Characterizing Scalability Issues in Spreadsheet Software using Online Forums

Kelly Mack, John Lee, Kevin Chang,
Karrie Karahalios, Aditya Parameswaran
Department of Computer Science
University of Illinois, Urbana-Champaign
knmack2, lee98, kcchang, kkarhal, adityagp@illinois.edu

ABSTRACT

In traditional usability studies, researchers talk to users of tools to understand their needs and challenges. Insights gained via such interviews offer context, detail, and background. Due to costs in time and money, we are beginning to see a new form of tool interrogation that prioritizes scale, cost, and breadth by utilizing existing data from online forums. In this case study, we set out to apply this method of using online forum data to a specific issue—challenges that users face with Excel spreadsheets. Spreadsheets are a versatile and powerful processing tool if used properly. However, with versatility and power come errors, from both users and the software, which make using spreadsheets less effective. By scraping posts from the website Reddit, we collected a dataset of questions and complaints about Excel. Specifically, we explored and characterized the issues users were facing with spreadsheet software in general, and in particular, as resulting from a large amount of data in their spreadsheets. We discuss the implications of our findings on the design of next-generation spreadsheet software.

INTRODUCTION

Spreadsheets are widely used by professionals in a variety of disciplines. Even in their most basic form, they act as a useful organizational aid, allowing users to record information in a two-dimensional tabular layout. With capabilities like formulae and macros to support complex calculations or automate processes, spreadsheets become indispensable as a comprehensive medium for data management and analysis [30].

At the same time, spreadsheets have problems: studies have shown that they foster errors [29, 27], they are not able to easily express useful data operations, such as joins or filters [11, 10], and they are known to be sluggish or crash when operating on very large datasets [30]. The last class of problems are what we call *scalability* problems, i.e., situations where users are unable to complete standard operations in a reasonable amount of time when the scale or complexity of the dataset increases [7]. With the increasing availability of large volumes of data in most domains, we suspect that scalability problems are only going to become more common—as an example, our biologist collaborators routinely use spreadsheets, and can now generate large genomic (VCF) files, but are unable to open these files even simply to verify correctness in spreadsheet software [12].

We conduct a case study of problems in spreadsheet software, with a special emphasis on scalability problems. Prior studies on spreadsheet software has largely focused on the ability of users to use a spreadsheet effectively, as opposed to spreadsheet software problems, including understanding users' conceptual models [23, 27, 25], and detecting and categorizing user errors [9, 20, 28, 13]. In contrast, we focus on settings where spreadsheet software does not meet expectations. For example, one user performed a copy/paste operation on a few hundred thousand rows of data; they waited two and a half hours and still could not accomplish their task. This behavior— where expectations and models of use are clear—is relatively unexplored.

A second difference with prior studies on understanding the use of spreadsheet software is our methodology. Our first attempt at characterizing spreadsheet problems was via interviews with current or aspiring spreadsheet users, such as our biologist collaborators, teaching faculty who managed MOOC student grades in spreadsheets, business analysts wanting to use spreadsheets to study ad campaigns, among others. While these interviews were illuminating in their depth of detail, they were lacking breadth—we wanted to survey a broad population of spreadsheet users, specifically Microsoft Excel users, to understand their problems. Thus, we turned to online forums, specifically Reddit, that captured users' self-reported opinions, questions, and frustrations on the Microsoft Excel spreadsheet subreddit[1].

To the best of our knowledge, online forums have not been used to understand the spectrum of problems in spreadsheet software. Online forums have been used for other goals, such as understanding the factors that (1) influence perception— e.g., the discovery of credibility factors in medical crowd-funding sites using Reddit [18], (2) describe use and behavior practices—e.g., physical interaction of users with disabilities using YouTube videos [8], and (3) highlight areas for improvement via "app mining"—prioritizing application bug fixes using Google Play Store reviews [17]. While our work is similar in its categorization of improvement areas to [17], they focus on improving an existing code/review pipeline to determine where to modify code, whereas we investigate the

---

[1]http://reddit.com/r/excel

scalability landscape surrounding multiple existing use cases in spreadsheets.

To accomplish this, we collected over 700 posts from an Excel Reddit forum. While analyzing the posts, four main types of operations emerged that posed problems, both in scalability, and otherwise—importing, managing, querying, and presenting data. We dig deeper to characterize these problem areas to highlight concrete areas of improvement for spreadsheet software, with an eye towards expanding the reach and usability of spreadsheets, especially for very large and complex datasets.

The contributions of this paper are (1) a mapping of challenges users face using spreadsheets in general, as well as (2) how they pertain to scalability, and (3) a methodology for broad evaluation of problems in features, capabilities, and intent specification for end-user software, and (4) a discussion of how to fix these problems, as a means towards building a more robust, scalable, powerful spreadsheet tool.

RELATED WORK

Our work builds on prior work in two different research areas, (1) spreadsheets and (2) the use of online community discourse as a primary data source for study.

Prior work in spreadsheets

Previous research in spreadsheets have focused on improving spreadsheets through understanding existing problems via user studies and analyzing existing spreadsheets. Researchers have conducted user studies to understand users' conceptual models of spreadsheets to identify how the cognitive process can affect error rates [27], to see how users navigate large spreadsheets [23], to evaluate how multiple users interact with a single spreadsheet [25], and to characterize the strengths and weaknesses of spreadsheets [26, 14]. Other studies have focused on errors: Powell et al. explore different types of errors that occur and how they can be minimized [28], while others study real spreadsheets to discover errors [9, 20]. Our approach is instead to identify scalability problems in spreadsheets by exploring troubleshooting posts on an online forum.

