

**A peer-reviewed version of this preprint was published in PeerJ on 29 July 2015.**

[View the peer-reviewed version](https://doi.org/10.7717/peerj-cs.12) (peerj.com/articles/cs-12), which is the preferred citable publication unless you specifically need to cite this preprint.

Borges HS, Valente MT. 2015. Mining usage patterns for the Android API. PeerJ Computer Science 1:e12 <https://doi.org/10.7717/peerj-cs.12>



# Mining usage patterns for the Android API

Hudson Silva, Marco Tulio Valente

API methods are not used alone, but in groups and following patterns. However, despite being a key information for API users, most usage patterns are not described in official API documents. In this article, we report a study that evaluates the feasibility of automatically enriching API documents with information on usage patterns. For this purpose, we mine and analyze 1,952 usage patterns, from a set of 396 Android applications. As part of our findings, we report that the Android API has many undocumented and non-trivial usage patterns, which can be inferred using association rule mining algorithms. We also describe a field study where a version of the original Android documentation is instrumented with the extracted usage patterns. During 17 months, this documentation received 77,863 visits from professional Android developers.



# Mining Usage Patterns for the Android API

Hudson Borges and Marco Tulio Valente

Department of Computer Science, UFMG, Brazil

{hsborges,mtov}@dcc.ufmg.br

## ABSTRACT

API methods are not used alone, but in groups and following some patterns. However, despite being a key information for API users, most usage patterns are not explicitly described in official API documents. In this article, we report a study to evaluate the feasibility of automatically enriching API documents with information on usage patterns. For this purpose, we extract and analyze 1,952 usage patterns, from a set of 396 Android applications. As part of our findings, we report that the Android API has many undocumented and non-trivial usage patterns, which can be inferred using association rule mining algorithms. We also describe a field study where a version of the original Android documentation was instrumented with the extracted usage patterns and source code examples. During 17 months, this documentation received 77,863 visits from real Android developers.

**Keywords:** Application Programming Interfaces, Usage Patterns, Android

## INTRODUCTION

Methods in modern APIs are not used independently of each other, but according to some patterns (Robillard et al., 2013; Long et al., 2009). For example, the Android JavaDoc page that documents the `beginTransaction` method explicitly reports that it is usually used together with `setTransactionSuccessful` and `endTransaction`. However, this page is an exception and most usage patterns are not documented at all. We reach this conclusion after inspecting the Android documentation, searching for 100 popular usage patterns, mined from a dataset of 396 applications. We found that only 12 patterns are somehow documented.

This article reports the first (to the best of our knowledge) large-scale field study on the instrumentation of API documents with usage patterns. The study is based on the Android API, which is selected due to its complexity, size, and relevance to Android developers (Syer et al., 2011; Ruiz et al., 2014). We consider an API usage pattern as set of API methods that are used together with a certain frequency (Robillard et al., 2013). We extend a tool, called APIMiner (Montandon et al., 2013), to mine usage patterns from a dataset of Android open-source applications. This tool also instruments the original API documents with information on the extracted usage patterns.

The study reported in this paper is divided in three parts:

- First, we report a characterization study on the usage of the Android API by client applications. Our central goal is to check whether the Android API and



the proposed dataset of Android clients are indeed interesting objects of study, considering our central goal.

- Next, we describe the methodology followed to extract usage patterns for the Android API and we characterize such patterns, in terms of their representativeness and complexity.
- Finally, we report a field study, when our version of the Android API instrumented with usage patterns and associated examples was made available to public access. During 17 months, it received 77,863 visits, coming from 160 countries.

## MATERIALS & METHODS

In this section, we present the tool we used to enhance API documents with information on usage patterns, the dataset used to mine these patterns, and the support and confidence thresholds used by the mining algorithm.

### APIMiner

APIMiner (<http://apiminer.org>) is a tool that instruments JavaDocs with code examples, extracted from API clients (Montandon et al., 2013). As illustrated in Figure 1, an “Example Button” is included in the original documentation, before the signature of each API method. By clicking on these buttons, developers are presented with source code examples for the documented API methods. A detailed presentation of the algorithms used by APIMiner to extract, summarize, and rank examples is out of the scope of this paper and we refer the interested reader to our previous work (Montandon et al., 2013).

| Public Methods         |                  |   |
|------------------------|------------------|---|
| Example<br>9 Examples  | abstract void    | <code>cancel()</code><br>Turn the vibrator off.   |
| Example<br>4 Examples  | abstract boolean | <code>hasVibrator()</code><br>Check whether the hardware has a vibrator.                        |
| Example<br>26 Examples | abstract void    | <code>vibrate(long[] pattern, int repeat)</code><br>Vibrate with a given pattern.               |
| Example<br>54 Examples | abstract void    | <code>vibrate(long milliseconds)</code><br>Vibrate constantly for the specified period of time. |

