



# From Federated Learning to Federated Analytics

Peter Kairouz  
[kairouz@](mailto:kairouz@)


Building on the work of many



# Data is born at the edge

Billions of phones & IoT devices constantly generate data

Data enables better products and smarter models




# Can data live at the edge?

Data processing is moving on device:

- Improved latency
- Works offline
- Better battery life
- Privacy advantages

E.g., on-device inference for mobile keyboards and cameras.




# What is federated learning?

# Federated learning

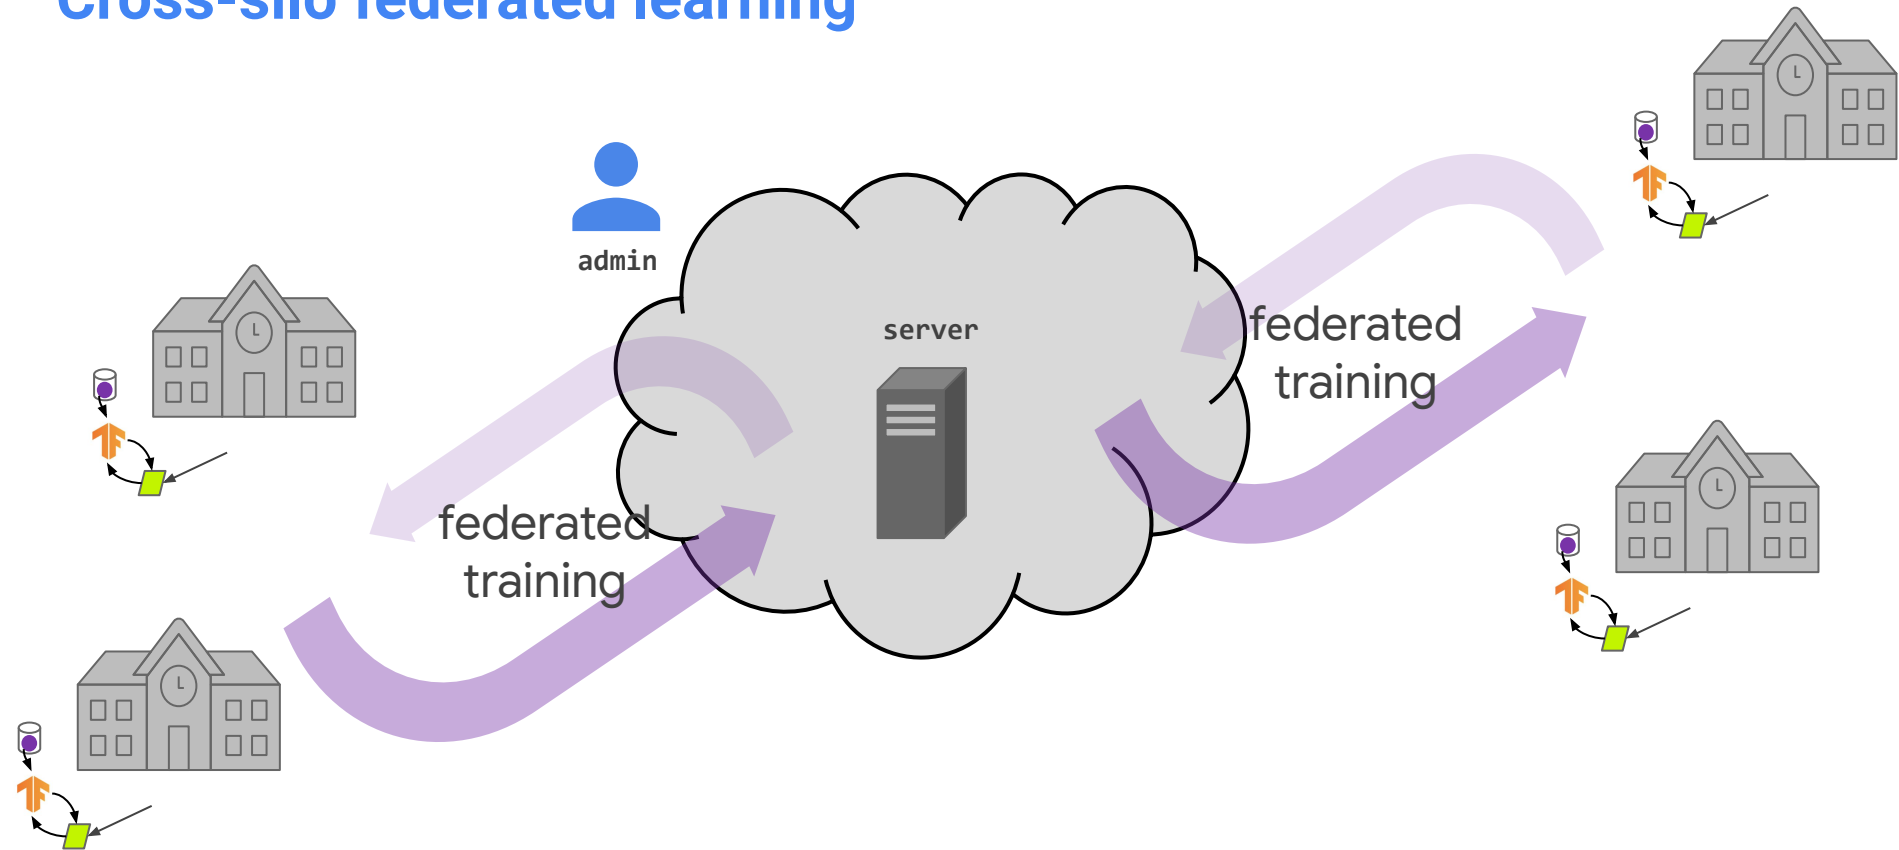
**Federated learning** is a machine learning setting where multiple entities (clients) collaborate in solving a machine learning problem, under the coordination of a central server or service provider. Each client's raw data is stored locally and not exchanged or transferred; instead, focused updates intended for immediate aggregation are used to achieve the learning objective.

