanation (e.g., include model) to include more of the ideas from [7] and other sources. Perhaps include example such as the SurveyDSL design from [9].

In the SCV analysis, we must seek to identify all the implicit assumptions in the application domain. These implicit assumptions need to be made explicit in the DSL's design and implementation.

We use the SCV analysis to guide our choices for elements of the DSL design [9]. The scope focuses our attention on what we are trying to accomplish. The terminology and commonalities suggest the DSL statements and constructs. The commonalities also suggest the semantics of the constructs and the nature of the

underlying computational model. The variabilities represent syntactic elements to which the DSL programmer can assign values.

1.8.2 DSL design guidelines

TODO: Better integrate the Karsai, Freeman, and other lists of guidelines.

Karsai et al. [41] identifies 26 guidelines important for DSL design, grouping them into 5 categories. The paper focuses on design of external DSLs. Here we expand their guidelines to include internal DSLs.

A. Language purpose guidelines (i.e., what purposes to satisfy)

1. *Identify language uses early.*

In Fowler’s terminology, we must identify the domain and what uses the language will have within the domain. We must carefully define the scope within the SCV analysis.

2. *Ask questions.*

What group of users will write the DSL programs? will read them? will deploy them for execution? Etc. What are each group’s purposes? What does each group need to be able to understand and use the DSL successfully? Can we simplify the DSL further?

3. *Make the language consistent.*

Avoid surprises. Keep the DSL narrowly focused on its purposes. A DSL feature should contribute to the purposes or be omitted. All features should be based on a cohesive set of concepts.

B. Language realization guidelines (i.e., how to implement)

4. *Decide carefully whether to use a text-based external DSL, a graphical external DSL, or an internal DSL hosted in some particular language.*

For the DSL’s identified domain, uses, and user groups, what are the advantages and disadvantages of each approach? What tools are available to support the DSL design, implementation, and use.

5. *Compose existing languages where possible.*

Can parts of the new DSL’s uses be handled by already implemented languages and tools? If so, we can avoid the time-consuming and error-prone work of designing and implementing a whole new DSL. We can combine the existing languages, embed them within a new “glue” language, extend an existing language with a few new features, etc.

In the “Little Languages” paper, Bentley describes how the processor for the external DSL `chem` generates a program in the DSL `pic`. The processor for `pic` itself generates a program in the DSL `troff`. These

“filter” programs are then connected using pipes in the Unix shell (a DSL). Furthermore, the implementation of `pic` specifies the lexical analysis and parsing phases using the DSLs `lex` and `yacc` and defines the overall build process using the DSL `make`.

Consider an internal DSL such as the Survey DSL. It includes many existing features of the host language (Ruby) in the new DSL.

6. *Reuse existing language definitions.*

Even if the implementation of an existing language cannot be reused, we can consider reusing its definition. This saves effort in language design and it may leverage the users’ knowledge of the existing language.

For example, the Pandoc Markdown dialect embeds LaTeX’s widely-known mathematical notation to specify mathematical symbols and expressions within the text.

Sometimes building a DSL entails embedding an existing API within a fluent interface to implement an internal DSL or devising a similar textual notation to form an external DSL [52]. Or perhaps we reenvision a textual or internal DSL as a graphical DSL.

7. *Reuse existing type systems.*

A language’s type system is probably the most difficult to design well and implement robustly. A new type system can also be difficult for users to learn to use effectively. Thus, by reusing an existing type system that the users may know, DSL developers can both make their work more efficient and the new DSL easier to understand and use.

An internal DSL selects from and builds on the host language’s existing type system.

C. Language content guidelines (i.e., what features to include)

8. *Reflect only necessary domain concepts.*

Which artifacts (or objects) from the domain must we capture to satisfy the DSL purposes? Which properties of those artifacts? Can we leave out the other artifacts and purposes? We should discuss possible designs with users and incorporate their feedback.

9. *Keep the DSL simple.*

Make the DSL easy for users to learn and use effectively. If the language is too complex, users may just ignore it. Keep the language as simple as possible to ensure effective use.

Simplicity may be especially difficult to achieve in internal DSLs because the boundary between the DSL features and the host language may not be clear. Seek to crisply define this boundary.

The next three guidelines help us achieve simplicity.

10. *Avoid unnecessary generality.*

This is an aspect of keeping the DSL simple. Design only what is necessary to solve the problem. Avoid excessive concern about generalizing and parameterizing the language beyond what is needed initially. It is difficult to predict what generalizations will actually prove useful.

