USENIX

THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

Characterizing and Detecting Propaganda-Spreading Accounts on Telegram

Klim Kireev, *EPFL, MPI-SP Max Plank Institute for Security and Privacy;* Yevhen Mykhno, *unaffiliated;* Carmela Troncoso, *EPFL, MPI-SP Max Plank Institute for Security and Privacy;* Rebekah Overdorf, *Ruhr University Bochum* (*RUB*), *Research Center Trustworthy Data Science and Security in University Alliance Ruhr, University of Lausanne*

https://www.usenix.org/conference/usenixsecurity25/presentation/kireev

This artifact appendix is included in the Artifact Appendices to the Proceedings of the 34th USENIX Security Symposium and appends to the paper of the same name that appears in the Proceedings of the 34th USENIX Security Symposium.

August 13–15, 2025 • Seattle, WA, USA

978-1-939133-52-6

Open access to the Artifact Appendices to the Proceedings of the 34th USENIX Security Symposium is sponsored by USENIX.

USENIX Security '25 Artifact Appendix: Characterizing and Detecting Propaganda-Spreading Accounts on Telegram

Klim Kireev[§†], Yevhen Mykhno, Carmela Troncoso[§†], Rebekah Overdorf[‡⋆]

§ EPFL, † MPI-SP Max Plank Institute for Security and Privacy

‡ Ruhr University Bochum (RUB), Research Center Trustworthy Data Science and Security in University Alliance Ruhr

⋆ University of Lausanne

A Artifact Appendix

A.1 Abstract

The present submission contains the dataset used in the paper, embeddings computed on this dataset, model checkpoints required for the functionality analysis, and source code both for the inference of the attached models and for training new ones if such training is useful for future research. In addition to that, we also provide the source code in the form of the Jupiter notebook, which can be used to build most figures present in the paper.

A.2 Description & Requirements

A.2.1 Security, privacy, and ethical concerns

We do not see any risk for evaluators coming from this artifact evaluation procedure.

A.2.2 How to access

Our artefact can be downloaded from Zenodo *https://zenodo.org/records/14736756*

A.2.3 Hardware dependencies

Most of the code needed to run the machine learning models requires the CUDA-enabled GPU. In particular, to train the attached models, we used NVIDIA RTX 4070. However, less powerful NVIDIA GPUs (any NVIDIA GPU with CUDA support and >2GB of VRAM) can also be used in this evaluation. Besides that, we recommend using a machine with >=32GB of RAM available in order to minimize page swapping. Finally, approx. 15GB of free disk space is required.

A.2.4 Software dependencies

The evaluation requires GNU/Linux operating system (we used Ubuntu 22.04), with CUDA 12.2 installed (our nvidia driver version: 535.183.01). The installation of the Python environment used in the report is done via the conda package manager (we used Miniconda 4.13.0).

A.2.5 Benchmarks

The dataset used in the evaluation is part of the artifact.

A.3 Set-up

A.3.1 Installation

First, evaluators need to download the artifact from Zenodo and unzip all of them in a single folder (we recommend creating a separate folder); source code files should be in the same folder as the dataset CSV file.

After this step, the main Python dependencies must be installed via commands listed in the INSTRUCTIONS.md file (conda ... , pip ...).

A.3.2 Basic Test

The functionality test can be run via: `python model_eval.py` and the output should be:

```
The dataset is loaded
Overall Accuracy: 83.5
New Topics Accuracy: 86.842
Validation Accuracy: 74.0
```

The script execution time can be 1-10 min.

A.4 Evaluation workflow

A.4.1 Major Claims

Major claims made in the paper:

(C1): Propaganda accounts and user accounts demonstrate different properties in terms of message length, lifespan, number of active channels, and number of messages (Section 3.3, Figures 4-7). However, regarding the number of replies per message sent, user and propaganda accounts are similar (Figure 8), demonstrating that propaganda

accounts are efficient in invoking users' reactions in the form of replies. These properties are demonstrated in the experiment E1.

(C2): The Trigger-Propaganda embeddings detector outperforms other detection methods present in the paper with respect to performance on the new topics, clean accuracy, and performance on the validation propaganda network (Table 3, Section 4.3). It is proven by the experiment E2.

A.4.2 Experiments

(E1): *[Data Analysis] [30-60 human-minutes + 1 compute-hour + 10GB disk + 20-32 GB RAM]*: The aim of the data analysis is to generate the figures present in the paper.

Preparation: First, the jupyter notebook server must be launched via command: `jupyter notebook`, then select the file `figures.ipynb` in the opened window (click the URL in the output if the window in the browser does not open)

Execution: Run all the cells in the notebook using the button "Run all cells"

Results: Resulting figures should match figures present in the paper

(E2): *[Detection performance] [10 human-minutes + 0.5-1 compute-hour]*: This experiment shows that the provided models function and produce the results reported in Table 3.

Preparation: Open the terminal in the folder with the dataset and the source code.

Execution: Run `python models_eval.py`, followed by `python handcrafted_eval.py`

Results: Resulting numbers should match the numbers reported in Table 3.

A.5 Version

Based on the LaTeX template for Artifact Evaluation V20231005. Submission, reviewing and badging methodology followed for the evaluation of this artifact can be found at https://secartifacts.github.io/usenixsec2025/.