working definition proposed in  
*Advances and Open Problems in Federated Learning* ([arxiv/1912.04977](https://arxiv.org/abs/1912.04977))

# Cross-device federated learning



# Cross-silo federated learning



# Characteristics of the federated learning setting

	<b>Datacenter distributed learning</b>	<b>Cross-silo federated learning</b>	<b>Cross-device federated learning</b>
<b>Setting</b>	Training a model on a large but "flat" dataset. Clients are compute nodes in a single cluster or datacenter.	Training a model on siloed data. Clients are different organizations (e.g., medical or financial) or datacenters in different geographical regions.	The clients are a very large number of mobile or IoT devices.
<b>Data distribution</b>	Data is centrally stored, so it can be shuffled and balanced across clients. Any client can read any part of the dataset.	<b>Data is generated locally and remains decentralized.</b> Each client stores its own data and cannot read the data of other clients. Data is not independently or identically distributed.	
<b>Orchestration</b>	Centrally orchestrated.	A central orchestration server/service organizes the training, but never sees raw data.	
<b>Wide-area communication</b>	None (fully connected clients in one datacenter/cluster).	Hub-and-spoke topology, with the hub representing a coordinating service provider (typically without data) and the spokes connecting to clients.	
<b>Data availability</b>	All clients are almost always available.		Only a fraction of clients are available at any one time, often with diurnal and other variations.
<b>Distribution scale</b>	Typically 1 - 1000 clients.	Typically 2 - 100 clients.	Massively parallel, up to $10^{10}$ clients.



