

Google COVID-19 Community Mobility Reports: Anonymization Process Description (version 1.0)

Ahmet Aktay, Shailesh Bavadekar, Gwen Cossoul, John Davis, Damien Desfontaines, Alex Fabrikant, Evgeniy Gabrilovich, Krishna Gadepalli, Bryant Gipson, Miguel Guevara, Chaitanya Kamath, Mansi Kansal, Ali Lange, Chinmoy Mandayam, Andrew Oplinger, Christopher Pluntke, Thomas Roessler, Arran Schlosberg, Tomer Shekel, Swapnil Vispute, Mia Vu, Gregory Wellenius, Brian Williams, and Royce J Wilson

community-mobility@google.com

April 13, 2020

Abstract

This document describes the aggregation and anonymization process applied to the initial version of Google COVID-19 Community Mobility Reports (published at <http://google.com/covid19/mobility> on April 2, 2020), a publicly available resource intended to help public health authorities understand what has changed in response to work-from-home, shelter-in-place, and other recommended policies aimed at flattening the curve of the COVID-19 pandemic. Our anonymization process is designed to ensure that no personal data, including an individual’s location, movement, or contacts, can be derived from the resulting metrics.

The high-level description of the procedure is as follows: we first generate a set of anonymized metrics from the data of Google users who opted in to Location History. Then, we compute percentage changes of these metrics from a baseline based on the historical part of the anonymized metrics. We then discard a subset which does not meet our bar for statistical reliability, and release the rest publicly in a format that compares the result to the private baseline.

COVID-19 Community Mobility Reports provide insights into changes in mobility patterns. These reports use anonymized, aggregated data to chart movement trends over time by geography, as well as by place categories, showing trends over several weeks. This works in a similar way to existing Google products and features. For example, Google Maps uses aggregated, anonymized data to show how busy certain types of places are, including when a local business tends to be the most crowded. Public health officials have suggested this same type of aggregated, anonymized data could also be helpful as they make critical decisions to combat COVID-19.

The COVID-19 Community Mobility Reports provide insights into what has changed in response to work-from-home, stay-at-home, and other recommended policies aimed at flattening the curve of the COVID-19 pandemic. They analyze trends in visits made to high-level categories of places, including workplaces, retail and recreational venues, groceries and pharmacies, parks, transit centers, and places of residence. Each version of the report will show trends over several weeks, with the most recent data representing 48 hours prior.

As explained in greater technical detail below, the anonymization process for these reports includes differential privacy [1], which is well-suited to produce analytics in contexts where the categories of data are known in advance. Our rigorous approach intentionally adds random noise to metrics in a way that maintains both users’ privacy and the overall accuracy of the aggregated data.

This paper is structured as follows: we introduce our method to produce anonymized metrics with differential privacy. We then explain how we post-process the anonymized metrics to generate the reports.

1 Definitions

Location History users The metrics in these reports are based on the data of Google users who have opted in to Location History [2], (“LH users”), a feature which is off by default.

Differential Privacy [3] Let ϵ be a positive real number and A be a randomized algorithm that computes a metric. In the context of this report, A is considered ϵ -differentially private if for all input datasets D_1 and D_2 such that D_2 can be obtained from D_1 by adding or removing a single user, and for all subsets of $S \in \text{im}A$:

$$Pr[A(D_1) \in S] \leq \exp(\epsilon) \cdot Pr[A(D_2) \in S].$$

Granularity levels The metrics are aggregated per day and per geographic area. There are three levels of geographic areas; in this paper, we call these granularity levels.

- Granularity level 0 corresponds to metrics aggregated by country / region.
- Granularity level 1 corresponds to metrics aggregated by top-level geopolitical subdivisions (e.g. US states).
- Granularity level 2 corresponds to metrics aggregated by higher-resolution granularity (e.g. U.S. Counties).

Granularity levels 1 and 2 are defined differently in different countries, to account for knowledge of local public-health needs. Note that in general, the geographic area represented gets smaller as the Granularity number increases. No metrics are published for geographic regions smaller than 3km².

