

Seven Habits of a Highly Effective Smell Detector

Emerson Murphy-Hill and Andrew P. Black
Portland State University
emerson@cs.pdx.edu, black@cs.pdx.edu

ABSTRACT

The process of refactoring code — changing its structure while preserving its meaning — has been identified as an important way of maintaining code quality over time. However, it is sometimes difficult for programmers to identify which pieces of code are in need of refactoring. “Smell detectors” are designed to help programmers in this task, but most smell detectors do not mesh well with “floss refactoring,” the recommended tactic in which refactoring and programming are finely interleaved. In this paper we present a smell detector that we have built with floss refactoring in mind by combining seven habits that we postulate are important to consider when designing usable smell detectors. We hope that this combination can help the designers of future smell detectors build tools that align with the way that programmers refactor.

Categories and Subject Descriptors

D.2.3 [Software Engineering]: Coding Tools and Techniques
; D.2.6 [Software Engineering]: Programming Environments

General Terms

Design, Human Factors

Keywords

smells, refactoring, tools

1. INTRODUCTION

Refactoring is the practice of restructuring code without changing its externally observable behavior. In his influential book, Fowler advocates frequent refactoring because it helps programmers understand code, find bugs, and add new features [3]. These recommendations do indeed seem to produce benefits in practice; several studies have shown quantitative improvements to code bases as a result of refactoring [6, 4, 8] and others have confirmed that it is a frequent practice [17, 18].

But before a programmer can refactor code, she must recognize the need for refactoring. For example, consider the following code snippet:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

RSSE '08, November 10, Atlanta, Georgia, USA
Copyright 2008 ACM 978-1-60558-228-3 ...\$5.00.

```
class TrainStation{
    int lengthOf(Train t) {
        return t.locomotiveCount() +
            t.boxcarCount() +
            1; //the caboose
    }
    ...
}
```

The method `lengthOf` exhibits the FEATURE ENVY smell, because the method sends several messages to a `Train` object, but sends no messages to itself. FEATURE ENVY is a problem that can make software more difficult to change because a class's responsibilities are contained not only in the class itself, but also spread throughout envious classes that access the class's members. Table 1 describes several other code smells.

This smell can be alleviated by delegating the functionality to the `Train` class by sequencing three smaller refactorings: Extract Method, Move Method, and Rename Method. However, the mechanics of these refactorings are not the topic of this paper; instead we focus on *how* the programmer recognizes that code needs to be refactored.

Until recently, programmers have been forced to locate smells manually, which can be difficult for two reasons. First, novice programmers sometimes cannot locate smells as proficiently as more experienced programmers, as Mäntylä has shown in an experiment [9]. Second, it is burdensome for a programmer, even an expert programmer, to inspect every encountered piece of code for every possible smell (22 cataloged in Fowler's book alone [3]). Fortunately, many smells can be detected automatically by tools called *smell detectors*. So how can the smell detector give the programmer the information that she needs without interfering with her primary task?

2. THE SEVEN HABITS

To postulate what makes an effective smell detector, we examine two sources of evidence: the refactoring task during which programmers find smells and the nature of how smells appear in program code. The name of each habit is **bolded**.

2.1 Habits from the Refactoring Task

As we have suggested previously [10], it appears that programmers refactor frequently to maintain healthy code, interleaving refactoring with other tasks such as adding features. We call this *floss refactoring*, and contrast it with *root canal refactoring*, where a programmer refactors code intensively and exclusively once it has become unhealthy. If a programmer were to use a smell detector during floss refactoring, she would need to run it frequently, interleaved with program modifications. Thus, we postulate that availability is an important habit of smell detectors:

Availability. *Rather than forcing the programmer to frequently go through a series of steps in order to see if a tool finds any code smells, a smell detector should make smell information as available as soon as possible, with little effort on the part of the programmer.*

Likewise, due to interleaving of coding and refactoring during floss refactoring, a smell detector should be unobtrusive:

Unobtrusiveness. *A smell detection tool should not block the programmer while gathering, analyzing, and displaying information about smells.*

Because the programmer performs floss refactorings only if they help accomplish an immediate programming goal, the programmer would be most interested in smells related to the code currently being worked on:

Context-Sensitivity. *A smell detector should first and foremost point out smells relevant to the current programming context. Fixing smells in a context-insensitive manner may be a premature optimization.*

Not every kind of smell that a tool can detect has equal value to the programmer. This is because some smells are obvious to the naked eye (e.g., LONG METHOD), while others are difficult for a programmer to find (e.g., FEATURE ENVY) [9]. Thus:

Relativity. *A tool should place more emphasis on smells that are more difficult to recognize without a tool.*

2.2 Habits from the Nature of Smells

In addition to the task of finding smells, we can learn how to make an effective detector by examining the qualities of smells themselves. Code smells can emanate from many pieces of code, and the same code can give off several smells. Consider again the `TrainStation` example. While it is a relatively small method, it contains at least three code smells: FEATURE ENVY, MAGIC NUMBER, and COMMENTS. Depending on the rest of the program, other smells like REFUSED BEQUEST may be present in the snippet as well. Indeed, nearly *all* the code in this method smells! Underlining everything that contains even a whiff of a code smell could quickly overwhelm the programmer, making the detector useless. Thus:

Scalability. *A proliferation of smells in code should not cause the tool to overload the programmer with smell information.*

Several smells do not emanate from a single point in the code, but instead speak about relationships between several program elements. Like compilation errors, code smells can relate to several program elements that may be distributed in many places across the program text. For instance, REFUSED BEQUEST is not simply a problem with an overriding method; it is a problem with that method, its overridden superclass method, and potentially sibling methods that override that superclass method but do not call `super`. Thus:

Relationality. *A smell detection tool should be capable of showing relationships between code fragments that give rise to smells.*

DATA CLUMPS	A group of data objects that are duplicated across code [3]
FEATURE ENVY	Code that uses many features from classes other than its own [3]
REFUSED BEQUEST	A method that overrides a superclass method, but does not use the super method's functionality [3]
SWITCH STATEMENT	A switch statement, typically duplicated across code [3]
MESSAGE CHAIN	A series of method calls to "drill down" to a desired object [3]
TYPECAST	Changing an object from a one type to another type [2]
INSTANCEOF	An operator that introspects on the type of an object [2]
MAGIC NUMBER	A hard-coded value that is poorly documented [3]
LONG METHOD	A method with too much code [3]
LARGE CLASS	A class with too much code [3]
COMMENTS	Comments denote code that is not self-explanatory [3]

Table 1: Some smell names and descriptions

Likewise, smells can be complex and difficult to understand. Smells are potentially worse than compilation errors because a piece of code either generates a compilation error or it does not, but an instance of code smell may be subtle or flagrant, widespread or centralized, or anywhere in between. Thus:

Expressiveness. *A smell detector should go further than simply telling the programmer that a smell exists; it should help the programmer find the source(s) of the problem by explaining why the smell exists.*

3. PRIOR SMELL DETECTORS

Several smell detectors have been proposed, but none exhibits all of the habits that we postulated in the last section. In this section, we discuss how our habits (in **bold**) relate to previously proposed tools. Smell detectors are typically one of two types. The first, visualizations, provide a separate view from the source code, where smells are represented graphically. The second, editor annotations, layer graphics or text on top of the programmer's editor to convey smell information.

3.1 Visualizations

The Crocodile tool displays code smells using a 3-dimensional visualization where the distance between objects represents some smell [14]. While highly **relational**, the tool is also *exclusively* so; the visualization technique does not appear generalizable to smells that are sometimes non-relational, such as MESSAGE CHAIN. The Noseprints tool displays smells using a 2-dimensional, full-screen visualization [13]. As with Crocodile, the programmer invokes the tool once for each smell she wishes to detect, and thus neither tool is **context-sensitive** nor as **available** as they might be otherwise. In jCosmo, the tool analyzes the *entire program* and displays a graph; the size and color of the graph nodes show which parts of

the system are affected by which code smells [2]. jCosmo is not **context-sensitive** because it displays smells for the entire program. Moreover, it is not always **available** because it is a separate tool from the development environment and takes considerable time to analyze code.

