

Teddy: Automatic Recommendation of Pythonic Idiom Usage For Pull-Based Software Projects

Purit Phan-udom*, Naruedon Wattanakul*, Tattiya Sakulniwat*,
Chaiyong Ragkhitwetsagul*, Thanwadee Sunetnanta*, Morakot Choetkiertikul*, Raula Gaikovina Kula†

*Faculty of Information and Communication Technology (ICT), Mahidol University

†Nara Institute of Science and Technology (NAIST)

{purit.pha, naruedonw, tattiya.sakul}@gmail.com, {chaiyong.rag, thanwadee.sun, morakot.cho}@mahidol.ac.th,
raula-k@is.naist.jp

Abstract—Pythonic code is idiomatic code that follows guiding principles and practices within the Python community. Offering performance and readability benefits, Pythonic code is claimed to be widely adopted by experienced Python developers, but can be a learning curve to novice programmers. To aid with Pythonic learning, we create an automated tool, called Teddy, that can help checking the Pythonic idiom usage. The tool offers a *prevention mode* with Just-In-Time analysis to recommend the use of Pythonic idiom during code review and a *detection mode* with historical analysis to run a thorough scan of idiomatic and non-idiomatic code. In this paper, we first describe our tool and an evaluation of its performance. Furthermore, we present a case study that demonstrates how to use Teddy in a real-life scenario on an Open Source project. An evaluation shows that Teddy has high precision for detecting Pythonic idiom and non-Pythonic code. Using interactive visualizations, we demonstrate how novice programmers can navigate and identify Pythonic idiom and non-Pythonic code in their projects. Our video demo with the full interactive visualizations is available at <https://youtu.be/vOCQReSvBxA>.

Index Terms—Pythonic Idioms, Code Review, Program Analysis

I. INTRODUCTION

According to the TIOBE index¹ and the GitHub 2019 annual report², Python is in the top three of the most popular programming languages. Python emphasizes being elegant and easy to read³. Similar to any other programming language, Python code also has an idiomatic “Pythonic” way of writing code to solve a particular problem.

In the age of social coding, such as Open Source GitHub projects, programmers find themselves in situations where they would start contributing to a project that involves Python programming. Learning and using Pythonic idioms can be a challenging task for novices who join the project [1]. There exists literature that provides examples on Pythonic idioms [2], however, most of the learning seems to be from the communities through experience and usage examples available on forums such as Stack Overflow [1]. Compared to non-Pythonic code, Pythonic idioms are well-accepted form of Python programming [3] due to many benefits. For example, Sakulniwat et al. [4] showed that adopting the `with open`

Pythonic idiom gives benefits not only in terms of readability, but also prevents memory leaks by automatically closing the file after finishing the task.

In this paper, we present *Teddy*, an automated tool for recommending Pythonic idiom usage for pull-based development software projects. Based on two common types of code contributions [5], Teddy offers two modes to recommend opportunities for novices to make their code more idiomatic. The first is *prevention mode* with Just-In-Time (JIT) analysis. In this mode, Teddy searches for development activities that are related to incorporation of new features and implementation (i.e., forward engineering). We identify non-Pythonic usage during the code review process by analyzing any submitted patches in pull request’s commits. The tool then recommends a Pythonic idiom counterpart of doing the same task to the developers. The second is *detection mode* with historical analysis. Teddy analyses the whole commit history of a given project and looks for both occurrences of Pythonic idiom (Py) and non-Pythonic (NPy) code snippets during software maintenance (i.e., re-engineering, corrective and management).

Teddy produces an interactive visualization to show the usage of Pythonic idioms over time based on the detected Py and NPy code in all the commits. To demonstrate our tool, we present a case study to demonstrate Teddy in a real-life scenario. An evaluation shows that Teddy has high precision for detecting idiomatic and non-idiomatic Python code. Using visualizations, we demonstrate how novice programmers navigate and identify Py or NPy code.

II. ARCHITECTURAL DESIGN

Teddy’s architecture has two aspects: Just-In-Time Pythonic idiom analysis (prevention mode) as shown in Fig. 1 and Pythonic idiom evolving in the history (detection mode) in Fig. 2. The prevention mode aims to give early feedback to code reviewers by identifying Py and NPy code snippets from GitHub pull requests and the results are delivered as pull request comments. The detection mode provides the Pythonic idiom visualization from the past commits which software practitioners can oversee the Pythonic adoption.

