On the Refinement of Spreadsheet Smells by means of Structure Information

Patrick Koch¹, Birgit Hofer², Franz Wotawa²

Abstract

Spreadsheet users are often unaware of the risks imposed by poorly designed spreadsheets. One way to assess spreadsheet quality is to detect smells which attempt to identify parts of spreadsheets that are hard to comprehend or maintain and which are more likely to be the root source of bugs. Unfortunately, current spreadsheet smell detection techniques suffer from a number of drawbacks that lead to incorrect or redundant smell reports. For example, the same quality issue is often reported for every copy of a cell, which may overwhelm users. To deal with these issues, we propose to refine spreadsheet smells by exploiting inferred structural information for smell detection. We therefore first provide a detailed description of our static analysis approach to infer clusters and blocks of related cells. We then elaborate on how to improve existing smells by providing three example refinements of existing smells that incorporate information about cell groups and computation blocks. Furthermore, we propose three novel smell detection techniques that make use of the inferred spreadsheet structures. Empirical evaluation of the proposed techniques suggests that the refinements successfully reduce the number of incorrectly and redundantly reported smells, and novel deficits are revealed by the newly introduced smells.

Keywords: Spreadsheets, Code smells, Static analysis

1. Introduction

End-user programming has gained a lot of attention. Without doubt, spreadsheets are the most prominent example of end-user programs. They are intuitive to use and widespread. Nearly everybody has access to them via spreadsheet environments like Microsoft Excel, Google Sheets, and Numbers. Professionals, managers, sales people, and administrators use spreadsheets exhaustively [1], and often base important decisions on them. Their easy usage is one of the main reasons why spreadsheets are so popular.

Preprint submitted to JSS, Special Issue on Program Debugging

Email addresses: Patrick.Koch@aau.at (Patrick Koch), bhofer@ist.tugraz.at (Birgit Hofer), wotawa@ist.tugraz.at (Franz Wotawa)

¹Klagenfurt University, Austria

²Graz University of Technology, Austria

On the downside, there is a lack of quality awareness. The ease of use often prevents people from taking training courses and lets them handle spreadsheets in a learning by doing style. Therefore, many people are not aware of the risks that come with spreadsheets. The list of horror stories³ where spreadsheet errors caused significant reputation loss or immense financial damages is long, as the following three examples illustrate: (1) JP Morgan lost \$400 million because of a fault in their value-at-risk model spreadsheet [2]. (2) The US-village West Baraboo has to pay more than \$400,000 in additional interests because of a spreadsheet error [3]. (3) The Canadian power generation company TransAlta lost \$24 million because of a copy-paste error in a spreadsheet that led them to buy more US power transmission hedging contracts than necessary [4]. One might think that these examples are just extraordinary exceptions, but field studies [5, 6] confirm a consistent lack of quality in business spreadsheets.

To address such quality issues, we are working among a large number of fellow researchers to develop quality assurance (QA) techniques for spreadsheets (see Jannach et al. [7] for an overview). As spreadsheets can be regarded as a form of software, as pointed out by Hermans et al. [8], a number of QA techniques for general software development were adopted for the domain. One such technique is the detection of spreadsheet smells that were adopted from code smells [9]. A code smell indicates lacks in the quality of a part of source code. This part could either be hard to comprehend, hard to maintain, or Spreadsheet smells similarly attempt to infer deficient parts in error-prone. spreadsheets [10]. There are three basic types of spreadsheet smells: (1) input smells [11] indicate, for example, missing input values, wrong numeric values, and mistyped strings; (2) formula smells [12] relate to complex formulas, long calculation chains, and duplicated formulas; and (3) inter-worksheet smells [13] point out cells and worksheets that excessively refer to other worksheets, or which are excessively referred to by other worksheets.

In this work, we regard spreadsheet smells to be used as part of a checking regime for spreadsheets. In a survey, Mireille Ducassé [14] classified debugging approaches where one method is checking. In checking, patterns in programs are used for identifying potential bug locations. Such patterns are code smells, which directly point to suspicious parts of the program and therefore potentially locate bugs. Debugging can be seen as an activity comprising fault detection, fault localization, and fault correction to eliminate bugs in software. In checking, fault detection and fault localization are preformed at the same time using patterns that are classifying a code smell. Similarly, we regard the use of spreadsheet smells to point out suspicious parts of a spreadsheet that a user should examine preferrentially when checking for correct functionality of said sheet.

We would however like to point out, that we currently lack substantive empirical evidence that clearly links spreadsheet smells to fault proneness, and we thus regard this as an opportunity for future work. Nevertheless, we still suspect our perception to be valid, as a recent case study linked a reduction in

³see http://www.eusprig.org/horror-stories.htm

spreadsheet smell detections with a subjective impression of heightened quality and maintainability of spreadsheets [10], and a number of previous works successfully used spreadsheet smells in the context of spreadsheet debugging, including the FaultySheet Detective tool introduced by Abreu *et al.* [15], as well as Table Clones and other approaches discussed by Dou *et al.* [16, 17].

Unfortunately, spreadsheet smell detection suffers from certain weaknesses and blind spots. For example, duplication is a common problem: if a smelly cell of a spreadsheet is copied, its copies are often also reported as being smelly, and seeing the same smell being reported over and over again can be very frustrating for a user. Other weaknesses affect only certain smells. For example, the input smell 'Pattern Finder' [18, 11] detects deviating cell types for cells in the same row or column. However, the reported smelly cells often turn out to be false positives, because not all cells in a row or column serve the same purpose.

We propose to compensate existing shortcomings by incorporating structural information in the detection procedures of spreadsheet smells. As shown by a study of Cunha *et al.* [19], providing abstract structure information can enhance the ability of end-users to understand and efficiently use spreadsheets. The detection of spreadsheet structures follows the approach we outlined in previous work [20], where we explained and evaluated the static analysis process. As the targeted smell refinements require detailed information about the inferred structures, the current work describes the structure analysis in greater detail, including a description of the handling of cell references and area references, as well as the theory of absolute versus relative references.

Inferred structure information (input groups, formula groups, computation blocks, and headers of such blocks) can be applied in several ways for smell refinement, as we demonstrate in three examples: (1) for the *Pattern Finder* input smell, we check for deviations only within groups instead of checking complete rows and columns; (2) for the *Long Calculation Chain* formula smell, we report a smell only once per group instead of reporting it for each individual cell; and (3) for *Feature Envy* inter-worksheet smell, we count the connections between worksheets for each group instead of for each cell. Moreover, structural information can be used to check for novel quality issues. We demonstrate this by introducing three new spreadsheet smell detection techniques which are based on the results of the structural analysis: (1) *Overburdened Worksheet* indicates worksheets that contain too much functionality; (2) *Inconsistent Formula Group Reference* signals inconsistencies within the references of groups of formulas; and (3) *Missing Header* indicates gaps in series of header cells.

We investigated the performance of the improved and new smells using an empirical evaluation on a well known dataset in combination with a manual investigation of detected smells on a selected subset of spreadsheets. This investigation revealed that the improved smell versions are successful in limiting false positive and duplicate detections, and the new smells can be used in combination with existing techniques to point out novel quality issues.

The remainder of this paper is organized as follows: Section 2 pictures a motivating example for the application of the proposed refinements, and Section 3 defines the most important concepts of spreadsheets used in this paper. Section 4 explains how our structural analysis approach works. Section 5 shows how three existing smell detection techniques can be improved by means of structure information and introduces three novel smell detection techniques that make use of structures. Section 6 evaluates the improved and new techniques, Section 7 discusses related work, and Section 8 concludes this paper.

2. Motivating Example

The spreadsheet illustrated in Figure 1 compares purchase options for a new office coffee machine. The coffee machine should be chosen from one of three alternatives (capsule, drip, automatic) in such a way that the total costs over a period of three years are the lowest. For this, the coffee consumption per department is captured in worksheets Department1/2/3 (see Figure 1a, the non-illustrated worksheets are similar), summed up in worksheet Total (see Figure 1b), and the resulting cost of the three alternatives compared in worksheet Investment (see Figure 1c).

Assume a smell detection technique reports cell E9 of the Investment worksheet as smelly, because E9's formula is too complex. Since the cells E10 and E11 have the same complex formula, they would also be reported as smelly. When only these three smells are reported, then a user would probably immediately see that these cells suffer from the same smell. However, when faced with an abundance of smell reports in a larger spreadsheet, the connection between related smelly cells is harder to comprehend for a user. In contrast, our proposed smell refinement detects and indicates related smelly cells as one smelly unit instead, providing concise feedback to the user.

3. Preliminaries

A spreadsheet consists of a set of worksheets; a worksheet consists of set of cells. A cell **c** has a content and can be uniquely identified by its coordinates and the worksheet it belongs to: The function $x(cell \mathbf{c})$ returns the column index of cell **c** within a worksheet; $y(cell \mathbf{c})$ returns its row index [21]. We can use the cells' coordinates to determine neighboring cells.

Definition 1 (Neighbors). Two cells, c and c', are neighbors if their Manhattan distance (simple sum of the differences of the respective column and row indices) is one. The function neighbors(cell c) returns c's neighbors [20]: neighbors(c) = { $c' \in cells(ws(c)) : |x(c') - x(c)| + |y(c') - y(c)| = 1$ }.

Example 1. Cell B5 of the worksheet from Figure 1c has as neighbors the cells B4, B6, A5, and C5.

Definition 2 (Connected). The function connected({cell} C) checks whether a set of cells C is completely connected by neighbourhood relations; it returns true, if $\{|C| = 1 \lor (\exists c \in C : \exists c' \in C \text{ with } c' \in neighbors(c) \land connected(C \setminus c))\}$.

2	A	В	С	D	E	F
1	Coffee Con	sumption: D	epartment	1		
2		Employee	-			
3	Quarter	Anderson	Bourne	Connor	Dredd	Total
4	1	35	172	80	98	=SUM(B4:E4)
5	2	47	200	51	66	=SUM(B5:E5)
6	3	46	139	46	39	=SUM(B6:E6)
7	4	74	155	57	43	=SUM(B7:E7)
8	Total	=SUM(B4:B7)	=SUM(C4:C7)	=SUM(D4:D7)	=SUM(E4:E7)	=SUM(F4:F7)
			(a) Depar	tment1		
1	A	В	() = ·F ···	C	D	E
1	Coffee Cons	sumption: All	Department	-		
2	conce cons	Department	Department	.5		
2	Quarter	1	2	3		Total
4	1	=Department1!F	4 =Departm	ent2!F4 =De	partment3!F4	=SUM(B4:D4
5	2	=Department1!F	5 =Departm	ent2!F5 =De	partment3!F5	=SUM(B5:D5
6	3	=Department1!F6		=Department2!F6 =Dep		=SUM(B6:D6
7	4	=Department1!F	7 =Departm	ent2!F7 =De	partment3!F7	=SUM(B7:D7
В	Total	=Department1!F	8 =Departm	ent2!F8 =De	partment3!F8	=SUM(B8:D8
			(b) Te	otal		
J.	A	В	C	D		E
1	Investmen	t Ontions				
2	Input	options				
2	Coffee / year	-TotallE8				
1	Service years	2				
5	Coffee total	-R2*R/				
6	conce total	-05 04				
7	Comparison	8				
8	Type	Investment	Coffee price	Yearly service	e T	otal
9	Capsule	100	0.3	100	=B9+C9*\$B\$	5+D9*\$B\$4
10	Drip	250	0.15	100	=B10+C10*\$F	3\$5+D10*\$R\$4
11	Automatic	800	0.1	200	=B11+C11*SE	3\$5+D11*\$B\$4
0	Drip Automatic	250 800	0,15 0,1	100 200	=B10+C10*\$E =B11+C11*\$E	3\$5+D10*\$B\$ 3\$5+D11*\$B\$

(c) Investment

Figure 1: Formula views of the different worksheets of our running example

All cells in a set of connected cells must reach every other cell in the set by following only neighborhood relations of cells within the set.

Beside its coordinates, each cell has a type depending on its content. This can be either formula, Boolean, numeric, string, error, or empty. 'Formula' is a dominant type: formula cells compute to other cell types, however, the cell's type is still 'formula'. The function $type(cell \mathbf{c})$ returns the type of a cell.

Example 2. Although cell F4 of worksheet Department1 (see Figure 1a) computes a numerical value, its type is 'formula'.

Formula cells might reference other cells. References can either be relative or absolute: a relative reference determines the referenced cell in relation to the position of the referencing cell; an absolute reference determines the referenced cell by indicating its position within the worksheet independent of the position of the referencing cell. There are two notations for referencing cells, the A1 and the R1C1 notation. Both support relative and absolute references.

In the A1 notation, a cell **c** references another cell **c'** by indicating the absolute position of **c'**, i.e., $x(\mathbf{c'})$ and $y(\mathbf{c'})$. Absolute references place a \$ sign preceding the coordinate position. Absolute references in A1 notation are relevant when copying cells. For example, if the relative cell reference 'A1' is copied from cell B1 to cell B2, the reference will change, according to the shift of the cell position, to 'A2'. In contrast, an absolute cell reference '\$A\$1', when copied from any cell to any other cell, will always refer to the coordinates A1. Cell references can be absolute (e.g., '\$A\$1'), relative (e.g., 'A1'), or mixed (e.g., '\$A1', 'A\$1'). Names of worksheets are only included when the worksheet of the referenced cell differs from that of the referencing cell.