However, a successful DSL is seldom static. It likely must evolve to meet the changes in the domain and the expectations of its users. Design the DSL so that it can be extended with new capabilities in the future.

An SCV analysis can reveal important “variabilities” (hot spots) that either should be incorporated in the initial design or that may become important in the future.

11. *Limit the number of language elements.*

This is a second aspect of keeping the DSL simple. A language with many elements is difficult to learn, use, and implement. It is better to have a few elements that can be combined flexibly.

If the domain and purposes are complex, look for ways to break them into a set of smaller problems, solve each subproblem by designing a sublanguage, and then combine the sublanguages to solve the larger problem. Users can focus on the sublanguages they need to carry out their specific tasks.

Alternatively, determine whether some of the more complex elements can be moved from the DSL’s core and to a “library” that can be accessed by DSL programs. Enable users to store their own language extensions in the library. This approach enables us to extend the functionality of the DSL without changing the core language’s structure.

For example, languages like Pascal included I/O statements as a language construct. Later languages such as C and Java moved I/O to a library.

12. *Avoid conceptual redundancy.*

This is a third aspect of keeping the DSL simple. If there are many possible ways to express the same concept in the DSL, users will likely become confused and may not use the DSL effectively. Avoid unnecessary redundancy.

13. *Avoid inefficient language elements.*

A DSL raises the level of abstraction and usually obscures the details of how a DSL program is actually executed. However, the DSL designers and implementers must ensure that the DSL programs execute with acceptable efficiency. Moreover, the user of the DSL must be able to understand which DSL programs are more efficient than others.

D. Concrete syntax guidelines (i.e., how to make the DSL readable)

14. *Adapt existing notations domain experts use.*

Where possible, build on the formal notation that the domain experts already know and use. If this needs to be modified or extended, keep the changes close to the style of the existing notation. Familiarity will make the DSL easier for domain experts to learn and

Of course, DSL designers may need to formalize the syntax and semantics of informal terminology, notation, and processes to enable them to be included in an automated DSL. Tools for processing the new DSL may also provide capabilities not present in the current practice.

15. *Use descriptive notations.*

Where possible, choose terms and symbols that suggest the intended meaning. Avoid using them in ways that differ significantly from the way they are used in the domain or in the general public.

For example, the symbol + usually denotes addition or some similar operation; to use it to denote multiplication would likely introduce confusion and make the DSL difficult to learn. Similarly, using the keyword if to denote something other than a conditional would be confusing.

16. *Make elements distinguishable.*

A DSL will be read by humans more frequently than written. So, when the needs of readers and writers conflict, favor the readers over the writers.

Make different elements appear differently in the displayed form. Avoid using a subtle difference in notation, spelling, location, size, font, or color as the sole way to distinguish between different elements. In most cases, make the DSL accessible to those with impaired vision.

17. *Use syntactic sugar appropriately.*

Syntactic sugar refers to elements of a language that do not add to the expressiveness of the DSL but which may make it easier to read and perhaps easier to parse.

If used in moderation, syntactic sugar makes the language more palatable. But overuse may just make the language fat -- more

verbose and confusing.

This guideline conflicts with the “avoid conceptual redundancy” guideline above. Designers must balance these concerns.

18. *Permit comments.*

Although comments do not make the DSL more semantically expressive, comments enable those writing DSL programs to explain their design decisions to others (or their future selves) who need to understand and modify the DSL program.

In addition, comments can allow information to be passed to tools that generate structured documentation on the DSL programs. (Consider JavaDoc.)

19. *Provide organizational structures for DSL programs*

To manage a DSL program that grows large and complex, we need to be able to break it up into subprograms. The subprograms may need to be organized into a graph structure and managed as a group in an archive.

The DSL and the tools that manage DSL programs should allow a group of subprograms to be organized into “packages” that can be selectively included in other DSL programs.

This solution seeks to manage complexity of DSL programs by breaking them into smaller subprograms. The “library” solution suggested for the “limit the number of language elements” guideline seeks to manage the complexity of the language itself.

20. *Balance compactness and comprehensibility.*

Compact notation is generally efficient to write and process. However, it may not be comprehensible to the reader.

Syntactic sugar can make a DSL more comprehensible, but make it more difficult to write. And, if overused, it can make the DSL confusingly verbose and thus less comprehensible.

So we must balance among the factors. As noted above, we should normally favor the human reader over the writer.

