

Google COVID-19 Search Trends Symptoms Dataset: Anonymization Process Description (version 1.0)

Shailesh Bavadekar, Andrew Dai, John Davis, Damien Desfontaines, Ilya Eckstein, Katie Everett, Alex Fabrikant, Gerardo Flores, Evgeniy Gabrilovich, Krishna Gadepalli, Shane Glass, Rayman Huang, Chaitanya Kamath, Dennis Kraft, Akim Kumok, Hinali Marfatia, Yael Mayer, Benjamin Miller, Adam Pearce, Irippuge Milinda Perera, Venky Ramachandran, Karthik Raman, Thomas Roessler, Izhak Shafran, Tomer Shekel, Charlotte Stanton, Jacob Stimes, Mimi Sun, Gregory Wellenius, and Masrouf Zoghi

covid-19-search-trends-feedback@google.com

September 2, 2020

Abstract

This report describes the aggregation and anonymization process applied to the initial version of COVID-19 Search Trends symptoms dataset [1], a publicly available dataset that shows aggregated, anonymized trends in Google searches for symptoms (and some related topics). The anonymization process is designed to protect the daily symptom search activity of every user with ϵ -differential privacy for $\epsilon = 1.68$.

1 Introduction

The COVID-19 Search Trends symptoms dataset is a publicly available dataset that shows aggregated, anonymized trends in Google searches for symptoms (and some related topics). The dataset provides a daily or weekly time series for each region showing the relative volume of searches for each symptom. Its goal is to help researchers, public health experts and data analysts better understand the impact of COVID-19, while protecting our users' privacy. For this purpose, the dataset covers approximately 400 symptoms, such as cough, fever and difficulty breathing. To ensure strict privacy standards, all published data is aggregated and anonymized, protecting each user's symptom search activity on a given day using a differentially private mechanism. No personal data or individual searches are included in the dataset.

The data of the symptoms dataset reflect the volume of Google searches that may be associated with a particular medical symptom. For this purpose, we count the daily number of searches relating to that symptom within a given geographic region and normalize this number based on the total search activity in that region. The resulting dataset is either a daily or a weekly series showing the relative frequency of searches for a particular symptom in a particular region. The dataset covers the recent period and we'll gradually expand its range as part of regular updates.

Similar to the Google COVID-19 Community Mobility Reports [2, 3] and as explained in greater technical detail below, the anonymization process of the symptoms dataset is based on differential privacy [4], which is a well established concept for producing data that satisfy formal privacy guarantees. For this purpose, we intentionally perturb our data by adding random noise and drop data that is deemed to be unreliable. The symptoms dataset is designed to maintain the privacy of our users while releasing aggregated and anonymized data that is as accurate and useful as possible.

The remainder of this report is structured as follows: First, we introduce some basic concepts and terminology of the symptoms dataset. We then explain how differential privacy is used to produce anonymized aggregates. Finally, we elaborate on how the published data is built from anonymized aggregates.

2 Definitions

The following definitions explain some common terms and concepts that we use throughout this report.

User A user who did a web search on Google.

Symptom Search A symptom search is a Google search query issued by a search user that relates to a particular medical symptom or health condition. The symptoms we are considering are compiled in a predefined list including medical issues such as cough, fever and difficulty breathing.

Geographic Granularity The COVID-19 Search Trends symptoms dataset is aggregated per geographic region. Similarly to the Google COVID-19 Community Mobility Reports [2, 3], we distinguish between three levels of geographic regions, which we call granularity levels:

- Granularity level 0 corresponds to data aggregated by country.
- Granularity level 1 corresponds to data aggregated by top-level geopolitical subdivisions (e.g., US states or equivalent geographical regions in other countries).
- Granularity level 2 corresponds to data aggregated by higher-resolution granularity (e.g., US counties or equivalent geographical regions in other countries).

Granularity levels 1 and 2 are defined differently in different countries to account for the differences between countries’ public health systems. Note that in general, the area of a geographic region gets smaller as the granularity number increases. All regions of the symptoms dataset have an area of at least 3km².