Definition 3 (R1C1 notation). In the R1C1 notation, a cell references another cell by indicating the relative position of the referenced cell. The relative position of cell c' with respect to c is indicated as R[y(c') - y(c)]C[x(c') - x(c)]. Absolute references to a cell c' are written as Ry(c')Cx(c'). The names of worksheets are only indicated when the worksheet of the referenced cell differs from that of the referencing cell. [20]

If a cell refers to a cell in the same row or column, the '0' for indicating the row or column index can be left out, e.g. 'RC[-1]' for indicating the cell above and 'R[-1]C' for indicating the left cell. References to absolute row or column positions omit the square brackets, e.g. 'R1C1' indicates cell A1, the first cell in the first column.

Definition 4 (Formula). The function formula (cell c) returns the formula of cell c in R1C1 notation if c is a formula cell; otherwise it returns void.

Example 3. Cell B5's formula of the worksheet Investment (Figure 1c) is '=B3*B4' in A1 notation and '=R[-2]C*R[-1]C' in R1C1 notation. Cell B3's formula of the same worksheet is '=Total!E8' in A1 notation and '=Total!R[5]C[3]' in R1C1 notation. Frequently, formulas are copied to other cells. Many spreadsheet environments support copying by drag&drop. Even though the copied cells reference other cells than the original cells, they are semantically equal. We identify semantically equal cells by means of copy-equivalence:

Definition 5 (Copy-equivalence). Two cells c, c' are copy-equivalent if their formulas in R1C1 notation are identical (formula(c) \equiv formula(c')).

Example 4. For the spreadsheet from Figure 1, we identify the cells E9, E10, and E11 of the worksheet Investment as copy-equivalent since they have the same formula ('=RC[-3]+RC[-2]*R5C2+RC[-1]*R4C2') in R1C1 notation. The other copy-equivalent cells are

- B8, C8, D8, E8, and F8 of the worksheets Department1, Department2 and Department3 ('=SUM(R[-4]C:R[-1]C)'),
- F4, F5, F6, and F7 of Department1, Department2 and Department3 (=SUM(RC[-4]:RC[-1]))
- B4, B5, B6, B7, and B8 of worksheet Total ('=Department1 RC[4]')
- C4, C5, C6, C7, and C8 of worksheet Total ('=Department2 RC[3]')
- D4, D5, D6, D7, and D8 of worksheet Total ('=Department3 RC[2]')
- E4, E5, E6, E7, and E8 of worksheet Total ('=SUM(RC[-3]:RC[-1])')

B3 and B5 of worksheet Investment do not have any copy-equivalent cells.

Since references can be absolute, relative or mixed, we separately consider the x and y coordinates when determining the position of a referenced cell:

Definition 6 (Coordinate reference). A coordinate reference \mathbf{r}_c represents either the x or the y coordinate of cell c and consists of a value and a type. The function absolute(coordinate reference \mathbf{r}_c) returns true if \mathbf{r}_c is an absolute reference; it returns false if \mathbf{r}_c is a relative reference. The function value(coordinate reference \mathbf{r}_c) returns for an absolute reference, the absolute index of an absolute reference; for a relative reference, it returns the deviation of the reference from a base index.

Example 5. The reference 'R4C[2]' features two coordinate references. Its row reference, 'R4', is absolute and refers to the fourth row of the worksheet. Its column reference, 'C[2]', is relative and refers to the column two positions right of the column of the cell containing the reference.

Definition 7 (Dereferencing coordinates). The function derefCoordinate(coordinate reference r_c , index coordinate) returns the dereferenced index of a coordinate reference r_c in the formula of the cell at index coordinate:

$$derefCoordinate(r_c, \ coordinate) = \begin{cases} value(r_c) & if \ absolute(r_c) \\ coordinate + value(r_c) & otherwise. \end{cases}$$

Example 6. Cell E9's formula of the worksheet Investment (Figure 1c) is '=B9+C9*\$B\$5+D9*\$B\$4' in A1 notation and '=RC[-3]+RC[-2]*R5C2+RC[-1]* R4C2' in R1C1 notation. Its first reference, 'RC[-3]' ('B9' in A1 notation), is a combination of a relative column-reference with value -3 and a relative rowreference with value 0. Its third reference, 'R5C2' ('\$B\$5' in A1 notation), is a combination of an absolute column-reference and an absolute row-reference.

Formulas may contain two different types of references: cell references and area references. Cell references, hereinafter referred to as "references", point to individual cells; area references point to sets of cells. Cell references consist of two parts: (1) the target worksheet of the reference (provided by the function $ws(reference \mathbf{r})$) and (2) the coordinate references in column and row orientation (provided by the functions $x(reference \mathbf{r})$ and $y(reference \mathbf{r})$).

Definition 8. The function deref(reference \mathbf{r} , cell \mathbf{c}) resolves a reference \mathbf{r} contained in the formula of cell \mathbf{c} , returning the referred cell \mathbf{c} ':

$$deref(r,c) = c' \in cells(ws(r)): \ x(c') \equiv derefCoordinate(x(r), x(c)) \land y(c') \equiv derefCoordinate(y(r), y(c))$$

The function refs (formula f) returns the set of all cell references of formula f.

In contrast to cell references, area references refer to a set of cells in a rectangular area of a worksheet. The rectangular area is described by its top-left start coordinates and its bottom-right end coordinates. Area references consist of three parts: (1) the worksheet to which the reference \mathbf{r}_a refers to (accessibly by the function $ws(area \ reference \ \mathbf{r}_a)$), (2) the x and y coordinate references of the start coordinates of the area reference (accessible by the functions $x_1(area \ reference \ \mathbf{r}_a)$ and $y_1(area \ reference \ \mathbf{r}_a)$), and (3) the x and y coordinate references of the end coordinates of the area reference (accessible by the functions $x_2(area \ reference \ \mathbf{r}_a)$ and $y_2(area \ reference \ \mathbf{r}_a)$).

Definition 9. Function deref(area reference r_a , cell c) resolves an area reference r_a contained in the formula of cell c, returning the set of referred cells:

$$deref(r_a, c) = \begin{cases} c' \in cells(ws(r_a)) : & derefCoordinate(x_1(r_a), x(c)) \leq x(c') \land \\ derefCoordinate(x_2(r_a), x(c)) \geq x(c') \land \\ derefCoordinate(y_1(r_a), y(c)) \leq y(c') \land \\ derefCoordinate(y_2(r_a), y(c)) \geq y(c') \end{cases} \end{cases}$$

Function refs_a (formula f) returns the set of all area references of formula f.

Definition 10 (Referenced cells). The function $\rho(\text{cell } c)$ returns the set of cells that are referenced in c's formula f:

$$\rho(c) = \bigcup_{r_c \in \operatorname{refs}(f)} \{\operatorname{deref}(r_c, c)\} \cup \bigcup_{r_a \in \operatorname{refs}_a(f)} \operatorname{deref}(r_a, c).$$

Example 7. For the cell E11 in the worksheet Investment, we compute the following reference information: refs(Investment!E11) = {RC[-3], RC[-2], R5C2, RC[-1], R4C2}, refs_a(Investment!E11) = {}, and ρ (Investment!E11) = {B11, C11, B5, D11, D4}. For the cell Total!E8, we compute the following reference information: refs(Total!E8) = {}, refs_a(Total!E8) = {RC[-3]:RC[-1]}, and ρ (Total!E8)={B8, C8, D8}.

4. Structure Analysis

This section describes our approach to infer structure information from a spreadsheet, that is then used to improve and create spreadsheet smells. The general approach was already presented in our IWPD paper [20]. In the present work, we describe the analysis process in greater detail. We provide examples for each analysis step by referring to the running example in Section 2, and also elaborate on previously unattended topics that are relevant for the definition of smell refinements and novel, structure-based smells (e.g. handling of cell and area references).

The goal of the structural analysis is to identify input groups, formula groups, computation blocks, and headers. Groups are one-dimensional areas in a worksheet. Cells belonging to the same group share the same purpose, e.g. summing up data, even though they might operate on different input data. Cells of an input group provide their value to formula groups, but do not themselves refer to any other group. Cells of formula groups process the same calculation on different input data and refer to input groups or intermediate formula groups. Computation blocks are rectangular areas containing related input groups, formula groups, and empty cells. Headers serve as labels for rows and columns of computation blocks.

Example 8. For the worksheet Department1 from Figure 1a, the structural analysis process will detect input groups for the area B4:E7, two formula groups (B3:F3 and F4:F7), and one computation block (B4:F3). The cells in the area A4:A8 are the row headers for this block and the cells B3:F3 are the column headers. A3 and B2 are meta headers. While all cells in the area A4:E7 are numerical values, the cells A4:A7 have a different purpose then the cells B4:E7.

Figure 2 illustrates the overall analysis process: (1) *Grouping* finds group of related cells, based on various criteria; (2) *Blocking* combines groups into cohesive blocks; and (3) *Header Assignment* relates header cells to blocks and the contained groups. Each step in the process infers specific structural properties of the input sheet, and includes information from preceding process steps for in the analysis.

4.1. Grouping

In the grouping step, we infer different groups the cells in each worksheet, according to their type (string, numeric, formula, ...), content (in case of formula cells), and position within the worksheet. The first action within the Grouping step is to determine *type-based groups* within each worksheet.



Figure 2: Block creation. Gray cells represent grouped cells, red cells represent non-blockable cells (i.e., non-empty cells that are not part of any group), and frames represent blocks.

Definition 11 (Type-based group). A type-based group C consists of a set of cells which (1) have the same cell type, (2) are connected (see Def. 2), and (3) feature the same formula in R1C1 notation (if the cells are of type formula):

$$typeBasedGroup(C) = \begin{cases} true & if (\{\forall c, c' \in C : type(c) = type(c') \\ \land formula(c) \equiv formula(c')\} \\ \land connected(C)) \\ false & otherwise. \end{cases}$$

Example 9. Worksheet Department1 (Figure 1) contains the following typebased groups: $\{A1\}, \{B2; A3:F3\}, \{A4:E7\}, \{F4:F7\}, \{A8\}, and \{B8:F8\}.$ Worksheet Total has $\{\{A1\}, \{B2\}, \{A3\}, \{B3:D3\}, \{E3\}, \{A4:A7\}, \{A8\}, \{B4:B8\}, \{C4:C8\}, \{D4:D8\}, \{E4:E8\}\}$ as type-based groups. Worksheet Investment has $\{\{A1:A5\}, \{B3\}, \{B4\}, \{B5\}, \{A7:A11;B8:E8\}, \{B9:D11\}, \{E9:E11\}\}$ as typebased groups.

In the example, the headers in the cells A4:A7 of the worksheet Department1 are in the same group as the input cells B4:E7. This renders header inference problematic. The header and input cells will be separated in the subsequent analysis, by examining which cells are referenced by formula cells.

The type-based groups mainly serve as starting point for further processing

of related formula cells. In the next step, we synthesize the formula cells of each worksheet into cohesive units. For doing so, we define *formula groups*.

Definition 12 (Formula group). A type-based group C is a formula group if the group consists of formula cells:

 $formulaGroup(C) = \begin{cases} true & if typeBasedGroup(C) = true \\ & \land (\forall c \in C : type(c) = `formula') \\ false & otherwise. \end{cases}$

The function formulaGroups(worksheet w) returns for a worksheet w the set of all formula groups contained in w. The function formula(formula group C) returns the formula of group C.

Example 10. The function formulaGroups(Department1) returns $\{\{B8:F8\}, \{F4:F7\}\}$, formulaGroups(Total) returns $\{\{B4:B8\}, \{C4:C8\}, \{D4:D8\}, \{E4:E8\}\}$, and formulaGroups(Investment) returns $\{\{B3\}, \{B5\}, \{E9:E11\}\}$.

Where necessary, formula groups are then divided into one-dimensional partitions: *partitioned formula groups* are one-dimensional areas, i.e., they either have a row-wise or column-wise orientation. When partitioning a formula group, we choose the orientation of the partitions in such a way that the smallest number of partitioned formula groups is created. In case of a draw, we opt for column-oriented groups because our analysis of a public spreadsheet corpus (EUSES, see Section 6 for further information) concluded that 34 % of worksheets contain references to column-oriented areas, whereas only 13 % feature references to row-oriented areas. Thus, it is more likely that a column-oriented group was applied by the spreadsheet's author.

Definition 13 (Partitioned formula group). A formula group C is a partitioned formula group g, if its cells form a one-dimensional group:

$$PFGroup(C) = \begin{cases} true & if formulaGroup(C) \land \{\{\forall c, c' \in C : x(c) = x(c')\} \lor \\ \{\forall c, c' \in C : y(c) = y(c')\}\} \\ false & otherwise. \end{cases}$$

The function partFormulaGroups(worksheet \boldsymbol{w}) returns all partitioned formula groups of worksheet \boldsymbol{w} .

Example 11. Since all formula groups in Example 10 are one-dimensional, the partitioned formula groups are the same.

After processing all partitioned formula groups of a spreadsheet, we first determine which *referred formula groups* are connected to each partitioned formula group via references from the cells within a group. **Definition 14 (Referred formula groups).** A partitioned formula group g has a group reference to another partitioned formula group g', if any of its cells refer to any cell within g':

$$refersTo(g,g') = \begin{cases} true & if \exists c : c \in g \land \rho(c) \cap g' \neq \emptyset\\ false & otherwise. \end{cases}$$

The function referred Formula Groups (Partitioned formula group g) returns the set of all partitioned formula groups to which g refers to.

Example 12. The Department1 worksheet has two partitioned formula groups: B8:F8 and F4:F7. Cell F8 in group B8:F8 refers to cells in group F4:F7; the function refers To(B8:F8,F4:F7) returns true. Hence, the function referred-FormulaGroups(B8:F8) returns {F4:F7}. Group F4:F7 does not refer to cells of any other partitioned formula group (referredFormulaGroups(F4:F7)= \emptyset).

Next, we establish *reference-based groups* using the references of each formula group. We thus identify connected areas of cells that serve as input for a specific calculation in the spreadsheet (a partitioned formula group).

