



USENIX Security '25 Artifact Appendix: Big Help or Big Brother?

Auditing Tracking, Profiling, and Personalization in Generative AI Assistants

Yash Vekaria
UC Davis

Aurelio Loris Canino
UNIRC

Jonathan Levitsky
UC Davis

Alex Ciechonski
UCL

Patricia Callejo
UC3M

Anna Maria Mandalari
UCL

Zubair Shafiq
UC Davis

A Artifact Appendix

A.1 Abstract

This artifact provides a framework for capturing and analyzing the network traffic generated while auditing Generative AI (GenAI) browser assistants. It is designed to assess the tracking, profiling, and personalization in Generative AI (GenAI) browser assistants. These assistants, implemented as browser extensions, offer functionalities such as search-based integration, querying, and page summarization. Leveraging powerful large language models (LLMs), they can access and process a wide range of user data in their browser. We use a MITM-based infrastructure to record network traffic in the form of '.flow' files and then a parser for further analysis of data collection and sharing practices of these assistants.

A.2 Description & Requirements

Our framework comprises capturing and analyzing the network traffic behaviour when using a GenAI extension. The artifact repository consists of the following components:

- **Flows/**: Directory contains the network traffic captured via Mitmproxy in the form of .flow files, generated while interacting with GenAI browser extensions. This directory is located in the current working directory and contains sub-directories labeled with extension name being tested (e.g., Merlin).
- **Output/**: Directory contains CSV files generated by processing the captured .flow files. Sample CSV files for different extensions generated during the experiments are available in the OSF repository.
- **parse_flows_to_csvs.py**: Python script that parses .flow files into summarized CSV files, extracting key details about relevant flows. This includes timestamp, request domain, payload, response type, cookies, domain's

parent organization and disconnect list based tracker's category (see Section A.2.5).

- **generate_entity_domain_mapping.py**: Python script that parses individual entity-domain mapping files for each domain from DuckDuckGo's tracker-radar repository and generates a consolidated ddg.json.
- **disconnect.json**: Disconnect's tracking protection list used to map domains to predefined tracker categories such as Advertising, Analytics, and Social.
- **ddg.json**: DuckDuckGo's entity-domain list maps the domains to corresponding parent organizations. ddg.json is an aggregated single file generated by parsing individual mapping files using generate_entity_domain_mapping.py.
- **requirements.txt**: Lists Python package dependencies required to run the analysis pipeline.
- **README.md**: Detailed instructions for setting up Mitmproxy, configuring it in Google Chrome, capturing network traffic, and executing the analysis pipeline.

A.2.1 Security, privacy, and ethical concerns

This artifact involves intercepting HTTPS traffic using a custom root Mitmproxy certificate. Some risk considerations are as follows:

- All HTTPS traffic from the configured browser instance may be logged in a decrypted form by Mitmproxy.
- Evaluators must be careful in choosing what type of online personal or private spaces they login to. This is because in presence of GenAI assistants, user's sensitive information can be potentially extracted and shared with assistant's server or LLM model's server.
- It is recommended to use a fresh Chrome profile and remove the Mitmproxy certificate after experiments.

For more information on ethical considerations, we encourage to read the corresponding section in the main paper.

A.2.2 How to access

The artifact can be accessed via Zenodo at the following link: <https://doi.org/10.5281/zenodo.15530229>. Additionally, code is also hosted on GitHub repository: <https://github.com/Yash-Vekaria/genai-assistants/>.

A.2.3 Hardware dependencies

No specialized hardware is required; the artifact can be run on standard consumer machines such as a laptop or a desktop.

A.2.4 Software dependencies

The artifact is compatible with any operating system such as Windows, Linux, and macOS. It requires Python 3.8+, mitmproxy for traffic interception, and relevant Python packages as listed in `requirements.txt`. Detailed setup instructions are provided in the README file.

A.2.5 Benchmarks

We use two external resources in our framework:

- **Disconnect List:** allows categorizing tracking domains into one of the 11 categories such as Advertising, Analytics, and Social. It is available at <https://github.com/disconnectme/disconnect-tracking-protection/blob/master/services.json>.
- **DuckDuckGo’s Entity-Domain Mapping:** allows mapping a domain to its parent organization or owner. It is openly available at <https://github.com/duckduckgo/tracker-radar/tree/main/domains/US>.

A.3 Set-up

The steps listed here allow setting up the necessary environment for evaluating the artifact.

A.3.1 Installation

- ❶ **MITM Proxy Installation:** Primarily, install mitmproxy following the official guide at (<https://mitmproxy.org/>).
- ❷ **Instantiate Chrome Profile:** Create a new Google Chrome profile and configure it to route its traffic through the mitmproxy server linked to (localhost:8080). First, visit <chrome://version/> in your Google Chrome browser to figure out path to Chrome executable and/or user data directory. Open Chrome via CLI by specifying a custom user directory.
- ❸ **Install MITM Certificate:** Now, in the instantiated browser, install the mitmproxy root certificate by visiting mitm.it in the configured browser and following the prompts to add it to the system’s trusted authorities. This step is required only once and ensures that HTTPS traffic can be intercepted.

- ❹ **Validation:** Once the browser has been configured with mitmproxy, validate it by launching the configured Chrome profile, and confirming that intercepted traffic is saved in a .flow file.

- ❺ **Install Python Dependencies:** For the analysis phase, set up a Python virtual environment and install the required Python dependencies.