Our design is divided into three main components: the Pythonic idiom and non-Pythonic code database, the idiom matching, and the idiom usage interactive visualization.

¹<https://www.tiobe.com/tiobe-index/>

²<https://octoverse.github.com/>

³Zen of Python at <https://www.python.org/dev/peps/pep-0020/>

TABLE I
EXAMPLES OF THE PYTHONIC IDIOM (PY) AND NON-PYTHONIC (NPY)
CODE STUDIED IN THIS PROJECT

Type	Name	Description
Py	enumerate	for-loop iteration using <code>enumerate</code> function
Py	file reading statement	Using <code>with open()</code> as <code>...</code> to open a file
NPy	variable swapping	Using a temporary variable to swap two variables' values
NPy	code formatting	Using <code>';</code> to put more than one statement in a single line

1. *The Py/NPy database.* Table I shows four examples of Py and NPy cases taken from our built database. In total, our database contains 113 code snippets of which 55 are Pys and 58 are NPy. The complete list of our labelled types, with their actual code snippets is available from our website at <https://muict-seru.github.io/icsme20-teddy-tooldemo>.

Our database is based on Python programming community's best practices [2, 6]. We identified and labelled ten different types of Pythonic idioms to be included in our database. The ten types are based on their functionalities and include *dictionary comprehension*, *enumerate*, *file reading statement*, *list comprehension*, *if statement*, *string formatting*, *set*, *tuple*, *variable swapping*, and *code formatting*, where each of them has both Py and NPy code snippets in our database. We thus have 20 original code snippets (10 Pys and 10 NPy). We then modify these original code snippets to create a variety of patterns, i.e., data augmentation, in each type, such as renaming identifiers and changing data types.

2. *Idiom matching based on clone search.* The Py/NPy database serves as a corpus in the idiom matching. We employ a code clone search technique, called Siamese [7], to determine whether code snippets from a software project similar to Py or NPy code snippets in our database. Siamese allows flexible matching at different levels of code abstraction using its multiple code representation technique. Thus, it offers a capability of detecting Pys and NPy with several modifications (i.e., Type-2 and Type-3 clones).

3. *Idiom usage interactive visualization.* To manage large software projects, we employ an interactive visualization to show the identified instances of Pythonic idiom and Non-Pythonic usage. Our tool allows a user to pan and zoom over many files to inspect Py and NPy usage and their evolution over commits in a project.

III. IMPLEMENTATION

Prevention mode. In this mode, Teddy is implemented as an automated bot during GitHub pull requests (Figure 1). The non-Pythonic code snippets in the Py/NPy database are stored in Siamese's code search index (using Elasticsearch). Then, Teddy connects with GitHub through Sme.io⁴, a webhook service, and uses the code changes in a pull request's commits as Siamese's search queries. The queries are matched with the NPy in the code search index using the idiom matching

⁴<https://smee.io>

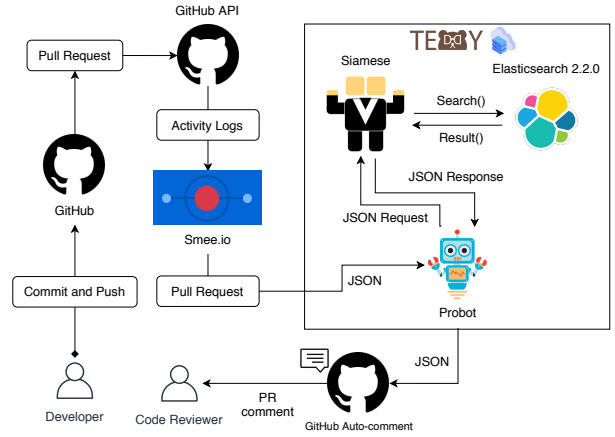


Fig. 1. Architecture of the prevention mode (JIT analysis)

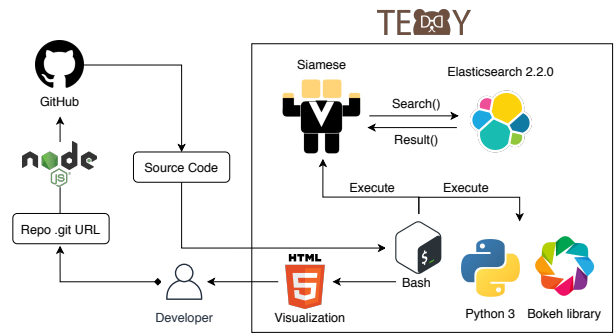


Fig. 2. Architecture of the detection mode (historical analysis)

component. A query that is a clone to any NPy is returned with a recommendation of its Py counterpart from the Py/NPy database. The overall results are formatted into a friendly message (i.e., with an explanation of a recommended Py snippet to replace the NPy one). Teddy then performs an auto-comment in the pull request by the Probot⁵ service.