Definition 15 (Reference-based group). A reference-based group g_r is a set of cells that are referred to by a partitioned formula group g. Each group g_r can be attributed to either a specific reference r or a specific area reference r_a of g. In case of a reference, g_r is the collection of all cells that are referred to by any cell of g via the reference r. In case of an area reference, g_r is a collection of all cells that are referred to by a specific cell of g via the area reference r_a . The function reference Groups (partitioned formula group g) returns the set of reference-based groups $\{g_r\}$ for a partitioned formula group g.

 $referenceGroups(g) = refGroups(g) \cup areaRefGroups(g)$

$$refGroups(g) = \bigcup_{r \in refs(f)} \left\{ \bigcup_{c \in g} \{deref(r, c)\} \right\}$$
$$areaRefGroups(g) = \bigcup_{r_a \in refs_a(f)} \left\{ \bigcup_{c \in g} deref(r_a, c) \right\}$$

where f = formula(g). The function referenceGroups(worksheet \boldsymbol{w}) returns for a worksheet \boldsymbol{w} the set of all reference-based groups of \boldsymbol{w} .

Example 13. The formula of the partitioned formula group Investment! E9:E11 (Figure 1c) has five cell references (refs(Investment!E9:E11)=RC[-3], RC[-2], R5C2, RC[-1], R4C2). Therefore, it has five reference-based groups, namely {B4}, {B5}, {B9:B11}, {C9:C11}, and {D9:D11}.

The formula of the partitioned formula group Department1!F4:F7 (Figure 1a) has one area reference (refs_a (Department1!F4:F7)={RC[-3]:RC[-1]}). Therefore, the group refers to four reference-based groups, one for each of the four cells in the group: {B4:E4}, {B5:E5}, {B6:E6}, and {B7:E7}.

To allow for concise further processing of higher-level structures, we then merge all overlapping reference-based groups that have the same orientation, i.e. each cell is at most part of one reference-based group in vertical and one group in horizontal orientation. The remaining non-empty cells that are contained neither in any partitioned formula group nor in any reference-based group are designated as *non-blockable cells*, in preparation of the following Blocking step of the analysis.

Definition 16 (Non-blockable cell). The function non-blockable(cell c) returns true, if c is non-blockable:

$$non-blockable(c) = \begin{cases} true & if \ type(c) \neq empty \land \{ \nexists g \in G : c \in g \} \\ false & otherwise. \end{cases}$$

where $G = partFormulaGroups(ws(c)) \cup referenceGroups(ws(c))$. The function non-blockables(worksheet w) returns the set of all non-blockable cells contained in w (non-blockables(w) = { $c \in cells(w) : non-blockable(c)$ }).

Example 14. The function non-blockables(Department1) returns {A1; B2; A3:F3; A4:A8}, non-blockables(Total) returns {A1; B2; A3:F3; A4:A8}, and non-blockables(Investment) returns {A1:A5; A7:A11; B8:E8}.

4.2. Blocking

In the blocking step, the formula groups and reference-based groups are aggregated to rectangular *areas*. These areas are called *blocks*, and each block contains input cells, formula cells, and empty cells, but no header cells.

Definition 17 (Area). Function area({cell} C, worksheet w) returns all cells contained in the rectangle spanned by the non-empty set of cells $C = \{c_1, \ldots, c_n\}$:

$$area(C, w) = \{c \in cells(ws(c_1)) : x_1 \le x(c) \le x_2 \land y_1 \le y(c) \le y_2\}$$

where $x_1 = \min\{x(c_1), \ldots, x(c_n)\}, y_1 = \min\{y(c_1), \ldots, y(c_n)\}, x_2 = \max\{x(c_1), \ldots, x(c_n)\}, and y_2 = \max\{y(c_1), \ldots, y(c_n)\}.$

Example 15. Area({A3, B4, C2}) computes $x_1 = 1$, $y_1 = 2$, $x_2 = 3$, and $y_2 = 4$ and returns as result {A2,A3,A4,B2,B3,B4,C2,C3,C4}.

Definition 18 (Block). A non-empty set of cells $C = \{c_1, \ldots, c_n\}$ forms a block if the area spanned by the cells C contains no non-blockable cells:

$$block(C) = \begin{cases} true & if \ \nexists c \in area(C) : non-blockable(c) \\ false & otherwise. \end{cases}$$

Example 16. For worksheet Departement1 (Figure 1a), $block(\{B4, D4\})$ returns true, but $block(\{A4, D4\})$ returns false, as cell A4 is non-blockable.

Blocks are established by expansion operations that add *block neighbors*, physically close groups, to an existing block. A group is regarded as physically close to a block if there is at most one column or row between the block and the group which should be added. We regard both, partitioned formula groups as well as reference-based groups, for block expansion.

Definition 19 (Block Neighbors [20]). The function neighbor (block b, group g) returns true if there is at most one row or column between b and g:

$$neighbor(b,g) = \begin{cases} true & if \ \exists c \in g, \ \exists c' \in b: |x(c) - x(c')| \le 2 \land y(c) = y(c') \\ true & if \ \exists c \in g, \ \exists c' \in b: x(c) = x(c') \land |y(c) - y(c')| \le 2 \\ false & otherwise. \end{cases}$$

Block creation for a worksheet follows the procedure of Algorithm 1): First, the set of blocks is initialized (Line 2) and the partitioned formula groups and the reference-based groups are computed (Lines 3-4). The sets G and G' are both initialized with the union of the previously computed groups (Line 5). While G' will be reduced in size during the computation of the blocks, G never changes. In the outer loop (Lines 6-13), new blocks are created as long as there are groups which are not yet part of any block. In the inner loop (Lines 9-12), groups are added to the blocks if they fulfil two criteria: (1) They must be neighboring to the block, and (2) the block and the group must form a valid block. Lastly, we return the inferred set of blocks for the given worksheet \mathbf{w} .

Algorithm 1 BLOCK CREATION

Re	quire: worksheet w
En	sure: set of blocks B
1:	procedure BLOCKS(worksheet w)
2:	$B \leftarrow \emptyset$
3:	$G_p \leftarrow partFormulaGroups(w)$
4:	$G_r \leftarrow referenceGroups(w)$
5:	$G, G' \leftarrow G_p \cup G_r$
6:	$\mathbf{while} \; \exists g \in G' \; \mathbf{do}$
7:	$b \leftarrow g$
8:	$G' \leftarrow G' \setminus g$
9:	for all $g' \in G$ do
10:	${f if} \ neighbor(b,g') \wedge block(b \cup g') \ {f then}$
11:	$b \leftarrow b \cup g'$
12:	$G' \leftarrow G' \setminus g'$
13:	$B \leftarrow B \cup \{b\}$
14:	return B

Figure 3 illustrates scenarios that might occur during the block creation. In the sub-figures (a), (b), (c), all groups can be merged to one large block. Sub-figures (d) and (e) additionally contain non-blockable cells. These cells prevent

us from building a single large block. In both cases, a second block (green hatched border) has to be created for the group to the right. We allow groups to belong to several blocks. Hence, groups within the first block (blue solid border) might be added to the second block (green hatched border). In sub-figure (d), the groups of the first block have row-orientation. Hence, adding any of these groups to the second block would violate the block criteria according to Def. 18. In contrast, all groups in sub-figure (e) are column-oriented. Therefore, some of the groups of the first block can also be added to the second one.



Figure 3: Block creation. Gray cells represent grouped cells, red cells represent non-blockable cells (i.e., non-empty cells that are not part of any group), and frames represent blocks.

Example 17. For the individual worksheets of our running example, the following blocks are computed: Department1!B4:F8, Total!B4:E8, Investment! B3:B5, and Investment!B9:E11.

4.3. Header Assignment

In the third and final part of the structural analysis process, headers are assigned to blocks. Headers are non-empty cells which are not part of a block, i.e., they are elements of the set of non-blockable cells.

The position of headers depends on the writing system: The Left-To-Right (LTR) system, used in western countries, places headers left to and/or above blocks; the Right-To-Left (RTL) system, used in Arabic countries, places headers right to and/or above blocks. In the LTR system, the left-most cells have the lowest column index; in the RTL system, the right-most cells have the lowest column index. Therefore, we assume that headers of blocks have lower row- or column indices than the cells of the block. In the following, we use the word 'left' to express that a cell has a lower column index than another cell, i.e., we focus on the LTR system, but the approach works similar in the RTL system.

Two types of cells can be located between a header cell h and its block b: empty cells and other header cells. If a header cell is between h and b, then his a higher level header, also called meta-header. If there are no other header cells between h and b, then h is a low-level header.

We call areas which contain headers layers. There are row- and column layers. Row layers are vertical groups which are located left to a block; they have the same number of rows as the block. Column layers are horizontal groups which are located above a block; they have the same number of columns as the underlying block. Figure 4 illustrates the position and shape of layers.



Figure 4: Column- and row layers for a block. The dark-shaded layers contain low-level headers; the light-colored layers contain higher-level headers.

We identify the column layers for a block b by investigating the horizontal areas in the row above b. If at least one of the cells of this area contains a non-blockable cell, a new layer has been detected. In succession, we check the area above this layer. We repeat this process until we reach the last row of potential headers or another block. We analogously detect the row header layers.

Definition 20 (Layers). The function columnLayers(block b) returns the set of detected column header layers of block b. Each header layer is described by the

set of cells within its area. Likewise, the function rowLayers(block b) returns the set of detected column header layers of block **b**.

Example 18. For our running example, we detect the following layers:

Block	Column layers	$Row \ layers$
Department1!B4:F8	B3:F3, and B2:F2	A4:A8
Total!B4:E8	B3:E3 and B2:E2	A4:A8
Investment!B3:B5	-	A3:A5
Investment!B9:E11	B8:E8	A9:A11

The layers Department1!B2:F2 and Total!B2:E2 contain higher-level headers; the other layers contain low-level headers.

In the next step, we check the remaining non-blockable cells for being metaheaders: If a non-blockable cell c is left of a column layer or above a row layer, c is a meta-header and has to be linked to the corresponding layer. If c is left to a column layer as well as above a row layer, we assign c to the row layer, because row headers typically act as identifiers for single items of a set while column headers typically act as descriptive headers for different characteristics that are recorded for the set. Hence, c is more likely to provide a common category for the underlying set of item-labels than a common category for the set of neighboring category labels.

Example 19. The following non-empty cells neither belong to a block nor to a layer: for the worksheets Department1 and Total, the cells A1 and A3, and for the worksheet Investment, the cells A1, A2, A7, and A8. We are able to assign some of these cells to layers:

Cell	Layer
Department1!A3	A4:A8
Total!A3	A4:A8
Investment!A2	A3:A5
Investment!A8	A9:A11

4.4. Comparison with previous work

We are not the first dealing with cell classification and header assignment for spreadsheets. As mentioned in Section 7, Abraham and Erwig [22] have identified header cells and core cells as part of the UCheck approach. UCheck assigns one of four roles to each cell: (1) header (i.e. labels), (2) footer (i.e. formula cells at the end of rows and/or columns which aggregate information), (3) core (i.e. data cells), and (4) filler (i.e. empty or particularly formatted cells which separate tables within spreadsheets). Thereby, UCheck uses different techniques to assign these roles, e.g. fence identification, content-based cell classification, and region-based cell classification. Based on the results of the cell classification, UCheck assigns first-level headers to core and footer cells and higher-level headers to header cells. Our analysis process is, in principle, based on the ideas of the cell classification and header assignment rules of UCheck. However, UCheck is not capable of identifying all header cells correctly. For example, UCheck fails to identify the headers for the quarters (cells A4:A7 of worksheet Total) and the headers for the departments (cells B2:E2 of worksheet Total) of the running example as headers. Moreover, our analysis not only identifies cell roles, but also provides information about cohesive structures (groups/blocks) within a worksheet.

5. Improved and New Spreadsheet Smell Detection Techniques

In this section, we demonstrate how the detected high-level structures can improve existing spreadsheet smell detection techniques, and how new smell detection techniques can be derived from the structure information.

5.1. Improved Smell Detection Techniques

The detected spreadsheet structures provide a number of opportunities to improve existing smell detection techniques. Table 1 presents basic refinement ideas, for smells presented by Cunha *et al.* [11], and Hermans *et al.* [12, 13].

Name	IS	\mathbf{FS}	IWS	Refinement Suggestion
Std. Deviation	•			Compare within group instead of column/row
Empty Cell	•			Report connected vacant areas in blocks
Pattern Finder	•			Compare within group instead of column/row
String Distance	•			Compare within group/block instead of sheet
Ref. to Empty Cells	•			Highlight group instead of individual cells
QFD	•			Compare within block instead of column/row
Multiple Operations		•		Report group instead of individual cells
Multiple References		•		Count group references instead of cell references
multiple itelefences		•		Report group instead of individual cells
Cond. Complexity		•		Report group instead of individual cells
Long Calc. Chain		•		Count group references instead of cell references
Long Cale. Chain		•		Report group instead of individual cells
Duplicated Formulas		•		Detect duplicated formula groups
Duplicated Formulas		•		Report group instead of individual cells
Inappr. Intimacy			•	Count group references instead of cell references
Feature Envy			•	Count group references instead of cell references
Middle Man			•	Report group instead of individual cells
Shotgun Surgery			•	Count changing groups instead of formulas

Table 1: Refinement suggestions for smells presented in the literature, categorized in input smells (IS), formula smells (FS), and inter-worksheet smells (IWS).

