

Docio: Documenting API Input/Output Examples

Siyuan Jiang, Ameer Armaly, Collin McMillan, Qiyu Zhi, and Ronald Metoyer

Department of Computer Science and Engineering
University of Notre Dame, Notre Dame, IN, USA
{sjiang1, aarmaly, cmc, qzhi, rmetoyer}@nd.edu

Abstract—When learning to use an Application Programming Interface (API), programmers need to understand the inputs and outputs (I/O) of the API functions. Current documentation tools automatically document the static information of I/O, such as parameter types and names. What is missing from these tools is dynamic information, such as I/O examples—actual valid values of inputs that produce certain outputs.

In this paper, we demonstrate Docio, a *prototype* toolset we built to generate I/O examples. Docio logs I/O values when API functions are executed, for example in running test suites. Then, Docio puts I/O values into API documents as I/O examples.

Docio has three programs: 1) `funcWatch`, which collects I/O values when API developers run test suites, 2) `ioSelect`, which selects one I/O example from a set of I/O values, and 3) `ioPresent`, which embeds the I/O examples into documents. In a preliminary evaluation, we used Docio to generate four hundred I/O examples for three C libraries: `ffmpeg`, `libssh`, and `protobuf-c`. Docio is open-source and available at: <http://www3.nd.edu/~sjiang1/docio/>

I. INTRODUCTION

Recent studies by Duala-Ekoko and Robillard [1] and by Sillito *et al.* [2] found that when learning Application Programming Interfaces (API), programmers look for input/output examples, **actual values** of a function’s input/output, in order to understand the function. In particular, the programmers asked “How does this data structure look at runtime?” [2] and “we have a `newInstance(String)` method that takes a `String` argument and I have no idea what this `String` is supposed to be” [1].

Programmers often refer to API documents to answer such questions [3], [4]. However, a study by Maalej and Robillard shows that most of the API documents do not contain input/output (I/O) examples [5]. One reason for the documents lacking I/O examples is that API developers have little tool support for documenting I/O examples. Current documentation tools, such as `doxygen` [6], `JavaDoc` [7], automate the process of converting comments and source code into structured html pages. If API developers write I/O examples in comments, the tools can present the I/O examples in html pages. Otherwise, the tools have no I/O examples to present.

When there is little help from API documents, programmers (the users of API libraries) often try a variety of approaches to get input/output examples [8]–[10]. For instance, programmers 1) search or ask a question online (e.g., `stackoverflow.com`), 2) look for existing code that use the API (e.g., open source repositories), or 3) write code to do experiments about the API. If the API documents have input/output examples, programmers may not need these approaches.

To help API developers add I/O examples into API documents, we built Docio, which is a prototype toolset that (1) logs actual I/O values when API functions are executed, (2) selects an example from the logged values for each API function, and (3) visualizes the examples to API document.

In the rest of the paper, we will describe what Docio adds to API documents, how API developers can use Docio, how Docio creates input/output examples, the preliminary results, the related work, and the future work.

II. DOCIO IN A NUTSHELL

Docio is for developers of C API libraries to create API documents with I/O examples—actual values of input/output—of API functions. Docio has three programs: `funcWatch`, `ioSelect` and `ioPresent`. First, `funcWatch` logs I/O values of a function when the function is executed (in tests or in other executable programs). Often the function is executed multiple times. Then, `ioSelect` selects one of the calls. Finally, `ioPresent` adds the I/O values of the selected call into the API document. `ioPresent` is a patch for the documentation tool, `doxygen` [6]. The patch enables `doxygen` to visualize I/O examples into html pages. Our online appendix is at: <http://www3.nd.edu/~sjiang1/docio/>

This website contains a virtual machine that is set up for running Docio, the videos about Docio, the installation guides, and three API documents generated by Docio.

III. MOTIVATIONAL EXAMPLE

In this section, we will show what Docio adds to API documents. Figure 1 shows a document generated by Docio and `doxygen`. The document is about the function `av_bprint_channel_layout` in `ffmpeg`.

The API document is composed of three parts. Part 1 is the declaration of `av_bprint_channel_layout`. Part 2 is a sentence describing the function. Parts 1 and 2 are generated by `doxygen`. Part 3 is what Docio adds to the document, which is an I/O example. In the rest of the section, we will walk through the three parts of the document from the perspective of a user who wants to learn `av_bprint_channel_layout`.