3.2 Editor Annotations

CodeNose addresses the limitations of jCosmo by presenting programmers with a representation of detected smells inside the editor [15]. Built on top of the Eclipse programming environment (<http://www.eclipse.org/>), CodeNose underlines locations in the program text where smells have been detected, much like Eclipse’s standard compilation warnings:

```
(Class<Long>) Class.getPrimitiveClass("long");
```

A similar line-based indicator for smell detectors has been independently proposed by Hayashi and colleagues [5], Bisanz [1], and Tsantalis and colleagues [16]. Indeed, the presentation of smells in the same manner as compilation warnings is intuitively appealing for several reasons. First, both smells and warnings tell the programmer about a “problem” with the source code, so it makes sense to report them similarly. Second, underlining is the default user interface mechanism supported by static analysis tools such as the Eclipse Test and Performance Tools Platform (<http://www.eclipse.org/tppt/>), where programmers can easily build their own custom smell detection tools. Third, underlining allows the programmer to quickly see smells while editing code, making the presentation apparently suitable for floss refactoring. However, underlining treats all smells equally, not **relatively**; it typically displays a problem at a single point, not **relationally**; and highlights may overlap or cover much of the program text, and thus is not **scalable**.

4. A SMELL DETECTOR THAT GRATIFIES THE SEVEN HABITS

We have built a prototype smell detector that provides three views of code smells. By default, Ambient View shows an smell visualization while the developer is working with code. Then, if the programmer notices something unusual, she mouses over the visualization to identify specific smells in Active View. If she desires details, she requests more information by mouse clicking to reveal details in Explanation View. While we describe the tool textually in this section, it is more useful to see it in action; a series of short screencasts can be found at <http://multiview.cs.pdx.edu/refactoring/smells>. We show how our tool exhibits the habits discussed in Section 2 in **bold**.

4.1 Ambient View

The initial view of the smell detector is ambient [12], where a visual representation of contextually relevant smells is displayed in the program editor, behind the program text (Fig. 1, top). By running the underlying smell detection engine in the background, (some-what like Eclipse’s incremental compilation), this visualization is intended to be visible and **available** at all times during code editing and navigating, light enough as to be **unobtrusive**.

The visualization is composed of circular sectors, which we call wedges, radiating from a central point in a half-circle. Each wedge represents a code smell, and the radius of each wedge represents the degree to which the current programming context exhibits that smell. For example, at the top of Figure 1, the southernmost wedge indicates the strongest smell while the northernmost wedge indicates the weakest smell. As the programmer navigates through the