Since similar smell detection techniques often can benefit from structure information in a similar way, we propose exemplary improvements for one representative of each smell group: (1) for *input smells* we improve the *Pattern Finder* smell, (2) for *formula smells* we refine the *Long Calculation Chain* smell, and (3) for *inter-worksheet smells* we improve the *Feature Envy* smell. For each investigated smell detection technique, we first discuss the original technique and its deficits. We then explain how to improve the smell detection process using structural information, and lastly discuss the benefits and drawbacks of the proposed improvements.

5.1.1. Pattern Finder

Cunha *et al.* [11] proposed the original Pattern Finder smell detection technique. The smell detects cells that break a pattern that holds for the other cells of the same row or column, e.g., a constant in a row of formula cells or a string in a column of numerical values. According to Cunha *et al.* [11], this smell is detected by checking in windows of four cells if one of these cells has a different type than the other cells. The authors provide an implementation of this technique in the *FaultySheet Detective* tool [15]. Examination of the tool⁴ furthermore provided the following insights: (1) Patterns are detected in column orientation only. (2) A cell's type refers to the type of its value, i.e., a cell with a constant numeric value within a series of formula cells which evaluate to a numeric value would not be indicated as smelly. (3) A broken pattern can only be detected if no other cell in a 5-cell-distance above or below the cell features the same evaluated type. (4) The smell is not detected for cells within the top or bottom five rows of a worksheet.

Pattern Finder attempts to establish patterns for entire columns of a worksheet. However, columns do not necessarily feature uniform content. For example, string cells are widely used to provide header information of a column. Similarly, cells at the bottom of a computation block might aggregate the data entered above or perform simple checks w.r.t. the validity of the above data. This explains why the top and bottom 5 rows are excluded. However, if a worksheet consists of several computation blocks, this workaround does not work.

We propose to focus smell detection on reference-based groups instead of generic sliding windows. This improves over the current detection process by allowing row-wise pattern detection and by extending the smell detection to the top and bottom five rows. Moreover, we base the smell detection on nonevaluated cell types. This allows for the detection of instances where formulas and values are mixed in the same group.

Algorithm 2 describes the updated Pattern Finder smell detection process. It detects reference-based groups whose cells feature more than one cell type. We iterate over the reference-based groups of the worksheet, checking for the smell (Lines 3 to 8). In the inner loop (Lines 5 to 8), we check whether group G contains the smell: If one of the cells has a different type, we add the group to the set of afflicted groups. The algorithm returns this set as result.

Example 20. The modified Total worksheet (Figure 5) illustrates the Improved Pattern Finder smell detection technique: Cell D4 has been changed from a formula to a fixed value. The original Pattern Finder does not indicate D4 as smelly, because every cell in column D evaluates to a number, but our improved Pattern Finder does. Indicating a constant within a group of formula cells as smelly helps users to detect formula cells which have been accidentally overwritten with constant values.

⁴Version 1.1 from http://ssaapp.di.uminho.pt/twiki/bin/view/Main/Software

Algorithm 2 DETECTPATTERNFINDER

Red	quire: worksheet w	
Ens	sure: reference groups afflicted by the Pattern Finder smell	
1:	procedure PatternFinderGroups(worksheet w)	
2:	$AfflictedGroups = \emptyset$	
3:	for all $G \in referenceGroups(w)$ do	
4:	Type \leftarrow type of first cell in G	\triangleright init
5:	for all $c \in cells(G)$ do	
6:	if Type $\neq type(c)$ then	
7:	$AfflictedGroups = AfflictedGroups \cup \{G\}$	
8:	break	
9:	return AfflictedGroups	

	А	В	С	D	E						
1	Coffee Consumption: All Departments										
2		Department									
3	Quarter	1	2	3	Total						
4	1	='Department 1'!F4	='Department 2'!F4	='Department 3'!F4	=SUM(B4:D4)						
5	2	='Department 1'!F5	='Department 2'!F5	='Department 3'!F5	=SUM(B5:D5)						
6	3	='Department 1'!F6	='Department 2'!F6	='Department 3'!F6	=SUM(B6:D6)						
7	4	='Department 1'!F7	='Department 2'!F7	468	=SUM(B7:D7)						
8	Total	='Department 1'!F8	='Department 2'!F8	='Department 3'!F8	=SUM(B8:D8)						

Figure 5: Example of Improved Pattern Finder smell detection technique

The focus on reference-based groups provides three key benefits in comparison to the original smell detection process: First, the algorithm compares cells within well-defined borders. This prevents the smell to be accidentally detected for column headers. Hence, the group-based approach reduces the number of false positives. Second, location limitations of the current detection approach do not apply for the updated algorithm. Groups in horizontal orientation and border areas of the worksheet are eligible for smell detection. Third, the updated algorithm only checks each reference-based group once, instead of checking every possible position of a sliding window, which makes the smell detection faster.

Setting the focus of the detection process on reference-based groups also introduces some drawbacks. First, smell detection is only applied to areas which are used as input values for calculations. Hence, the smell cannot be detected for areas containing output formulas, non-computational values, and labels. Further, the smell assumes that every referenced cell should contain a value at the time of smell detection. Consequently, blank spots which are reserved to be filled by a spreadsheet's user (as often used in form spreadsheets) will wrongly be indicated as smelly. However, if all cells in the same reference group are empty, they feature the same type. In such a case, no smell is reported.

5.1.2. Long Calculation Chain

Hermans *et al.* [12] proposed the original Long Calculation Chain smell detection technique, which detects formulas referring to a long chain of formulas, because long calculation chains are difficult to follow and to understand. This smell is detected by computing the maximum number of references which need to be followed when evaluating a formula. If this number exceeds a certain threshold, the formula cell is indicated as smelly.

The main drawback of this smell detection technique is its tendency to cause redundant calculations and detection notifications. Neighboring cells with the same R1C1 formula usually share the same preceding calculations. Hence, detectable issues for these cells can be traced back to one and the same general structural flaw. Computing the length of the calculation chain for each cell individually is inefficient. Moreover, each afflicted cell is reported individually. As a remedy, we propose to apply smell-detection on formula groups and inter-group dependencies instead of individual cells and cell references:

Definition 21. The function longestChain(partitioned formula group g) calculates the longest chain of group g as follows:

$$longestChain(g) = \begin{cases} longestChain(G) + 1 & if G \neq \emptyset \\ 1 & else \ if G_r \neq \emptyset \\ 0 & otherwise \end{cases}$$

where $longestChain(G) = max\{l | \forall g \in G : l = longestChain(g)\},\ G = referredFormulaGroups(g), and G_r = referenceGroups(g).$

We compute the length of the longest formula group chain of formula group \mathbf{g} by adding 1 to the longest chain of the formula groups to which \mathbf{g} refers to. The chain of a formula group that has only references to input cells has a length of 1. The chain of a formula group that has no references has a length of 0.

Like the original smell detection technique, the detection function of the group-based Long Calculation Chain smell returns a metric value. To decide on whether any given partitioned formula group is smelly, a threshold for the calculated metric is required. Groups whose calculated metric exceeds this threshold are indicated as smelly. Hermans *et al.* proposed a threshold of 4 to indicate a small risk, and a threshold of 7 to indicate a high risk.

Example 21. The Investment worksheet (Figure 1c) illustrates the benefits of the Improved Long Calculation Chain: Cells E9, E10, and E11 are part of the partitioned formula group E9:E11. Each cell has a longest calculation chain of length 7. The specific references for each cell differ. However, the overall contextual structure of the calculation is shared among all cells. A specific reference chain path for cell E9 is Department1!B4 \rightarrow Department1!F4 \rightarrow Department1!F8 \rightarrow Total!B8 \rightarrow Total!E8 \rightarrow Investment!B3 \rightarrow Investment! B5 \rightarrow Investment!E9. Similar paths can be reported for cells E10 and E11. Alternatively, the Improved Long Calculation Chain reports only one calculation path for the entire group Investment!E9:E11. One possible group reference chain is Department1!B4:E4 \rightarrow Department1!F4:F7 \rightarrow Department1!B8:F8 \rightarrow Total!B4:B8 \rightarrow Total!E4:E8 \rightarrow Investment!B3:B3 \rightarrow Investment!B5:B5 \rightarrow Investment!E9:E11.

Checking for long calculation chains on a per-group basis has several benefits: First, associated smells can be reported once per group instead of individually for each cell. This helps to reduce the number of reported smells. Second, per-group detection provides users with additional context to facilitate understanding of the overall calculation structure of the spreadsheet. A better mental model of a spreadsheet enables users to introduce well-considered changes. Third, group-wise smell detection implies that each reference is computed only once for an entire group of related formula cells. Hence, the detection approach may require substantially less individual references to be checked.

Setting the focus to group-based references introduces one key flaw: inconsistent groups. Inter-group references are not necessarily consistent for each individual cell within a group. Hence, the smell might be reported for a cell within a group even though this cell is not affected by the smell on a per-cell basis. As a remedy, only the partition of a formula group which is affected by the smell could be reported instead of the entire group. However, this would require per-cell detection of the smell in combination with per-group detection, neutralizing the calculation performance benefit of the improvement. Moreover, although usual spreadsheet programs prohibit circular references on a per-cell reference basis, inconsistent group references might introduce circular reference paths in-between formula groups. Such instances need to be handled correctly when calculating the metric's value.

5.1.3. Feature Envy

Hermans et al. [13] proposed the Feature Envy smell. This smell detects worksheets which contain a large number of references to other worksheets. It is difficult to follow many different relations to other worksheets when trying to understand or debug a spreadsheet. Feature Envy reports worksheets for excesses in the number of individual connections to other worksheets. However, even a limited number of semantically different connections to other worksheets can render a worksheet difficult to comprehend and, thus, should be reported as smelly. Moreover, advanced tasks, e.g., elaborate data analysis, require a greater number of processing steps. Spreadsheet creators fulfil such tasks in two different ways: (1) They add more functionality into individual formulas, or add more formula cells to the worksheets. (2) They add new worksheets that refer to interim results. The first way increases either the complexity of the individual formula cells or the size of the worksheet; both consequences make a worksheet more difficult to understand. Therefore, the second way becomes the preferred option at some point. Consequently, we argue that a high number of semantically equivalent connections to the same worksheet should not necessarily be indicated as smelly. To allow for a more purposeful smell detection, we propose to base the smell detection on references of partitioned formula groups instead of individual formula connections.

Algorithm 3 describes the updated Feature Envy detection process. The function $ws(reference-based group \mathbf{g}_r)$ permits access to the worksheet of group g_r . The function countFeatureEnvyConnections(worksheet \mathbf{w}) in Line 1 counts the number of total smell occurrences within worksheet w. It first initializes the variable Count with the value 0. It then iterates the partitioned formula groups of worksheet w. For each group g, the function iterates the set of reference-based groups to which g refers to. For each referred group g_{ref} , the function increments Count if g_{ref} 's worksheet differs from worksheet w. Lastly, the function returns Count, the number of group references to other worksheets.

Algorithm 3 DETECTFEATUREENVY

Require: worksheet w

Ensure: number of connections from formula groups in w to other worksheets 1: **procedure** COUNTFEATUREENVYCONNECTIONS(worksheet w)

 2:
 $Count \leftarrow 0$

 3:
 for all $g \in partFormulaGroups(w)$ do

 4:
 for all $g_{ref} \in referenceGroups(g)$ do

 5:
 if $ws(g_{ref}) \neq w$ then

 6:
 $Count \leftarrow Count + 1$

 7:
 return Count

Example 22. The cells in the area B4:D8 of worksheet Total (Figure 1b) feature 12 references to other worksheets. However, the cells in each column can be grouped into the partitioned formula groups B4:B8, C4:C8, and D4:D8. Hence, the Improved Feature Envy only reports three inter-worksheet connections.

The detection function of the group-based Feature Envy smell returns a value. Worksheets whose calculated metric value exceeds a certain threshold are indicated as smelly. Hermans *et al.* proposed a threshold of 3 to indicate a small risk, and a threshold of 7 to indicate a high risk.

The proposed improvement offers two main benefits: First, applying group connections for the calculation of the detection metric provides users with more meaningful feedback in regard to the overall quality of the connection structure of the spreadsheet. This supports users in making high-level structural improvements, eliminating the root cause of indicated smells instead of alleviating its effects. Second, group-based smell detection only requires to check for interworksheet connections of each partitioned formula group, instead of checking the connections of all cells of a worksheet. Hence, this detection approach may require substantially less individual references to be checked.

The drawback of using group connections is a potential loss of information. While a high number of semantically similar inter-worksheet connections might be a necessary design, reporting the circumstance might still offer an opportunity for improvement. As a remedy, the cardinality of processed group connections might be introduced into the detection metric, be reported to the user as contextual information, or both.

5.2. New smell detection techniques

We elaborated different approaches to formulate new smell detection methods that are based on the structural information. From this list of fundamental smell ideas, we present three smells that showcase utilization of different aspects of the available structure information: The Overburdened Worksheet smell indicates that a worksheet contains too much functionality. The Inconsistent Formula Group Reference smell signals inconsistencies occurring within the group-resolving step of the analysis process. The Missing Header smell indicates gaps in the headers of a block. For each of the introduced smells, we first outline its basic concept. We then present the smell detection process and provide an example. Lastly, we highlight benefits and possible limitations of the new smell, and explain its significance related to the overall quality of a spreadsheet. For doing so, we will refer to the ISO/IEC 25010:2011 International Standard for System and Software Quality Models [23] and the quality model [24] for spreadsheets that is based on a predecessor of this standard.

5.2.1. Overburdened Worksheet

Each part of a spreadsheet serves a specific purpose: a reference group provides a set of common input data for further calculations; a formula group performs a calculation on sets of input data; a block distinguishes functionally enclosed areas of a worksheet. The Overburdened Worksheet smell indicates worksheets which feature an excessive number of any structure type, e.g., blocks:

Definition 22 (Overburened Worksheet). The function overburdened Worksheet (worksheet w) returns the detection metric for the smell:

overburdenedWorksheet(w) = |blocks(w)|.