21. *Use the same style everywhere.*

If a language has several sublanguages, make all the sublanguages similar in style. This makes the various languages easier to understand and use as a group.

For example, it would be confusing for one sublanguage to group items with { and } and another to use `begin` and `end` for a similar purpose.

Of course, this guideline must be balanced with the above guidelines suggesting we should “compose existing languages.” “reuse existing language definitions,” and “adapt existing notations domain experts use.”

22. *Identify usage conventions.*

To keep the DSL definition simple, we generally should not rigidly enforce minor syntactic issues such as layout. However, good DSL programming style can make a program more comprehensible (and more pleasant to read).

In parallel with the language definition, we should describe good style means for layout, naming conventions, order of elements, commenting, etc.

E. Abstract syntax guidelines (i.e., how to represent the DSL internally)

23. *Align abstract and concrete syntax*

This means that:

- elements that differ in concrete syntax should differ in internal representations
- elements that are similar in *meaning* should have similar internal representations (e.g., be subclasses of same base class)
- the internal representation of an element should not be dependent upon the context in which the element appears

The goal is to be able to map the human readable representation of the DSL to the internal representation. This makes the execution easier to understand and debug. It may enable the generation of runtime error messages that tie back to the DSL program. (Remember the discussion of shallow versus deep embedding of DSLs.)

24. *Prefer layout that does not affect translation from concrete to abstract syntax.*

This guideline suggests that issues like indentation should not affect the semantics. Not all language designers (e.g., of Python and Haskell) agree with this guideline.

25. *Enable modularity.*

To manage the complexity of the language, the “limit the number of language elements” guideline suggests providing a library of non-primitive elements that extend the core language.

To manage the complexity of large DSL programs, the “organizational structure for DSL programs” guideline suggests breaking a program into a group of related pieces and storing the pieces in a library of packages. This can be done at the DSL source code level.

This “enable modularity” guideline suggests the capability to build the internal representation of a DSL program by composing separately compiled “modules” incrementally.

Perhaps all of these mechanisms are implemented by a single module mechanism or there may separate mechanisms for composing precompiled modules, DSL source code packages, and DSL extensions.

26. *Introduce interfaces.*

An *interface* usually defines a set of operation signatures (the name, parameters and their types, and return value type) and perhaps constraints on the operation’s execution (preconditions and postconditions).

If the modules of a large DSL program have well-defined interfaces, then the DSL’s module mechanism can check whether the modules conform with each other’s expectations.

Interfaces and modules together support an information-hiding approach to program development. Each module “hides” its internal details from the other modules in the system.

1.8.3 More guidelines for internal DSL design

TODO: Better integrate the Karsai, Freeman, and other lists of guidelines.

Drawing on the experience in designing, implementing, and evolving the JMock internal DSL (which provides mock objects for testing [25]), Freeman and Pryce make four recommendations for constructing an internal DSL [25]. To some extent, these recommendations apply to design of all DSLs.

27. *Separate syntax and semantics (interpretation) into separate layers.*

The concern of the syntax layer is to provide a fluent, readable language for users familiar with the application domain. These may not be programmers—or at least not programmers who are experts in the host language.

The concern of the semantic layer is to provide a correct, efficient, and maintainable interpreter for the language. The developers and maintainers of this layer are typically experts in the host language who can take advantage of the implementation language’s capabilities and idioms.

As with most nontrivial software development tasks, mixing these concerns can make the implementation difficult to develop, understand, and maintain. It is better to translate the syntax to to an appropriate semantic model (e.g., an abstract syntax tree) for processing, hiding the details of each behind well-designed interfaces.

28. *Use, and perhaps abuse, the host language and its conventions to enable the writing of readable DSL programs.*

For internal DSLs, the syntax layer may need to violate the conventional programming styles and naming conventions of the host language to achieve the desired readability and fluency.

For example, DSLs in object-oriented languages may use *method chaining* [24,102] extensively to achieve the desired fluency of the interface [21,101]. These practices usually discouraged in the usual programming practices.

29. *Don't trap the user in the internal DSL.*

Note: This concern of this guideline is similar to discussion of libraries, packages, and modules in the Karsai et al guidelines in a previous subsection.

The DSL should encapsulate its internal implementation details to avoid unexpected dependence on implementation features that might change over time and to enable naive users to use the DSL safely.

However, it is difficult to anticipate all possible uses of a DSL over time. Thus it is helpful to enable expert users of the host language to extend the DSL by providing alternative implementations of key abstractions.