**Figure 1.** JavaDoc instrumented by APIMiner Montandon et al. (2013)

For this article, we extend APIMiner with a capability to provide examples for API methods that are often called together. An API usage pattern has the following form:

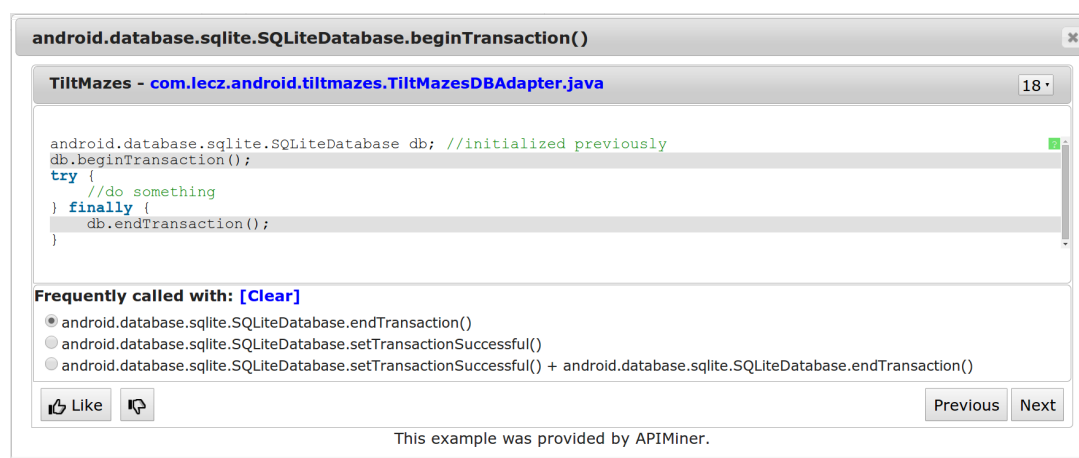
$$M \Rightarrow M_1, M_2, \dots, M_n$$

where  $M$  and  $M_i$ ,  $1 \leq i \leq n$ , are methods from the API of interest. This pattern expresses that when the method  $M$  (antecedent term) is called by a given client method  $\mathcal{M}$ , methods  $M_1, M_2, \dots, M_n$  (consequent terms) are usually called by  $\mathcal{M}$ , not necessarily in this order.



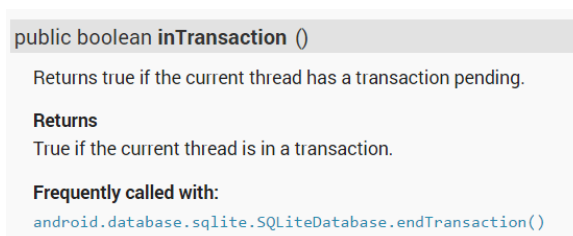
53 The extension adds a new module to APIMiner for inferring usage patterns. This  
54 module relies on a standard association rules mining algorithm (Agrawal and Srikant,  
55 1994), which retrieves association rules with the same syntax and semantics of the  
56 proposed usage patterns. The transactions used as input to this algorithm are the  
57 methods of the client systems and the API methods called by them (all calls are statically  
58 extracted, based on the static types of the target objects). To evaluate the significance  
59 of the extracted patterns we rely on two measures: support (number of transactions  
60 that include the methods in the pattern) and confidence (the probability of finding  
61 the methods from the consequent term in the subset of the transactions including the  
62 antecedent method).

63 We also extend the JavaDoc interface to show examples for usage patterns, as  
64 presented in Figure 2. In the usage scenario illustrated by this figure, the API user  
65 initially requested examples for the `beginTransaction()` method, using the original  
66 interface provided by APIMiner. The window presenting the examples has a bottom  
67 panel with the usage patterns that have `beginTransaction` as the antecedent term. For  
68 instance, after selecting the first pattern the user sees examples that include not only  
69 `beginTransaction()` but also `endTransaction()`.



**Figure 2.** Interface for presenting examples for API usage patterns

70 We also instrumented the JavaDoc sections that document the API methods with  
71 information on other methods they are frequently called with (if any), as presented in  
72 Figure 3.



**Figure 3.** Usage patterns are presented in the detailed documentation of methods



## Mining Dataset

Different from work that process Android bytecode downloaded from Google Play (Ruiz et al., 2014), in our case it is important to have the original source code, to guarantee the extraction of examples with a minimal level of legibility. For this reason, we relied on GitHub to construct a dataset with the source code of Android systems, which we use for mining the usage patterns considered in this paper. We downloaded Android projects from GitHub that have at least 50 commits, to restrict the analysis to projects with a minimal level of activity, and that are not forks of other projects, to avoid many similar projects in the dataset. By considering these requirements, we create a mining dataset with 396 projects, including well-known applications, such as WordPress, Astrid, K9, and ConnectBot. Considering all projects, the dataset includes 57,658 classes and 450,762 methods.