Detection mode. In this mode, Teddy first obtains the source code from a given GitHub URL by cloning the repository (Figure 2). Then, the tool iterates through the commit versions of the cloned repository, from the first to the last commit. Siamese creates a code search index from the project source code in each commit, and injects the Py/NPy code snippets from the Py/NPy database as search queries. The results are combined and fed into the visualization component, using the Bokeh⁶ library, where a scatter-plot graph is generated.

IV. EVALUATION

Evaluation Metrics. We prepared the Siamese search index with a ground-truth dataset containing Python code snippets that are labelled as Py, NPy, and normal code. We performed the idiom matching by sending 113 queries in our Py/NPy database to Siamese and checked how many Py/NPy code snippets were correctly retrieved. From the results of 113

⁵<https://github.com/probot/probot>

⁶<https://github.com/bokeh/bokeh>

queries, we computed mean average precision (MAP), mean reciprocal rank (MRR), query recall (QR) and overall recall (OR). The calculation of MAP and MRR follow the definitions in information retrieval [8]. For recall, we defined two types of recall, QR and OR, as follows:

Query Recall (QR) determines how complete is the retrieved items with respect to the subset of queries with *non-empty results* (i.e., returned queries). QR for a set of returned queries r is defined as: $QR = \frac{1}{|r|} \sum_{i=1}^{|r|} \frac{|R_i|}{|A_i|}$ where R_i is the set of retrieved relevant items and A_i the set of all relevant items for the i -th returned query, respectively.

Overall Recall (OR) evaluates how complete is the number of relevant items retrieved with respect to the entire query set. In contrast to QR, OR also considers the queries with empty results returned. OR for a query set R is defined as $OR = \frac{1}{|R|} \sum_{i=1}^{|R|} \frac{|R_i|}{|A_i|}$ where R_i is the set of retrieved relevant items and A_i the set of all relevant items for the i -th query.

Ground-Truth. We collected Python source code files, containing Py and NPy snippets inside, from external sources, mainly from three GitHub repositories of Flask, Tensorflow TFX, Ipyparallel, and programming websites⁷. First, we searched for Py and NPy code snippets within the three GitHub repositories by using CCGrep [9], a code clone detector that can detect clones by giving regular expressions. Table II shows the three groups of Python code files. The first group, the normal code, are Python code snippets that do not contain Py nor NPy code snippets inside. The second group, the Py group, are files that contain a Pythonic idiom. The third group, the NPy group, are files that contain a non-Pythonic code snippet.

TABLE II
THE GROUND-TRUTH DATASET

File group	Description	Files
Normal	Python code without Py or NPy	30
Py	Pythonic idiom code snippets	20
NPy	Non-Pythonic code snippets	20
Total		70

Accuracy Results. Table III presents four configurations C_1 , C_2 , C_3 , and C_4 where we varied two of Siamese’s parameters including the similarity measure (n-gram token ratio and fuzzywuzzy) and the similarity thresholds (T_0 , T_1 , T_2 , T_3) of the four code representations in Siamese. We found that different similarity measures offer different performance on the idiom matching. In general, idiom matching using n-gram token ratio gains higher MAP and MRR values comparing to those of fuzzywuzzy. The similarity threshold also plays an important role in the accuracy of idiom matching. Choosing low similarity thresholds (i.e., C_1 and C_2) resulted in high recall but low MAP, which mean a large number of false positives were retrieved. In contrast, choosing higher similarity

⁷geeksforgeeks.org, programiz.com, tutorialpoint.com, journaldev.com, beginnersbook.com, stackoverflow.com, and note.nkmk.me

TABLE III
CONFIGURATIONS OF SIAMESE EVALUATED IN THE EXPERIMENT

	Sim.	Multi-representation similarity threshold				Error Measures			
		T_0	T_1	T_2	T_3	MAP	QR	OR	MRR
C_1	FWZ	0	0	0	0	0.32	1.00	1.00	0.44
C_2	NTR	0	0	0	0	0.35	0.89	0.88	0.45
C_3	FWZ	40	40	40	40	0.38	0.47	0.46	0.41
C_4	NTR	40	40	40	40	0.89	0.50	0.04	0.83