Using online communities as a data source

The availability of diverse and large amounts of online community data has led researchers to mine this data to answer research questions [15, 8, 21, 18, 16]. While this method may bias the user sample to users with internet access and a level of technology savviness, prior works have successfully created rich characterizations of users via this approach. In addition to the papers mentioned in the introduction [8, 18, 17], Kulshrestha et al. [21] measured the political bias of an individual Twitter search result by extracting features from the Twitter user's account, while Keelan et al. [16] extracted Youtube videos to measure the sentiment (positive/negative) surrounding immunization.

CHARACTERIZING REDDIT

Reddit is a website that hosts a variety of forums called subreddits. Each subreddit is a forum dedicated to a specific topic and is named /r/topic_name. Within these forums, users can post questions or notes and can comment on each others' posts,

leading to discussions. Reddit has an API [3] to access data from the site, making it ideal to use for those interested in automating the scraping and categorizing of data.

The Forum

The /r/excel forum has been around for over 8 years, and as of September 2017, had over 70,000 subscribers [4]. Note that a subscriber is someone who subscribes to (follows) the forum; anyone with a Reddit account can read, post, or comment in the /r/excel forum without following it. Reddit recently removed its statistics regarding the traffic of subreddits to improve user privacy [5], resulting in the number of subscribers being the most valuable statistic publicly available. Note also that this forum is monitored (meaning a moderator can take down posts as they see fit), and one of the moderators of the forum disclosed in a post entitled "Please welcome our new Corporate Overlords" that the subreddit is involved in the Excel Influencer Program [2].

The Users of the Forum

Reddit does not collect any information from the user other than a username, password, and email, so it is difficult to characterize the type of users that frequent Reddit, never mind those that visit the /r/excel subreddit. However, in their posts and comments, some users share their experience level with Excel. Often users who ask questions begin with a statement of their unfamiliarity with Excel. Others state that they use Excel frequently for work but that they need help performing a new and/or complex operation. The users who answer questions typically do not state their credentials, but often their complex solutions indicate a high level of experience with Excel.

In some cases, there are tools other than Excel that are useful for managing (particularly large amounts of) data like Microsoft's Access software [22] or relational databases. Users of the forum (both post makers and commenters) showed varying degrees of knowledge about these alternatives. In some cases, the user wanted to know if a tool was more appropriate than Excel. In others, the users specifically said they knew other programs would perform better, but they were forced to use Excel. Sometimes the user did not mention Access or databases, but was recommended by commenters to use a database as opposed to a spreadsheet. Out of the 712 posts we collected, 89 mentioned one or more of the terms *"Access", "SQL",* or *"database"* in the post body or comments.

The Uses of Excel

In the 712 posts, the uses of Excel seemed to fall into two overarching categories: Excel for personal use, and Excel for professional use. Traditionally, in both of these areas, we think of spreadsheets as being used for record keeping of data like addresses and emails, time trackers and schedules, or financial information and budgets. There were many posts regarding these topics in both personal and professional settings. However, the more unusual uses of Excel were impressive, detailed next.

Regarding personal uses, several users asked about how to keep track of and calculate *sports statistics*; fantasy football was particularly popular. One user wanted to design a spreadsheet that automatically organizes a table tennis tournament.

Another wanted to create an ELO rating system for professional tennis players. The games mentioned also included virtual ones. Several people asked questions about keeping track of *statistics for video games* including Pokemon and League of Legends. Some users went even further and *created fully functional versions of popular games* like Tic-Tac-Toe or Blackjack in Excel. Another popular personal use was *keeping track of health information* like food consumption and body measurements in Excel.

Professionals from a variety of industries used Excel spreadsheets for work. Unsurprisingly, many users' questions related to time tracking and inventory or office management. Several posts dealt with *financial information*, including two people who were pulling live stock statistics from the web. Manufacturing industry professionals created spreadsheets that serve as *specialized calculators*. One user created a sheet that accepts input from a bar code scanner and organizes data accordingly. *Real estate industry* professionals used Excel to track renovation progress at tens of locations or to modeling rental income. Several types of *scientists utilize Excel's data analytics features*. Two biologists asked questions regarding spreadsheets they used to organize animal tracking data. Other scientists used Excel to create models of hydraulics and analyze a set of lunar data. The medical industry uses Excel not just to *keep track of things like patient diagnoses and diets*, but also for *medical research*. One user in particular created graphs in Excel that analyzed a set of brain wave data.

Overall, we found posts that touched on the following users/industries: real estate, finance, manufacturing, accountants, biologists, doctors and medical staff, teachers, researching students and professors, family management, and quantified self. Note that this is information from users who identified the context of their work. The majority of posts simply listed their problem without giving context as to where the spreadsheet is used.

METHODOLOGY

One goal of this work is to understand the breadth of issues users experience while using Excel, specifically issues related to scalability. To accomplish this, we chose to analyze posts from the Excel forum "/r/excel" on Reddit. There are other spreadsheet related forums such as /r/spreadsheets and /r/sheets. However, as Microsoft Excel is one of the largest players in the spreadsheet industry, we limited our scope to the Excel forum.