To decide on whether any given worksheet is smelly, a threshold for the calculated value is required. Worksheets whose calculated value exceeds this threshold are indicated as smelly.

Example 23. The Investment worksheet (Figure 1c) has two calculation blocks: B3:B5 and B9:E11. If we set the threshold of this smell to the extremely low level of two, the worksheet would be indicated as smelly.

As worksheet *Investment* contains a minimal example, it remains comprehensible despite featuring multiple calculation blocks. However, in general, multiple blocks in a worksheet indicate a suboptimal spreadsheet structure, which can be easily resolved by moving some blocks to a new worksheet.

The provided detection function utilizes the number of calculation blocks per spreadsheet as significance metric. However, other structure information may be employed, as well. We included the number of formula groups per worksheet as an additional metric in our evaluation. Another possibility would be to count the reference-based groups of a worksheet, or the number of intra-worksheet connections in-between formula groups. When relying on basic spreadsheet information, the number of cells or the number of formulas might also be employed to indicate an overburdened worksheet.

The Overburdened Worksheet smell provides a natural counter-balance to existing inter-worksheet smells indicating worksheets that refer too abundantly to other worksheets. The goal for an overall, optimal spreadsheet structure is then a balanced partition of the required functionality over a number of worksheets which neither overburdens any individual sheet nor renders any sheet overly reliant on inter-worksheet connections.

The quality and success of the Overburdened Worksheet's smell detection process highly depends on the success of the previous structural analysis process. Inconclusive structural analysis might result in an excessive number of small structures. In such a case, the size-based metrics provide misleading information. Smell metrics which combine the quantity of structures with their respective size might lessen the influence of ambiguous structure analysis.

While an overburdened worksheet does not influence the functionality and security of a spreadsheet, it influences the maintainability (in particular the subcatetory analyzability, see ISO/IEC 25010:2011 standard [23]) and usability (subcategory understandability): A spreadsheet that has a clear modular structure with worksheets as modules is easier to understand and maintain than a spreadsheet that contains all information in a single worksheet.

5.2.2. Inconsistent Formula Group Reference

The Inconsistent Formula Group Reference smell highlights an inconsistency that becomes apparent during the structural analysis process. Common spreadsheet programs already point out inconsistencies regarding the formulas of groups of related cells. Inconsistent Formula Group Reference points out inconsistencies regarding references to individual cells of such groups.

More elaborate worksheets depend on reference chains, linking sequential calculations. Structural analysis enables tracking of references in between formula cells, as well as references in between formula groups. However, references in between formula groups are not always concise. For example, one group might refer only to a subset of the cells of another group. Inconsistent Formula Group Reference points out those instances.

Definition 23 (Inconsistent group reference). The function

inconsGroupRef(partitioned formula group g, partitioned formula group g') identifies inconsistent group references:

$$inconsGroupRef(g,g') = \begin{cases} true & if \nexists g_r \in referenceGroups(g) : g_r \equiv g' \\ & \land \exists g_r \in referenceGroups(g) : g_r \cap g' \neq \emptyset \\ false & otherwise. \end{cases}$$

Example 24. Our running example in Figure 1 contains an occurrence of the Inconsistent Formula Group Reference smell. In the Investment worksheet, cell B3 creates a formula group that refers to the single cell E8 of the Total worksheet. However, the cell Total!E8 is part of the formula group Total!E4:E8. Hence, Investment!B3:B3 inconsistently refers to Total!E4:E8.

The smell points to inconsistencies within reference chains in between formula groups. Such inconsistencies may be introduced during the creation or expansion of the spreadsheet, e.g., a newly created set of calculations mistakenly refers to only a part of preceding formulas, or an inner part of a calculation chain is expanded, but successive calculations are not updated accordingly.

As demonstrated in the example, referring only to a part of a formula group may not always indicate an error, but be the intended behavior. However, even in those cases, a spreadsheet may be restructured to remove the inconsistency. Moreover, inconsistency detection depends on the success of the previous structural analysis process. Incorrect grouping, poor partitioning of formula groups, or resolving of group references might lead to false positive detection.

A spreadsheet with inconsistent references is difficult to analyze. Inconsistent references might point to faults caused by expanding a spreadsheet. While some parts of a spreadsheet are updated, others might be forgotten to be updated. When the forgotten updates lead to errors, the spreadsheet does not provide the intended functionality. Hence, the Inconsistent Formula Group Reference influences the overall quality of a spreadsheet with respect to its analyzability (subcategory of maintainability) and to a certain extent to its functionality (see ISO/IEC 25010:2011 standard [23]).

5.2.3. Missing Header

Headers are not always provided for each column and/or row of a block. This results in empty spots within the header layers of affected blocks. The Missing Header smell reports cases of such vacant spots.

Definition 24 (Missing headers). The function missingHeaders(block b) returns the set of missing header cells of block b:

$$missingHeaders(b) = \{ c \in (\bigcup_{l \in L_{col}} l) \cup (\bigcup_{l \in L_{row}} l) : type(c) = `empty' \}$$

where $L_{col} = columnLayers(b)$ and $L_{row} = rowLayers(b)$.

Example 25. Figure 6 illustrates the Missing Header smell. It depicts a revised version of the Department1 worksheet of our running example. For demonstration purposes, we have removed the label of cell D3. Structural analysis detects a block in area B4:F8. Column-headers for this block are available in row 3. However, one spot in the header layer of the block, cell D3, is vacant.

The Missing Header smell gives feedback about the quality of non-calculation parts of worksheets. Missing headers impair comprehensibility of worksheets, as calculation relations do not necessarily provide contextual information.

	A	В	С	D	E	F					
1	Coffee Consumption: Department 1										
2		Employee									
3	Quarter	Anderson	Bourne		Dredd	Total					
4	1	35	172	80	98	=SUM(B4:E4)					
5	2	47	200	51	66	=SUM(B5:E5)					
6	3	46	139	46	39	=SUM(B6:E6)					
7	4	74	155	57	43	=SUM(B7:E7)					
8	Total	=SUM(B4:B7)	=SUM(C4:C7)	=SUM(D4:D7)	=SUM(E4:E7)	=SUM(F4:F7)					

Figure 6: Example of Missing Header smell

Inference of headers and header layers is directly dependent on the preceding block detection results. Hence, header detection carries on the same limitations as affected the previous analysis steps. For example, the current blocking approach does not compute blocks for tables that only collect data, but do not process it. Thus, no headers can currently be inferred for such tables. Another drawback of the proposed analysis method is a high likelihood of false positives in higher-order header layers. Items of such layers usually provide contextual information to underlying header layers; they are, therefore, usually not completely filled in. This is an intended behavior. Nevertheless, the current definition would count such instances as missing headers.

Moreover, conflicts may occur, whereby meta-headers can be assigned to both, underlying column- and row header layers. Following the current header assignment process, such conflicts are resolved by a static default decision. To attain more reliable results for higher-order headers, structural analysis would benefit from a more elaborate approach to correctly decide ties.

Badly or even undocumented spreadsheets are obviously difficult to understand. Hence, missing headers influence the overall quality of a spreadsheet by a reduced understandability (subcategory of the usability characteristic in the ISO/IEC 25010:2011 standard [23]).

6. Empirical Evaluation

In this section, we evaluate the performance of the improved and new smell detection techniques. We first outline our study design. We then present the details and results and lastly, we discuss the presented results.

6.1. Study Design

Study Rationale. The rationale of this study is to evaluate whether structural information improves the detection of spreadsheet smells. The improvement potential for smells varies, based on the source smell and on how structure information can be applied. In general, we assume that the improvement of the detection techniques results in a reduction of false positives, and limits the number of redundant smell detections. Newly introduced smell detection techniques are expected to perform similar to already established ones. The performance analysis of the underlying structural detection process is out of the focus of this empirical evaluation. We refer the interested reader to our IWPD paper [20] for a detailed evaluation of the detection performance for blocks, groups, and headers.

Objective/Units of study & Context. The objective of this particular evaluation is to investigate the detection performances of the improved and new spreadsheet smell detection techniques. The context of our study is the EUSES corpus, a publicly available collection of spreadsheets. Units of analysis are the sets of spreadsheet smells: original smells, improved smells, and new smells. For each of these sets, we record the respective detection metrics when applied to spreadsheets of the EUSES [25] corpus.

Research questions.

- 1. Can existing smell detection techniques be improved by applying structural information in the smell's detection process?
- 2. Are novel spreadsheet smell detection techniques that are based on structural information able to detect new quality issues, and do they perform similar to traditional ones?

Concepts & Measures. As established by previous research [12][13], we use detection metrics of smells as basis of analysis and comparison; to determine whether a given entity is smelly, every smell detection technique calculates a metric value for the entity. The target entities for baseline techniques are either cells or worksheets, whereas entities for the improved and new techniques are either groups, blocks or worksheets. To allow for comparison of the approaches, we thus regard a cell as well as a group and a block as individual entities of a spreadsheet which a user has to check, if indicated as smelly. Metrics that are recorded per worksheet are directly comparable. Following common practice, we aggregate these metrics into quartile plots. Quartile plots aggregate the results of individual entities, illustrating the percentage \mathbf{x} of the analyzed entities that feature a metric value of \mathbf{y} or lower. To incorporate a wider range of metric values, we use a logarithmic scale for the y-axis. This allows us to compare our results with previous work.

Data Collection. Data collection is based on our own test implementation, $FRITZ^5$. The 'evaluation' command of FRITZ analyzes a supplied spreadsheet corpus and writes a selectable set of calculated metrics to CSV files for further processing. The evaluation uses files of the EUSES [25] spreadsheet corpus. The corpus can be downloaded from the tera-PROMISE Repository⁶. It contains 4490 files in 11 categories. Not all files of this corpus are fit for evaluation with FRITZ. In a preprocessing step, we exclude files which (1) are not readable by external library components used by the evaluation tool, (2) are not processable due to limitations of the evaluation tool, or (3) do not contain any formulas. This preprocessing operation is provided by FRITZ, using the

⁵available at http://spreadsheets.ist.tugraz.at/index.php/software/

⁶http://openscience.us/repo/spreadsheet/euses.html, last visited 2017-01-31

command 'preprocess'. This command supports different filtering options, e.g., the 'complete' option which applies all the previously mentioned filter criteria. The resulting filtered EUSES corpus consists of 1735 files in 10 categories. We then run the automatic 'evaluation' command offered by FRITZ, using the evaluation option 'SMELLS_COMPLETE'. FRITZ has a 5 minute timeout limit per file. Three files (*personal/FindFunction.xls, inventory/in_emit99.xls*, and *grades/PregnancyDiet.XLS*) exceed this limit. The stated results refer to the 1732 files that are fit for evaluation and do not cause a timeout.

Data Analysis. In order to compare our improvements to the original smell detection techniques, we have to introduce a common baseline. To establish this baseline, we collect data on the detection performance of our own re-implementations of the basic smell detection approaches. For our data analysis, we first compare the new baseline with the results of the smells' original authors where available. We then compare the collected data for the improved smell detection techniques to our baseline results. For our set of novel smells, we analyze the smell detection performance following the established analysis approach for spreadsheet smells and provide general remarks.

Case & Data Selection. The general case of the study is provided by the spreadsheets of the EUSES corpus. Units of analysis within the study are the sets of smells: baseline, improved, new. The data resulting of analyzing all eligible spreadsheets within the corpus is used for each analysis unit. Eligible spreadsheets are determined in a preprocessing step, using the FRITZ tool.

Replication. The focus of the present study is to build up on, rather than to replicate existing work. However, to allow for comparison, we had to replicate some parts of related publications, as the tools and data of those studies are no longer available. In order to guarantee that our work is replicable, we provide references to the used dataset and tools.

6.2. Smell Detection Improvements

Since the tools that were used for evaluating the original smell detection process are either not publicly available or not designed to support automatic evaluation using a spreadsheet corpus, we implemented the baseline smell detection techniques in our own evaluation tool, FRITZ. For each improved smell, we first compare our baseline implementation with the smell's original evaluation results (using it's original evaluation dataset). We then compare this baseline implementation with our improved variant, using the EUSES corpus as dataset. As we want to highlight the reduction in total detections caused by avoiding redundancies, the improved versions of the *Long Calculation Chain* and *Feature Envy* smells use the same thresholds for detection as the baseline techniques.

6.2.1. Pattern Finder

Cunha *et al.* implemented this technique in the FaultySheet Detective tool, and evaluated it using 180 selected spreadsheets of the EUSES corpus. To enable comparison, we reproduced the subset of spreadsheets after consultation of the authors. Since the FaultySheet detective tool does not support a batch-mode analysis, the manual execution of the tool for each spreadsheet and evaluation of the results is very time consuming. Hence, we chose to limit the comparison of our implementation with the author's original results using the FaultySheet Detective⁷ to the category "homework" (16 spreadsheets). However, we provide corpus archive as download⁸, to allow for validation of our results. Table 2 illustrates the analysis results using the following metrics:

- FaultySheet: smell instances detected by FaultySheet Detective.
- Relevant: amount of FaultySheet detections we regard as relevant, based on manual inspection. A relevant detection indicates a cell which discerns from the obvious intention of the spreadsheet's author (e.g., a number instead of a date). Other detections (e.g., labels and descriptions within a table) are not regarded as relevant.
- Cols: smell instances detected by FRITZ using column-oriented windows. This corresponds to FaultySheet Detective's detection approach.
- Rows: smell instances detected by FRITZ using row-oriented windows.