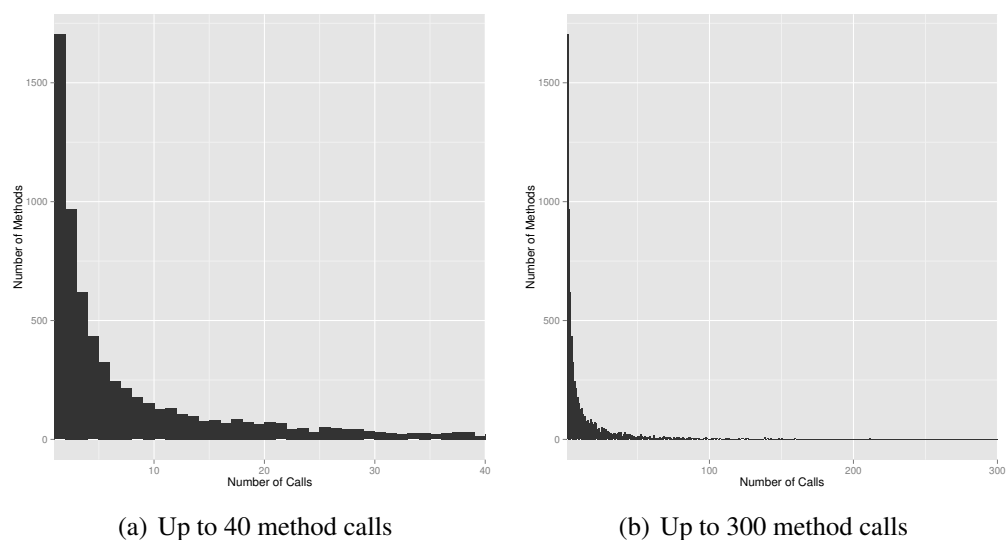
For this study, we use the version 4.1.2-r1 of the Android API. We initially evaluate the usage of the API by the systems in the proposed dataset. Considering just the methods that call the API, 59% call a second API method, which shows the feasibility of searching for usage patterns. In the dataset, 40% of the public or protected methods from the Android API are never called. Table 1 presents the ten packages with the highest percentage of unused classes, i.e., classes that do not have any method call in the dataset. The packages with the highest percentage of unused classes are `android.provider` and `android.drm`. In both packages, 81% of the classes are never used.

**Table 1.** Packages with the highest percentage of unused classes

| Package                           | # Classes | % Unused Classes |
|-----------------------------------|-----------|------------------|
| <code>android.provider</code>     | 177       | 81               |
| <code>android.drm</code>          | 26        | 81               |
| <code>android.mtp</code>          | 5         | 80               |
| <code>android.sax</code>          | 7         | 71               |
| <code>android.media</code>        | 131       | 69               |
| <code>android.security</code>     | 3         | 67               |
| <code>android.test</code>         | 39        | 64               |
| <code>android.renderscript</code> | 55        | 64               |
| <code>android.widget</code>       | 222       | 56               |
| <code>android.net</code>          | 86        | 54               |

We also analyzed the distribution of the number of API calls in our dataset. Figure 4 shows two histograms with the frequency of the number of calls. For each value  $n$  in the x-axis, the y-axis represents the number of methods with exactly  $n$  calls. To ease visualization, the first histogram shows methods with at most 40 calls; and the second histogram with at most 300 calls (in the full dataset, the number of calls ranges from 1 to 6,729). The histogram is right-skewed, meaning that while most methods are called few times (median equal to 5), we also have high and very high values. This finding implies that centrality and dispersion statistics measures (e.g., mean and standard deviation) should not be used to describe our empirical data on number of calls. Instead, the histogram suggests the data best fits a heavy-tailed distribution, possibly a power law. To check this possibility, we use the statistical framework proposed by Clauset





**Figure 4.** Histogram of number of calls

et al. (Clauset et al., 2009) for discerning power-law behavior in empirical data. By following this framework, we reject the (null) hypothesis that power-laws are a plausible explanation for our data, for a significance level of 5% ( $p$ -value= 0.00). However, this is not the same as concluding that the number of calls do not match a heavy-tailed distribution. In fact, Figure 4 suggests a heavy-tailed behavior, possibly matching an alternative distribution (e.g., stretched exponential or log-normal).