NTR = N-gram Token Ratio, FWZ = Fuzzywuzzy

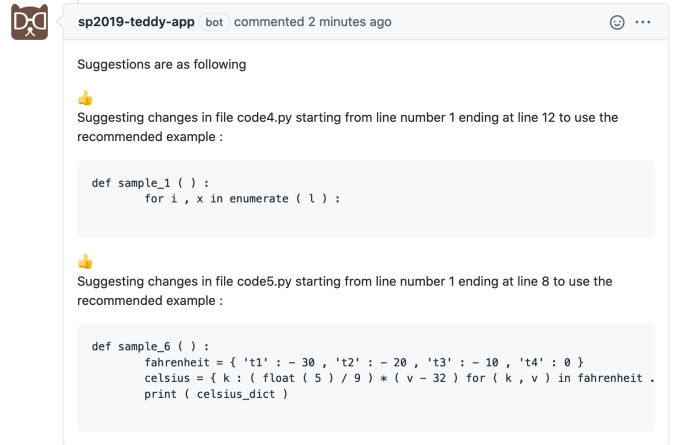


Fig. 3. Result of prevention mode: Recommending Pythonic idioms in a pull request (available at <https://github.com/MUICT-SERU/flask/pull/2>).

thresholds (i.e., C_3 , and C_4) gave high MAP, but we needed to sacrifice the recall.

As a result, we found that the configuration C_4 gives the best idiom matching performance by offering MAP of 0.89 and MRR of 0.83. This configuration, however, has a relatively low recall. The QR value of 0.5 shows that, for all the queries that returned results, we retrieved half of the Py and NPy code snippets. The low OR value of 0.04 came from a low number of queries that returned the results. The modified variants that we added into the Py/NPy database made them matched with only very specific patterns of Py and NPy not available in the ground-truth dataset. In the end, we prefer to get high precision to avoid presenting false positives to users.

Case Study on the Flask Project. We performed a case study on Flask, a popular web application framework written in Python, with 50.9k stars, 13.6k forks, and more than 3.9k commits on GitHub.

In the prevention mode, we needed to simulate the situation of creating a pull request. This could not be done on the actual Flask project due to the permission to install the Teddy bot and to create a pull request. Thus, we forked the Flask project into our GitHub account first and installed the Teddy bot on the repository. Then we created a new pull request with 5 commits. Each commit contained one Python file that performs a specific task. There were two files that contained the non-Pythonic code that could be replaced with the Pythonic