Constructing and coding a Dataset from Reddit

We collected two datasets of posts from the /r/excel forum using the Reddit API: one dataset to understand the general scope of Excel challenges and a second dataset to understand scalability issues specifically.

The random dataset
Collection. To collect our first dataset, we used the Reddit API's "random" function, fetching 50-100 posts with each call and collected 278 posts in May of 2017 over the period of two weeks. With this data we aimed to understand the breadth of the issues users encountered.

Coding. We used open coding to code our random posts. Two authors individually examined 50-100 fetched posts at a time and developed criteria to separate the posts into themes (and then sub-themes). The two then compared their themes to create a codebook to create a common set of themes. They repeated the process of fetching and categorizing posts until no new themes emerged.

Most of these posts related to a user's uncertainty of how to perform operations (e.g., how can one convert columns into rows) or described a behavior not meeting expectations (e.g., "AND" operation did not join two clauses in an array formula correctly, resulting in no data returned). Two posts described scalability issues, where the dataset size caused undesired behavior: in one, a formula calculated for 30 minutes before crashing, and in the second, Excel froze every time the screen shifted. We were intrigued by the scalability issues because (1) to date, they have not been chronicled in existing work, and (2) the outcomes of anecdotal cases did not have solutions. We, therefore collected a second dataset to explore scalability (see below).

Establishing Reliability. For both this and the following set, after creating the categories, two of the authors coded sets of 10% of the posts from each dataset separately until they achieved a Cohen's kappa of at least .78. They then proceeded to code the entirety of the dataset, comparing 5% of all post codes every 85 posts (20% of all posts), ensuring that they maintained a Cohen's kappa of at least .7.

The search term dataset
Collection. We used the Reddit API and searched for scalability posts via keywords within the same /r/excel forum. Adapting the methodology used in [19, 21, 18], we finalized the set of search terms in the following manner: (1) From the randomly collected posts, we identified scalability posts and their keywords. (2) Using the keywords found in (1), we searched for additional scalability posts. We then extracted new scalability keywords from these posts. We repeated this process until no new keywords/phrases were identified. The final list of scalability keywords was *"big", "crash", "forever", "freeze", "lag", "large", "long time",* and *"slow"*.

Using our final set of eight keywords/phrases, we created a dataset of 434 posts by collecting the first 100 posts returned by each keyword or until the API call ceased to return results, whichever happened first, from July through August of 2017. Typically, approximately 70 posts were returned, but some keywords returned as few as nine.

While this approach was designed to filter scalability issues from the forum as a whole, we found that people also addressed Excel "crashing" or "freezing" in other situations (e.g., "freeze panes" or "crash course in Excel"). To capture posts that addressed the scalability issues, we created a codebook to separate scalability issues from non-scalability issues. We stopped collecting data when no new themes emerged in our coding of problem area themes.

Coding for Theme. We adapted the coding process we used for the random dataset to code this search term dataset. We ran each search term separately and coded resulting posts in

batches as they arrived. Two of the authors developed criteria to separate the posts into themes (and then sub-themes). The two then compared their themes to create a codebook to create a common set of themes. They repeated the process of fetching and categorizing posts until no new themes emerged. While we started with a blank slate after coding the random dataset, we found that the posts essentially fell into the same themes as the randomly fetched posts. The posts that did not fall into the four main themes were placed in a *miscellaneous* category.

Coding for Scalability. As described earlier, we separated the scalability related posts from the non-scalability related posts. The authors created the following criteria for this coding: a scalability issue is defined as *poor* behavior within an Excel workbook that contains a *large amount of data*. Specifically, poor behavior is characterized as Excel lagging, not responding, crashing, or taking long enough to respond that the user declared the spreadsheet unusable. The inclusion criteria for scalability posts were: (1) the user specified observing thousands of rows of data, (2) the user's file size was 5MB [2] or larger, (3) the user did not quantitatively specify how large their data was, but they expressed suspicion that the issues were due to the dataset size, (4) a user could successfully import a small amount of data, but failed when the data size increased.

We coded the same dataset thematically and for scale. For both codings of the search term dataset, we established coding reliability in the same manner as we described in the random dataset coding section. From the 434 search term posts, 81 posts addressed scalability concerns, bringing the combined number of posts dealing with scalability issues to 83 out of 712.

THE THEMES
We made two observations which laid out the architecture for our thematic coding. First, the posts were overwhelmingly questions, as opposed to statuses or updates. Second, the majority of the questions dealt with *four themes: importing data, managing data, querying data, and presenting data.* While the majority of posts fell into these four themes, the remainder were placed in a Miscellaneous theme. These posts generally asked broader, questions than other posts, as we discuss further below.

Our four main themes were further categorized into sub-themes. We introduce the themes and sub-themes below with a brief description of the posts they contain, followed by a table presenting paraphrased quotes (to preserve anonymity) of posts from a selection of sub-themes. Each post received one of the following labels as well as a yes or no label indicating whether it related to scalability to better understand both the post issue whether the issue was affected by scalability. The pair of labels allowed us to see the operation areas affected by scalability. To illustrate how the "scalability" code changes the nature of the problems users encountered, we present several posts in Table 1 with their scalability code.

---

[2]This is the size where users began experiencing "lag" and "freezing" issues.