Finally, we correlated the project's size (in terms of number of methods) and the usage of the API (in terms of API calls/method). Initially, we checked whether this data follows a normal distribution, which was not the case. For this reason, we performed the Kendall Tau rank correlation test to measure the correlation between the defined variables. The result was a correlation coefficient of -0.43, which indicates that smaller projects depend more on the Android API than large ones.

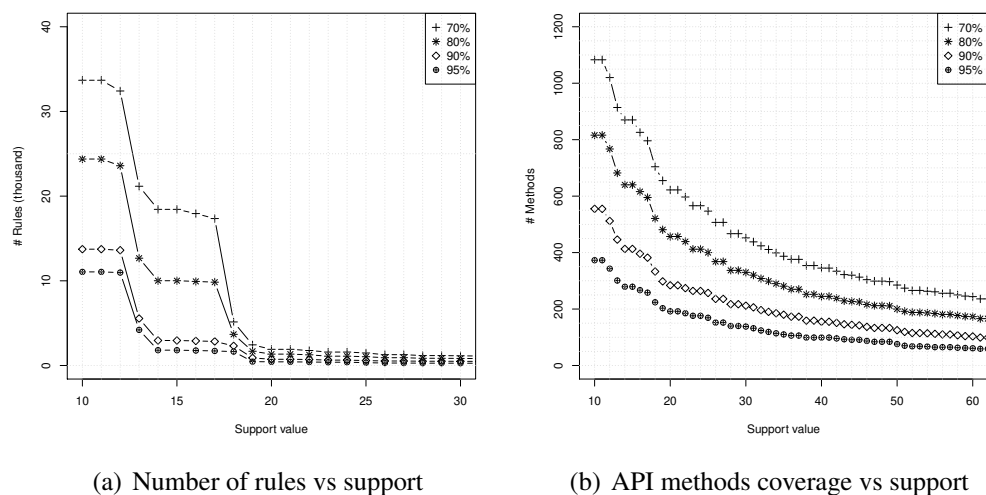
## Mining Usage Patterns

The first step for mining usage patterns is to set up the support and confidence thresholds. Figure 5a shows the number of association rules by varying the support values. We restrict the confidence thresholds to 70%, 80%, 90%, and 95% to keep a minimum quality in the rules. We can observe that small support values generate too many rules. Fixing a support of 10 transactions, we have 11,054 rules for a confidence of 95% and 33,685 rules for a confidence of 70%. In fact, the number of association rules starts to grow very rapidly for support values less or equal to 20. For this reason, we decide to use a support of 21 transactions.

Figure 5b shows the number of methods covered by the extracted rules, considering just rules with a single method in the antecedent term. For a support of 21 transactions, the coverage ranges from 192 methods (confidence of 95%) to 624 methods (confidence of 70%). To increase API's coverage, we decide to fix a confidence of 70%.

By using the proposed thresholds, 1,952 usage patterns are mined, covering 624 API methods, which represent 5% of the public and protected methods in the Android API





**Figure 5.** Setting up the support and confidence thresholds

and 8% of the methods called in our dataset.

## RESULTS

Initially, this section provides examples of usage patterns and also a classification of patterns in two categories, intra and inter-class. We then report an evaluation with developers, to validate whether the patterns really include methods that are called together. Finally, we report a field study, when we made our tool available to public usage.

### Examples and Types of Patterns


Table 2 presents the usage patterns with the highest support. The most common usage pattern—found in 1,362 transactions with a confidence of 75%—models the computation required to create a focused UI window, which is called an Activity by the Android API. The pattern expresses that client methods that set an Activity's view by calling `Activity setContentView(View)` usually call `Activity.onCreate(Bundle)` to initialize the view.

In Table 2, for nine patterns the methods in the antecedent and in the consequent term come from the same class. We refer to such patterns as intra-class usage patterns. It is more simple to discover these patterns without tool support, since their documentation is restricted to a single JavaDoc page. However, considering the 1,952 usage patterns, there is an almost equal distribution between intra-class and inter-class usage patterns. Specifically, 50.3% of the usage patterns are inter-class, i.e., they have methods coming from more than one class. Figure 6 shows boxplots describing the distribution of the support values, regarding intra-class and inter-class usage patterns. The first and second quartiles of both distributions are very similar. Relevant differences start on the third quartile (63 methods for intra-class patterns vs 47 methods for inter-class patterns). For this reason, among the ten usage patterns in Table 2, only one is inter-class.