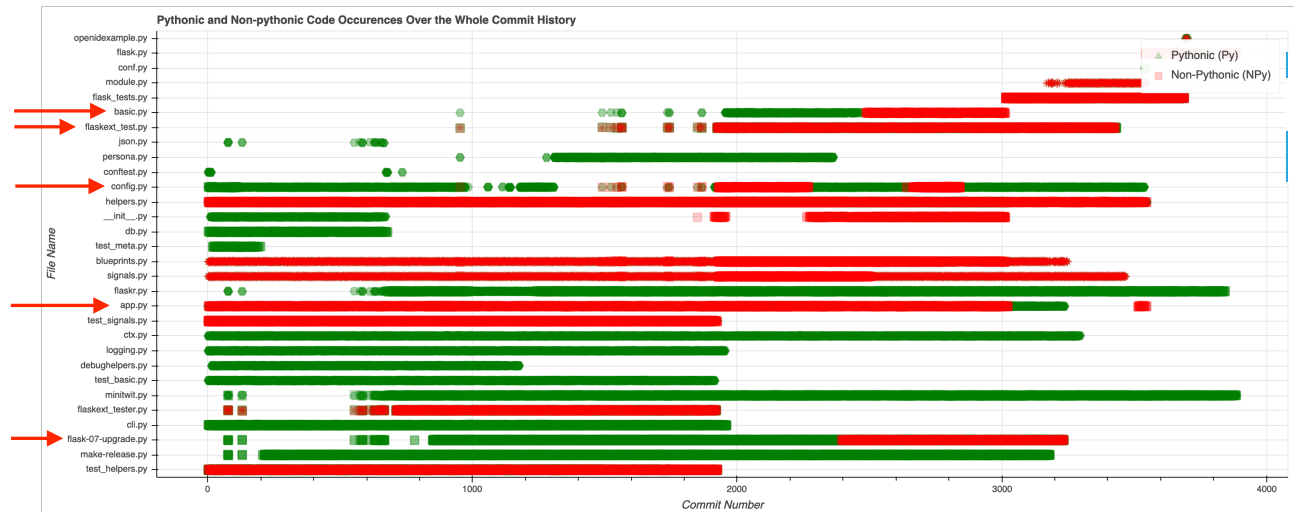


Fig. 4. Result of detection mode: An interactive visualization for Py and NPy usage over the whole commit history. A full interactive sample of the visualization, which includes the pan and zoom functions, is available on the study website at <https://muict-seru.github.io/icsme20-teddy-tooldemo/flask.html>.

idiom of *enumeration* and *dictionary comprehension*. In the detection mode, we provided the GitHub repository URL of the Flask project to Teddy. Then, the Teddy tool cloned the source code of the project into its system and performed the idiom matching and visualization.

Teddy Outputs. Figure 3 shows that Teddy bot detected the two non-Pythonic code snippets in the pull request and gave the recommendations for their Pythonic idiom usages. We can see that Teddy suggested a Py example (i.e., `sample_1` snippet) based on the code snippet in the Py/NPy database to replace the detected NPy in `code4.py`.

Figure 4 shows a visualization of the occurrences of Pythonic idioms and non-Pythonic code over time. The x-axis displays the number of commits while the y-axis displays the files containing the detected Py and NPy. Each dot represents one occurrence of Py and NPy in the project, which can be hovered on to show details of the type of Py and NPy and the line numbers (start, end) that it appears in the file.

From our visualization, a user is able to see changes of the Pythonic idioms over time. For instance, the figure shows that about half of the files detected by Teddy in the project contain non-Pythonic code (red lines) that never be replaced by Pythonic idiom. In contrast, the other half of the detected files contain Pythonic idioms (green lines) that also never change to non-Pythonic code. There are a few files that show a switch from non-Pythonic code to Pythonic idioms (i.e., changing from red to green) such as `app.py` and `flaskext_test.py`. The other direction (Py to NPy or from green to red lines) can also be observed in `basic.py` and `flask-07-upgrade.py`. Lastly, there is a file, `config.py`, that show multiple changes between Py and NPy. These files are pointed to by an arrow in the figure.

V. CONCLUSION

Teddy promotes the usage of Pythonic idioms by giving recommendations to developers. Our evaluation shows that

Teddy accurately detects Pythonic idioms and non-Pythonic code with MAP of 0.89 and MRR of 0.83. We demonstrate the tool’s usefulness by a case study using Flask GitHub project.

ACKNOWLEDGEMENT

This work is supported by Japanese Society for the Promotion of Science (JSPS) KAKENHI Grant Numbers 18H04094 and 20K19774.

REFERENCES

- [1] C. Alexandru, J. Merchante, S. Panichella, S. Proksch, H. C. Gall, and G. Robles, “On the usage of pythonic idioms,” in *Onward! 2018*, 2018, pp. 1–11.
- [2] J. Knupp, *Writing Idiomatic Python 3.3*, 2013.
- [3] P. S. Foundation, “Glossary - python 3.8.0 documentation,” <https://docs.python.org/3/glossary.html>, 2020, online; accessed 15-June-2020.
- [4] T. Sakulniwat, R. Kula, C. Ragkhitwetsagul, M. Choetkier-tikul, T. Sunetnanta, D. Wang, T. Ishio, and K. Matsumoto, “Visualizing the Usage of Pythonic Idioms over Time: A Case Study of the with open Idiom,” *IWESEP '19*, pp. 43–48, 2019.
- [5] L. P. Hattori and M. Lanza, “On the nature of commits,” in *ASE '08*, 2008, p. III–63–III–71.
- [6] “Pythonic,” <https://pythonic-examples.github.io>, 2020, online; accessed 15-June-2020.
- [7] C. Ragkhitwetsagul and J. Krinke, “Siamese: scalable and incremental code clone search via multiple code representations,” *Empirical Software Engineering*, pp. 1–49, Mar 2019.
- [8] C. Manning, P. Raghavan, and H. Schütze, *An Introduction to Information Retrieval*. Cambridge University Press, 2009, vol. 21.
- [9] K. Inoue, Y. Miyamoto, D. M. German, and T. Ishio, “Code clone matching: A practical and effective approach to find code snippets,” 2020.