Reproducibility Companion Paper: Describing Subjective Experiment Consistency by p -Value P-P Plot

Jakub Nawala
Lucjan Janowski
jnawala@agh.edu.pl
AGH University of Science and
Technology, Institute of
Telecommunications
Kraków, Poland

Bogdan Ćmiel
AGH University of Science and
Technology, Department of
Mathematical Analysis,
Computational Mathematics and
Probability Methods
Kraków, Poland

Krzysztof Rusek
AGH University of Science and
Technology, Institute of
Telecommunications
Kraków, Poland

Marc A. Kastner
National Institute of Informatics
Tokyo, Japan

Jan Zahálka
Czech Technical University in Prague
Prague, Czech Republic

ABSTRACT

In this paper we reproduce experimental results presented in our earlier work titled “Describing Subjective Experiment Consistency by p -Value P-P Plot” that was presented in the course of the 28th ACM International Conference on Multimedia. The paper aims at verifying the soundness of our prior results and helping others understand our software framework. We present artifacts that help reproduce tables, figures and all the data derived from raw subjective responses that were included in our earlier work. Using the artifacts we show that our results are reproducible. We invite everyone to use our software framework for subjective responses analyses going beyond reproducibility efforts.

CCS CONCEPTS

• **Human-centered computing** → **User models; User studies;** •
Social and professional topics → *Quality assurance.*

KEYWORDS

Quality of Experience; Subjective Experiment; Consistency; Reproducibility; P-P Plot; Subjective Data

ACM Reference Format:

Jakub Nawala, Lucjan Janowski, Bogdan Ćmiel, Krzysztof Rusek, Marc A. Kastner, and Jan Zahálka. 2021. Reproducibility Companion Paper: Describing Subjective Experiment Consistency by p -Value P-P Plot. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21)*, October 20–24, 2021, Virtual Event, China. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3474085.3477935>

1 ARTIFACTS ORGANISATION

The artifacts are available for download from the following GitHub repository: <https://github.com/Qub3k/subjective-exp-consistency-check> [2]. Its file structure is presented in Fig. 1.

This work is licensed under a Creative Commons Attribution International 4.0 License.

MM '21, October 20–24, 2021, Virtual Event, China.
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8651-7/21/10.
<https://doi.org/10.1145/3474085.3477935>

The `reproduce.py` script is the entry point, serving as a master script governing the reproducibility process. The script’s help message is designed so as to provide sufficient information necessary to understand framework’s operation. The following snippet shows how to invoke the `reproduce.py`’s help message using the command-line interface (CLI).

```
$ python3 reproduce.py -h
```

Should the help message be insufficient, the reader is encouraged to take a look at the `README.md` file. More specifically, its “Reproducibility” section provides further guidance on how to use the framework. At last, since the framework is entirely open-sourced, its operation can be investigated by looking at the source code.

```
.
├── .gitignore
├── G_test_on_real_data.py
├── G_test_results.csv
├── LICENSE
├── README.md
├── _logger.py
├── bootstrap.py
├── figures/
├── friendly_gsd.py
├── gsd.py
├── gsd_prob_grid.pkl
├── hdtv1_exp1_scores_pp_plot_ready.csv
├── p_value_pp_plot_in-depth_tutorial.ipynb
├── probability_grid_estimation.py
├── qnormal.py
├── qnormal_prob_grid.pkl
├── reproduce.py
├── reproducibility/
├── requirements.txt
├── subjective_quality_datasets.csv
└── typical_vs_atypical.csv
```

Figure 1: Structure of the GitHub repository.

Another two important files in the repo are: (i) `subjective_quality_datasets.csv` and (ii) `G_test_results.csv`. The former one includes raw subjective data that is processed in the original paper [1]. The most important output of this processing is the `G_test_results.csv` file. It includes results of running computationally intensive bootstrapped version of the G-test of goodness-of-fit (cf. Fig. 2 from [1]). Effectively, recreating these results is the most significant part of the reproducibility efforts.