The implementation of the DSL should itself be approached as a software family. Although the DSL syntax may look different than the host language, the DSL implementation should seek to work seamlessly with other host language programs.

We should track the changes that expert users make. These help identify possible future enhancements of the DSL and its implementation.

30. *Map error reports to the syntax layer rather than to the semantics layer, which is hidden from the DSL user.*

Note: The concern of this guideline is similar to that of Karsai et al's "align concrete and abstract syntax" guideline in a previous subsection.

Good error reports are critical to a successful DSL. As much as possible, errors in both syntax and semantics should be stated in terms of the syntactic structure of the specific DSL program. This is often difficult to accomplish, but it is important because it is unlikely that the DSL's users will be familiar with the internal details of the implementation.

Deep embedding of DSLs can make it easier to trace errors back to recognizable syntactic structures.

1.9 What Next?

TODO

TODO: Is a source code section needed?

1.10 Exercises

TODO

1.11 Acknowledgements

These notes evolved during the 2014-2018 period. They draw ideas from (in alphabetical order) Bentley [3], Buck [6], Coplien [7], Fowler [18,19,22–24], Freeze [26], Hudak [35,36], Karsai et al. [41], Mernik et al. [52], Thompson [73], Freeman and Pryce [25], van Deursen et al. [74], my own research [9] related to the Ruby-based Survey DSL, my experiences reimplementing several of Fowler’s DSLs in Scala, Lua, or Python, various Wikipedia articles, and likely other sources.

In Fall 2014, I wrote the first short version of these notes to accompany the Haskell-based Sandwich DSL project I assigned to my CSci 450 (Organization of Programming Languages) class.

In Spring 2016, I updated these notes slightly to accompany the Scala-based Sandwich DSL project I assigned to my CSci 555 (Functional Programming) class.

In Fall 2016, I adapted and revised much of this document for possible use in CSci 450, but I did not use it that semester.

In Spring 2017 for a Haskell-based offering of CSci 556 (Multiparadigm Programming), I continued to develop these notes, adding discussion of the boundaries between DSLs and other computing artifacts, of the Expression Language (later called ELI Calculator language) case study, and of my use of DSLs in the workflow for producing this document.

In Summer and Fall 2017, I updated these notes slightly for possible use in CSci 450, but did not use it that semester.

In Spring 2018, I added discussion of Fowler’s definition of DSLs [24], advantages and disadvantages of DSLs [24], and Karsai’s DSL design guidelines [41] and modified the document for use in CSci 658 (Software Language Engineering).

I retired from the full-time faculty in May 2019. As one of my post-retirement projects, I am continuing work on possible textbooks based on the course materials I had developed during my three decades as a faculty member. In January 2022, I began refining the existing content, integrating separately developed materials together, reformatting the documents, constructing a unified bibliography (e.g., using citeproc), and improving my build workflow and use of Pandoc.

I maintain this chapter as text in Pandoc’s dialect of Markdown using embedded LaTeX markup for the mathematical formulas and then translate the document to HTML, PDF, and other forms as needed.

1.12 Concepts and Terms

TODO: Update this list

Domain-specific languages (DSLs); language nature; fluency; limited expressiveness; domain focus; DSLs versus general-purpose programming languages; DSLs versus APIs; external versus internal DSLs; shallow versus deep embedding of internal DSLs; use of algebraic data types to implement DSLs

1.13 References

- [1] Alfred V. Aho, Brian W. Kernighan, and Peter J. Weinberger. 1979. Awk—a pattern scanning and processing language. *Software: Practice and Experience* 9, 4 (1979), 267–279.
- [2] Conrad Barski. 2011. *Land of Lisp: Learn to program in Lisp, one game at a time*. No Starch Press, San Francisco, California, USA.
- [3] Jon Bentley. 1986. Programming pearls: Little languages. *Communications of the ACM* 29, 8 (1986), 711–721.
- [4] Frederick P. Brooks, Jr. 1987. No silver bullet: Essence and accident in software engineering. *IEEE Computer* 20, 4 (1987), 10–19.
- [5] Frederick P. Brooks, Jr. 1995. 'No Silver Bullet' refired. In *The mythical man month* (Anniversary). Addison-Wesley, Boston, Massachusetts, USA.
- [6] Jamis Buck. 2006. Writing domain specific languages (blog post). Retrieved from <https://weblog.jamisbuck.org/2006/4/20/writing-domain-specific-languages.html>
- [7] James Coplien, Daniel Hoffman, and David Weiss. 1998. Commonality and variability in software engineering. *IEEE Software* 15, 6 (1998), 37–45.
- [8] Douglas Crockford. 2022. *Introducing JSON*. Retrieved from <https://www.json.org/json-en.html>
- [9] H. Conrad Cunningham. 2008. A little language for surveys: Constructing an internal DSL in Ruby. In *Proceedings of the ACM SouthEast conference*, Auburn, Alabama, USA, 282–287.
- [10] H. Conrad Cunningham. 2022. *Exploring programming languages with interpreters and functional programming (ELIFP)*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from <https://john.cs.olemiss.edu/~hcc/docs/ELIFP/ELIFP.pdf>
- [11] Ecma International. 2017. ECMA-404 standard: The JSON data interchange format. Retrieved from <https://www.ecma-international.org/publications-and-standards/standards/ecma-404/>
- [12] Ecma International. 2021. ECMA-376 Standard: Office open XML file formats. Retrieved from <https://www.ecma-international.org/publications-and-standards/standards/ecma-376/>