Importing Data
This theme includes posts related to importing data from sources such as CSV files, PDFs, websites, and other Excel workbooks.

Managing Data
This theme emerged from posts that discussed cleaning and updating data after it has been imported. This theme included the three sub-themes below.

Manipulating Data. This sub-theme includes questions about actions required to transform data into a analyzable state. (e.g., cleaning data after import and changing the format of data for a non-aesthetic purpose–changing the format of a date so it is consistent with existing data and formulas in the sheet.)

Deleting Data. This sub-theme includes questions that ask how to delete or hide data from view.

Updating Data. This sub-theme includes questions asking how to update existing entries in a workbook or how to add more data to a workbook, e.g., adding a fixed value to all of the entries in a column or adding new entries to a list of records in a sheet.

Querying Data
This theme includes posts related to extracting information from the data in a workbook.

Calculations. This sub-theme includes questions on how to create new values from existing data in the spreadsheet. Note that simply selecting existing data from the sheet does *not* qualify as a calculation; *new* data is created.

Selections. This sub-theme includes posts about finding and/or selecting pieces of information from a workbook (e.g., questions whose main topic is VLOOKUP). Other common examples discuss operations such as finding duplicates or finding which records fall between two dates.

Aggregations. This sub-theme includes posts that ask questions about grouping all records of a sheet into themes, counting entities (often via COUNTIF), summing entities, or using pivot tables. The main distinction between *Aggregations* and *Selections* is that *Selections* deal with selecting data that meets certain conditions while disregarding the remainder of the data, whereas *Aggregations* deal with organizing/categorizing *all* of the data in a sheet. A *selection* functions as a drop down filter. An *Aggregation* functions similarly to a SQL groupby statement, as it organizes all of the data in a table.

Presenting Data
This theme includes posts relating to how the data is presented to the Excel user (and not necessarily for the purpose of creating a formal presentation).

Visualizations. This sub-theme includes posts that relate to the creation of charts, graphs, and tables. The questions address creating templates, changing the style of a chart, or selecting the correct data to be used in a chart.

Style. This sub-theme includes posts that deal with the style of the Excel sheet or the style of the data within the sheet.
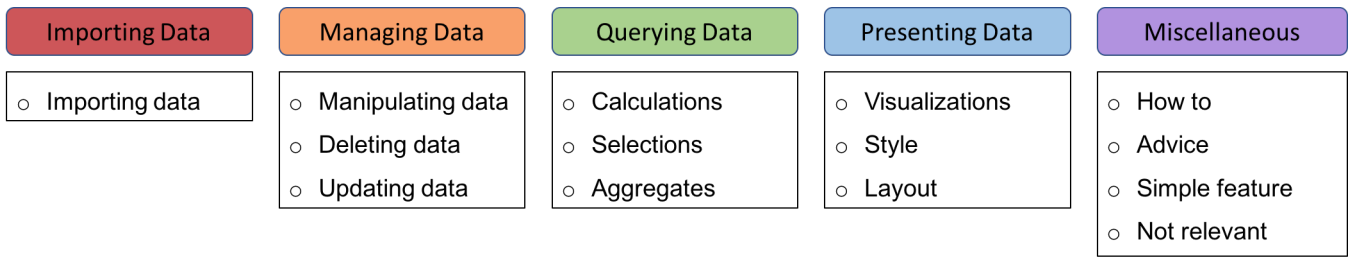
Figure 1. A hierarchical representation of the themes and their sub-themes.

Examples of style actions in the sheet include highlighting, outlining cells, and changing row/column widths. Actions that deal with the style of the data include changing font properties or the number of decimal places displayed. One important distinction with the style of the data is that it is done for an aesthetic purpose; it does not change the traits of the data that would influence queries.

Layout. This sub-theme includes posts that deal with the position or layout of data within the workbook. Questions asking how to move data *within the same workbook* also fall into this theme. Other common questions include moving all data in a certain manner (e.g., up one row or converting rows to columns).

Miscellaneous
These posts were significant enough to warrant coding, but did not fall into one of the aforementioned themes.

How-To or Isn't Working. This sub-theme included posts that either ask how to do a task (not present in the previous four themes) or ask why a simple feature in Excel does not meet user expectations. Often these formats of questions are combined into one—"X is not working, how would I solve this problem". Creative uses of Excel (e.g., creating games within a spreadsheet) fell in this sub-theme.

Advice. This sub-theme includes posts that ask for general advice about Excel functionality. If a post was specifically asking about aggregating data, for example, we included it in the *aggregate* theme. However this theme captured questions that were broad in scope—for example, "how do I structure a large amount of data in a spreadsheet".

Simple Expected Features. This sub-theme includes posts that contain descriptions of simple features that did not function as expected. We define simple features as those required for Excel to properly function as data storage mechanism (e.g., opening/closing the application, adding a row to a sheet, saving the data, and scrolling).

Not Relevant. This sub-theme includes posts that were not relevant to our case study addressing spreadsheet needs and desires. These posts included advertisements or requests for software alternatives to Excel.