First, the user reads Part 1 and gets some basic information: (1) the name of the function, (2) the return type, which is `void`, and (3) the types and the names of the parameters. According to the previous studies [1], [2], the user may have the following questions after reading Part 1:

- Q1. What does `av_bprint_channel_layout` do?
- Q2. What is `struct AVBPrint`?
- Q3. What is `nb_channels`?

Second, the user reads Part 2, which answers Q1: the function “append a description of a channel layout to a bprint buffer.” From this description, the user may expect an input representing “a description of a channel layout”, and an output representing “a bprint buffer”. The names of the parameters indicate that the parameter `channel_layout` is a channel layout and `bp` is a bprint. But the user still does not know where the description of the channel layout is, and where the buffer of the bprint is. In other words, the user may have the questions:

- Q4. Where is the description of a channel layout?
 Q5. Where is the buffer of the bprint?

Part 3, the I/O example, provides some hints for some unanswered questions. The I/O example shows that `bp->str` changed from an empty string to “stereo”. So the reader can guess that `bp->str` is the buffer of the bprint (answers Q5).

The I/O example also shows that `nb_channels` can be “-1”, which is related to Q3. For Q4, the user still does not know where the description of a channel layout is. What the reader can know from the I/O example is that when `channel_layout` is three, the description is “stereo”.

In summary, all the three parts have useful information that helps the user to understand `av_bprint_channel_layout`. Part 1 has static information extracted from source code. Part 2 has the high-level description that only API developers can provide. Docio adds Part 3, the actual, concrete input/output values of a function.

IV. USAGE OF DOCIO

The usage of Docio has three steps: 1) run `funcWatch`, 2) run `ioSelect`, and 3) run `ioPresent`. If an API developer wants to get I/O examples for all the functions in the library, the developer needs to have a list of the function names for running `funcWatch` (so that `funcWatch` knows which function to log). To help developers with this situation, we modified `doxygen` so that `doxygen` writes all the function names in a file when `doxygen` is executed.

We put a demo video in our online appendix (Section II) in which we followed these steps and generated nearly two hundred I/O examples for the `ffmpeg` library.

V. ENVISIONED USERS

The envisioned users of Docio are the developers of API libraries. The developers can use Docio to create API documents with I/O examples. Although the developers are the envisioned users of Docio, the API users (the programmers who use and learn API libraries) are the ones that read the documents and benefit from the I/O examples.

Additionally, the programs in Docio can be used independently for other purposes. For example, `funcWatch` can be used to detect whether a parameter always takes the same value.

VI. DESIGN OF DOCIO

In this section, we will describe the design of Docio. Figure 2 shows the architecture of Docio, which illustrates the relationship among the three programs: `funcWatch`, `ioSelect`,

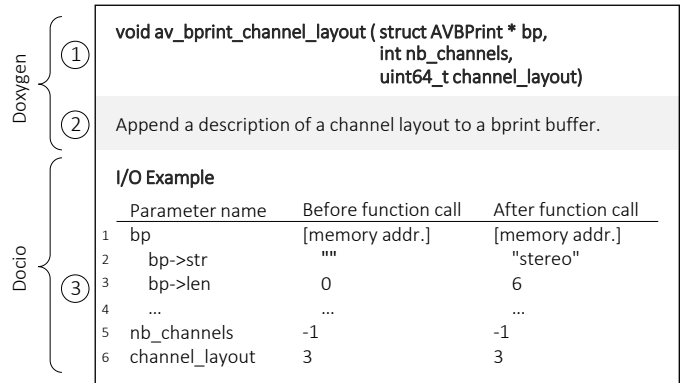


Fig. 1: The document of function `av_bprint_channel_layout` in `ffmpeg`. The document has three parts. Part 1 is the function declaration. Part 2 is a description of the function. Parts 1 and 2 are generated by `doxygen`. Part 3 is an input/output example created by Docio. The input/output example is a table with three columns: “Parameter name”, “Before function call”, and “After function call”. The first column lists the names of parameters. The second column lists the values of the parameters before `av_bprint_channel_layout` is called, i.e. input values. The third column lists the values of the parameters after `av_bprint_channel_layout` is called, i.e. output values.

Each row in the table represents a pair of input/output values. For example, the first row represents the values of the first parameter, `bp`, which is pointer-type. Because the actual values of pointers—memory addresses—are not meaningful, Docio does not put actual values in the table but “[memory addr.]” to indicate the values are memory addresses.

The value that `bp` points to is a struct variable, which is composed of the fields of the struct. The second, third and fourth rows represent the fields. The three dots in the fourth row does not represent a field but represent three more fields of the struct variable. Docio presents the three fields in the actual document, but we omit them due to page limits. Section III discussed this figure in detail.