FaultySheet and Relevant numbers result from manual inspections of the resulting sheets by one researcher. For each metric, we provide the total number of detections for all analyzed worksheets (*Cells Total*), as well as the average and median number of detections per worksheet (*Cells Average* and *Cells Median*). Moreover, we state the number and percentage of worksheets that feature any detection (*Worksheets* (>0) and % *Worksheets* (>0)), as well as the average number of detections for these worksheets (*Cells Average* (>0)).

Metric	FaultySheet	$\mathbf{Relevant}$	\mathbf{Cols}	Rows
Cells Total	180	20	181	129
Cells Average	6.6	0.7	6.6	4.8
Cells Median	0.5	0	1	0
Worksheets (>0)	14	4	15	9
% Worksheets (>0)	50%	14%	54%	32%
Cells Average (>0)	12.9	5.0	12.1	14.3

Table 2: Pattern Finder detection metrics based on the homework folder of Cunha *et al.*'s evaluation set. Cell count average and median are calculated on a per-worksheet basis.

Our recorded detection numbers for the FaultySheet tool diverge from the numbers stated by the smell's authors [11]. Cunha *et al.* reported 58 detected Pattern Finder smells; 56 of these they categorized as "no smells", leaving two genuine smell detections. In contrast, we counted 180 smelly cells as detected by the FaultySheet Detective tool. Manual inspection categorized 20 of these as relevant detections, for example the number 38412 in a column labelled "Target

⁷Version 1.1 from http://ssaapp.di.uminho.pt/twiki/bin/view/Main/Software

 $^{^{8} \}tt http://spreadsheets.ist.tugraz.at/wp-content/uploads/EUSES_small.zip$

date for next steps" that otherwise contained proper date values. In comparison, our analysis tool managed to find the same 180 smell instances using column-oriented windows. Hence, the baseline implementation of the technique adequately reproduces the performance of the original approach. One additional instance was detected due to technical specifics of a utilized library component. When using row-based windows, FRITZ detected 129 smelly cells.

We thus proved the adequacy of our implementation when compared to the original tool. However, during this evaluation we also revealed a significant shortcoming in terms of relevant detections. We assume that specific implementation details are to blame for these shortcomings. To check how specific implementation choices influence the results, we devised a number of different interpretations of the base technique, and evaluated the approach on the EUSES corpus. Figure 7 illustrates the results of the evaluation of various interpretations of the Pattern Finder smell as implemented in FRITZ. Pattern Finder Column and Pattern Finder Row only use detection windows in the respective orientation and exclude detections in the first and last five columns and rows. The Pattern Finder Column -border and Pattern Finder Row -border metrics work the same, but suspend the border restrictions. The Pattern Finder Combined metric reports cases where a cell is indicated as smelly by both, a horizontal and a vertical detection window, and excludes detections in border areas. The Pattern Finder Combined -border metric works the same way, but also allows detections in the border areas. Outliers that exceed metric values of 100 are not depicted.



Figure 7: Quartile plot of various interpretations of the original Pattern Finder smell, evaluated on EUSES. See Section 6.1 Concepts & Measures for a description of quartile plots.

Pattern Finder Column detects more smell instances than its row-based counterpart. The detection numbers of Pattern Finder Combined is significantly lower than both, indicating that only a limited overlap exists between column-

based and row-based smell detections. When allowed to detect smells in border areas (-border), the detection rates of all approaches increase noticeably. The rise is especially significant for the *Pattern Finder Row* -border metric, where at least one smell instance is detected for more than 90% of worksheets. Borders of worksheets usually contain a high number of string cells that are used as descriptive headers and footers. Hence, we suspect that the additional detections in border areas include a high number of such cells, and can thus be regarded as false positive detections. Pattern Finder Combined -border also registers a moderate increase in its detection rate in comparison to the results of it's borderexcluding counterpart. Nevertheless, it identifies fewer smelly cells than the comparable row- and column-based approaches. We argue that the combination of row-based and column-based detection windows reduces the likelihood of false positive detections by counteracting the tendency of spurious detections in border areas. We consequently regard Pattern Finder Combined -border as the most comprehensible of the analyzed approaches, and add this measure as additional baseline for comparison with our *Improved Pattern Finder* approach.

Figure 8 compares two interpretations of our improved smell detection approach with the original technique, *Cunha Pattern Finder*, and the selected baseline variant, *Combined Pattern Finder -border*. The metric *Group Pattern Finder* illustrates the result of our implementation based on Algorithm 2. The *Group Evaluated Pattern Finder* first evaluates formula cells and uses the evaluated cell types for comparison with other cells. In total, *Cunha Pattern Finder Finder Combined -border* used as baseline, identified 8 003 smelly cells (1.3 detections per worksheet). In comparison, *Group Pattern Finder Finder* counted 49 667 group detections, (8.2 per worksheet), and *Group Evaluated Pattern Finder Combined -border* (7.3 per worksheet).



Figure 8: Quartile plot comparing the original Pattern Finder, our selected baseline variant, and two improved smell variants on the EUSES dataset.

As expected, Cunha Pattern Finder reports a significantly larger number of smell detections than the other approaches, detecting at least one smelly cell within 40 % of analyzed worksheets. Group Evaluated Pattern Finder's result is similar to the chosen baseline. Both evaluate formula cells before comparing the types, hence both approaches are likely to detect similar cases of the smell. The higher number of detections of the Group Evaluated Pattern Finder is likely due to cases where a pattern is broken only in one orientation, but not the other. The results of Group Pattern Finder follow the same overall trend as those of Group Evaluated Pattern Finder. Due to cases whereby the evaluation of formula cells conceals otherwise mismatching cell types, the number of individual detections is higher. As such instances might indicate genuine deficits, we suggest using Group Pattern Finder over the Group Evaluated Pattern Finder approach.

6.2.2. Long Calculation Chain

Figure 9 compares our evaluation results with the author's initial results. *Hermans Chain Length* indicates the evaluation result as published by Hermans *et al.* [12]. *Baseline Chain Length* is the result of our interpretation of the smell using FRITZ. *Group Chain Length* refers to the results of the improved smell detection version, as described in Section 5.1.2.



Figure 9: Quartile plot of Long Calculation Chain metrics

Baseline Chain Length computes an average chain length of 24.5, while Group Chain Length computes an average chain length of 2.1. When applying the threshold of 7 for a smell detection, the baseline approach identifies 84 031 formula cells as smelly, whereas the group-based approach identifies only 4879 groups as smelly. Consultation of the smell's original authors revealed that detailed numbers for Hermans Chain Length are no longer available. The presented numbers are thus extracted from the result plots given in [12]. When comparing significant features of the graph, our baseline result indicates a higher proportion of calculation chains as smelly than was indicated by the original results. Moreover, *Baseline Chain Length* features a greater proportion of formulas that feature chain lengths significantly longer than 10 than both the original and group-based smell versions. These deviations might be caused by specific design choices regarding the compared smell implementations, or by distinctions in the evaluation datasets introduced by pre-processing operations. The graph of *Group Chain Length* follows a similar progression as the graph of the baseline implementation. The number of individual detections, however, is noticeably lower. This is the expected result, as the same structural flaws are detected by both approaches. Nevertheless, the baseline implementation has a higher individual metric number due to redundant smell detections. In terms of detection rates, the baseline approach results in a substantially larger number of individual detections than the group-based approach. *Group Chain Length*, therefore, offers a more concise way of communicating these flaws.

6.2.3. Feature Envy

Figure 10 compares the results of the improved Feature Envy detection metric to the smell's baseline implementation and the author's original results. *Hermans Feature Envy* illustrates the evaluation result of the original smell detection technique as published by Hermans *et al.* [13]. *Baseline Feature Envy* is the result of our baseline interpretation of the smell in FRITZ. *Group Feature Envy* depicts the result of the improved smell detection process as described in Section 5.1.3.



Figure 10: Quartile plot of Feature Envy metrics

In terms of numbers, *Baseline Feature Envy* reports 3 058 758 inter-worksheet connections (505.4 connections per worksheet). In comparison, *Group Feature Envy* counts 102 694 inter-worksheet group-connections (16.9 connections per

worksheet). When applying the threshold of 7 for a high risk smell detection, the baseline approach detects 737 worksheets as smelly, whereas the group-based approach only results in 533 smell detections. Consultation of the smell's original authors revealed that detailed numbers for *Hermans Feature Envy* are no longer available. The presented numbers are thus extracted from the result plots given in [13].

When comparing significant features of the graph, Hermans Feature Envy detects noticeably more worksheets with a significant number of inter-worksheet connections than the other approaches. Based on Hermans et al.'s evaluation, about 70 % of the worksheets fall in this category, whereas both Baseline Feature Envy and Group Feature Envy start reporting significant detections at the 85 % threshold. Hermans Feature Envy also features wider plateaus of percentage-areas where worksheets contain the same number of Feature Envy connections, whereas the other approaches do not report similar features. Starting at the 90 % mark, the results of all approaches follow a similar trend. However, the maximal values of the approaches are noticeably different: Our implementations detect a number of worksheets with more than 1000 connections, but Hermans et al. reported no findings of this magnitude.

In terms of detection rates, *Baseline Feature Envy* results in a larger number of individual detections than *Group Feature Envy*. However, the group-based approach is based on the number of semantically different connections, instead of the total number of connections. It is therefore likely that the group-based approach prevents reporting of false positive smell detections.

6.3. New Smell Detection Techniques

Figure 11 compares the results of our novel smell detection techniques. Overburdened Worksheet Blocks and Overburdened Worksheet Groups are smell detection metrics as proposed in Section 5.2.1, with the former counting the number of calculation blocks, and the latter counting the number of formula groups per worksheet. Inconsistent Formula Group Reference illustrates the results of the corresponding smell as described in Section 5.2.2. Missing Header counts the number of missing headers as described in Section 5.2.3.

Overburdened Worksheet Blocks counts 24 006 blocks (3.9 counted per worksheet). In comparison, Overburdened Worksheet Groups counts 105 414 groups for the same worksheets (17.4 groups per worksheet). Moreover, worksheets feature 102 281 Inconsistent Formula Group References and 98 686 Missing Headers. The results of all novel smells follow a power law like distribution; each individual metric curve has a gentle slope at first and most of its variability on the tail. This is the usual case for smell metrics, as demonstrated by Hermans et al. [12][13].

The values of the metrics for *Overburdened Worksheet Groups* are significantly higher than their block-based counterparts. This is expected, as blocks aggregate multiple groups. For smell detection, both versions of the *Overburdened Worksheet* smell require a detection threshold. Worksheets are indicated as smelly only if this threshold is exceeded by the recorded smell metric. Following Hermans *et al.* 's recommendation, we provide threshold values for the 70 %,



Figure 11: Quartile plot of novel smell metrics

80%, and 90% marks of the smell metric curves in Table 3. Worksheets whose metrics surpass these thresholds are regarded as featuring low risk, medium risk, and high risk respectively of being affected by the related smell.

Smell Detection Technique	70 %	80%	90%
Overburdened Worksheet Blocks	4	5	9
Overburdened Worksheet Groups	11	19	37

Table 3: Overburdened Worksheet detection thresholds.

Inconsistent Formula Group Reference and Missing Header smells are reported for each individual instance that is detected within a worksheet, hence no detection thresholds have to be provided. Inconsistent Formula Group Reference has one or fewer detections for about 70% of the worksheets; Missing Header for about 63%. Both metrics surpass 10 or fewer detections per worksheet at 80% to 85%. Hence, about 15% of the worksheets have more than 10 individual detections of the smells. The upper ends of both metric curves exceed 100 detections per worksheet. In case of Inconsistent Formula Group Reference, these results are probably caused by VLOOKUP and similar spreadsheet functions which refer to entire areas of worksheets instead of individual groups, introducing a large number of inconsistent references. In case of Missing Header, the high metric numbers are likely caused by limitations in the header detection process and the currently applied focus of the header detection method.

6.4. Manual Investigation

To complement the empirical study we described above, we also conducted a manual investigation of detected smells. For this investigation, we again employed the *homework* category of spreadsheets of Cunha*et al.*'s evaluation dataset that we previously used to compare the results of the *Pattern Finder* smell in Section 6.2.1. For each of the basic, improved, and novel smells, we applied the smell detection techniques to all sheets in the collection using the suggested thresholds. We then tallied how many of each smell instance were detected, how many of the detections were relevant (signalled a genuine issue), and for the basic and improved techniques, how many detections were missing that were successfully indicated by the respective other technique.

Table 4 summarizes the results of this investigation. In general, significantly less detections were recorded for the improved and novel techniques. *Pattern Finder* was the most detected smell. However, many of the detections of the basic smell were spurious, and many of relevant cases that are detected of the improved version were missed by the basic implementation. The 9 *Missing* cases for the *Improved Pattern Finder* were multiplicities of the same structural issue of one worksheet. The *Long Calculation Chain* smell was detected in one spreadsheet only (see example below). Also, we found no detections for both versions of the *Feature Envy* smell, as the spreadsheets in this category do not sufficiently rely on inter-worksheet references. For the *New Techniques*, almost all detections were relevant and pointed out novel issues.

Smell Detection Technique	Detected	Relevant	Missing				
Basic Techniques							
Pattern Finder	181	20	80				
Long Calculation Chain	99	99	0				
Feature Envy	0	0	0				
Improved Techniques							
Pattern Finder	9	9	9				
Long Calculation Chain	0	0	0				
Feature Envy	0	0	0				
Novel Techniques							
Overburdened Worksheet	12	9	-				
Incons. Ref.	12	12	-				
Missing Header	5	5	-				

Table 4: Manual evaluation of detected smells based on the homework folder of Cunha $et\ al.$'s evaluation set.