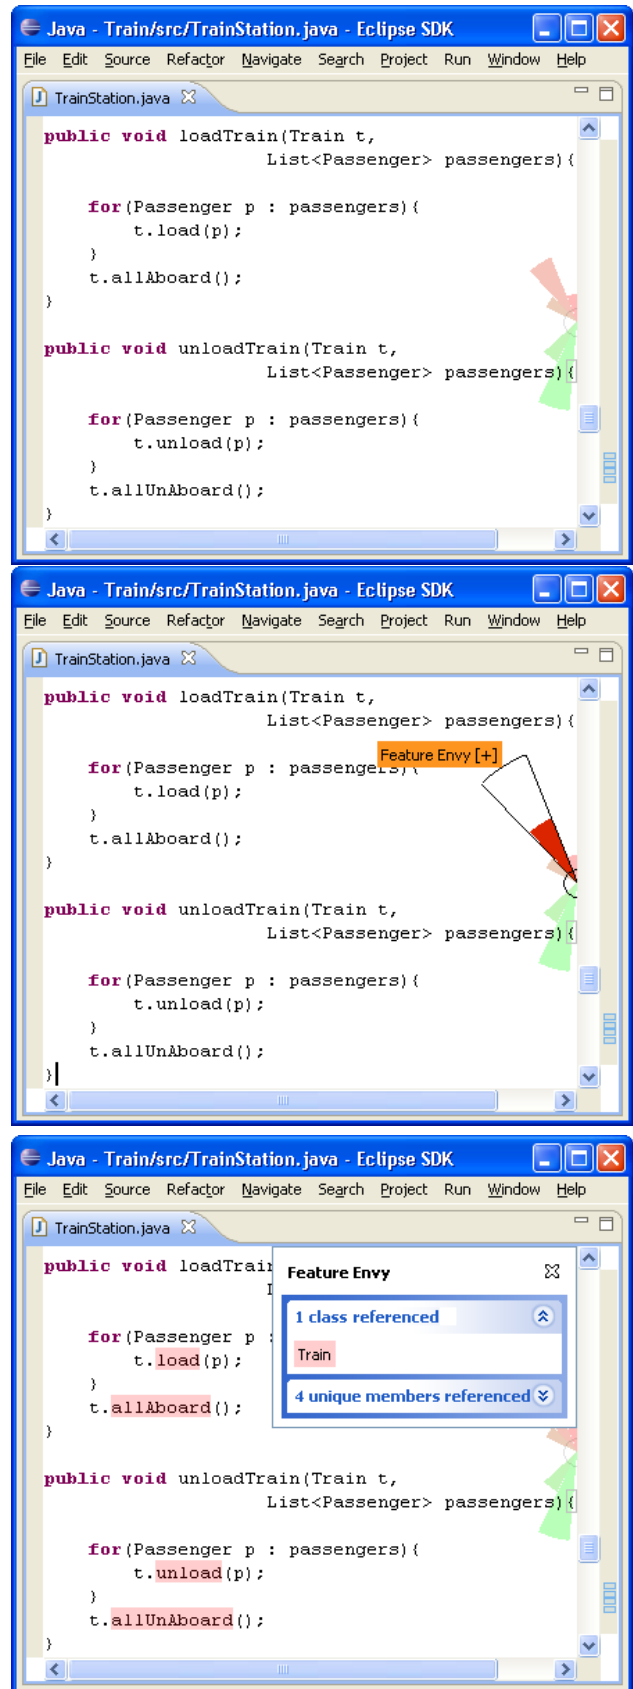


Figure 1: Three progressively more informative views of the code smells in program code.

code, the program text flows in front of the visualization, but the radius of each wedge changes as the programmer's context changes. The radius of each wedge is controlled by a smell analyzer that evaluates a smell in the programmer's **context**. However, the maximum screen area available for each wedge is bounded, and thus the visualization is designed to **scale** as the number of smells increases.

Wedges are colored from red-to-green, north-to-south. Smells are assigned to wedges such that the southernmost (and greenest) wedge is the most obvious smell and the northernmost (and reddest) wedge is the least obvious smell. Our subjective obviousness ordering is reflected in Table 1, where the least obvious smell appears at the top, as it does in Ambient View¹. For example, at the top of Figure 1, the view indicates that there is a strong unobvious smell (in this case, `FEATURE ENVY`, although the smell names are intentionally omitted from this view) as well as a strong obvious smell (`LARGE CLASS`). Because more obvious smells are distinguished from less obvious smells, both spatially and chromatically, the programmer can judge the **relative** importance of the displayed smells.

The purpose of this view is to allow the programmer to occasionally glance at the visualization to determine if there are any strong, relevant code smells and to give a rough estimate of their extent. There is very little burden or commitment on the part of a programmer to determine whether a smell exists; she only needs to glance at the visualization, unlike with most existing smell detection tools which require the programmer to activate the tool and inspect the results. Indeed, such batch-smell-detection tools are applicable in only 1 of 3 refactoring tasks (code inspection) [3, pp.58-59], and are less useful for floss refactoring.

4.2 Active View

If the programmer observes something interesting or unusual in the Ambient View, she can then mouse-over a wedge to reveal the name of that wedge's associated smell. Furthermore, the wedge's color is darkened to bring it into the programmer's focus. In the middle of Figure 1, we have moused-over the second smell from the top. If the programmer is interested in further details of the detected smell, she can click on the smell label to activate the Explanation View.

The purpose of the Active View is to provide a little more information than the Ambient View, and to help the programmer transition to the Explanation View. Again, the transition from view to view in order to reveal more information was designed to be as fast and as painless as possible.

4.3 Explanation View

The Explanation View provides detailed information about a particular smell. In essence, this view was designed to explain *why* the smell's wedge has the displayed radius in the most **expressive** way possible. Each smell is displayed using different visual elements, but smells in the Explanation View typically have two common components.