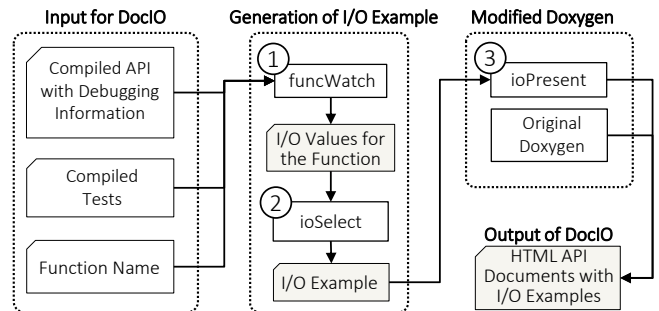


Fig. 2: The architecture of Docio. The inputs of Docio are (1) compiled API, (2) compiled API tests, (3) an API function name. The three programs in Docio are labeled with numbers: (1) `funcWatch`, (2) `ioSelect`, and (3) `ioPresent`. `funcWatch` runs tests and code examples, and logs input/output values of the function specified by the input. `ioSelect` chooses an example from all the logged I/O values. `ioPresent` adds the example to html documents.

and `ioPresent`. In the following subsections, we will describe how we design and implement each program.

A. `funcWatch`: Logging I/O values

The input of `funcWatch` is an API function name and a test. `funcWatch` runs and monitors the test, and, during the execution of the test, `funcWatch` logs the I/O values of the given API function.

We built `funcWatch` based on the code of `Flashback` [11], which uses `libdwarf` [12], a library that provides access to Dwarf debugging information. By using `libdwarf`, `funcWatch`

inserts breakpoints at the entry and the exits of the given function (a C function may have multiple exits).

Then, funcWatch runs the given test with the breakpoints. After a breakpoint is reached and the test paused, funcWatch collects the parameters/return values of the API function. Furthermore, funcWatch is able to dereference pointers at runtime, so funcWatch can collect not only pointers’ values, but also the values that the pointers point to. FuncWatch also calculates the addresses of a struct variable’s fields, so that funcWatch can log the values of the fields. Similarly, funcWatch also work for union and enum types.

In C, arrays are often passed without size information to functions. While a common practice is to pass the size of an array as a separate parameter, the naming conventions differ across different projects and developers. So there is no simple way to automatically detect the size parameters. In this case, funcWatch logs only the first item of an array.

In a test, the given function can be called multiple times. FuncWatch logs a call id for every parameter/return value in order to distinguish the values collected from different calls.

B. ioSelect: Choosing I/O examples

The input of ioSelect is a set of I/O values of an API function. Each value is labeled with a call id logged by funcWatch. IoSelect selects a call id randomly, and all the values with the same call id become the I/O example for the function. Because ioSelect selects I/O examples randomly, for the same API function and the same test, Docio may create different I/O examples. We implemented ioSelect as a standalone program so that we can easily add new algorithms for selecting I/O examples in the future.

C. ioPresent: Adding I/O examples to documents

The input of ioPresent is a set of I/O examples. Each API function has only one I/O example. The output of ioPresent is html API documents with the I/O examples.

The main functionality of ioPresent is to visualize an I/O example into a html table like the table in Figure 1. For example, in Figure 1, ioPresent automatically detects `bp->str` is a field of a struct variable that `bp` points to. Then, ioPresent automatically indents the row of `bp->str` and makes the row of `bp` collapsible (so that all the values that `bp` points to can be hidden if needed).

We implemented ioPresent as a patch of doxygen. Specifically, we modified two parts of doxygen: *Data Organizer* and *Html Generator*.

Data Organizer builds data structures to represent elements in the API libraries (that are going to be documented), such as classes and functions. First, we added a new class to represent a new type of element—the I/O examples. Second, we added I/O examples into the data structure that represents a function. We also modified *Html Generator*, in which we implemented the process of visualizing I/O examples in html format.

VII. PRELIMINARY EVALUATION

To assess the ability of Docio to generate I/O examples for real-world libraries, we did a preliminary evaluation to answer

TABLE I
THE NUMBER OF FUNCTIONS THAT
HAVE I/O EXAMPLES GENERATED BY DOCIO

API library	# of functions in source code	# of functions in api doc.	# of functions with I/O examples
ffmpeg	20,347	625	191
libssh	1,473	519	202
protobuf-c	3,171	18	9

the question: how many functions will have I/O examples by using Docio?