Figure 12 illustrates an excerpt of the finalGrades.xls that was part of the manual investigation. This example contains three deficits, that were successfully indicated by spreadsheet smells: (1) The student numbers in Column A, after the first entry, are created by successive, self referencing formulas, instead of usually employed continuous numbers. The issue is highlighted by many *Long Calculation Chain* detections, as each formula after the 7th in the chain is considered smelly. It is also indicated by one instance of the *Group Pattern Finder* smell for the column. The basic *Pattern Finder* smell, in contrast, does not detect this issues. (2) Many cells in Column K are empty, but are referred to by subsequent calculations, that do not properly take missing values into account. This issue is, again, highlighted by *Group Pattern Finder* smell, detected for the

column, but was not detected by the basic smell version. (3) The calculations in Column L use different formulas, dividing the total in Column N bei either 7 or 8. The exact number to divide by should likely depend on the optional entry in Column K. However, closer investigation reveals that this is not properly implemented in the sheet. The issue is revealed by the *Inconsistent Formula Group Reference* smell, as the percentage calculation in *Column M* refers to a number of smaller *formula groups* in Column L that are created for the different formulas. Moreover, the issue is also detected by the *Overburdened Worksheet* smell, as many small formula groups are calculated in Column L, which exceeds the threshold for this smell.

	A	В	0	:	D	J		К	L	M	N
1		CM 290 ·	- 02 Ho	mew	ork					11.12.2003	
2		Fall 2003					Op	ional *			
3				0	Coleman	Amelco (2) How	ard (2)	AVE.	7.	total
4		Student Na	me Code	Name							
-5	#										
6	1		Doc		3	2			1,4	48%	10
7	2		Bugs		3	3			2,7	90%	19
8	3		Smee		3	2		2	2,4	81%	17
9	4		Minnie		3	21			1,9	62%	13
10	5		Beauty	, []	3	3		3	2,0	67%	16
11	6		Grump	9	3	2			2,7	90%	19
					(a) Va	lue view					
	Α	В	С	D	J	К	L	M		N	
1		CM 290 - 0						37966			
2		Fall 2003				Optional*					
3		Curdent Name	Code Name	Colemar	n Amelco (2) Howard (2)	AVE.	~		total	
5	#	Student Name	Lode Name		-						
6	1		Doc	3	2		=N6/7	=L6/3	=D6+E6+F6+G6+H6+I6+J6+K6		6
7	=A6+1		Bugs	3	3		=N7/7	=L7/3	=D7+E7+F7+G7+H7+I7+J7+K7		
8	=A7+1		Smee	3	2	2	=N8/7	=L8/3	=D8+E8+F8+	-G8+H8+l8+J8+k	8
9	=A8+1		Minnie	3	1		=N9/7	=L9/3	=D9+E9+F9+	-G9+H9+I9+J9+K	9

	(1)		•
- (b) Formul	a view
- 1	~	/ _	

=D10+E10+F10+G10+H10+I10+J10+K10

=N11/7 =L11/3 =D11+E11+F11+G11+H11+I11+J11+K11

=L10/3

=N10/8

Figure 12: Spreadsheet finalGRADES.xls of Cunha et al.'s evaluation set.

6.5. Discussion

Beauty

Grumpy

3

10 =A9+1

11 =A10+1

In the empiric evaluation of the *Feature Envy* smell, the baseline interpretation features significantly lower individual detection numbers. Indeed, it counts 8003 of 45010 detections. The results from the manual investigation suggest that this drop in the detection rate likely excludes a significant portion of the previously detected false positive smell instances. Further, the *Improved Pattern Finder* detection reveals more smell instances than our chosen baseline (i.e., the non-evaluated version reports 49667 detections). This increase of the detection rate is likely attributable to genuine smell detections. This is also in line with the results from our manual investigation, where the improved technique revealed additional issues that were not detected by the original technique.

When comparing the results of improved Long Calculation Chain and Feature Envy smells with their baseline implementations, the corresponding metric curves in the empiric evaluation consistently follow a similar trend. However, the improved versions of the smells indicate overall lower individual metric values. Indeed, when applying the suggested threshold values, the improved versions of the smells report a substantially lower number of smell detections than their baseline counterparts (4879 instead of 84 031 for Long Calculation Chain, and 533 instead of 737 for Feature Envy). However, genuine detections that would be indicated by the baseline smell are still reported by the improved version. Our manual investigation also indicates that the improved techniques are successful in limiting the number of superficial detections. The proposed improvements thus are successful in reducing the number of individual entities that are indicated and have to be checked by a user.

Revisiting the first research question (**RQ1**): "Can existing smell detection techniques be improved by applying structural information in the smell's detection process?", we conclude that our proposed improvements provide a clear benefit over the original smell detection techniques. The improved version of the Pattern Finder detection process decreases the amount of false positives and reveals new issues. The improved Long Calculation Chain and Feature Envy detection techniques limit duplicate reports, while still including genuine cases.

The metric results of the novel smell detection techniques follow the expected trend set by previous spreadsheet smell evaluations. The Overburdened Worksheet variants require a threshold to indicate worksheets as smelly. We provided the cut-off values of four for block-based detection and eleven for group-based detection of the smell, each indicating a low risk for the corresponding worksheet. Occurrences of inconsistent formula group references and missing headers can directly be reported as smell detections. As illustrated by our manual investigation, all three smell types could be used alongside the established set of smells and were able to identify novel issues.

Looking at the second research question (**RQ2**): "Are novel spreadsheet smell detection techniques that are based on structural information able to detect new quality issues, and do they perform similar to traditional ones?", we conclude that the newly introduced smell detection techniques indeed perform similar to traditional spreadsheet ones. Detection values and rates follow the same distribution trend, a power law like distribution which is the usual case for smell metrics (see Hermans *et al.* [12][13]). They are moreover mechanically similar to the existing smells, and successfully point out novel issues, as demonstrated by the manual investigation. Consequently, each of the newly introduced smells is apt to be used alongside the currently established smell catalogue.

Lastly, detection of structure refined smells is dependent upon successful inference of structure information. However, due to erroneous or unexpected spreadsheet layouts, the proposed structure analysis approach might lead to incomplete or misleading results. Fortunately, cases of "improper" spreadsheet structuring usually also cause issues that are detected by structure aware smells. For example *Inconsistent Formula Group Reference* highlights any structurally unsound modification of a formula group, if the initial group as referenced by another formula group, or if the initial group referred to any other formula group. A list of examples for such detections is given in previous work [26]. If one of these smells identifies and reports the initial structural issue, the user is able to apply adequate fixes and refactorings. This allows for a bootstrap approach of iterative cycles of structure inference, smell detection, and refactoring, until a sound spreadsheet structure is accomplished.

6.6. Threats to Validity

A threat to the external validity of our evaluation is the representativeness of the EUSES corpus for the overall population of spreadsheets. However, the corpus consists of 4490 spreadsheets in 11 categories, providing an adequate variety of samples. Moreover, the corpus has already been extensively used for empirical evaluations, providing further credit as adequate evaluation baseline.

A further threat to the external validity of our results are the preprocessing operations we applied to the corpus before the actual evaluation took place. However, this operation only affects the comparison between the results of our baseline implementation and the respective smell's original evaluation. Both the evaluation of our baseline implementation as well as the evaluation of improved and new detection techniques are based on the same, preprocessed set of spreadsheets.

A threat to the internal validity of our results might concern the baseline smell detection implementations in the FRITZ tool. Abstract smell definitions, provided by each smell's original author, leave room for interpretation for a concrete implementation. Hence, our specific design decisions when implementing the baseline smells might affect the related evaluation results.

Another threat with respect to the internal validity is related to the correctness of our tool implementation, FRITZ, providing spreadsheet file handling, abstraction, structural analysis, smell detection, and automatic evaluation. We minimized this risk by manual testing, sanity checking of evaluation results, and comparison of the results with the original evaluation results. Moreover, the tool is publicly available. This allows other researchers to replicate our results.

7. Related Work

Hermans et al. [13] were among the first to define smells for spreadsheets. They adapted Fowler's inter-class smells [9] from object-oriented software to spreadsheets by treating worksheets as classes: When two or more worksheets have a strong connection, the spreadsheet is difficult to understand and to maintain; changes to a worksheet might also have impacts on other worksheets. In their paper, they redefined well-known smells like *Inappropriate Intimacy*, *Feature Envy*, and *Shotgun Surgery*. In an ensuing work [12], they deal with intraworksheet smells and propose smells like *Multiple Operations* (derived from *Long Method* smell), *Multiple References* (from *Long Parameter List*), *Conditional Complexity*, and *Long Calculation Chain* for spreadsheets. In more recent work, Hermans *et al.* propose refactorings for formula smells [27]. They indicate the need for refactoring by shading smelly cells and adding comments. These comments mention the proposed refactoring action (e.g. "Common subformula can be extracted"). Hermans and Dig [28] provide tool support for refactoring intra-formula smells. Unfortunately, support for automated refactoring of inter-formula smells is not offered yet, meaning the user has to manually change the spreadsheet when inter-formula smells like long calculation chains have been detected.

Cunha et al. [18, 11] focus on input cells and identify, e.g., outliers of numerical values (Standard Deviation smell), typos (String Distance smell), references to empty cells, mixed use of strings and numerical values in a column (Pattern Finder smell), and deviations in data entries (Quasi-Functional Dependency smell). To provide a better overview of existing smells, they aggregate Hermans et al.'s smells and their own smells in a catalog and provide a tool named SmellSheet Detective⁹ which implements all of these smells. Abreu et al. [29] improve the functionality of the SmellSheet Detective by combining smell detection with spectrum-based fault localization. They provide an implementation of their approach in the tool FaultySheet Detective⁹.

Several approaches detect faults by identifying structures in spreadsheets. UCheck [22] uses header cells as unit information for input and formula cells. A unit can be a simple unit like 'Employee' or a dependent unit like 'Employee[Anderson]'. Units can be combined using the &-operator, e.g., 'Employee[Anderson]&Quarter[1]'. Formula cells inherit their unit from referenced cells. Since a formula might reference several cells, the resulting units are a combination of the referenced cells' units. All units must be well-formed. Violations occur, for example, if two dependent units with the same base unit are combined using an &-operator, e.g. 'Employee[Anderson]&Employee[Bourne]'. Im more recent work [30], Cunha *et al.* extended their approach to automatically infer rational schemas from spreadsheets, and to map these schemas to ClassSheets, object-oriented models for spreadsheets that were previously introduced by Engels and Erwig [31]. They also provided and evaluated a catalogue of refactorings for said models, and showed a positive effect on end-users' productivity via an empirical evaluation [32].

Dimension [33] derives dimension information from the headers (e.g., length, time, and speed) and uses the corresponding units (e.g., meter, second, and meter/second) as units for the input cells. Formula cells inherit their units from the input cells to which they refer to. Invalid operations (e.g., adding meter and decimeter or meter and meter/second) are reported as errors.

AmCheck [16] identifies cell arrays and detects smells based on these arrays. There are two types of smells that can be detected by AmCheck: the Missing Formula smell and the Inconsistent Formula smell. The Missing Formula smell occurs in cells which have a constant input value instead of a formula; the Inconsistent Formula smell occurs in cells whose formulas differ from the formulas

⁹download via http://ssaapp.di.uminho.pt/twiki/bin/view/Main/Software

of the other cells in the same cell array.

Zhang *et al.* [34] have empirically evaluated UCheck, Dimension, and Am-Check with respect to precision, recall, efficiency, scope, and limitations. This evaluation shows that AmCheck has the best precision and recall rate and that UCheck and Dimension find different faults compared to AmCheck.

CACheck [17] improves over AmCheck by additionally detecting inhomogeneous cell arrays. A row-/column-based cell array is inhomogeneous if it contains a formula cell that references cells in a different column/row. Furthermore, CACheck removes invalid cell arrays by means of filtering rules.

CUSTODES [35] clusters cells by means of strong and weak features. Strong features are, for example, copy-equivalent formulas and cell dependency patterns; weak features are the position of the cell within a worksheet, the cell's labels, and the cell's style. Outlier cells in the individual clusters are identified and classified either as Missing Formula smell or as Dissimilar Formula smell.

Koci *et al.* [36] propose a machine learning approach that classifies cells into five categories: headers, attributes (i.e., row headers), meta data (i.e., captions), data, and derived (i.e., content that is derived from the actual data). Their approach extracts features (e.g., cell type, color, alignment, font type, font style, column- and row index) from the cells and applies different supervised learning techniques (e.g., Random Forest) on them. In a post-processing step, they detect and repair misclassified cells by means of predefined patterns.

TableCheck [37] detects table clones, i.e., two rectangular blocks of cells which have the same labels. Table clones are problematic, as they might become inconsistent when a spreadsheet evolves. TableCheck also reports when detected clones contain inconsistencies like missing and inconsistent formulas. TableCheck differs from AmCheck, CACheck and CUSTODES, as it detects inconsistencies between blocks rather than smells within a block.

Amalfitano *et al.* [38] propose a reverse engineering process for automatically retrieving data models from spreadsheets. This process is a top-town approach, meaning that a spreadsheet is decomposed into several worksheets which contain several areas, subareas and sub-subareas. The decomposition process makes use of the cells' formatting properties for refining the model. The derived model is visualized as a UML class diagram. Another interesting work of Amafitano *et al.* [39] presents a tool that helps to analyze connections of cells and VBA code.

The range of research publications regarding the overall topic of spreadsheet quality assurance is considerable. Therefore, we have focused on papers that are closely related to ours: The last mentioned papers [22, 33, 16, 17, 35, 36, 37, 38] deal with the identification of spreadsheet computation structures. Our structural analysis process builds upon the ideas of UCheck [22] and is explained in detail in Section 4. The first mentioned papers [13, 12, 27, 18, 11] deal with spreadsheet smells. We discuss how these spreadsheet smells can benefit from the structural analysis in Section 5. Since we have limited the discussion of related work to smells and structural analysis, we refer the interested reader to Jannach *et al.*'s overview paper [7] for a general overview of quality assurance techniques for spreadsheets.