RESULTS AND DISCUSSION
The main goal of this work is to present the landscape of scalability issues present in Excel as captured in a popular

Category	Scalability	Quote
Import	No	"Is there a way I can automatically pull reports from Google, Facebook, and email into an Excel file?"
Import	Yes	"I have over 300 Excel files. I would like to bring them into one file, but it is tens of MB's large and it does not conatin them all yet. Should I be using something other than Excel?"
Calculation	No	"I want to add up how many hours I've worked sporadically throught the day to get a day total."
Calculation	Yes	"In my 100,000 row spreadsheet, I have many formulas to calculate ratios I need for a report. However, now the file is large and extremely slow."
Layout	No	"How do I freeze both columns and rows."
Layout	Yes	"I wrote some VBA code to help copy and post a large dataset of 5K rows. I'm only trying it on a couple hundred rows and it already crashes Excel."

Table 1. Excerpts from collected posts.

Figure 2. The legend for the following charts.

subreddit forum. We first present characteristics of the posts before describing distribution of results in categories.

Figure 3 presents the distribution of the set of 278 random posts, Figure 4 presents the distribution of the 434 search term posts, and Figure 5 shows the distribution of the posts we categorized as relating to scalability. 83 of the 712 posts related to scalability. 81 came from the search terms posts, while the other two came from the random posts. We now discuss the characteristics of the posts and the conclusions we made.

Results Regarding All Posts
We collected 712 posts in total. Of these posts, 644 posts had unique authors. 21 posts had no author [3]. We were interested to see if authors of comments were similarly diverse, or if there were a few users that commented on many posts. We

---

[3]We have not found a source for why a post can have no author. We assume that this occurs when a user makes a post and then deletes their account.
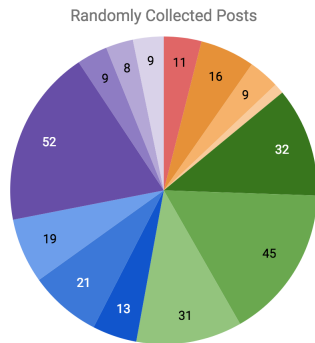
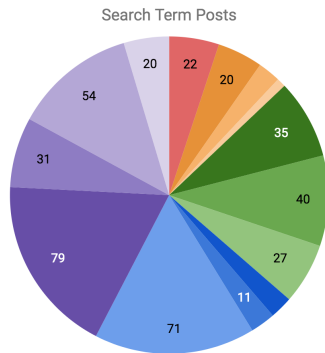**Figure 3. The distribution of the random posts (legend in figure Figure 2).**

Figure 4. The distribution of the search term posts (legend in Figure 2).

found that there were 315 unique comment authors– one of which was "no author". For these 712 posts, the most unique posts that one user commented on was 8. The most comments that one user made was 18 across a total of 7 different posts. Again, there does not appear to be a small number of users that dominate the comments.

The average length of a post (not including the comments) was 135 words, and the average number of comments was 7– note that this includes all comments posted by users and bots. However, 42 posts received no comments. On this subreddit, users can specify that their post is a question which starts out in the "unresolved state". The post maker can then wait for other users to comment and suggest solutions. If a solution works,
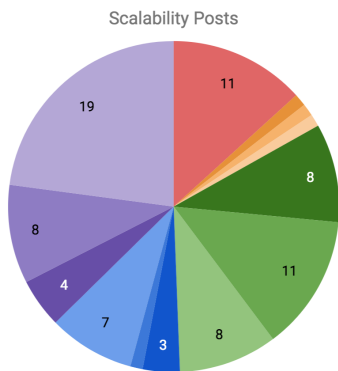

Figure 5. The distribution of the scalability posts (legend in Figure 2).

they comment "solution verified" and the post is marked as resolved. If the user fails to mark the post as resolved within a certain amount of time, a bot comments reminding them to mark a solution as verified. 91 posts had the "solution verified" comment.

Reddit allows users to "upvote" and "downvote" posts. On the /r/excel subreddit, when a user tries to downvote a post, a box pops up indicating that the user should comment why they are downvoting and says specifically "this is NOT a disagree button". This comment likely indicates that the forum encourages the number of upvotes and downvotes to reflect the quality of the post as opposed to if they like/agree with it. The average upvote to downvotes ratio of the 712 posts was 84.2%, indicating that the majority of posts users believe to be of high quality.

Traditional methods of determining quality of posts and discussions include looking at the thread depth of posts [31]. Here, we measure thread depth as the height of the comment tree, where the original post has thread depth 0, a direct reply to the original post has depth 1, and so on. The average thread depth of the posts in the forum was 3, with the maximum and minimum thread depths being 10 and 1. One of the posts with a thread depth of 1 stated that Excel crashed every time it was closed and listed the solutions that had already been tried. A post that had a thread depth of 10 asked a more general question. For example, one post that had a thread depth of 10 asked how to best encrypt a spreadsheet.

Results Specific to Scalability Posts
As we specifically wanted to study the circumstances that lead to scalability issues, we analyzed the post content and comments of scalability posts in depth.

File Size
We were curious as to what sized file are susceptible to scalability issues. Microsoft specifies data limits in the Excel documentation [1]. Specifically, Excel can handle no more than 1,048,576 rows by 16,384 columns. We expected that worksheets with over 1 million rows would start to have scalability issues. Indeed, the largest spreadsheet we saw in our scalability posts contained about 5 million rows, and Excel was slow to partially open it. The smallest sheet that we classified as having scalability issues had 18 thousand rows. This user wanted to perform multiple conditional if statements on the data, but needed a more time efficient solution. As the solution of using conditional if statements would have worked on a smaller dataset, we classified it as a scalability issue.