In this evaluation, we used Docio on three API libraries: ffmpeg, libssh, and protobuf-c. For each library, we executed funcWatch on every function (found by doxygen) and every test in the test suites. Then, we executed ioSelect and ioPresent to generate API documents with I/O examples.

In total, Docio generates I/O examples for 402 functions, 191 functions in ffmpeg, 202 in libssh, and 9 in protobuf-c. To put the numbers in context, Table I shows the number of functions in source code, the number of functions in API documents, along with the number of functions that have I/O examples. The API documents have much fewer functions than source code, because doxygen includes only the functions that are commented in a predefined format. This default setting helps doxygen exclude the internal functions, which should not be seen by API users.

Docio did not generate I/O examples for all the functions that are in API documents, because not all the functions were executed in the test suites. For example, ffmpeg has more than twenty thousand functions defined in source code. Only about six hundred functions are documented. Docio created I/O examples for 30% of the documented functions.

VIII. VISUALIZING ALL I/O VALUES

Sometimes, it is useful to see all the possible values of the parameters and return variables of an API function. So we built a prototype to visualize all the values logged by funcWatch for an API function. For each parameter in a function, we draw a bar chart of all the possible values of the parameter, where the height of each bar represents the frequency of that value. For example, in Figure 3, function `av_gcd` has three bar charts. Each bar chart represents a parameter or a return variable. Parameter `a` has five different values: 0, 1, 3, 4, 25 (logged by funcWatch). The most frequent value for `a` is 0, which occurs in more than 140 calls.

The bar charts are interactive to show the relationships between the parameters and return variables. For example, in Figure 3, the cursor is hovered over bar 25 in bar chart `a`. The other two bar charts are updated to reflect the values collected when `a` was 25. In other words, Figure 3 shows that there are about 30 calls where `a` is 25. In these calls, `b` is either 1 or 25. And all these calls return 1.

IX. RELATED WORK

In this section, we will present two categories of related work. **The first category** is about introducing different types

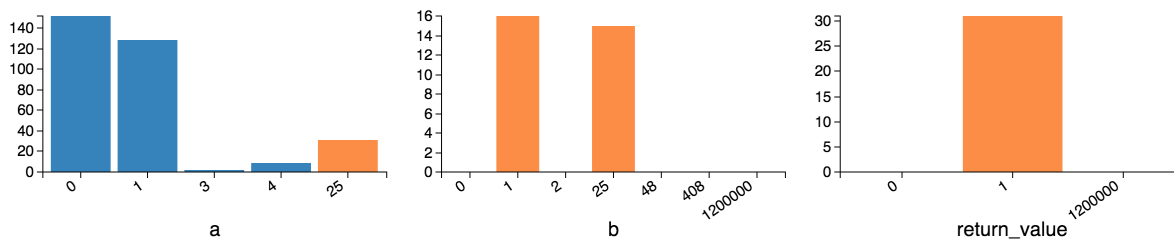


Fig. 3: All the logged I/O values for function `av_gcd`. Three bar charts correspond to three variables: `a`, `b`, and `return`. In this figure, the user has moused over the possible value “25” in `a`, and the corresponding values for `b` and `return` are updated and highlighted.

of contents into API documents, such as information from StackOverflow posts [13], usage information [14], frequently-asked-questions from the Web [15], and code example summaries mined from the Web [16], [17]. Different from these approaches, Docio does not mine the Web, and Docio adds I/O values instead of source code examples into API documents.

The second category is about collecting and analyzing dynamic information, which can be useful in API documentation. For example, *Daikon* analyzes dynamic information to report likely program invariants [18]. The reports can be useful in API documents. We plan to embed this type of reports to API documents in the future in order to study whether and how the reports can help programmers use API functions.

X. FUTURE WORK AND LIMITATIONS

Docio is a prototype toolset, which serves as our first step towards investigating what information is useful for API users. Previous studies have shown that programmers ask questions about dynamic information [1], [2], such as I/O examples. What is not shown is whether and how this information helps programmers. First, we will conduct an empirical study to assess the usefulness of I/O examples. In this study, we plan to hire programmers to complete certain program tasks with and without I/O examples in documents. By comparing the programmers’ performance with and without I/O examples, we can see how I/O examples can help programmers.

Furthermore, we plan to design and implement different selection strategies for `ioSelect`. Our assumption is that some I/O examples are more useful than others. We plan to investigate some potential characteristics of useful I/O examples. For example, the most common I/O values may be more useful.