- [13] Elixir Team. 2022. Elixir. Retrieved from <https://elixir-lang.org>
- [14] Stuart I. Feldman. 1979. Make—a program for maintaining computer programs. *Software: Practice and Experience* 9, 4 (1979), 255–265.
- [15] Matthias Felleisen, David Van Horn, and Conrad Barski. 2013. *Realm of Racket: Learn to program, one game at a time!* No Starch Press, San Francisco, California, USA.
- [16] Michael Fogus and Chris Houser. 2011. *The joy of Clojure*. Manning, Shelter Island, New York, USA.
- [17] International Digital Publishing Forum. 2022. EPUB 3.2. Retrieved from <https://www.w3.org/publishing/epub32/epub-spec.html>
- [18] Martin Fowler. 2005. Language workbenches: The killer-app for domain specific languages? (Blog post). Retrieved from <http://www.martinfowler.com/articles/languageWorkbench.html>
- [19] Martin Fowler. 2005. Generating code for DSLs (blog post). Retrieved from <https://www.martinfowler.com/articles/codeGenDsl.html>
- [20] Martin Fowler. 2005. CommandQuerySeparation (blog post). Retrieved from <https://www.martinfowler.com/bliki/CommandQuerySeparation.html>
- [21] Martin Fowler. 2005. FluentInterface (blog post). Retrieved from <https://martinfowler.com/bliki/FluentInterface.html>
- [22] Martin Fowler. 2008. One lair and twenty Ruby DSLs. In *The Thought-Works anthology: Essays on software technology and innovation*, Thought-Works, Inc. (ed.). Pragmatic Bookshelf, Raleigh, North Carolina, USA. Retrieved from https://media.pragprog.com/titles/twa/martin_fowler.pdf
- [23] Martin Fowler. 2008. DomainSpecificLanguage (blog post). Retrieved from <https://www.martinfowler.com/bliki/DomainSpecificLanguage.html>
- [24] Martin Fowler and Rebecca Parsons. 2010. *Domain specific languages*. Addison-Wesley, Boston, Massachusetts, USA.
- [25] Steve Freeman and Nat Pryce. 2006. Evolving an embedded domain-specific language in Java. In *Companion to the conference on object-oriented programming languages, systems, and applications*, Portland, Oregon, USA, 855–865. Retrieved from <http://jmock.org/oopsla2006.pdf>
- [26] Jim Freeze. 2006. Creating DSLs with Ruby. Retrieved from <https://www.artima.com/articles/creating-dsls-with-ruby>
- [27] Daniel P. Friedman and Matthias Felleisen. 1995. *The little schemer* (Fourth ed.). MIT Press, Cambridge, Massachusetts, USA.
- [28] Daniel P. Friedman and Matthias Felleisen. 1995. *The seasoned schemer* (Second ed.). MIT Press, Cambridge, Massachusetts, USA.