Temporal Granularity For each geographic region and symptom, the data is released either as daily or weekly aggregates. The particular temporal granularity depends on the quality of the data. In general, we try to provide daily aggregates whenever possible. However, if the accuracy of our data is affected too much by our privacy protections, we may choose to release weekly aggregates instead. Weekly aggregates help to improve accuracy while maintaining the same level of privacy. Since weekly aggregates are based on more data points, the noise added for differential privacy introduces less relative error.

The decision between daily and weekly aggregates is made per symptom and region based on the data available at the time of the initial release of the symptoms dataset, and is made in a differentially private manner (more details can be found in the final paragraph of this report). We keep the temporal granularity fixed for a given region and symptom throughout the full time range of the dataset release.

Differential Privacy Let A be a randomized algorithm for computing some metric over a given dataset. In the context of this report, we consider a pair of datasets D_1 and D_2 neighboring if D_2 can be obtained from D_1 by adding or removing a single user’s search activity on a given day. We then call A ϵ -differentially private if for any pair of neighboring datasets D_1 and D_2 and for all sets S of the possible outputs of A ¹:

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

3 Anonymizing User Data

The COVID-19 Search Trends symptoms dataset is based on anonymized counts designed to protect personal data. More precisely, we protect every user’s symptom search activity on a particular day with ϵ -differential privacy for $\epsilon = 1.68$. The majority of this ϵ , 97.5%, is used to anonymize symptom search counts, while 2.5% is used to anonymize general search activity for normalization purposes. See Figure 1 for a system diagram of the anonymization process. All counts are anonymized by adding appropriately scaled Laplace noise [5]. We generate this noise using our open-source differential privacy library [6].

¹We generate a fixed set of metrics (all <day, symptom, region> combinations) and add Laplace noise to all of them independent of whether they contain user data or not. Thus, the released metrics satisfy (ϵ, δ) -differentially privacy, with $\delta = 0$.

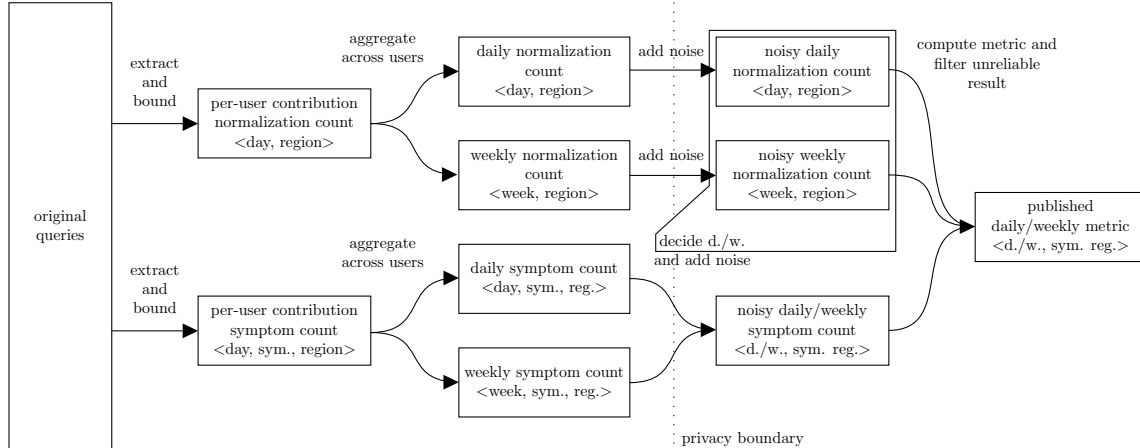


Figure 1: System diagram of the data generation and anonymization process

3.1 Symptom Search Counts

To generate the symptoms dataset, we count the number of searches that relate to a particular medical symptom and group them by date and geographical region. As a result, each count corresponds to a triplet of $\langle \text{day, symptom, region} \rangle$. For each day and geographic granularity level, a user can contribute at most once to any given count (per-symptom bound) and to no more than three counts in total (cross-symptom bound). Similar to the process described by Wilson et al. [7], we arbitrarily discard symptom searches of a user that exceed their contribution bound. Approximately 75% of users search at most three symptoms per day, so the amount of data dropped is relatively small.