# Characteristics of the federated learning setting

	<b>Datacenter distributed learning</b>	<b>Cross-silo federated learning</b>	<b>Cross-device federated learning</b>
<b>Addressability</b>	Each client has an identity or name that allows the system to access it specifically.		Clients cannot be indexed directly (i.e., no use of client identifiers)
<b>Client statefulness</b>	Stateful --- each client may participate in each round of the computation, carrying state from round to round.		Generally stateless --- each client will likely participate only once in a task, so generally we assume a fresh sample of never before seen clients in each round of computation.
<b>Primary bottleneck</b>	Computation is more often the bottleneck in the datacenter, where very fast networks can be assumed.	Might be computation or communication.	Communication is often the primary bottleneck, though it depends on the task. Generally, federated computations uses wi-fi or slower connections.
<b>Reliability of clients</b>	Relatively few failures.		Highly unreliable --- 5% or more of the clients participating in a round of computation are expected to fail or drop out (e.g., because the device becomes ineligible when battery, network, or idleness requirements for training/computation are violated).
<b>Data partition axis</b>	Data can be partitioned / re-partitioned arbitrarily across clients.	Partition is fixed. Could be example-partitioned (horizontal) or feature-partitioned (vertical).	Fixed partitioning by example (horizontal).

# Federated learning vs fully decentralized learning

	<b>Federated learning</b>	<b>Fully decentralized (peer-to-peer) learning</b>
<b>Orchestration</b>	A central orchestration server/service organizes the training, but never sees raw data.	No centralized orchestration.
<b>Wide-area communication pattern</b>	Hub-and-spoke topology, with the hub representing a coordinating service provider (typically without data) and the spokes connecting to clients.	Peer-to-peer topology.

# Federated beyond learning

# Beyond learning: federated analytics

**Federated analytics** is the practice of applying data science methods to the analysis of raw data that is stored locally on users' devices. Like federated learning, it works by running local computations over each device's data, and only making the aggregated results — and never any data from a particular device — available to product engineers. Unlike federated learning, however, federated analytics aims to support basic data science needs.

*definition proposed in*

<https://ai.googleblog.com/2020/05/federated-analytics-collaborative-data.html>

# Federated analytics

- Federated histograms over closed sets
- Federated heavy hitters discovery over open sets
- Federated density of vector spaces
- Federated selection of random data subsets
- Federated SQL?
- etc...

# Federated heavy hitters (frequent item) discovery



= “The **moon** is full, the sky full of stars.”



= “The full **moon** is two days before Halloween this month.”



= “You see the **moon** instead of how dark the night is.”

We are going to focus on words and assume each device has a single word - both assumptions can be relaxed