- [29] Jeremy Gibbons and Nicolas Wu. 2014. Folding domain-specific languages: Deep and shallow embeddings (functional pearl). In *Proceedings of the 19th ACM SIGPLAN international conference on functional programming*, Gothenburg, Sweden, 339–347. Retrieved from <https://www.cs.ox.ac.uk/jeremy.gibbons/publications/embedding.pdf>
- [30] David Goodger. 2022. *reStructuredText: Markup syntax and parser component of Docutils*. Docutils Project. Retrieved from <https://docutils.sourceforge.io/rst.html>
- [31] John Gruber. 2022. Markdown. Retrieved from <https://daringfireball.net/projects/markdown/>
- [32] David Hagen. 2022. Parsita. Retrieved from <https://pypi.org/project/parsita/>
- [33] Haskell Organization. 2022. Parsec: Monadic parser combinators. Retrieved from <https://hackage.haskell.org/package/parsec>
- [34] Rich Hickey. 2020. A history of Clojure. *Proceedings of the ACM on Programming Languages* 4, HOPL, Article 71 (2020), 1–46.
- [35] Paul Hudak. 1996. Building domain-specific embedded languages. *ACM Computing Surveys* 28, 4 (1996), 196–201. Retrieved from <https://dl.acm.org/doi/fullHtml/10.1145/242224.242477>
- [36] Paul Hudak. 1998. Modular domain specific languages and tools. In *Proceeding of the 5th international conference on software reuse (ICSR'98)*, IEEE, Victoria, British Columbia, Canada, 134–142.
- [37] Paul Hudak, John Peterson, Joseph H. Fasel, and Reuben Thomas. 2000. A gentle introduction to Haskell 98. Retrieved from <https://www.haskell.org/tutorial/>
- [38] Roberto Ierusalimschy. 2009. A text pattern-matching tool based on parsing expression grammars. *Software: Practice and Experience* 39, 3 (2009), 221–258. Retrieved from <http://www.inf.puc-rio.br/~roberto/lpeg/lpeg.html>
- [39] Roberto Ierusalimschy. 2022. LPeg: Parsing expression grammars for Lua, version 1.0. Retrieved from <http://www.inf.puc-rio.br/~roberto/lpeg/>
- [40] json-schema.org. 2022. JSON Schema. Retrieved from <https://json-schema.org>
- [41] Gabor Karsai, Holger Krahn, Claas Pinkernell, Bernhard Rumpe, Martin Schindler, and Steven Voelkel. 2009. Design guidelines for domain specific languages. In *Proceedings of OOPSLA workshop on domain-specific modeling*, Vancouver, British Columbia, Canada. Retrieved from <https://arxiv.org/pdf/1409.2378>
- [42] Brian W. Kernighan. 1984. *PIC—a graphics language for typesetting, revised user manual*. Bell Laboratories, Computing Science, Murray Hill, New Jersey, USA. Retrieved from <http://doc.cat-v.org/unix/v8/picmemo.pdf>

- [43] Michael E. Lesk and Eric Schmidt. 1975. *Lex: A lexical analyzer generator*. Bell Laboratories, Murray Hill, New Jersey, USA. Retrieved from <http://mbzz.de/unix/ultrix/etc/lex.pdf>
- [44] John R. Levine and Tony Mason and Doug Brown. 1992. *Lex & yacc* (Second ed.). O'Reilly Media, Sebastopol, California, USA.
- [45] Peter Linz. 2011. *Formal languages and automata* (Fifth ed.). Jones & Bartlett, Burlington, Massachusetts, USA.
- [46] John MacFarlane. 2022. Pandoc: A universal document converter. Retrieved from <https://www.pandoc.org>
- [47] Simon Marlow. 2022. Alwx: A lexical analyzer generator for Haskell. Retrieved from <https://www.haskell.org/alex/>
- [48] Simon Marlow. 2022. Happy: The parser generator for Haskell. Retrieved from <https://www.haskell.org/happy/>
- [49] John McCarthy. 1978. History of LISP. *ACM SIGPLAN Notices* 8 (1978), 217–223. Retrieved from <https://pages.cs.wisc.edu/~horwitz/CS704-NOTES/PAPERS/mccarthy.pdf>
- [50] John McCarthy, Paul W. Abrahams, Daniel J. Edwards, Timothy P. Hart, and Michael I. Levin. 1962. *LISP 1.5 programmer's manual*. MIT Press, Cambridge, Massachusetts, USA. Retrieved from <https://apps.dtic.mil/sti/pdfs/AD0406138.pdf>
- [51] Robert Mecklenburg. 2004. *Managing projects with GNU Make: The power of GNU Make for building anything* (Third ed.). O'Reilly Media.
- [52] Marjan Mernik, Jan Heering, and Anthony M. Sloane. 2005. When and how to develop domain specific languages. *ACM Computing Surveys* 37, 4 (2005), 316–344. Retrieved from <https://pdfs.semanticscholar.org/bd06/01088d5f217dc136a898f577763df92891cb.pdf>
- [53] Bertrand Meyer. 1997. *Object-oriented program construction* (Second ed.). Prentice Hall, Englewood Cliffs, New Jersey, USA.
- [54] Martin Odersky, Lex Spoon, and Bill Venner. 2021. *Programming in Scala* (Fifth ed.). Artima, Inc., Walnut Creek, California, USA.
- [55] Terence Parr. 2009. *Language implementation patterns: Create your own domain-specific and general programming languages*. Pragmatic Bookshelf, Raleigh, North Carolina, USA.
- [56] Terence Parr. 2022. ANTRL. Retrieved from <https://wwwantlr.org/>
- [57] PLT Inc. 2022. Racket. Retrieved from <https://www.racket-lang.org>
- [58] Christian Queinnec. 2003. *Lisp in small pieces*. Cambridge University Press, Cambridge, UK.
- [59] Eric S. Raymond. 1995. Making pictures with GNU PIC. Retrieved from <https://lists.gnu.org/r/groff/2011-08/pdfbLauVhlfQs.pdf>