2 SETUP AND EXECUTION

We recommend to create a separate Python virtual environment and install there all the dependencies listed in the `requirements.txt` file. Importantly, Python version 3.7 or newer is required. When the required packages are installed the `reproduce.py` script can be run through the CLI, as shown below.

```
$ python3 reproduce.py [-h] [-n N] {scenario}
```

The `{scenario}` place-holder identifies the execution scenario. As of the time of writing this paper there are five such scenarios. They are identified by subsequent integers from the range 1–5. The following list provides details on each scenario.

- (1) Reproduce the original experiments using existing G-test results (i.e., the `G_test_results.csv` file). This executes immediately, but in principle does not reproduce the most important piece of the data presented in the original paper.
- (2) Reproduce only these G-test results that are necessary to draw Fig. 3 from the original paper. This already takes a significant amount of time to run (approximately nine days according to our internal tests). This scenario makes sense since Fig. 3 is a central part of the discussion presented in the original paper.
- (3) Reproduce all G-test results. This scenario reproduces from the scratch all the results presented in the original paper. As such this is the most important scenario of the reproducibility framework. However, according to our internal tests, this takes around 21 days to run.
- (4) Run the G-test for randomly selected N stimuli (cf. the `-n N` argument of the call to `reproduce.py` above). This scenario can be used to quickly check the correctness of N randomly selected results from the `G_test_results.csv` file. We note that according to our internal tests processing one stimulus takes about four to seven minutes.
- (5) Reproduce the probability grids of the Generalised Score Distribution (GSD) and Quantized Normal (QNormal) models (cf. Sec. GSD Parameters Estimation in [1]). Those grids are internally used when running the G-test.

Since our framework is implemented in Python in its entirety, it can be run on any platform.¹ All execution times mentioned in this paper were measured using the following hardware setup: Intel Core i3-8130U CPU, 16 GB of 2400 MHz RAM and 256 GB SSD disk (Lenovo LENSE30256GMSP34MEAT3TA).

¹We confirmed framework's operation on the three popular operating systems: Windows 10, Mac OS 10.15 and Ubuntu Linux 18.04.05.

2.1 Batch Processing Capability

Reproducing complete G-test results takes a significant amount of time when done on a single machine. Thus, we make available the batch processing friendly variation of the G-test running framework. It can be used to run multiple parallel instances of the G-test, each running on a different chunk of the input data. Crucially, the `reproduce.py` script does not support batch processing, as this would greatly complicate its structure. Instead, the `friendly_gsd.py` script must be used. Still, both `reproduce.py` and `friendly_gsd.py` scripts will produce the same output. For more information we refer the reader to the “Batch Processing” section of the `README.md` file.

3 EXPERIMENTS

Execution scenarios 1, 2 and 3 (cf. Sec 2) reproduce all tables and figures presented in the original paper. However, only scenario 3 reproduces entirely the data used to generate figures and tables. The other two scenarios either use the data processing outputs from the original paper (scenario 1) or reproduce only a part of the data (scenario 2).

Since the G-test running framework internally uses pre-calculated probability grid of the GSD model, to achieve the complete reproducibility (i.e., being able to achieve the same results when being provided only with raw subjective data contained in the `subjective_quality_datasets.csv` file) one has to run execution scenario 5 as well.

All in all, both scenario 3 and scenario 5 must be executed to check results reproducibility. The snippet below shows two calls to the `reproduce.py` script that fulfil this goal.

```
$ python3 reproduce.py 3
$ python3 reproduce.py 5
```

The first call produces four types of output: (i) CSV files with reproduced tables contents, (ii) PDF files with reproduced figures, (iii) a CSV file with G-test results and (iv) tables contents written to the standard output. Fig. 2 shows files that are created as a result of this call. Significantly, the reproduced figures have the same formatting as the one used in the original paper (cf. Fig. 3).