Suggested Solutions
Every scalability post included a question and most received more than one suggested solution from other Reddit users. We analyzed the solutions that were frequently recommended in each of the importing data, managing data, querying data, and presenting data themes. The following solutions were common across all of these four themes:

1. **Use a database.** This common solution was suggested by multiple people in 38 of the 83 scalability posts. Reddit users commented that, once a large enough amount of data

is stored in the spreadsheet, a database is the best way to ensure that one can still manage and query the data in a reasonable amount of time. Though commenters indicated using a database is the best solution, they offered other suggestions, if using Excel is unavoidable. For example, PowerQuery and PowerPivot, Excel extensions that are meant to process large amounts of data more easily, are not as powerful as a database, but are still better than using basic Excel. Otherwise, users are recommended to perform tedious operations like breaking up the data into chunks and processing the smaller chunks one at a time so as not to overwhelm the system.

2. **Turn of automatic calculations.** This was a suggested solution for 12 of the 83 scalability posts. Excel is set by default so that once a user updates a cell c, the system will recalculate any cells that depend on c immediately. This can be very time consuming, especially when an update to one cell triggers a chain of recalculations. Turning off automatic calculations results in Excel waiting to recalculate formulas after an update until the user manually selects to recalculate. This does not necessarily make the calculations faster, but will allow the spreadsheet to remain interactive while changes are made.

3. **Save Excel files as .XLSB files.** This was a suggested solution for 12 of the 83 scalability posts. By default, Excel files are saved as .XLS or .XLSX files. If instead a spreadsheet is saved as a .XLSB file, it is stored in binary format. The effects of the binary format include a smaller file size and shorter loading and saving times.

4. **Use as little conditional formatting as possible.** This was a suggested solution for 10 of the 83 scalability posts. Conditional formatting is the adjusting the formatting of a cell based on a conditional statement. For instance, highlight each cell in column C if its value is less than 50. Multiple users report that conditional formatting slows down an Excel file, particularly if it affects many rows.

5. **Avoid using volatile functions.** This was a suggested solution for 10 of the 83 scalability posts. Functions like OFFSET are considered volatile, meaning that any update to any cell in the spreadsheet triggers a recalculation of this formula. Obviously, this can become very expensive unnecessarily.

6. **Avoid using computationally intensive functions.** This was a suggested solution for 9 of the 83 scalability posts. Computationally intensive functions include VLOOKUP and COUNTIF. Users indicate that these functions, though sometimes necessary, oftentimes slow a spreadsheet down, particularly if the sheet has a lot of data. One alternative to VLOOKUP is INDEX MATCH. INDEX MATCH is preferred to VLOOKUP as INDEX MATCH often examines fewer cells in total than VLOOKUP.

We drew two main conclusions:

1. Users understand the capabilities of Excel, but not how to operationalize them. As a whole, the results stress the importance of Excel as a data storage tool and as a data processing engine (these posts account for 63.2% of all posts collected). Note the large difference between the number of posts that fell in the *querying data* (29.5% of posts) theme compared to *importing* and *manipulating data* combined (13.3% of posts). This indicates that users are generally comfortable with using Excel as a tool for storage, but they struggle when it comes to applying functions which allow them to extract useful information from their data.

In the majority of posts, users state what they would like Excel to do, and ask how to obtain this result. People do not underestimate what can be accomplished with Excel. The questions span the whole gamut of data analysis, from data collection (crawling the web with VBA code) to consumption of information (visualizations) and everything in between. Many of the posts in the *How to* category deal with people trying to accomplish creative tasks with Excel. One user was trying to implement Tic-Tac-Toe. Another was trying complete their RSA encryption homework in a spreadsheet.

These results together indicate that we need to make these complex and powerful capabilities more accessible to users. The technical knowledge of post makers ranged from people who were not familiar with VBA and wanted to avoid a solution using VBA to those who used it frequently. To cater to both types of users, perhaps more research can be given into spreadsheet tutorials and documentation.

Many of the posts tell a story of an often complex task the user wants to complete. Context is given, a goal is explained, and ideas of implementation are shared. Often these posts reflect an understanding of how to perform operations in Excel, but express confusion with stringing multiple operations together to perform a complete task. While documentation is very useful for understanding how a specific function works, it is not much help in understanding how to build a solution to a large problem out of basic functions. More research in the area of helping users understand how to integrate functions together would be of great use to many of the post makers on the /r/excel subreddit.

2. Scalability issues affect a wide range of operations in spreadsheets. Every theme and sub-theme contains posts relating to scalability. Analyzing this distribution can help prioritize which operations are the most important to adapt to handle large quantities of data. For example, a large amount of the scalability issues deal with either importing data or querying data. These two areas in particular require further attention.

Notice that scalability issues cause a large number of the issues in the *simple expected feature* posts. This indicates that as data sets become larger, even the simplest of features in Excel cease to respond in a timely manner. Without these features, as the post makers comment, Excel is essentially useless. Users turn to other options like databases, but some are reluctant to use them because they are not familiar with SQL. Ideally, these users who are not familiar with databases should be able to use the spreadsheet interface they are comfortable with, while utilizing the processing power and scalability of a database.

Moreover, from an expressiveness point of view, many of the users' issues could be solved by a simple SQL statement.