**Table 2.** Usage patterns with highest support

| Usage Pattern   | Support | Confidence |
|---|---------|------------|
| Activity setContentView(View) ⇒<br>Activity.onCreate(Bundle)                  | 1,362   | 75         |
| Toast.show() ⇒<br>Toast.makeText(...)   | 1,133   | 86         |
| ViewGroup.getChildCount() ⇒<br>ViewGroup.getChildAt(int)                      | 1,077   | 75         |
| ViewGroup.getChildAt(int) ⇒<br>ViewGroup.getChildCount()                      | 1,077   | 74         |
| AlertDialog.Builder.setTitle(...) ⇒<br>AlertDialog.Builder.Builder(Context)   | 1,019   | 98         |
| ContentValues.put(String,String) ⇒<br>ContentValues.ContentValues()           | 973     | 80         |
| AlertDialog.Builder.create() ⇒<br>AlertDialog.Builder.Builder(Context)        | 765     | 94         |
| AlertDialog.Builder.setMessage(...) ⇒<br>AlertDialog.Builder.Builder(Context) | 736     | 96         |
| LayoutInflater.inflate(...) ⇒<br>View.findViewById(int)                       | 697     | 72         |
| ContentValues.put(String,Integer) ⇒<br>ContentValues.ContentValues()          | 632     | 77         |



**Figure 6.** Support values for intra-class and inter-class usage patterns



## 156 Evaluation by Developers

157 Association rules can generate meaningless patterns, due to random relations that exist  
 158 in most datasets. To validate the mined relations, we randomly selected a sample of  
 159 45 usage patterns, among the “less reliable” ones. By less reliable we refer to patterns  
 160 with the lowest support and confidence. This sample was selected among 477 usage  
 161 patterns ( $\approx 25\%$ ) with support less than 50 and confidence less than 80%. We call this  
 162 first sample the treatment group. Moreover, we randomly selected five sequences of  
 163 methods, among the existing methods in the Android API. We call this second sample  
 164 the control group. We presented each sequence of methods to two experienced Android  
 165 developers and asked them the following question:  
 166

167 *Do you recommend to refer to method Y when documenting method X,*  
 168 *since they are usually used together?*

169 Both developers provided negative answers for all patterns in the control group, as  
 170 expected. Table 3 summarizes the results for the 45 patterns in the treatment group. The  
 171 first developer provided positive answers for 33 usage patterns (73%) and the second one  
 172 for 37 patterns (82%). When combining the answers, 28 patterns (62%) were evaluated  
 173 positively by both developers and 3 patterns (7%) received two negative answers. 42  
 174 patterns (93%) received at least a positive recommendation and 17 patterns (38%)  
 175 received at least a negative recommendation. Therefore, according to the developers’  
 176 judgment, we can reach a precision of 62% or 93% depending on the classification  
 177 criteria (two positive answers vs at least one positive answer). This result reveals that  
 178 the mined patterns are not coincidences or spurious relations, specially considering that  
 179 we evaluated patterns in the first quartile of the support and confidence distributions.

**Table 3.** Usage patterns evaluation by developers

| Developer | Answers  |          |
|-----------|----------|----------|
|           | Yes      | No       |
| #1        | 33 (73%) | 12 (27%) |
| #2        | 37 (82%) | 8 (18%)  |
| #1 & #2   | 28 (62%) | 3 (7%)   |
| #1    #2  | 42 (93%) | 17 (38%) |

180 However, it is important to highlight that the developer’s judgment is influenced by  
 181 their personal experience with the API. For example, we discussed in details the patterns  
 182 with two negative answers. In common, the developers initially argued they did not  
 183 find reasons to the methods being called together. To clarify this common answer, we  
 184 analyzed the following pattern in more details:  
 185

186 `Fragment.onActivityCreated(android.os.Bundle)  $\Rightarrow$  Fragment.getActivity()`  
 187

188 One of the developers said this pattern does not make sense, “because the first  
 189 method is a callback, called after an Activity is created”. We then presented to this  
 190 developer many instances of client code calling both methods. Most examples are  
 191 subclasses of Fragment that override onActivityCreated and then call the superclass



192 method, using `super`. After that, `getActivity()` is invoked to access the current  
193 Activity context. The developer acknowledged that some Android developers may use  
194 the methods in this way, but he also highlighted that the same “behavior can be achieved  
195 by arranging the Fragment classes in other ways”, as he usually prefers to do.

## 196 **Field Study**

197 We conducted a field study with the purpose of assessing the importance that API users  
198 give to usage patterns. Particularly, the study was not designed to evaluate the usability  
199 of the provided source code examples, which it is always a challenge in the case of open  
200 field studies. Instead, we focused on the frequency that real users search for examples,  
201 including examples for usage patterns. We claim that answers for such questions—  
202 coming from the real usage of a complex API such as Android—are important for  
203 developers that want to instrument their API documents with usage patterns and to  
204 guide further research on code summarization techniques. In this way, the study aims to  
205 answer the following research questions:

- 206 • Do API users search for source code examples?
- 207 • Do API users search for examples for usage patterns?