- ❻ **Download Benchmarks:** For disconnect list, download services.json from here: <https://github.com/disconnectme/disconnect-tracking-protection/blob/master/services.json>, place the downloaded file inside the current working directory, and rename it to disconnect.json. For generating DuckDuckGo’s mapping, run generate_entity_domain_mapping.py and then place the generated ddg.json in the current working directory.

A.3.2 Basic Test

A basic functionality check can be performed to verify the artifact is operational. For the sake of this test, we will consider the browser extension Merlin. Execute the following steps:

First, start mitmweb in one terminal and instantiate a fresh or unused chrome profile in the second parallel terminal window. Once the Chrome browser opens, install the GenAI browser extension (in this case Merlin) from the Chrome Web Store. Next, login to the extension and pin it to your browser. Next, perform three interactions: *search*, *browse*, and *summarize*. For *search*, ask a question in Google search and let Merlin auto-generate the response in side space on the right. In case, it does not, but provides an option to generate, click on that option in the side space. Alternatively, extension icon can be clicked to ask a question. For *browse* scenario, visit any webpage and ask a question about the page to the Merlin by clicking on the extension icon. Lastly, for *summarize*, the same webpage can be summarized using one of the extension-provided options. Now close the browser window, and terminate the processes in both the terminal windows.

The successful capture of network traffic can be confirmed by verifying the creation of the .flow file in the working directory. The analysis script: parse_flows_to_csvs.py can be now run on the captured .flow file to generate a summary CSV of the network traffic in the Output folder. This can be used to verify different claims in our paper such as: “Merlin can be seen contacting google-analytics.com and sharing user’s query prompt to this domain in the payload.”

A.4 Evaluation workflow

A.4.1 Major Claims

- (C1): *We visit personal as well as private spaces online in presence of GenAI browser assistant and observe the assistant to collect varying granularity of data such as page’s textual content, partial content, form fields, and*

whole DOM depending on the extension as depicted in Table 2. This is demonstrated by experiment (E1) and discussed in Section 5.3 of the paper.

(C2): An audit of sharing of collected information with first- as well as third-parties is measured for different user-, chat-, and browser-specific identifiers as shown in Table 3 and Figure 3. Sider and Merlin share chat identifiers with google-analytics.com while TinaMind shares it with analytics.google.com. Merlin also shares user’s raw query prompts with Google Analytics. MaxAI and Harpa also send data to Mixpanel. This can be manually verified with the help of experiment (E2) and is discussed in Section 5.3 of the paper.

(C3): We show that some GenAI assistants profile and personalize more than others based on five user attributes: location, age, gender, income, and interests. More specifically, two extensions (Monica and Sider) profile and personalize in-context as well as out-of-context based on all five user attributes. In contrast, Perplexity and TinaMind showed no strong evidence of profiling or personalization, while Harpa exhibited only in-context personalization behavior. The results in Table 4 can be reproduced by performing experiment (E2) as discussed in Section 5.4 of the paper.

A.4.2 Experiments

Our auditing framework is semi-automated, with the data collection infrastructure being automated but the experiments being manual.

(E1): Auditing User Tracking (Data Collection) in GenAI Assistants [30 human-mins. for experimentation + 15 human-mins. for manual analysis per extension]:

Preparation: To reproduce results in Table 2, read Section 5.3 and refer to prompts in Table 6. Since the experiments were performed and analyzed manually, the evaluators can limit the evaluation to at most 20 scenarios, comprising of public as well as private spaces described in Table 2. Create fresh accounts for each of the 10 private spaces (wherever possible). Ensure that the crawling infrastructure is set-up and validated as explained in Section A.3 and account for the assistant being tested is pre-created. In ChatGPT for Google, it also needs to be linked to a chatgpt.com account.

Execution: To test the selected extension, for each space, follow Section A.3.2 to start mitmweb in one terminal window and initiate a new browser profile in the second terminal. Within the browser instance, install the extension, pin it, and login to extension. Now, visit the space being tested and open some sensitive or copyrighted content to understand collection practices of the assistant. If it’s a private space, first authenticate into the website and then navigate to an appropriate content

page to test for. Refer to Table 6 row corresponding to the selected space and ask those prompts in the sidebar chat of the assistant one after the other. Once the testing completes, close the browser instance and stop the data collection to save the .flow file.

Results: Manually analyze the flow file of each or a sample of evaluated spaces to validate the observed data sharing with respect to the results displayed in the Table 2 for the tested extension.

(E2): Auditing User Tracking (Data Sharing), Profiling, and Personalization in GenAI Assistants [30 human-mins. for experimentation + 15 human-mins. for manual analysis per extension]:

Preparation: Similar to E1, we manually perform three scenarios for each extension – search, browse, and summarize. Read Section 5.4 to understand how to perform experiments for each scenario using fresh accounts (wherever possible). If reusing the same account (e.g., Google account), ensure all data is deleted between subsequent tests and that each test is carried out in a fresh browser profile.

Execution: For each scenario per extension, use a distinct browser profile and start mitmweb and instance of the browser from two separate terminal windows, followed by installing the extension and performing the login. Next, perform each experimental scenario as per Section 5.4 of our paper using the prompts listed in Section 8.1. Table 5 (Section 8.3) lists the webpages to visit during the browse and summarize scenarios. After experimenting all scenarios of an extension, generate a summarized .csv file.

Results: The generated .csv file can be used to compare the data sharing practices of the assistant with the first- and third-parties by analyzing payload, request URL, and cookies against Table 3. Likewise, third-party sharing flows in Figure 3 can also be easily verified. Profiling and personalization of the selected assistant across the three scenarios can be validated against Table 4 by analyzing the responses returned by the assistant while performing experiments.

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/>.