Multiple times, we witnessed a user trying to create a complex formula to accomplish what could be done with a simple select statement. For example, one post maker asked "I want to be able to input two dates, and have shown the text and date for every date between those. I think i[sic] need to use some sort of array formula, but I'm not sure how (the exact formula, not array as a whole)". The suggested solution posted was to use the following formula:

```
=IFERROR(INDEX(B:B,SMALL(IF(($B$5:$B$80>=$C$34)
* ($B$5:$B$80<=$F$28), ROW($B$5:$B$80)),
ROWS($B$5:$B5))), "")
```

Not only is this formula hard to parse and understand, but also the post maker was not able to conceive of this formula themselves. They further could not alter the suggested formula for their task and needed assistance to organize the results. The solution to this problem is much simpler and cleaner in SQL:

```
SELECT text, date FROM table WHERE date >
date1 AND date < date2
```

Multiple users similarly try to create a complex formula using VLOOKUP in order to complete what at its basic form is a join. Similar observations could be made of many of the questions in the *querying data* theme—that they could be expressible as simple SQL queries (with joins, group-by, and selections) [10, 24, 6]. Using SQL queries would help users avoid using computationally intensive formulae (as was suggested by Reddit users) like COUNTIF, SUMIF, and VLOOKUP.

LIMITATIONS
Our methodology poses some limitations. Our collection approach using the "random" API call is described as a "the serendipity button," indicating that it returns a random post from the forum. However, verifying that it is truly random is challenging. In our results, we see a bias towards more recent posts. We also have a self-selection bias in the posters that contribute to the subreddit we studied.

CONCLUSIONS
In our case study, we set out to apply a method of utilizing online forum posts to understand usability issues people face with spreadsheets. The majority of posts in the subreddit were Excel users asking questions about the four main steps needed to process data: importing, managing, querying, and presenting—we found that all four are important. Furthermore, not only are scalability issues very real problems plaguing those who want to analyze large amounts of data, but these issues touched on all four areas of actions users take to process data. This motivates the need for developing a spreadsheet application which continues to perform as expected on very large volumes of data [12].

REFERENCES
1. 2017. Excel specifications and limits. `goo.gl/PiWptB`. (2017). Accessed: September 24, 2017.

2. 2017. Please welcome our new Corporate Overlords. `https://www.reddit.com/r/excel/comments/3lusom/please_welcome_our_new_corporate_overlords/`. (2017). Accessed: September 24, 2017.

Name	Num Posts	Percentage
Import	11	4.0
Manipulate	16	5.8
Delete	9	3.2
Update	3	1.1
Calculation	32	11.5
Selection	45	16.2
Aggregation	31	11.2
Visualization	13	4.7
Style	21	7.6
Layout	19	6.8
How to	52	18.7
Advice	9	3.2
Simple	8	2.9
Irrelevant	9	3.2

Table 2. Accessible representation of distribution of random posts (figure 3).

Name	Num Posts	Percentage
Import	2	5.1
Manipulate	20	4.6
Delete	10	2.3
Update	4	.9
Calculation	35	8.1
Selection	40	9.2
Aggregation	27	6.2
Visualization	10	2.3
Style	11	2.5
Layout	71	16.4
How to	79	18.2
Advice	31	7.1
Simple	54	12.4
Irrelevant	20	4.6

Table 3. Accessible representation of distribution of search term posts (figure 4).

Name	Num Posts	Percentage
Import	11	13.3
Manipulate	1	1.2
Delete	1	1.2
Update	1	1.2
Calculation	8	9.6
Selection	11	13.3
Aggregation	8	9.6
Visualization	3	3.6
Style	1	1.2
Layout	7	8.4
How to	4	18.7
Advice	8	9.6
Simple	19	22.9

Table 4. Accessible representation of distribution of scalability posts (figure 5).

3. 2017a. Reddit API Documentation. `https://www.reddit.com/dev/api/`. (2017). Accessed: September 15, 2017.

4. 2017. /r/excel metrics (Microsoft Excel | Help and Support with your Formula, Macro, and VBA problems | A Reddit Community). `http://redditmetrics.com/r/excel`. (2017). Accessed: September 24, 2017.

5. 2017b. Upcoming Changes: View counts, users here now and traffic pages. `https://www.reddit.com/r/ModSupport/comments/6atvgi/upcoming_changes_view_counts_users_here_now_and/`. (2017). Accessed: September 24, 2017.

6. Azza Abouzied, Joseph Hellerstein, and Avi Silberschatz. 2012. DataPlay: interactive tweaking and example-driven correction of graphical database queries. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*. ACM, 207–218.

7. Michael Alexander, Jared Decker, and Bernard Wehbe. 2014. *Microsoft business intelligence tools for Excel analysts*. John Wiley & Sons.

8. Lisa Anthony, YooJin Kim, and Leah Findlater. 2013. Analyzing user-generated youtube videos to understand touchscreen use by people with motor impairments. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1223–1232.

9. Salvatore Aurigemma and Ray Panko. 2014. Evaluating the Effectiveness of Static Analysis Programs Versus Manual Inspection in the Detection of Natural Spreadsheet Errors. *Journal of Organizational and End User Computing (JOEUC)* 26, 1 (2014), 47–65. DOI: `http://dx.doi.org/10.4018/joeuc.2014010103`

10. Eirik Bakke, David Karger, and Rob Miller. 2011. A spreadsheet-based user interface for managing plural relationships in structured data. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2541–2550. DOI: `http://dx.doi.org/10.1145/1978942.1979313`

11. Eirik Bakke and David R Karger. 2016. Expressive query construction through direct manipulation of nested relational results. In *Proceedings of the 2016 International Conference on Management of Data*. ACM, 1377–1392.