Additionally, the quality and the quantity of the I/O examples generated by Docio depend on the test suites of the API libraries. In the future work, we may consider code examples, tutorial examples, and real-world applications that use the API libraries in addition to the test suites to generate I/O examples.

Implementation-wise, Docio currently has three major limitations. First, `funcWatch` does not support 64-bit memory addressing so Docio can work only in 32-bit operating systems. Second, `funcWatch` collects only the values for statically-linked libraries. Third, `funcWatch` collects all the I/O values no matter whether a test passes or fails. If a test fails, the logged values may be invalid. The current workaround for this limitation is to detect the outputs of the tests in a script and discard the outputs of `funcWatch` for those failed tests.

XI. CONCLUSION

To conclude, we built and demonstrated a prototype toolset, Docio. The goal of Docio is to help API developers create and add I/O examples into API documents. Docio has three programs: `funcWatch`, `ioSelect`, and `ioPresent`. Our preliminary evaluation shows that Docio can generate four hundred I/O examples for three real-world C libraries.

ACKNOWLEDGMENT

We thank Douglas Smith for his work in visualization of all I/O values. This work was partially supported by the NSF CCF-1452959 and CNS-1510329 grants, and the Office of Naval Research grant N000141410037. Any opinions, findings, and conclusions expressed herein are the authors’ and do not necessarily reflect those of the sponsors.

REFERENCES

- [1] E. Duala-Ekoko and M. P. Robillard, “Asking and answering questions about unfamiliar apis: An exploratory study,” ser. ICSE ’12, 2012.
- [2] J. Sillito, G. C. Murphy, and K. D. Volder, “Asking and answering questions during a programming change task,” *IEEE Trans. Softw. Eng.*, vol. 34, no. 4, July 2008.
- [3] A. Forward and T. C. Lethbridge, “The relevance of software documentation, tools and technologies: A survey,” ser. DocEng ’02, 2002.
- [4] M. P. Robillard and R. DeLine, “A field study of api learning obstacles,” *Empirical Software Eng.*, vol. 16, no. 6, pp. 703–732, 2010.
- [5] W. Maalej and M. P. Robillard, “Patterns of knowledge in api reference documentation,” *IEEE Trans. Softw. Eng.*, vol. 39, no. 9, Sept. 2013.
- [6] (2015, Dec.) Doxygen. [Online]. Available: <http://www.doxygen.org>
- [7] (2016, May) Javadoc. [Online]. Available: <http://www.oracle.com/>
- [8] A. J. Ko, B. A. Myers, M. J. Coblenz, and H. H. Aung, “An exploratory study of how developers seek, relate, and collect relevant information during software maintenance tasks,” *IEEE TSE*, vol. 32, no. 12, 2006.
- [9] J. Lawrance, C. Bogart, M. Burnett, R. Bellamy, K. Rector, and S. D. Fleming, “How programmers debug, revisited: An information foraging theory perspective,” *IEEE Trans. Softw. Eng.*, vol. 39, no. 2, 2013.
- [10] T. Roehm, R. Tiarks, R. Koschke, and W. Maalej, “How do professional developers comprehend software?” ser. ICSE ’12, 2012.
- [11] A. Armaly and C. McMillan, “Pragmatic source code reuse via execution record and replay,” *Journal of Software: Evolution and Process*, 2016.
- [12] (2016, May) libdwarf. [Online]. Available: <https://www.prevanders.net>
- [13] C. Treude and M. P. Robillard, “Augmenting api documentation with insights from stack overflow,” ser. ICSE ’16, 2016.
- [14] J. Stylos, B. A. Myers, and Z. Yang, “Jadeite: Improving api documentation using usage information,” ser. CHI EA ’09, 2009.
- [15] C. Chen and K. Zhang, “Who asked what: Integrating crowdsourced faqs into api documentation,” ser. ICSE Companion 2014, 2014.
- [16] J. Kim, S. Lee, S.-W. Hwang, and S. Kim, “Enriching documents with examples: A corpus mining approach,” *ACM TOIS*, vol. 31, no. 1, 2013.
- [17] S. Subramanian, L. Inozemtseva, and R. Holmes, “Live api documentation,” ser. ICSE 2014, 2014.
- [18] M. D. Ernst, J. H. Perkins, P. J. Guo, S. McCamant, C. Pacheco, M. S. Tschantz, and C. Xiao, “The daikon system for dynamic detection of likely invariants,” *Science of Computer Programming*, vol. 69, 2007.