2 Generating anonymized metrics

We are releasing aggregated, anonymized data that is designed to ensure that no personal data, including an individual’s location, movement, or contacts, can be derived from the resulting metrics. To that end, we anonymize the statistics with differential privacy. We query the underlying data using our open-source differential privacy library [4], which adds Laplace noise [5] to protect each metric with differential privacy.

2.1 Daily visits in public places

We count the number of unique LH users who visited a public place of a given category in a given day at each granularity level. There are seven different categories derived from the data: retail, recreation, eateries (reported as part of “Retail & recreation”); groceries, pharmacies; transit; and parks. We add Laplace noise to each count according to the following table.

Granularity level	Scale of Laplace noise	Corresponding ϵ parameter
0	1/0.11 \approx 9.09 ($\sigma \approx 12.86$)	0.11
1	1/0.11 \approx 9.09 ($\sigma \approx 12.86$)	0.11
2	1/0.22 \approx 4.55 ($\sigma \approx 6.43$)	0.22

Table 1: Noise parameters used for the daily visits in public places metrics

For each location (at all geographic levels), each LH user can contribute at most once to each category. We also bound the contribution of each LH user to 4 $\langle \text{category}, \text{location} \rangle$ pairs per day and per geographic level, using a process similar to the one described in this paper [6]: if an LH user contributes to more than 4 pairs in a given day and given geographic level, we randomly select 4 of them, and discard the others.

For example, suppose that on the same day, an LH user goes to public places in all 7 categories in two distinct neighboring countries. This makes a total of 14 $\langle \text{category}, \text{location} \rangle$ pairs at country level. We would randomly discard 10 of these pairs when computing country-level statistics.

This process does not significantly affect data accuracy: in the US, at county level, 99% of LH users contribute 3 or fewer $\langle \text{category}, \text{place} \rangle$ pairs per day on average. Thus, each daily place visit is protected by differential privacy with $\epsilon = 0.44$, and the total daily contribution of each user with a maximum of $\epsilon = 1.76$.

2.2 Residential

For the purposes of this analysis, we use signals like relative frequency, time and duration of visits to calculate metrics related to places of residence. We calculate an average amount of time spent at places of residence for LH users in hours. This computation is performed for each day and geographic area, using the differentially private mean mechanism from our open-source library [7]. This mechanism works as follows:

- We compute the amount of time spent at place of residence in a given day and geographic area in hours by summing up the individual values per user offset by 12, so all individual values fall into the range $[-12; 12]$. We then add Laplace noise to this sum; the scale of the noise is indicated in the table below. We denote the real sum S , and noisy sum NS .
- We compute the count of unique users who spent any time at residences in a given day and geographic area. We then add Laplace noise to this count; the scale of the noise is indicated in the table below. We refer to the real count C , and the noisy count NC .
- Finally, we compute the ratio NS/NC for each day and each geographic area, add 12 as offset and clamp it to the range $[0, 24]$ hours/day.

For example, at county-level, NS is obtained by first sampling a random number from a Laplace distribution of scale 109.1, and then adding that number to S . In the table below, we also indicate the standard deviation σ of the noise added to each value.

Granularity level	Scale of Laplace noise: sum (total hours/day)	Scale of Laplace noise: count (number of users)	Corresponding ϵ parameter
0	$12/0.055 \approx 218.2$ ($\sigma \approx 308.6$)	$1/0.055 \approx 18.2$ ($\sigma \approx 25.71$)	0.11
1	$12/0.055 \approx 218.2$ ($\sigma \approx 308.6$)	$1/0.055 \approx 18.2$ ($\sigma \approx 25.71$)	0.11
2	$12/0.110 \approx 109.1$ ($\sigma \approx 154.3$)	$1/0.110 \approx 9.09$ ($\sigma \approx 12.86$)	0.22

Table 2: Noise parameters used for the residential metrics

Each user can contribute to at most one region per granularity level, which protects these metrics by differential privacy with $\epsilon = 0.44$ total budget across all granularities. A description of the differentially private mean mechanism implemented and a proof of its privacy guarantees is described in [8] (Algorithm 2.4).

2.3 Workplaces

For the purposes of this analysis, we use signals like relative frequency, time and duration of visits to calculate metrics related to places of residence and places of work of LH users. We calculate how many LH users spent more than 1 hour at their places of work. This computation is performed for each day and geographic area. Then, we add Laplace noise to each count according to the following table.