208 To answer such questions, we analyzed the accesses made to a public instance of  
209 APIMiner, available at <http://apiminer.org>. We collected access data to this  
210 site from May 13th, 2013 to October 14th, 2014 (17 months), using a private logging  
211 service and Google Analytics. During this time frame, APIMiner received a total of  
212 77,863 visits, which gives an average of 150 visits/day. These visits came from 160  
213 countries and the top three countries in number of visits were India (14.3%), United  
214 States (11.8%), and Brazil (4.9%). In total, 63,314 users visited the platform and 14,867  
215 visits (19.1%) were from returning visitors. Finally, 54,704 visits (70.2%) came from  
216 queries performed in search engines, mostly in Google. The visits generated 114,124  
217 page views. However, 60,029 page views (52.6%) have a very short duration, less than  
218 one second, or retrieved pages that are not instrumented by APIMiner (e.g., the site front  
219 page or several tutorials included in the Android documentation). Therefore, such page  
220 views were discarded from our analysis. Furthermore, when analyzing the requests for  
221 examples, we excluded requests to methods from the `SQLiteDatabase` class, because  
222 this class is used in the main page of the site to illustrate our usage patterns concept.

### 223 ***Do API users search for source code examples?***

224 The visits generated 14,402 requests for source code examples, i.e., clicks in the “Ex-  
225 ample Button”. Therefore, on average 26.6% of the considered page views included an  
226 “Example Button” click. During the field study, 35,596 examples were presented to the  
227 users, i.e., on average 2.47 examples were presented after each click in the “Example  
228 Button”. Therefore, the users navigated through the list of examples to see at least a  
229 second example.

230 Table 4 presents the top ten methods with more requests for examples. These  
231 methods, with the exception of `Toast.setGravity(...)`, do not have usage patterns.  
232 However, among the 933 methods that received requests for example, 125 methods  
233 (13.4%) have at least one usage pattern.



**Table 4.** Top ten methods with the highest number of requests for examples

| Method  | Req. |
|---|------|
| ViewPager.OnPageChangeListener.onPageSelected(int)              | 101  |
| SimpleCursorAdapter.SimpleCursorAdapter(Context,int,Cursor,...) | 91   |
| RectF.RectF(float,float,float,float)                            | 57   |
| Point.Point(int,int)  | 56   |
| ViewPager.OnPageChangeListener.onPageScrollStateChanged(int)    | 55   |
| SimpleCursorAdapter.setViewBinder(ViewBinder)                   | 50   |
| BitmapRegionDecoder.decodeRegion(Rect,BitmapFactory.Options)    | 50   |
| RectF.RectF()   | 44   |
| Toast.setGravity(int,int,int)                                   | 43   |
| ViewPager.OnPageChangeListener.onPageScrolled(int,float,int)    | 42   |

#### 234 ***Do API users search for examples for usage patterns?***

235 In our dataset, 624 methods have usage patterns. Considering these methods, 306 meth-  
 236 ods (49%) received at least one click in their respective “Example Button”. Computing  
 237 all clicks, these methods received 1,301 requests for examples. In other words, a win-  
 238 dow with source code examples including an option to present just examples for usage  
 239 patterns—such as the one presented in Figure 2—was activated 1,301 times during the  
 240 field study. In such situations, the users selected an option to just show examples for a  
 241 given usage pattern 399 times (30.6% of the window activations).

## 242 **DISCUSSION**

243 Our main findings on using APIMiner for extracting examples for the Android API are  
 244 as follows.

### 245 **Most Android API methods are underused**

246 Only 60.5% of the methods in the Android API are called by the systems in our dataset.  
 247 Therefore, even considering that this dataset might not represent an ideal sample of the  
 248 whole population of Android applications (e.g., it only includes open-source apps), this  
 249 result suggests that a considerable proportion of the Android API methods are rarely  
 250 used by real clients.

251 Therefore, API developers should monitor the usage of their API elements by real  
 252 client systems. As a result, it is likely to conclude that many elements are underused  
 253 or not used at all, suggesting a possible move to a streamlined API. In fact, a similar  
 254 coverage rate is reported in other studies on API usage. Ma, Amor, and Tempero found  
 255 that only 21% of the methods in the Java API are used in a corpus of open-source  
 256 software (Ma et al., 2006). A not fundamentally different coverage pattern happens to  
 257 appear on crowd-based Q&A forums. For example, Parnin et al. report that 13% of  
 258 the classes in the Android API do not have discussion threads at all at Stack Overflow  
 259 (Parnin et al., 2012).