8. Conclusion and Future Work

In this paper, we proposed to compensate present shortcomings of spreadsheet smells by refining smell detection procedures using structural information. To that end, we first presented an analysis process that infers structural information from spreadsheets. We then demonstrated smell refinements on the examples of the Pattern Finder, Long Calculation Chain, and Feature Envy smells. Furthermore, we introduced three new smells that make use of inferred structure information, namely Overburdened Worksheet, Inconsistent Formula Group Reference, and Missing Header. Empirical evaluation indicated that refined smells indeed have a positive effect on detected smells, and that novel smells are an adequate strategy to indicate further quality deficits.

The empirical evaluation shows that the use of structure information improves the performance of the smell detection techniques: (1) improved *Pattern Finder* reduces the number of incorrect detections while increasing the number of genuine detections; (2) improved *Long Calculation Chain* limits the number of redundant smell reports; and (3) improved *Feature Envy refines* the smell's detection focus, reducing the number of detections of permissible cases. The evaluation of the new smell detection techniques indicates their applicability along the current smell catalog to detect novel quality issues.

The proposed smell refinements alleviate a major drawback of existing smell detection processes: i.e., that the original approaches often highlight effects instead of causes. The refined smell detection approach condenses many related smell detections into one; the user gets a clearer picture of the overall issue and is less focused on an overwhelming number of problems regarding individual cells. Moreover, the new smells provide additional perspectives for users to assess the overall spreadsheet quality. In particular, the Overburdened Worksheet smell acts as a good counterbalance to existing formula- and inter-worksheet smells.

A major drawback of the proposed improved and new smell detection techniques is their reliance on a successful structural analysis process and the associated limited applicability to any given spreadsheet. Not all spreadsheets follow the same approach to general spreadsheet structuring, and not all spreadsheets contain applicable formulas, required as cues for the analysis. Hence, the structural analysis process in its current form might not always be be applicable and successful, leading to missing or false positive smell detections. However, detection of structure refined smells also allows a user to tackle this issue by fixing unsound spreadsheet structures using an iterative bootstrap process.

In future work, we want to examine how to best provide adequate representations of structures and related smells to users. This includes information about the inter-relations between different groups of a spreadsheet via group references, e.g. in form of a graph, which would be a valuable asset for spreadsheet comprehension. Moreover, smell detection in traditional software development usually is accompanied with a set refactorings, standard transformations of code that remove the indicated issue. To provide similar actions for spreadsheets, we currently investigate structure-based interactions in spreadsheets that allow us to formulate refactorings for structure-based smells. Lastly, the presented structure analysis process, as well as the derived smells pose opportunities for future work. For example, by extending the inference of groups to include non-formula related cells, and defining/evaluating further refined and novel smells.

Acknowledgment

The work described in this paper has been been funded by the Austrian Science Fund (FWF) project *DEbugging Of Spreadsheet programs (DEOS)* under contract number I2144 and the Deutsche Forschungsgemeinschaft (DFG) under contract number JA 2095/4-1.

References

References

- C. Scaffidi, M. Shaw, B. A. Myers, Estimating the numbers of end users and end user programmers, in: 2005 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC),, 2005, pp. 207–214. doi:10. 1109/VLHCC.2005.34.
- Report of JPMorgan Chase & Co. Management Task Force Regarding 2012 CIO Losses (January 2013).
- B. D. Bridgeford, W. Baraboo to pay more for borrowed money than believed, Barboo News Republic, last visited: July, 29th 2016.
 URL http://www.wiscnews.com/baraboonewsrepublic/news/local/ article_7672b6c6-22d5-11e1-8398-001871e3ce6c.html
- [4] J. Lee, Spreadsheet horror stories that will make you re-think your receivables management strategy, last visited: July, 29th 2016.
 URL http://blog.anytimecollect.com/
 5-receivables-management-spreadsheet-horror-stories/
- [5] R. Panko, What we don't know about spreadsheet errors today: The facts, why we don't believe them, and what we need to do, in: Proceedings of the 16th EuSpRIG Conference, 2016, pp. 79–93. URL http://arxiv.org/abs/1602.02601
- [6] S. G. Powell, K. R. Baker, B. Lawson, Impact of errors in operational spreadsheets, Decision Support Systems 47 (2) (2009) 126 – 132. doi: 10.1016/j.dss.2009.02.002.
- [7] D. Jannach, T. Schmitz, B. Hofer, F. Wotawa, Avoiding, finding and fixing spreadsheet errors - A survey of automated approaches for spreadsheet QA, Journal of Systems and Software 94 (2014) 129–150. doi:10.1016/j.jss. 2014.03.058.

- [8] F. Hermans, B. Jansen, S. Roy, E. Aivaloglou, A. Swidan, D. Hoepelman, Spreadsheets are code: An overview of software engineering approaches applied to spreadsheets, in: Software Analysis, Evolution, and Reengineering (SANER), 2016 IEEE 23rd International Conference on, Vol. 5, IEEE, 2016, pp. 56–65.
- [9] M. Fowler, Refactoring Improving the Design of Existing Code, Addison Wesley object technology series, Addison-Wesley, 1999.
- [10] B. Jansen, F. Hermans, Code smells in spreadsheet formulas revisited on an industrial dataset, in: Software Maintenance and Evolution (ICSME), 2015 IEEE International Conference on, IEEE, 2015, pp. 372–380.
- [11] J. Cunha, J. P. Fernandes, H. Ribeiro, J. Saraiva, Towards a catalog of spreadsheet smells, in: 12th International Conference on Computational Science and Its Applications (ICCSA), Vol. 7336 of Lecture Notes in Computer Science, 2012, pp. 202–216. doi:10.1007/978-3-642-31128-4_15.
- [12] F. Hermans, M. Pinzger, A. van Deursen, Detecting code smells in spreadsheet formulas, in: 28th IEEE International Conference on Software Maintenance (ICSM), 2012, pp. 409–418. doi:10.1109/ICSM.2012.6405300.
- [13] F. Hermans, M. Pinzger, A. van Deursen, Detecting and visualizing interworksheet smells in spreadsheets, in: 34th Int. Conference on Software Engineering (ICSE), 2012, pp. 441–451. doi:10.1109/ICSE.2012.6227171.
- [14] M. Ducassé, A pragmatic survey of automatic debugging, in: Proceedings of the 1st International Workshop on Automated and Algorithmic Debugging, AADEBUG '93, Springer LNCS 749, 1993, pp. 1–15.
- [15] R. Abreu, J. Cunha, J. P. Fernandes, P. Martins, A. Perez, J. Saraiva, Faultysheet detective: When smells meet fault localization, in: 2014 IEEE International Conference on Software Maintenance and Evolution (IC-SME), 2014, pp. 625–628. doi:10.1109/ICSME.2014.111.
- [16] W. Dou, S. Cheung, J. Wei, Is spreadsheet ambiguity harmful? Detecting and repairing spreadsheet smells due to ambiguous computation, in: 36th International Conference on Software Engineering, ICSE, 2014, pp. 848– 858. doi:10.1145/2568225.2568316.
- [17] W. Dou, C. Xu, S. C. Cheung, J. Wei, Cacheck: Detecting and repairing cell arrays in spreadsheets, IEEE Transactions on Software Engineering PP (99) (2016) 1–1. doi:10.1109/TSE.2016.2584059.
- [18] J. Cunha, J. P. Fernandes, P. Martins, J. Mendes, J. Saraiva, Smellsheet detective: A tool for detecting bad smells in spreadsheets, in: IEEE Symp. on Visual Languages and Human-Centric Computing, VL/HCC, 2012, pp. 243–244. doi:10.1109/VLHCC.2012.6344535.

- [19] J. Cunha, J. P. Fernandes, J. Mendes, J. Saraiva, Embedding, evolution, and validation of model-driven spreadsheets, IEEE Trans. Software Eng. 41 (3) (2015) 241-263. doi:10.1109/TSE.2014.2361141. URL https://doi.org/10.1109/TSE.2014.2361141
- [20] P. W. Koch, B. Hofer, F. Wotawa, Static spreadsheet analysis, in: 7th IEEE International Workshop on Program Debugging (IWPD), ISSRE Workshop Proceedings, 2016, pp. 167–174. doi:10.1109/ISSREW.2016.8.
- [21] B. Hofer, A. Riboira, F. Wotawa, R. Abreu, E. Getzner, On the Empirical Evaluation of Fault Localization Techniques for Spreadsheets, in: Fundamental Approaches to Software Engineering (FASE'13), Vol. 7793 of Lecture Notes in Computer Science, 2013, pp. 68–82. doi:10.1007/ 978-3-642-37057-1_6.
- [22] R. Abraham, M. Erwig, UCheck: A spreadsheet type checker for end users, Journal of Visual Languages and Computing 18 (1) (2007) 71–95. doi: 10.1016/j.jvlc.2006.06.001.
- [23] ISO/IEC 25010:2011 Systems and software engineering Systems and software Quality Requirements and evaluation (SQuaRE) – System and software quality models, International Organization for Standardization, Geneva, Switzerland (2011).
- [24] J. Cunha, J. P. Fernandes, C. Peixoto, J. Saraiva, A quality model for spreadsheets, in: 8th International Conference on the Quality of Information and Communications Technology, QUATIC, 2012, pp. 231–236. doi:10.1109/QUATIC.2012.16.
- [25] M. F. II, G. Rothermel, The EUSES spreadsheet corpus: a shared resource for supporting experimentation with spreadsheet dependability mechanisms, ACM SIGSOFT Software Engineering Notes 30 (4) (2005) 1–5. doi:10.1145/1082983.1083242.
- [26] P. Koch, Smelly spreadsheet structures: Structural analysis of spreadsheets to enhance smell detection, Master's thesis, Graz University of Technology (2016).
- [27] F. Hermans, M. Pinzger, A. van Deursen, Detecting and refactoring code smells in spreadsheet formulas, Empirical Software Engineering 20 (2) (2015) 549–575. doi:10.1007/s10664-013-9296-2.
- [28] F. Hermans, D. Dig, Bumblebee: a refactoring environment for spreadsheet formulas, in: Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, (FSE-22), 2014, pp. 747–750. doi:10.1145/2635868.2661673.
- [29] R. Abreu, J. Cunha, J. P. Fernandes, P. Martins, A. Perez, J. Saraiva, Smelling faults in spreadsheets, in: 30th IEEE International Conference on

Software Maintenance and Evolution (ICSME), 2014, pp. 111–120. doi: 10.1109/ICSME.2014.33.

- [30] J. Cunha, M. Erwig, J. Mendes, J. Saraiva, Model inference for spreadsheets, Autom. Softw. Eng. 23 (3) (2016) 361-392. doi:10.1007/ s10515-014-0167-x. URL https://doi.org/10.1007/s10515-014-0167-x
- [31] G. Engels, M. Erwig, Classsheets: automatic generation of spreadsheet applications from object-oriented specifications, in: Proceedings of the 20th IEEE/ACM international Conference on Automated software engineering, ACM, 2005, pp. 124–133.
- [32] J. Cunha, J. P. Fernandes, P. Martins, J. Mendes, R. Pereira, J. Saraiva, Evaluating refactorings for spreadsheet models, Journal of Systems and Software 118 (2016) 234–250. doi:10.1016/j.jss.2016.04.043. URL https://doi.org/10.1016/j.jss.2016.04.043
- [33] C. Chambers, M. Erwig, Automatic detection of dimension errors in spreadsheets, Journal of Visual Languages and Computing 20 (4) (2009) 269–283. doi:10.1016/j.jvlc.2009.04.002.
- [34] R. Zhang, C. Xu, S. Cheung, P. Yu, X. Ma, J. Lu, How effectively can spreadsheet anomalies be detected: An empirical study, Journal of Systems and Software 126 (2017) 87 – 100. doi:10.1016/j.jss.2016.03.061.
- [35] S.-C. Cheung, W. Chen, Y. Liu, C. Xu, CUSTODES: Automatic spreadsheet cell clustering and smell detection using strong and weak features, in: Proceedings of the 38th International Conference on Software Engineering, ICSE '16, ACM, 2016, pp. 464–475. doi:10.1145/2884781.2884796.
- [36] E. Koci, M. Thiele, O. Romero, W. Lehner, A machine learning approach for layout inference in spreadsheets, in: Proceedings of the 8th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K) - Volume 1, 2016, pp. 77–88. doi:10.5220/0006052200770088.
- [37] W. Dou, S.-C. Cheung, C. Gao, C. Xu, L. Xu, J. Wei, Detecting table clones and smells in spreadsheets, in: Proceedings of the 2016 24th ACM SIG-SOFT International Symposium on Foundations of Software Engineering, FSE 2016, ACM, 2016, pp. 787–798. doi:10.1145/2950290.2950359.
- [38] D. Amalfitano, A. R. Fasolino, P. Tramontana, V. D. Simone, G. D. Mare, S. Scala, A reverse engineering process for inferring data models from spreadsheet-based information systems: An automotive industrial experience, in: Data Management Technologies and Applications Third International Conference, DATA, 2014, pp. 136–153. doi:10.1007/978-3-319-25936-9_9.

[39] D. Amalfitano, V. D. Simone, A. R. Fasolino, P. Tramontana, EXACT: A tool for comprehending vba-based excel spreadsheet applications, Journal of Software: Evolution and Process 28 (6) (2016) 483–505. doi:10.1002/ smr.1787.