- [60] Arnold Robbins. 2015. *Effective AWK programming: : Universal text processing and pattern matching* (Fourth ed.). O’Reilly Media, Sebastopol, California, USA.
- [61] Ruby Community. 2022. Ruby: A programmer’s best friend. Retrieved from <https://www.ruby-lang.org>
- [62] Scala Language Organization. 2022. The Scala programming language. Retrieved from <https://www.scala-lang.org/>
- [63] Diomidis Spinellis. 2001. Notable design patterns for domain-specific languages. *Journal of Systems and Software* 56, 1 (2001), 91–99. Retrieved from <https://www.spinellis.gr/pubs/jrnl/2000-JSS-DSLPatterns/html/dslpat.html>
- [64] Richard M. Stallman and Zachary Weinberg. 2022. The c preprocessor. Retrieved from <https://gcc.gnu.org/onlinedocs/cpp/>
- [65] Johnson Stephen C. 1975. *Yacc: Yet another compiler-compiler*. Bell Laboratories, Murray Hill, New Jersey, USA. Retrieved from <http://dinosaur.compilertools.net/yacc/>
- [66] The Graphviz Authors. 2022. DOT Language. Retrieved from <https://graphviz.org/doc/info/lang.html>
- [67] The Graphviz Authors. 2022. Graphviz. Retrieved from <https://learnxinyminutes.com/docs/haskell/>
- [68] The LaTeX Project. 2022. LaTeX: A document preparation system. Retrieved from <https://www.latex-project.org/>
- [69] The R Foundation. 2022. R: The R project for statistical programming. Retrieved from <https://www.r-project.org/>
- [70] Scott Thibault, Renaud Marlet, and Charles Consel. 1999. Domain-specific languages: From design to implementation—Application to video device driver generation. *IEEE Transactions on Software Engineering* 25, 3 (1999), 363–377.
- [71] Dave Thomas. 2018. *Programming Elixir >= 1.6: Functional |> concurrent |> pragmatic |> fun*. Pragmatic Bookshelf, Raleigh, North Carolina, USA.
- [72] David Thomas, Chad Fowler, and Andrew Hunt. 2004. *Programming Ruby* (Second ed.). Pragmatic Bookshelf, Raleigh, North Carolina, USA.
- [73] Simon Thompson. 2011. *Haskell: The craft of programming* (Third ed.). Addison-Wesley, Boston, Massachusetts, USA.
- [74] Arie van Deursen, Paul Klint, and Joost Visser. 2000. Domain specific languages: An annotated bibliography. *ACM SIGPLAN Notices* (2000). Retrieved from <https://dl.acm.org/doi/pdf/10.1145/352029.352035>
- [75] W3C Math Working Group. 2022. W3C math home. Retrieved from <https://www.w3.org/Math/>

