



USENIX Security '26 Artifact Appendix: Interpolation-Based Optimization for Enforcing ℓ_p -Norm Metric Differential Privacy in Continuous and Fine-Grained Domains

Chenxi Qiu
University of North Texas

A Artifact Appendix

A.1 Abstract

This artifact accompanies the paper “Interpolation-Based Optimization for Enforcing ℓ_p -Norm Metric Differential Privacy in Continuous and Fine-Grained Domains” and provides the complete MATLAB implementation for the proposed *anchor-interpolated privacy optimization (AIPO)* framework.

The package contains end-to-end driver scripts to reproduce the paper’s main results (e.g., `main_2norm.m`), core algorithmic routines under `./functions/` (including optimization, log-convex interpolation, and violation evaluation), and implementations of representative baselines. It also includes processed OpenStreetMap road-network datasets (nodes/edges) for Rome, London, and New York City under `./datasets/`, plus cached intermediate files under `./intermediate/`. Running the provided scripts generates metrics such as utility loss, mDP violation ratio, and runtime, and writes structured outputs to subfolders under `./results/`. The artifact is released as a single ZIP archive on Zenodo (DOI: 10.5281/zenodo.17851733).

A.2 Description & Requirements

The artifact is distributed as a single ZIP archive. After extracting, the top-level directory contains:

- Primary driver scripts to reproduce tables/figures:
`main_2norm.m`,
`main_1norm.m`,
`main_granularity.m`,
`main_ablation_privacybudget.m`
- `functions/`: implementation of anchor-based optimization, log-convex interpolation, evaluation (utility/violation), and baseline methods.
- `datasets/`: processed road-network datasets (`nodes.csv`, `edges.csv`) for each city benchmark.
- `intermediate/`: cached intermediate files (e.g., pre-computed loss matrices) to reduce preprocessing overhead.

- `results/`: output directory populated by the scripts (`cost/violation/time/PPR` and related logs).

Configuration is centralized in `parameters.m` (e.g., number of repeats). Input data and/or download instructions are under `dataset/`. A paper copy (`paper.pdf`) is included for reference. To run the experiments, start MATLAB in the repository root, modify `parameters.m` if needed, and then execute `main_2norm.m` to reproduce the main results. To reproduce the supplementary (appendix) results, run `main_1norm_appendix.m`, `main_ablation_budget_appendix.m`, and `main_granularity_appendix.m`. This organization allows evaluators to quickly locate the README, source code, datasets, and run scripts without navigating extraneous files.

A.2.1 Security, privacy, and ethical concerns

Executing this artifact poses *low risk* to evaluator machine security and data privacy. The MATLAB scripts are intended to run locally in user space, require no administrator/root privileges, and do not ask evaluators to disable security mechanisms. The code writes outputs only to local folders (i.e., `./results/`) and is not designed to modify system files. The included datasets are road-network graphs (nodes/edges) and do not contain personal or sensitive information, and the artifact does not require uploading data or sending telemetry.

A.2.2 How to access

The artifact is archived on Zenodo. It can be downloaded from:

<https://doi.org/10.5281/zenodo.17851733>

(DOI: 10.5281/zenodo.17851733). The record provides `AIPO_artifact.zip`, which contains the full MATLAB code, datasets, and scripts needed to reproduce the results.

A.2.3 Hardware dependencies

- *CPU*. A commodity x86_64 machine; a multi-core processor is recommended (4+ cores).

- *Memory.* 32 GB RAM is recommended to run `main_2norm.m`, `main_1norm.m`, `main_granularity.m`, `main_ablation_privacybudget.m`. More memory can reduce runtime by avoiding swapping.
- *Accelerators.* Not required (no GPU/FPGA needed).
- *Network.* Not required for normal use.

A.2.4 Software dependencies

- *OS.* Windows 10/11, macOS, or Linux (any recent 64-bit distribution).
- *MATLAB.* MATLAB (R2022b or newer recommended).
- *Toolboxes.* The **Optimization Toolbox** is required for linear programming routines used by some baselines. The **Statistics and Machine Learning Toolbox** is used for random sampling (e.g., `randsample`).
- *Other dependencies.* No external solvers, GPUs, or compilers are required for the provided workflows. Network access is not needed for normal use.

A.2.5 Benchmarks

Included datasets. The artifact bundles all data needed to run the experiments entirely offline. The `datasets/` directory contains preprocessed road-network graphs for: **Rome**, **New York City**, and **London**.