Geographic level	Scale of Laplace noise	Corresponding ϵ parameter
0	$1/0.11 \approx 9.09$ ($\sigma \approx 12.86$)	0.11
1	$1/0.11 \approx 9.09$ ($\sigma \approx 12.86$)	0.11
2	$1/0.22 \approx 4.55$ ($\sigma \approx 6.43$)	0.22

Table 3: Noise parameters used for the work places metrics

The count is aggregated by places of residence of LH users. Since each user can contribute to at most one geographic area per granularity level, these metrics are protected by differential privacy with $\epsilon = 0.44$.

3 Generating the report from the anonymized metrics

The metrics described above are generated for each day, starting on 2020-01-01. They are then used to generate the percentage changes relative to day of the week published in the reports. All operations described below use only the output of the differentially private mechanisms described in the previous section; so they do not consume any privacy budget.

Additional privacy protections We discard all metrics for which the geographic region is smaller than 3km^2 , or for which the differentially private count of contributing users (after noise addition) is smaller than 100. Geographic regions smaller than 3km^2 may be merged such that the union of their area is above the 3km^2 threshold. This merging does not occur across country boundaries, except for the Vatican City and Italy.

3.1 Computing percentage changes from a baseline

For each individual metric generated using the mechanisms described above, we compute the ratio between the metric for a given day D and the same metric computed for the baseline period. The reference baseline is defined in the following way.

- We consider the 5-week range from 2020-01-03 through 2020-02-06.
- Within this 5-week range, we consider the 5 days with the same day of week as D. For example, if D is 2020-03-20, D is a Friday, so we consider the 5 Fridays in this 5-week range (Jan 3 to Jan 31, inclusive).
- We compute the median of the differentially private metrics for these 5 baseline days.
- This median metric is the baseline metric for D.

We then compute and publish the ratio between the metric for D and the baseline metric, as a percentage.

3.2 Removing unreliable metrics for Residential, Workplace, Transit, and Parks

In some regions, the noise added to obtain differential privacy can reduce the confidence that we are capturing a meaningful change, typically when there is not a lot of data for the metric. When, because of this uncertainty, the percentage change for one of these metrics has a 5% chance (or higher) of being wrong by more than ± 10 absolute percentage points, we do not publish it and instead include an asterisk denoting that there is not enough data available to present privacy-safe information. More precisely:

- Before releasing a ratio metric/baseline, we compute 97.5% confidence intervals for the metric and its baseline. Let us denote $[\text{metric_min}, \text{metric_max}]$ and $[\text{baseline_min}, \text{baseline_max}]$ these respective confidence intervals.
- We compute the ratios $\text{metric_min}/\text{baseline_max}$ and $\text{metric_max}/\text{baseline_min}$.
- If one of these ratios differs from the differentially private ratio by more than 10 absolute percentage points, we do not publish the corresponding percentage changes.

If the last condition is not satisfied, then the probability of being wrong by more than 10 absolute percentage points in each direction is lower than 2.5%. By union bound, this means that there is at most a 5% risk of being wrong by more than 10 absolute percentage points. Note that the confidence intervals are based on an already differentially private value and on public data (the scale and shape of the noise), so no privacy budget is consumed by this operation.

References

- [1] Enabling developers and organizations to use differential privacy. <https://developers.googleblog.com/2019/09/enabling-developers-and-organizations.html>.
- [2] Location History. <https://support.google.com/accounts/answer/3118687?hl=en>.
- [3] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.
- [4] Google’s C++ Differential Privacy Library. <https://github.com/google/differential-privacy>.
- [5] Laplace Distribution. https://en.wikipedia.org/wiki/Laplace_distribution.
- [6] Royce Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-Marengo, and Bryant Gipson. Differentially Private SQL with Bounded User Contribution. *To appear in Proceedings on Privacy Enhancing Technologies Symposium 2020*.
- [7] DP Library Bounded Mean. https://github.com/google/differential-privacy/blob/master/differential_privacy/algorithms/bounded-mean.h.
- [8] N. Li, M. Lyu, D. Su, and W. Yang. *Differential Privacy: From Theory to Practice*. Morgan & Claypool Publishers, 2016.