- [76] W3C SVG Working Group. 2022. Scalable vector graphics (SVG). Retrieved from <https://www.w3.org/Graphics/SVG/>
- [77] Martin P. Ward. 1994. Language-oriented programming. *Software–Concepts and Tools* 15, 4 (1994), 147–161. Retrieved from <http://www.gkc.org.uk/martin/papers/middle-out-t.pdf>
- [78] Web Hypertext Application Technology Working Group (WHATWG). 2022. HTML: Living standard. Retrieved from <https://html.spec.whatwg.org/>
- [79] Hadley Wickham. 2019. *Advanced R* (Second ed.). CRC Press. Retrieved from <https://adv-r.hadley.nz/index.html>
- [80] Wikipedia: The Free Encyclopedia. 2022. R (programming language). Retrieved from [https://en.wikipedia.org/wiki/R_\(programming_language\)](https://en.wikipedia.org/wiki/R_(programming_language))
- [81] Wikipedia: The Free Encyclopedia. 2022. Lisp (programming language). Retrieved from [https://en.wikipedia.org/wiki/Lisp_\(programming_language\)](https://en.wikipedia.org/wiki/Lisp_(programming_language))
- [82] Wikipedia: The Free Encyclopedia. 2022. Scheme (programming language). Retrieved from [https://en.wikipedia.org/wiki/Scheme_\(programming_language\)](https://en.wikipedia.org/wiki/Scheme_(programming_language))
- [83] Wikipedia: The Free Encyclopedia. 2022. Clojure. Retrieved from <https://en.wikipedia.org/wiki/Clojure>
- [84] Wikipedia: The Free Encyclopedia. 2022. Unix philosophy. Retrieved from https://en.wikipedia.org/wiki/Unix_philosophy
- [85] Wikipedia: The Free Encyclopedia. 2022. Turing completeness. Retrieved from https://en.wikipedia.org/wiki/Turing_completeness
- [86] Wikipedia: The Free Encyclopedia. 2022. General-purpose language. Retrieved from https://en.wikipedia.org/wiki/General-purpose_language
- [87] Wikipedia: The Free Encyclopedia. 2022. Domain-specific language. Retrieved from https://en.wikipedia.org/wiki/Domain-specific_language
- [88] Wikipedia: The Free Encyclopedia. 2022. PIC (markup lanugage). Retrieved from [https://en.wikipedia.org/wiki/PIC_\(markup_language\)](https://en.wikipedia.org/wiki/PIC_(markup_language))
- [89] Wikipedia: The Free Encyclopedia. 2022. AWK. Retrieved from <https://en.wikipedia.org/wiki/AWK>
- [90] Wikipedia: The Free Encyclopedia. 2022. Lex (software). Retrieved from [https://en.wikipedia.org/wiki/Lex_\(software\)](https://en.wikipedia.org/wiki/Lex_(software))
- [91] Wikipedia: The Free Encyclopedia. 2022. Yacc. Retrieved from <https://en.wikipedia.org/wiki/Yacc>
- [92] Wikipedia: The Free Encyclopedia. 2022. Make (software). Retrieved from [https://en.wikipedia.org/wiki/Make_\(software\)](https://en.wikipedia.org/wiki/Make_(software))
- [93] Wikipedia: The Free Encyclopedia. 2022. C preprocessor. Retrieved from https://en.wikipedia.org/wiki/C_preprocessor

- [94] Wikipedia: The Free Encyclopedia. 2022. DOT (graph description language). Retrieved from [https://en.wikipedia.org/wiki/DOT_\(graph_description_language\)](https://en.wikipedia.org/wiki/DOT_(graph_description_language))
- [95] Wikipedia: The Free Encyclopedia. 2022. HTML. Retrieved from <https://en.wikipedia.org/wiki/HTML>
- [96] Wikipedia: The Free Encyclopedia. 2022. LaTeX. Retrieved from <https://en.wikipedia.org/wiki/LaTeX>
- [97] Wikipedia: The Free Encyclopedia. 2022. Markdown. Retrieved from <https://en.wikipedia.org/wiki/Markdown>
- [98] Wikipedia: The Free Encyclopedia. 2022. reStructuredText. Retrieved from <https://en.wikipedia.org/wiki/ReStructuredText>
- [99] Wikipedia: The Free Encyclopedia. 2022. EPUB. Retrieved from <https://en.wikipedia.org/wiki/EPUB>
- [100] Wikipedia: The Free Encyclopedia. 2022. Office open XML (OOXML). Retrieved from https://en.wikipedia.org/wiki/Office_Open_XML
- [101] Wikipedia: The Free Encyclopedia. 2022. Fluent interface. Retrieved from https://en.wikipedia.org/wiki/Fluent_interface
- [102] Wikipedia: The Free Encyclopedia. 2022. Method chaining. Retrieved from https://en.wikipedia.org/wiki/Method_chaining
- [103] Wikipedia: The Free Encyclopedia. 2022. Command-query separation. Retrieved from https://en.wikipedia.org/wiki/Command-query_separation
- [104] Wikipedia: The Free Encyclopedia. 2022. API (application programming interface). Retrieved from <https://en.wikipedia.org/wiki/API>