We also note here that the G-test used in the framework has randomness built into it. Thus, reproduced results will not exactly match the ones generated for the purposes of the original paper. This is because we use the bootstrapped version of the G-test that internally generates 10,000 synthetic random samples based on each observed sample. We refer readers interested into more details on the topic to section “In-depth Tutorial about Generating p -Value P–P Plots for Your Subjective Data” of the `README.md` file in the GitHub repository.

Running scenario 5 (the second call from the snippet above) creates two files only: (i) `reproduced_gsd_prob_grid.pkl` and (ii) `reproduced_qnormal_prob_grid.pkl`. They are pickled Python objects and more specifically, pickled Pandas DataFrames. They can be manually compared with the corresponding pickle files from the original paper: `gsd_prob_grid.pkl` and `qnormal_prob_grid.pkl`.

p-value_pp-plot_fig_one_a.pdf
 p-value_pp-plot_fig_one_b.pdf
 table_one_score_distribution.csv
 table_two_pvals.csv
 p-value_pp-plot_HDTV1_fig_three.pdf
 p-value_pp-plot_ITS4S2_fig_three.pdf
 p-value_pp-plot_ITS4S_AGH_fig_three.pdf
 p-value_pp-plot_AGH_NTIA_fig_three.pdf
 table_three_five_lowest_pvalue_res_its4s_agh.csv
 table_four_five_lowest_pvalue_res_its4s2.csv
 G_test_on_subjective_quality_datasets_chunk000_of_001.csv

Figure 2: Files generated as a result of running execution scenario 3.

Figure 3: Reproduced Fig. 1b from the original paper. Please note that the formatting of the figure is the same as in the original paper.

We note that our framework reproduces Fig. 1a and 1b from the original paper using ready-made CSV files.² This is because the two plots use synthetic data. Differently put, the plots were not generated from subjective responses gathered during any real-life subjective experiment.

²The CSV files are available in the reproducibility folder in the GitHub repository.

4 REPRODUCIBILITY EFFORTS

The code is open-source, well readable, and sufficiently commented. All results of the original paper are easily reproducible, directly generating the figures used in the original paper.

Since the submitted software was of high quality in the first version already, the reproducibility review has mostly consisted of minor fixes and ease-of-use improvements. Firstly, the review process resulted in the batch processing mode of the software being more accessible to the user. This is essential, as the sequential mode runs for days to reproduce all results. Secondly, the authors have fixed the random stimuli scenario that did not fail gracefully when the $-n$ parameter was omitted (now the default is 3 stimuli). Finally, the authors have been very responsive not only to the reviewer comments, but also to general GitHub user comments.

All of the above aspects lead us to believe this is a software worthy of the reproducibility badge.

5 INVITATION

Although this paper focuses on reproducibility, our GitHub repository [2] was created to help others use our framework in future analyses as well. We invite everyone, who has at hand a data set of subjective responses, to use the framework. It can test how well the GSD models subjective responses distribution and provide insights into subjective experiment consistency. For more details we refer the reader to our original paper [1] and the README.md file in the repository.

ACKNOWLEDGMENTS

This work was supported by the Polish Ministry of Science and Higher Education with the subvention funds of the Faculty of Computer Science, Electronics and Telecommunications of AGH University and by the PL-Grid Infrastructure. Furthermore, the research leading to these results has received funding from the Norwegian Financial Mechanism 2014–2021 under project 2019/34/H/ST6/00599.

We also kindly acknowledge the help of Franz Hahn, who identified and helped us resolve an issue with our source code.

REFERENCES

- [1] Jakub Nawala, Lucjan Janowski, Bogdan Ćmiel, and Krzysztof Rusek. 2020. Describing Subjective Experiment Consistency by p-Value P-P Plot. In *Proceedings of the 28th ACM International Conference on Multimedia*. ACM, New York, NY, USA, 852–861. <https://doi.org/10.1145/3394171.3413749>
- [2] Jakub Nawala, Lucjan Janowski, Bogdan Ćmiel, and Krzysztof Rusek. 2021. *subjective-exp-consistency-check GitHub Repository*. <https://github.com/Qub3k/subjective-exp-consistency-check>