Summary Pane. In the bottom panel of Figure 1, a summary pane is displayed at the upper right of the editor. This pane is fixed relative to the program code (that is, it does not move when the programmer navigates away), but may be moved manually within the editor at the programmer's behest. Generally, this pane displays a summary of the data collected from the smell analyzer. In the Figure, the summary pane for the `FEATURE ENVY` smell shows that 4 members from the `Passenger` class are accessed; more information about these members can be accessed by clicking on the

¹We currently have implemented analyzers for all the smells in Table 1, with the exception of `REFUSED BEQUEST`, `MAGIC NUMBER`, and `COMMENTS`.

downward-pointing double-chevron.

Editor Annotations. The code editor is typically annotated to point out where smells originate from. For example, in the Figure you can see where the 4 members of the `Passenger` class are referenced in the editor. Each member reference to an external class is **related** together visually using the same color highlight.

Taken together, editor annotations and the summary pane were built to help the programmer not only understand *if* code smells, but *why* the code smells.

4.4 Tool Discussion

While we have outlined how our smell detection tool works in general, a number of details have significant bearing on the tool's practicality.

First, how does the tool determine the radius of each wedge in the Ambient View? The maximum size of each wedge is fixed, so that it does not monopolize the program editor. An individual smell analyzer is responsible for converting a smell in the programmer's working context to a scalar value between zero and the maximum radius of the wedge. While the formula for the radius is different for different analyzers, some formulas are more complex than others. For instance, `LARGE CLASS` is a relatively simple formula because the radius increases as the size of the class increases, while `FEATURE ENVY` incorporates the number of external classes referenced, the number of external members referenced, and whether internal members are referenced.

Second, how does the tool search for smells efficiently? Several smell analyzers require complex program analysis, so as the number and complexity of analyzers increase, the development environment may begin to respond more slowly. However, having detection run in a background thread and caching smell results for unchanged program elements have been important techniques for maintaining acceptable performance. Moreover, we hope that a more intelligent search strategy, starting in the programmer's current context and radiating outward to less contextually relevant code, will improve performance even further. Smell analysis may also be made more efficient by using heuristics to analyze smells for the Ambient View, but use full static analysis in the Explanation View. In this way, the development environment can remain responsive during normal development and be fully accurate during smell inspection.

Third, what constitutes the programmer's "current context?" In our implementation, we define current context as the the union of all methods, whole or partial, that are visible in the open editor. In the future, more sophisticated definitions of context may be used as well, such as task contexts used by the Mylyn tool [7] or Parnin and Görg's usage contexts [11].

5. PLANNED EVALUATION

We plan on completing an evaluation in the near future. We are primarily interested in showing that our "habits," as embodied by the tool, are useful in helping programmers do their jobs. To that end, we propose an experiment with three parts.

In the first part, we ask the programmer to inspect several pieces of source code, and then tell us when the detector is displaying an interesting smell. We will also purposely tell the programmer to pause at specific points in the source code, and tell us which smells and in which order they would ask the tool for more detailed information. The purpose of this part of the experiment is to determine if the tool conveys smell information to the programmer efficiently.

In the second part of the experiment, we ask programmers to use the Explanation View to learn more about one or two smells. Each time they use the Explanation View, we will ask the program-

mer about the extent of the smell, how likely she is to correct the smell, and whether the display tells her something surprising about the code. The purpose of this part of the experiment is to determine whether the Explanation View is sufficiently detailed to convey complex smell information.

We will repeat these two parts with each programmer, except that we will not give the programmer the aid of the tool. Our intent is to establish a baseline for judging the programmers' behaviors, and to allow programmers to respond in a more comparative manner in the final part of the experiment.

In the final part, we ask programmers to subjectively evaluate their experience with the tool. We will use a questionnaire where the programmers answer a variety of questions about whether the tool helped them quickly and efficiently find and understand smells in the code. Moreover, we will ask a series of questions regarding how the programmers might use the tool in their own code.

We hope that this proposed evaluation will help show that our smell detection tool, and thus our guidelines, are an effective approach to smell detection.

6. CONCLUSION

We have presented seven principles for the design of code smell detectors that fit into the typical floss refactoring workflow, and described a realization of these principles in the form of an ambient smell detector. While we feel that the seven principles are important to building usable smell detectors, we have not yet verified that this is the case, and plan to do so in a forthcoming experiment.

