

Differential Privacy

There are situations where Apple can improve the user experience by getting insight from what many of our users are doing, for example: What new words are trending and might make the most relevant suggestions? What websites have problems that could affect battery life? Which emoji are chosen most often? The challenge is that the data which could drive the answers to those questions—such as what the users type on their keyboards—is personal.

A privacy-preserving system

Apple has adopted and further developed a technique known in the academic world as *local differential privacy* to do something really exciting: gain insight into what many Apple users are doing, while helping to preserve the privacy of individual users. It is a technique that enables Apple to learn about the user community without learning about individuals in the community. Differential privacy transforms the information shared with Apple before it ever leaves the user's device such that Apple can never reproduce the true data.

The differential privacy technology used by Apple is rooted in the idea that statistical noise that is slightly biased can mask a user's individual data before it is shared with Apple. If many people are submitting the same data, the noise that has been added can average out over large numbers of data points, and Apple can see meaningful information emerge.

Differential privacy is used as the first step of a system for data analysis that includes robust privacy protections at every stage. The system is opt-in and designed to provide transparency to the user. The first step we take is to privatize the information using local differential privacy on the user's device. The purpose of privatization is to assure that Apple's servers don't receive clear data. Device identifiers are removed from the data, and it is transmitted to Apple over an encrypted channel. The Apple analysis system ingests the differentially private contributions, dropping IP addresses and other metadata. The final stage is aggregation, where the privatized records are processed to compute the relevant statistics and the aggregate statistics are then shared with relevant Apple teams. Both the ingestion and aggregation stages are performed in a restricted access environment so even the privatized data isn't broadly accessible to Apple employees.

Privacy budget

The Apple differential privacy implementation incorporates the concept of a per-donation *privacy budget* (quantified by the parameter epsilon), and sets a strict limit on the number of contributions from a user in order to preserve their privacy. The reason is that the slightly-biased noise used in differential privacy tends to average out over a large numbers of contributions, making it theoretically possible to determine information about a user's activity over a large number of observations from a single user (though it's important to note that Apple doesn't associate any identifiers with information collected using differential privacy).

Apple uses local differential privacy to help protect the privacy of user activity in a given time period, while still gaining insight that improves the intelligence and usability of such features as:


- QuickType suggestions
- Emoji suggestions
- Lookup Hints
- Safari Energy Draining Domains
- Safari Autoplay Intent Detection (macOS High Sierra)
- Safari Crashing Domains (iOS 11)
- Health Type Usage (iOS 10.2)

For each feature, Apple seeks to make the privacy budget small while still collecting enough data to to enable Apple to improve features. Apple retains the collected data for a maximum of three months. The donations do not include any identifier, and IP addresses are not stored.

For Lookup Hints, Apple uses a privacy budget with epsilon of 4, and limits user contributions to two donations per day. For emoji, Apple uses a privacy budget with epsilon of 4, and submits one donation per day. For QuickType, Apple uses a privacy budget with epsilon of 8, and submits two donations per day.

For Health types, Apple uses a privacy budget with epsilon of 2 and limits user contributions to one donation per day. The donations do not include health information itself, but rather which health data types are being edited by users.

For Safari, Apple limits user contributions to 2 donations per day. For Safari domains identified as causing high energy use or crashes, Apple uses a single privacy budget with epsilon of 4. For Safari Auto-play intent detection, Apple uses a privacy budget with epsilon of 8.



The Count Mean Sketch technique allows Apple to determine the most popular emoji to help design better ways to find and use our favorite emoji. The top emoji for US English speakers contained some surprising favorites.

Techniques

Local differential privacy guarantees that it is difficult to determine whether a certain user contributed to the computation of an aggregate by adding slightly biased noise to the data that is shared with Apple. But before adding this noise, it's necessary to define a data structure that captures a sketch of user input with a small number of bits. Apple currently makes use of two specific techniques:

Count Mean Sketch

In our use of the Count Mean Sketch technique for differential privacy, the original information being processed for sharing with Apple is encoded using a series of mathematical functions known as *hash functions*, making it easy to represent data of varying sizes in a matrix of fixed size.

The data is encoded using variations of a SHA-256 hash followed by a privatization step and then written into the sketch matrix with its values initialized to zero.

The noise injection step works as follows: After encoding the input as a vector using a hash function, each coordinate of the vector is then flipped (written as an incorrect value) with a probability of $1/(1 + e^{\epsilon/2})$, where ϵ is the privacy parameter. This assures that analysis of the collected data cannot distinguish actual values from flipped values, helping to assure the privacy of the shared information.

In order to stay within the privacy budget we do not send the entire sketch matrix to the server but only a random row of the matrix. When the information encoded in the sketch matrix is sent to Apple, the Apple server tallies the responses from all devices sharing information and outputs the mean value for each element of the array. Although each submission contains many randomized elements, the average value across large numbers of submissions gives Apple meaningful aggregate data.

Hadamard Count Mean Sketch

The Hadamard Count Mean-based Sketch technique uses a noise injection method similar to the one used in the Count Mean Sketch technique, but with an important difference: It applies a type of mathematical operation called a Hadamard basis transformation to the hashed encoding before performing the privatization step. Additionally, it samples only 1 bit at random to send instead of the entire row as in the Count Mean Sketch technique. This reduces communication cost to 1 bit at the expense of some accuracy.

Seeing user data

Users can examine the information being shared with Apple for the categories of data that are protected using Differential Privacy. In iOS, the information is visible under Settings > Privacy > Analytics > Analytics Data, in entries that begin with "DifferentialPrivacy." In macOS, users can launch the Console app and view the information under the Differential Privacy category of System Reports.

Controlling participation

The data-gathering features that use differential privacy are linked to the user setting for Device Analytics. Users are presented with the option of sending diagnostic information when they set up a device running macOS or iOS, and they can always change their choice later in System Preferences on macOS or the Settings app on iOS.

The beginning

Apple launched differential privacy for the first time in macOS Sierra and iOS 10. Since then, we have expanded to other use cases such as Safari and Health types. As Apple continues to refine differential privacy algorithms, we look forward to using them to improve user experience in other areas of our products, while continuing to work to protect our users' private information.