As an example, assume that on June 3, 2020 a user made two searches related to fever in Santa Clara county (CA), one search related to fever in San Bernardino county (CA), one search related to fever in Clark county (NV), and one search related to cough in Clark county (NV). The following table lists a possible configuration of the user’s contribution to the respective counts after contribution bounding.

level	count	contribution	bound type
0	$\langle 2020-06-03, \text{fever}, \text{United States} \rangle$	1 (originally 4)	per-symptom
0	$\langle 2020-06-03, \text{cough}, \text{United States} \rangle$	1	
1	$\langle 2020-06-03, \text{fever}, \text{California} \rangle$	1 (originally 3)	per-symptom
1	$\langle 2020-06-03, \text{fever}, \text{Nevada} \rangle$	1	
1	$\langle 2020-06-03, \text{cough}, \text{Nevada} \rangle$	1	
2	$\langle 2020-06-03, \text{fever}, \text{Santa Clara} \rangle$	1 (originally 2)	per-symptom
2	$\langle 2020-06-03, \text{fever}, \text{San Bernardino} \rangle$	1	
2	$\langle 2020-06-03, \text{fever}, \text{San Bernardino} \rangle$	1	
2	$\langle 2020-06-03, \text{fever}, \text{Clark} \rangle$	1	
2	$\langle 2020-06-03, \text{cough}, \text{Clark} \rangle$	0 (originally 1)	cross-symptom

Table 1: Example of per-symptom and cross-symptom contribution bounding

To protect a user’s daily symptom searches with 1.638-differential privacy (97.5% of the total ϵ), we then add appropriately scaled Laplace noise. In the case of daily granularity, the noise is added directly to the respective count. In the case of weekly granularity, we first sum the daily counts to obtain a weekly count and then add the noise to the weekly count. The following table lists the noise parameterization for each granularity level in terms of its scale b and standard deviation $\sigma = \sqrt{2}b$. The total ϵ sums up to 1.638, i.e., the privacy budget we spend on computing symptom searches.

level	noise added to daily or weekly count	ε per level
0	$b = 3/\varepsilon_0 \approx 17.857$ ($\sigma \approx 25.254$)	$\varepsilon_0 = 0.168$
1	$b = 3/\varepsilon_1 \approx 8.108$ ($\sigma \approx 11.467$)	$\varepsilon_1 = 0.37$
2	$b = 3/\varepsilon_2 \approx 2.727$ ($\sigma \approx 3.857$)	$\varepsilon_2 = 1.1$

Table 2: Noise parameters for daily and weekly symptom search counts

3.2 Normalization Counts

To normalize our data, we scale each daily or weekly symptom search count proportional to the total search activity in the given geographical region during the respective time period. We estimate the search activity based on the number of unique users who have issued a search query in a particular geographical region during a day. Consequently, each daily normalization count corresponds to a tuple of <day, region>.

For each day and each geographic granularity level, a user can contribute to at most one normalization count. Similarly to the symptom search counts, we arbitrarily discard contributions that exceed this contribution limit.

The daily and weekly normalization counts are protected with 0.021-differential privacy each, or 0.042-differential privacy in combination (2.5% of the total ε). In the case of daily counts, we add appropriately scaled Laplace noise to the respective count. In the case of weekly counts, we first sum the respective daily counts and then add Laplace noise to the sum. The following table lists the noise parameterization for each granularity level in terms of its scale b and standard deviation $\sigma = \sqrt{2}b$. Note that the total privacy budget sums up to 0.042.

level	noise added to daily and weekly count	ε per level
0	$b = 1/\varepsilon'_0 \approx 434.783$ ($\sigma \approx 614.875$)	$\varepsilon'_0 = 0.0023$
1	$b = 1/\varepsilon'_1 \approx 212.766$ ($\sigma \approx 300.897$)	$\varepsilon'_1 = 0.0047$
2	$b = 1/\varepsilon'_2 \approx 71.429$ ($\sigma \approx 101.015$)	$\varepsilon'_2 = 0.014$

Table 3: Noise parameters for daily and weekly normalization counts

4 Reporting the Metrics

The daily or weekly data published in the COVID-19 Search Trends symptoms dataset are based exclusively on the anonymized counts described in the previous section. Due to the post-processing property of differential privacy, the following process does not consume any privacy budget.