Contents and format. Each dataset provides a road graph within the target region (node/edge lists; `nodes.csv` and `edges.csv`). The provided driver scripts (e.g., `main_2norm.m`) load these files automatically.

Compared methods included. Implementations of the paper’s compared mechanisms are included in the artifact alongside AIPO (e.g., exponential-mechanism and LP-based baselines), exposed through the run scripts and the `functions/` directory. No external repositories or additional datasets/models are required.

No external downloads. All datasets are included with the package; internet access is not needed to reproduce results.

Data use and privacy. The datasets are public/aggregated research data and contain no personally identifiable information.

A.3 Set-up

A.3.1 Installation

Step 1: Download the artifact. Download the ZIP archive from the Zenodo DOI <https://doi.org/10.5281/zenodo.17851733> and extract it to a local folder.

Step 2: Install dependencies. Install MATLAB (R2022b or newer recommended) with the following toolboxes: (i) **Optimization Toolbox** (required for LP-based baselines), and (ii) **Statistics and Machine Learning Toolbox** (used for `randsample`).

A.3.2 Basic Test

This test verifies that (i) MATLAB can locate and execute the artifact code, (ii) datasets can be loaded, (iii) the core optimization/interpolation pipeline runs end-to-end, and (iv) result files are generated.

Prerequisites. MATLAB is launched in the artifact root directory and the code is on the path: `addpath(genpath(pwd))`;

Command. Run the MATLAB file `main_2norm.m`. No input parameters are required (the script uses the default settings shipped with the artifact).

Expected successful output. The test is successful if:

- MATLAB finishes without errors, and
- the `results/2norm/` directory is updated with output files produced by the run.

Troubleshooting. If MATLAB reports missing functions (e.g., `randsample` or LP solvers), verify that the required toolboxes are installed and that `addpath(genpath(pwd))` was executed in the artifact root directory.

A.4 Evaluation workflow

A.4.1 Major Claims

(C1): AIPO empirically satisfies (ϵ, d_p) -mDP on the evaluated city benchmarks, while discretization-based baselines (e.g., LP and COPT) may exhibit nonzero empirical violations. Supported by (E1) / Table 1.

(C2): Under the same privacy budgets, AIPO improves utility loss compared to predefined-noise mechanisms (Laplace, EM, TEM) and hybrid methods (COPT, RMP). Supported by (E2) / Table 2.

(C3): AIPO is more time-efficient than optimization-based baselines (LP, COPT). Supported by (E3) / Table 3.

A.4.2 Experiments

All three experiments (E1)–(E3) are executed together by running the MATLAB script `main_2norm.m`. With the default setting (`NT_TEST=1`), each dataset run typically takes about 5–6 hours, so running all three datasets takes about 15–18

hours (runtime may vary by machine). Increasing `NT_TEST` scales the runtime approximately linearly.

Preparation: Extract the artifact ZIP, start MATLAB in the artifact root, and run:

```
addpath(genpath(pwd));
```

To set the number of repeats, set value of "NT_TEST" in parameters.m.

Execution: Run:

```
main_2norm.m
```

(E1): (*Supports C1*) *mDP violation evaluation.* MATLAB prints a table titled: “[city] road map - Violation ratio”, as Fig. 1 shows:

```
rome road map - Violation ratio
Method e=0.2 e=0.4 e=0.6 e=0.8 e=1.0 e=1.2 e=1.4 e=1.6
Pre-defined Noise Distribution:
EM 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000
Laplace 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000
TEM 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000
Hybrid Method:
RMP 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000
COPT 1.97810.000 8.09710.000 6.29410.000 5.99810.000 5.72310.000 5.47610.000 5.16510.000 4.92910.000
LP 2.07610.000 2.10810.000 1.81710.000 1.97710.000 1.65210.000 1.36210.000 1.01710.000 0.83410.000
AIPO-R 8.2810.000 7.58210.000 5.01210.000 3.31410.000 2.09910.000 1.25010.000 0.89810.000 0.57310.000
AIPO* 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000 0.00010.000
```

Figure 1: Printed table for mDP violation ratio (Rome roadmap).

The results are also saved to:

```
results/violation/city/violation_method.mat
```

Verify that AIPO and predefined-noise mechanisms (EM, Laplace, TEM) to achieve 0% empirical violations, while relaxed/approximate baselines (COPT, LP, AIPO-R) may yield nonzero violation ratios, matching the results in Table 1.