# The TrieHH Algorithm

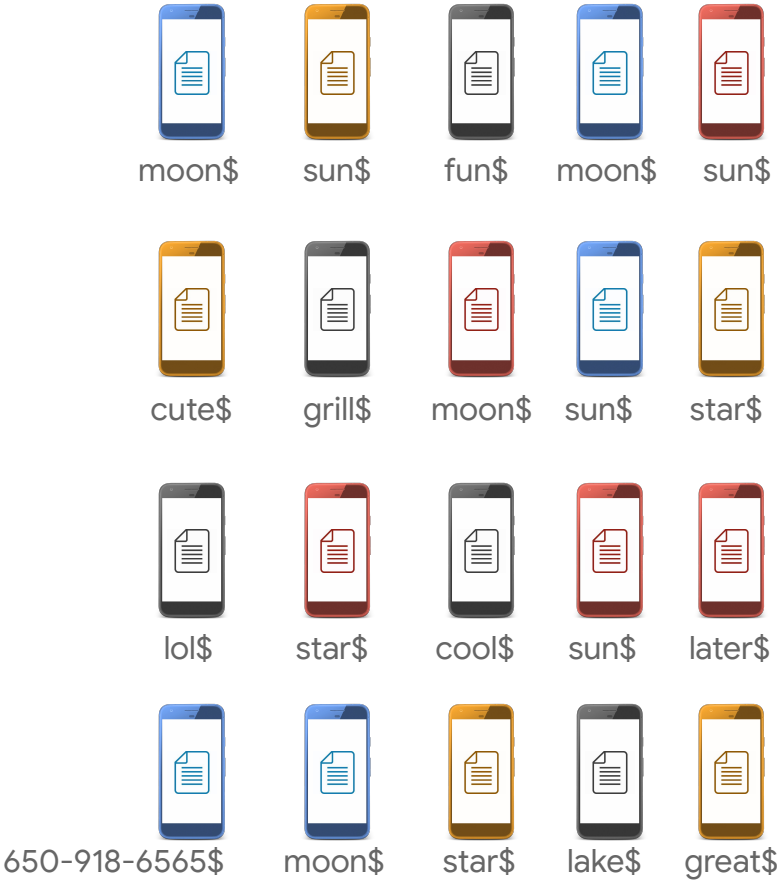
# Algorithm via example

n = 20, each user has a single word

“moon” and “sun” appear 4 times

“star” appears 3 times

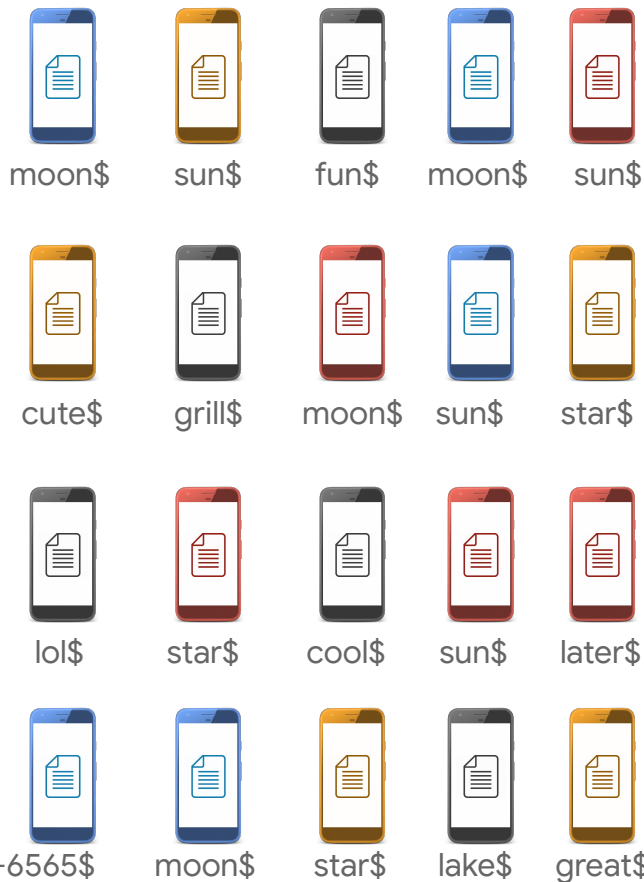
“\$” denotes end of word





# Round 1: random device selection

root



Randomly select  $m = 10$  devices at random

# Round 1: random device selection

root



# Round 1: voting stage

root

's': 3 votes