7. ACKNOWLEDGEMENTS

Thanks to Chris Parnin and Claudia Rocha for their helpful advice, and to the National Science Foundation for partially funding this research under CCF-0520346.

8. REFERENCES

- [1] M. Bisanz. Pattern-based smell detection in TTCN-3 test suites. Master's thesis, Georg-August-Universität Göttingen, Dec 2006.
- [2] E. v. Emden and L. Moonen. Java quality assurance by detecting code smells. In *Proceedings of the Ninth Working Conference on Reverse Engineering*, pages 97–106, Washington, DC, USA, 2002. IEEE Computer Society.
- [3] M. Fowler. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.
- [4] B. Geppert and F. Rosler. Effects of refactoring legacy protocol implementations: A case study. In *METRICS '04: Proceedings of the 10th International Symposium on Software Metrics*, pages 14–25, Washington, DC, USA, 2004. IEEE Computer Society.
- [5] S. Hayashi, M. Saeki, and M. Kurihara. Supporting refactoring activities using histories of program modification. *IEICE - Trans. Inf. Syst.*, E89-D(4):1403–1412, 2006.
- [6] Y. Kataoka, T. Imai, H. Andou, and T. Fukaya. A quantitative evaluation of maintainability enhancement by refactoring. In *ICSM '02: Proceedings of the International Conference on Software Maintenance*, page 576, Washington, DC, USA, 2002. IEEE Computer Society.
- [7] M. Kersten and G. C. Murphy. Mylar: a degree-of-interest model for IDEs. In *AOSD '05: Proceedings of the 4th International Conference on Aspect-Oriented Software Development*, pages 159–168, New York, NY, USA, 2005. ACM.
- [8] R. Kolb, D. Muthig, T. Patzke, and K. Yamauchi. A case study in refactoring a legacy component for reuse in a product line. In *ICSM '05: Proceedings of the 21st IEEE International Conference on Software Maintenance*, pages 369–378, Washington, DC, USA, 2005. IEEE Computer Society.
- [9] M. V. Mäntylä. An experiment on subjective evolvability evaluation of object-oriented software: explaining factors and interrater agreement. In *Proceedings of the International Symposium on Empirical Software Engineering*, pages 287–296, November 2005.
- [10] E. Murphy-Hill and A. P. Black. Refactoring tools: Fitness for purpose. *IEEE Software*, 25(5), September–October 2008.
- [11] C. Parnin and C. Görg. Building usage contexts during program comprehension. In *ICPC '06: Proceedings of the 14th IEEE International Conference on Program Comprehension*, pages 13–22, Washington, DC, USA, 2006. IEEE Computer Society.
- [12] C. Parnin and C. Gorg. Design guidelines for ambient software visualization in the workplace. *Visualizing Software for Understanding and Analysis, 2007. VISSOFT 2007. 4th IEEE International Workshop on*, pages 18–25, June 2007.
- [13] C. Parnin, C. Görg, and O. Nnadi. A catalogue of lightweight visualizations to support code smell inspection. In *Proceedings of the 2008 ACM Symposium on Software Visualization*, 2008.
- [14] F. Simon, F. Steinbrückner, and C. Lewerentz. Metrics based refactoring. In *Proceedings of the Fifth European Conference on Software Maintenance and Reengineering*, pages 30–38, Washington, DC, USA, 2001. IEEE Computer Society.
- [15] S. Slinger. Code smell detection in Eclipse. Master's thesis, Delft University of Technology, March 2005.
- [16] N. Tsantalis, T. Chaikalis, and A. Chatzigeorgiou. JDeodorant: Identification and removal of type-checking bad smells. In *CSMR*, pages 329–331. IEEE Computing Society, 2008.
- [17] P. Weißgerber and S. Diehl. Are refactorings less error-prone than other changes? In *MSR '06: Proceedings of the 2006 International Workshop on Mining Software Repositories*, pages 112–118, New York, NY, USA, 2006. ACM.
- [18] Z. Xing and E. Stroulia. Refactoring practice: How it is and how it should be supported — an Eclipse case study. In *ICSM '06: Proceedings of the 22nd IEEE International Conference on Software Maintenance*, pages 458–468, Washington, DC, USA, 2006. IEEE Computer Society.