## 260 **Less than 10% of the Android API methods called by clients have usage** 261 **patterns**

262 In our dataset, 8.1% of the Android methods have usage patterns. Even though this  
263 dataset does not include thousands of apps, as datasets based on Android bytecode (Ruiz  
264 et al., 2014), the considered support threshold was properly adapted to its size.

265 Therefore, APIMiner shows that it is feasible to instrument API documents in a  
266 seamless way both with source code examples and with usage patterns. The only re-  
267 quirement is to have a representative sample of client systems. However, API developers  
268 should expect to retrieve usage patterns for less than 10% of their API methods.

## 269 **In one out of three opportunities users search for usage patterns**

270 During the field study, when examples for methods with usage patterns were available,  
271 in 30.6% of the cases the users requested examples for the mined patterns. Therefore,  
272 by including examples for usage patterns, API developers can help such users on their  
273 specific needs when browsing API documents.

## 274 **Threats to Validity**

275 There are at least four threats that could undermine the validity of our results. First,  
276 although the Android API is a complex and popular API, we cannot claim that our  
277 findings apply univocally to other APIs, specially to APIs for other languages or targeting  
278 a different application domain. Second, even in the universe of Android applications and  
279 considering a mining dataset with 396 projects, we might have missed usage patterns  
280 that are common just among applications from a particular category (e.g., location-based  
281 applications). Third, the selection of the support and confidence thresholds, as usual  
282 in association rules mining, is to some extent a subjective decision. To control this  
283 threat, we experimented various threshold combinations aiming to balance coverage  
284 and representativeness. Despite that, we could not estimate the impact on our field  
285 study of a different threshold selection, specially a less strict one. Fourth, our field  
286 study is an observational study, i.e., the subjects are not divided in treated and control  
287 groups. Therefore, there is always a risk of selection bias, specially when compared with  
288 controlled experiments. However, controlled experiments with real software practitioners  
289 are very hard to conduct, specially in the area of API reuse. The main reason is that  
290 real-world development tasks may take hours or even days to be concluded .

## 291 **RELATED WORK**

292 Kagdi et al. compared frequent itemset and sequential-pattern matching and concluded  
293 that the latter is usually worth the additional cost (Kagdi et al., 2007). However, they  
294 do not aim the extraction of examples, when the order of the calls is immediately  
295 visible in the extracted code fragments. CodeWeb is a system that mines not only  
296 usage patterns regarding method calls, but also other forms of reuse, such as inheritance  
297 (Michail, 2000). PR-Miner (Li and Zhou, 2005) extracts general programming rules  
298 using frequent itemset mining, with focus on detecting buggy code. Code Recommenders  
299 (<http://www.eclipse.org/recommenders>) extends the Eclipse built-in Java  
300 API documentation with information such as the methods usually overridden when  
301 subclassing a selected type or the methods usually called on a selected object.



Systems such as Strathcona (Holmes et al., 2006) and MAPO (Zhong et al., 2009) explore the syntactic context provided by the IDE to recommend examples. However, the examples do not have documentation purposes, because they are highly dependent of a particular development context. In contrast, systems such as APIMiner (Montandon et al., 2013) and eXoaDocs (Kim et al., 2013) generate a new JavaDoc instrumented with source code examples. However, they do not provide support for API usage patterns. Altair is a tool that automatically generates API function cross-references, which are useful to populate *see also* sections in API documents (Long et al., 2009). The recommended functions are not computed using association rules, but based on their structural similarity with the function the cross-reference refers to.

Baker is a tool that links source code examples extracted from Q&A sites to API documentation. The tool relies on a constraint-based technique to uniquely identify fully qualified names in source code snippets (Subramanian et al., 2014). ExPort is a tool that detects complex API usages, which can for example crosscut function implementations (Moritz et al., 2013). Saied et al. propose a technique to detect multi-level usage patterns, which are API methods uniformly used across variable client programs, independently of usage context (Saied et al., 2015a). The authors later proposed a technique to infer API usage patterns using structural and semantic relationships mined in the own API source code, i.e., without requiring client programs (Saied et al., 2015b). Since our central goal was to evaluate usage patterns in the field, with real API users, we decided to conduct our study using the most established usage patterns mining technique (association rules). Further work can include a comparison with the aforementioned techniques.

## CONCLUSION

This paper provides a large-scale study of API usage patterns, including the extraction of patterns for a relevant API, an evaluation by expert developers, and a field study, when such patterns were presented to real users. For practitioners, specially API developers and maintainers, our study shows that with the wide availability of source code repositories, like GitHub, it is feasible to generate API documents instrumented with source code examples and usage patterns, both mined automatically. Moreover, the heavy-tailed behavior observed on the usage of API elements suggest to practitioners that most elements of their APIs may be underused or not used at all, and therefore it might be possible to evolve to a streamlined API. Our study also shows that it is possible to collect real data on the usage of research prototypes. Further research is also possible, specially on techniques for summarizing code examples and for mining usage patterns.