Computing the Reported Data Given a particular geographic region, symptom and time interval (either day or week), the normalized search count published in the COVID-19 Search Trends is computed as

$$c \cdot \max\{(A/B), 0\}$$

, where A is the noisy symptom search count, B is the normalization count and c is a scaling factor specific to the geographic region. For each geographic region, c is chosen in a way that maps the data in the initial release of the symptoms dataset to values between 0 and 100. Because we keep c fixed, future releases may contain metrics that exceed the value of 100. We also want to note that c is determined purely based on noisy symptom search counts and normalization counts, so no privacy budget is spent on its computation.

Removing Unreliable Data In some geographic regions, the noise added for differential privacy reasons can introduce a disproportionate amount of uncertainty to a metric. Typically, this happens when the respective symptom search count is empty or small. To address this uncertainty, we only keep a metric if it has a chance of 50% or more to be within 25% points of its raw value, i.e., the metric before adding noise.

More precisely, let A be some noisy symptom search count obtained after adding Laplace noise to the raw symptom search count a^* (note that a^* is not noisy but still subject to contribution bounding). Similarly,

let B be the corresponding noisy normalization count computed from the raw count b^* . To decide whether the data associated with A will be kept or dropped we:

- Compute a confidence interval $[l, r]$ based on A and B that contains a^*/b^* with a probability of at least 50%.
- Keep the data if $|A/B - l| \leq 0.25 \cdot A/B$ and $|A/B - r| \leq 0.25 \cdot A/B$. Otherwise drop it.

The confidence interval $[l, r]$ is entirely based on the anonymized counts A and B . As a result, no privacy budget is spent on removing unreliable data.

Deciding Between Daily and Weekly Granularity The decision whether a certain symptom is published at the daily or weekly granularity within a given geographic region is based on the amount of data for this symptom in that region. Intuitively, if metrics are available for more than half of the days within the sample period from February 2020 to July 2020 (i.e., not dropped as unreliable, as explained in the previous section), we want to produce data for this symptom at the daily resolution. Otherwise, if too many daily metrics are dropped, we opt for the weekly resolution, which is more reliable.

For the sake of usability, we make the decision between daily and weekly granularity once and stick to it over the course of the entire data release.

To avoid consuming additional privacy budget, we switch between daily and weekly granularity based on the anonymized volume of search activity. More precisely, for each symptom we make our decision according to the following process:

- All level 0 regions are published with daily granularity.
- We order all other regions that are of the same level and contained in the same level 0 region by the total search activity, which we approximate based on the anonymized normalization counts.
- Starting with the geographic region that has the highest search activity, we start publishing daily metrics according to the process described above.
- We continue down the list of regions, by decreasing volume of search activity. For each new region, we look at the last 20 regions we published². If 11 or more had more than 50% of metrics dropped in the time period from February 2020 to July 2020, we switch from daily to weekly granularity and publish weekly metrics for all regions from this point forward. Otherwise, we publish the current region with daily granularity, and repeat the process.

Taking a majority vote over the last 20 regions mitigates potential outliers in the ordering.

The key insight is that this process uses anonymised data only to order the regions, i.e., the daily normalization counts. Thus the privacy budget we spend on the ordering is already accounted for. Moreover, we never compute anonymized daily and weekly symptom search counts for the same symptom and region pair at the same time. This means we only spend privacy budget on one of the two counts, resulting in the promised $\epsilon = 1.68$.

References

- [1] Google covid-19 search trends symptoms dataset. <http://goo.gle/covid19symptomdataset>.
- [2] Google covid-19 community mobility reports. <https://www.google.com/covid19/mobility/>.
- [3] Ahmet Aktay, Shailesh Bavadekar, Gwen Cossoul, et al. Google covid-19 community mobility reports: Anonymization process description (version 1.0). Technical report, Google, 2020.
- [4] Differential privacy. https://en.wikipedia.org/wiki/Differential_privacy.
- [5] Laplace distribution. https://en.wikipedia.org/wiki/Laplace_distribution.

²If less than 20 regions were published, we only consider the regions published so far.

- [6] Google differential privacy library. <https://github.com/google/differential-privacy>.
- [7] Royce J. Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-Marengo, and Bryant Gipson. Differentially private sql with bounded user contribution. *Proceedings on Privacy Enhancing Technologies*, 2020(2):230 – 250, 2020.