12. Mangesh Bendre, Vipul Venkataraman, Xinyan Zhou, Kevin Chen-Chuan Chang, and Aditya Parameswaran. 2017. Towards a Holistic Integration of Spreadsheets with Databases: A Scalable Storage Engine for Presentational Data Management. *arXiv preprint arXiv:1708.06712* (2017).

13. Chris Chambers and Chris Scaffidi. 2010. Struggling to excel: A field study of challenges faced by spreadsheet users. In *Visual Languages and Human-Centric Computing (VL/HCC), 2010 IEEE Symposium on*. IEEE, 187–194. DOI:`http://dx.doi.org/10.1109/VLHCC.2010.33`

14. David G Hendry and Thomas RG Green. 1994. Creating, comprehending and explaining spreadsheets: a cognitive interpretation of what discretionary users think of the spreadsheet model. *International Journal of Human-Computer Studies* 40, 6 (1994), 1033–1065. DOI: `http://dx.doi.org/https://doi.org/10.1006/ijhc.1994.1047`

15. Sun Hee Jang. 2011. YouTube as an innovative resource for social science research. In *Australian Association for Research in Education Conference (AARE 2011 Conference)*. Citeseer, 1–16.

16. Jennifer Keelan, Vera Pavri-Garcia, George Tomlinson, and Kumanan Wilson. 2007. YouTube as a source of information on immunization: a content analysis. *jama* 298, 21 (2007), 2482–2484.

17. Swetha Keertipati, Bastin Tony Roy Savarimuthu, and Sherlock A. Licorish. 2016. Approaches for prioritizing feature improvements extracted from app reviews.. In *EASE*, Sarah Beecham, Barbara Kitchenham, and Stephen G. MacDonell (Eds.). ACM, 33:1–33:6. `http://dblp.uni-trier.de/db/conf/ease/ease2016.html#KeertipatiSL16`

18. Jennifer G Kim, Ha Kyung Kong, Karrie Karahalios, Wai-Tat Fu, and Hwajung Hong. 2016. The power of collective endorsements: credibility factors in medical crowdfunding campaigns. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 4538–4549.

19. Danai Koutra, Paul N Bennett, and Eric Horvitz. 2015. Events and controversies: Influences of a shocking news event on information seeking. In *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 614–624. DOI: `http://dx.doi.org/10.1145/2736277.2741099`

20. Susan E Kruck and Steven D Sheetz. 2001. Spreadsheet accuracy theory. *Journal of Information Systems Education* 12, 2 (2001), 93–108. DOI: `http://dx.doi.org/10.4018/joeuc.2014010103`

21. Juhi Kulshrestha, Motahhare Eslami, Johnnatan Messias, Muhammad Bilal Zafar, Saptarshi Ghosh, Krishna P Gummadi, and Karrie Karahalios. 2017. Quantifying search bias: Investigating sources of bias for political searches in social media. (2017), 417–432. DOI: `http://dx.doi.org/10.1145/2998181.2998321`

22. Microsoft. 2017. Access. (2017). `https://products.office.com/en-us/access`

23. Roland Mittermeir and Markus Clermont. 2002. Finding high-level structures in spreadsheet programs. In *Reverse Engineering, 2002. Proceedings. Ninth Working Conference on*. IEEE, 221–232. DOI: `http://dx.doi.org/10.1109/WCRE.2002.1173080`

24. Arnab Nandi and Michael Mandel. 2013. The interactive join: recognizing gestures for database queries. In

CHI'13 Extended Abstracts on Human Factors in Computing Systems. ACM, 1203–1208.

25. Bonnie A Nardi and James R Miller. 1990a. An ethnographic study of distributed problem solving in spreadsheet development. In *Proceedings of the 1990 ACM conference on Computer-supported cooperative work*. ACM, 197–208. DOI: http://dx.doi.org/10.1145/99332.99355

26. Bonnie A Nardi and James R Miller. 1990b. *The spreadsheet interface: A basis for end user programming*. Hewlett-Packard Laboratories.

27. Ray Panko. 2016. What We Don't Know About Spreadsheet Errors Today: The Facts, Why We Don't Believe Them, and What We Need to Do. *arXiv preprint arXiv:1602.02601* (2016).

28. Stephen G Powell, Kenneth R Baker, and Barry Lawson. 2008. A critical review of the literature on spreadsheet errors. *Decision Support Systems* 46, 1 (2008), 128–138. DOI:http://dx.doi.org/10.1016/j.dss.2008.06.001

29. Kamalasen Rajalingham, David R Chadwick, and Brian Knight. 2008. Classification of spreadsheet errors. *arXiv preprint arXiv:0805.4224* (2008).

30. Thomas Reschenhofer and Florian Matthes. 2015. An empirical study on spreadsheet shortcomings from an information systems perspective. In *International Conference on Business Information Systems*. Springer, 50–61.

31. Steve Whittaker, Loen Terveen, Will Hill, and Lynn Cherny. 2003. The dynamics of mass interaction. In *From Usenet to CoWebs*. Springer, 79–91.