## REFERENCES

- Agrawal, R. and Srikant, R. (1994). Fast algorithms for mining association rules in large databases. In *20th International Conference on Very Large Data Bases (VLDB)*, pages 487–499.
- Clauset, A., Shalizi, C. R., and Newman, M. E. J. (2009). Power-Law Distributions in Empirical Data. *Society for Industrial and Applied Mathematics Review*, 51(4):661–703.
- Holmes, R., Walker, R., and Murphy, G. (2006). Approximate structural context



- 344 matching: An approach to recommend relevant examples. *IEEE Transactions on*  
345 *Software Engineering*, 32(12):952–970.
- 346 Kagdi, H. H., Collard, M. L., and Maletic, J. I. (2007). Comparing approaches to mining  
347 source code for call-usage patterns. In *4th Workshop on Mining Software Repositories*  
348 *(MSR)*, page 20.
- 349 Kim, J., Lee, S., Hwang, S., and Kim, S. (2013). Enriching documents with examples:  
350 A corpus mining approach. *ACM Transactions on Information Systems*, 31(1):1.
- 351 Li, Z. and Zhou, Y. (2005). PR-Miner: automatically extracting implicit program-  
352 ming rules and detecting violations in large software code. In *13th Symposium on*  
353 *Foundations of Software Engineering (FSE)*, pages 306–315.
- 354 Long, F., Wang, X., and Cai, Y. (2009). API hyperlinking via structural overlap. In *17th*  
355 *Symposium on Foundations of Software Engineering (FSE)*, pages 203–212.
- 356 Ma, H., Amor, R., and Tempero, E. D. (2006). Usage patterns of the Java Standard API.  
357 In *13th Asia-Pacific Software Engineering Conference (APSEC)*, pages 342–352.
- 358 Michail, A. (2000). Data mining library reuse patterns using generalized association  
359 rules. In *22nd International Conference on Software Engineering (ICSE)*, pages  
360 167–176.
- 361 Montandon, J. E., Borges, H., Felix, D., and Valente, M. T. (2013). Documenting APIs  
362 with Examples: Lessons Learned with the APIMiner Platform. In *20th Working*  
363 *Conference on Reverse Engineering (WCRE)*, pages 401–408.
- 364 Moritz, E., Linares-Vasquez, M., Poshypanyk, D., Grechanik, M., McMillan, C., and  
365 Gethers, M. (2013). ExPort: Detecting and visualizing API usages in large source code  
366 repositories. In *28th International Conference on Automated Software Engineering*  
367 *(ASE)*, pages 646–651.
- 368 Parnin, C., Treude, C., Grammel, L., and Storey, M.-A. (2012). Crowd documentation:  
369 exploring the coverage and the dynamics of API discussions on Stack Overflow.  
370 Technical report, Georgia Tech, College of Computing.
- 371 Robillard, M. P., Bodden, E., Kawrykow, D., Mezini, M., and Ratchford, T. (2013).  
372 Automated API property inference techniques. *IEEE Transactions on Software*  
373 *Engineering*, 39(5):613–637.
- 374 Ruiz, I. J. M., Adams, B., Nagappan, M., Dienst, S., Berger, T., and Hassan, A. E.  
375 (2014). A large-scale empirical study on software reuse in mobile apps. *IEEE*  
376 *Software*, 31(2):78–86.
- 377 Saied, M., Benomar, O., Abdeen, H., and Sahraoui, H. (2015a). Mining multi-level API  
378 usage patterns. In *22nd International Conference on Software Analysis, Evolution*  
379 *and Reengineering (SANER)*, pages 23–32.
- 380 Saied, M. A., Abdeen, H., Benomar, O., and Sahraoui, H. (2015b). Could we infer API  
381 usage patterns only using the library source code? In *23rd International Conference*  
382 *on Program Comprehension (ICPC)*, pages 1–10.
- 383 Subramanian, S., Inozemtseva, L., and Holmes, R. (2014). Live API documentation. In  
384 *36th International Conference on Software Engineering (ICSE)*, pages 643–652.
- 385 Syer, M. D., Adams, B., Zou, Y., and Hassan, A. E. (2011). Exploring the development  
386 of micro-apps: A case study on the BlackBerry and Android platforms. In *11th IEEE*  
387 *Working Conference on Source Code Analysis and Manipulation (SCAM)*, pages  
388 55–64.
- 389 Zhong, H., Xie, T., Zhang, L., and Pei, J. (2009). MAPO: Mining and recommending



390 API usage patterns. *23rd European Conference on Object-Oriented Programming*  
391 (*ECOOP*), 5653:318–343.