(E2): (*Supports C2*) *Utility-loss comparison.* MATLAB prints a table titled: “[city] road map - Utility loss”, as Fig. 2 shows:

```
rome road map - Utility loss
Method e=0.2 e=0.4 e=0.6 e=0.8 e=1.0 e=1.2 e=1.4 e=1.6
Pre-defined Noise Distribution:
EM 8.5410.000 8.5110.000 8.5010.000 8.4910.000 8.4910.000 8.4910.000 8.4910.000 8.4810.000
Laplace 8.5410.000 8.5110.000 8.5010.000 8.4910.000 8.4910.000 8.4910.000 8.4910.000 8.4810.000
TEM 8.5910.000 8.5910.000 8.5910.000 8.5910.000 8.5910.000 8.5910.000 8.5910.000 8.5910.000
Hybrid Method:
RMP 5.8610.000 5.1910.000 5.0010.000 4.9210.000 4.8810.000 4.8510.000 4.8310.000 4.8210.000
COPT 7.7110.000 8.5110.000 8.4810.000 8.4710.000 8.4610.000 8.4410.000 8.4110.000 8.3710.000
LP 4.8910.000 3.1810.000 2.5610.000 2.2610.000 2.1110.000 2.0310.000 1.9910.000 1.9710.000
AIPO-R 3.9310.000 2.9910.000 2.6210.000 2.4610.000 2.3610.000 2.3210.000 2.2910.000 2.2710.000
LB 2.3210.000 1.7810.000 1.7410.000 1.7310.000 1.7310.000 1.7310.000 1.7310.000 1.7310.000
AIPO* 5.6410.000 4.5610.000 3.9010.000 3.4710.000 3.1810.000 2.9510.000 2.8210.000 2.7110.000
```

Figure 2: Printed table for utility loss (Rome roadmap).

The results are also saved to:

```
results/cost/city/loss_method.mat
```

Verify that the relative ordering among methods matches Table 2: AIPO achieves lower utility loss than predefined-noise mechanisms (Laplace, EM, TEM) and hybrid methods (COPT, RMP), while incurring higher utility loss than the lower bound (LB) and relaxed/discretization-based methods (AIPO-R, LP).

(E3): (*Supports C3*) *Computation time evaluation.* MATLAB prints a table titled: “[city] road map -

Computation time (s)”, as Fig. 3 shows:

```
rome road map - Computation time (s)
Method e=0.2 e=0.4 e=0.6 e=0.8 e=1.0 e=1.2 e=1.4 e=1.6
COPT 145.97210.000 152.04210.000 143.86810.000 143.13110.000 143.39010.000 141.73310.000 143.31910.000 142.01910.000
LP 237.31310.000 46.38310.000 262.01810.000 118.06610.000 89.10910.000 65.73710.000 82.72510.000 120.17510.000
AIPO* 15.45410.000 15.71110.000 14.77610.000 13.34710.000 12.79510.000 12.49610.000 12.33510.000 12.51110.000
```

Figure 3: Printed table for computation time (Rome roadmap).

The results are also saved to:

```
results/time/city/time_method.mat
```

Verify that AIPO is faster than LP and COPT as in Table 3 (absolute runtimes may vary by machine, but the ordering should be consistent).

A.5 Notes on Reusability

Scaling experiments up/down. The provided run scripts execute predefined parameter sweeps. To *scale down* for quick testing, reduce the number of privacy budgets (e.g., fewer ϵ values), restrict to a single dataset (e.g., Rome), or disable expensive LP-based baselines. To *scale up*, extend the parameter lists or enable additional baselines.

Running on a new road network. To evaluate on a new region, add a new folder under `datasets/` with `nodes.csv` and `edges.csv` in the same format as the included benchmarks.

Changing the metric or utility model. The code separates dataset loading, metric construction, and utility/loss computation. Users can replace or modify the distance computation (e.g., switch from Euclidean distance to Manhattan distance, or plug in a custom metric) and update the loss-matrix generation accordingly. Cached inputs under `intermediate/` can be deleted to force regeneration when changing these components.

Extending baselines. Baseline mechanisms are organized under `functions/` (including a `benchmarks/` subdirectory). New mechanisms can be added by implementing a perturbation-generator function with the same input/output interface as existing baselines and registering it in the run scripts.

A.6 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/usenixsec2026/>.