'm': 2 votes

'f', 'g', 'c', 'l', '6': 1 vote each

**Each device votes on one character**



# Round 1: vote thresholding stage

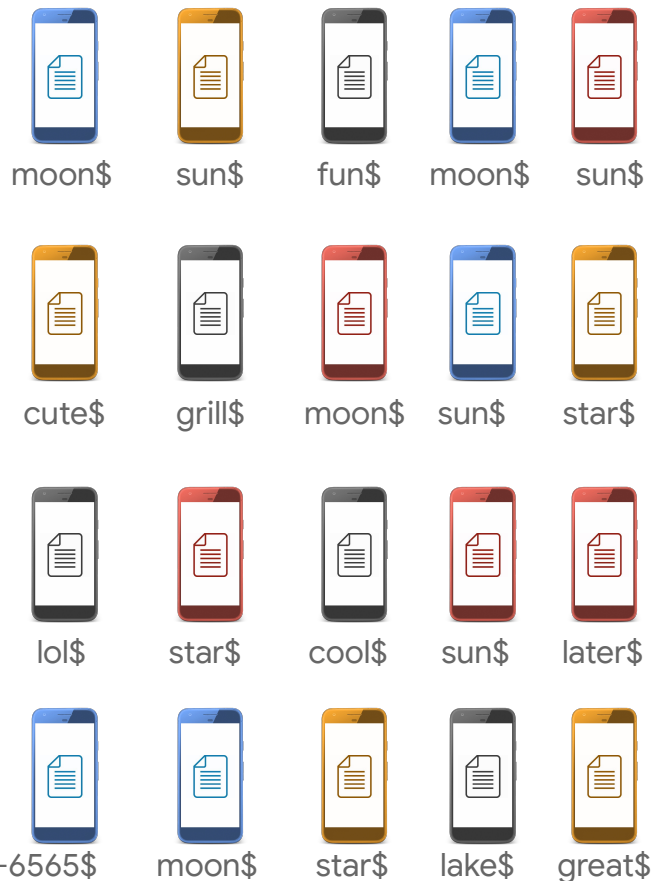


's': 3 votes

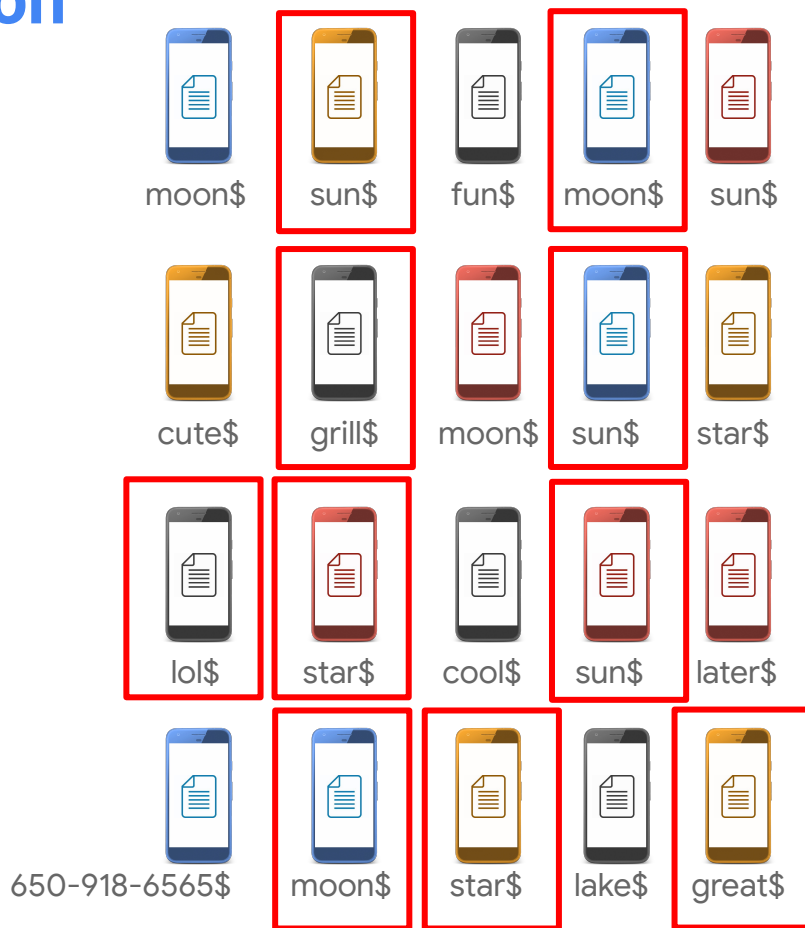

'm': 2 votes

'f', 'g', 'c', 'l', '6': 1 vote each

**Chars with votes  $\geq \theta = 2$  are added to trie**




# Round 2: random device selection



Randomly select  $m = 10$  devices at random

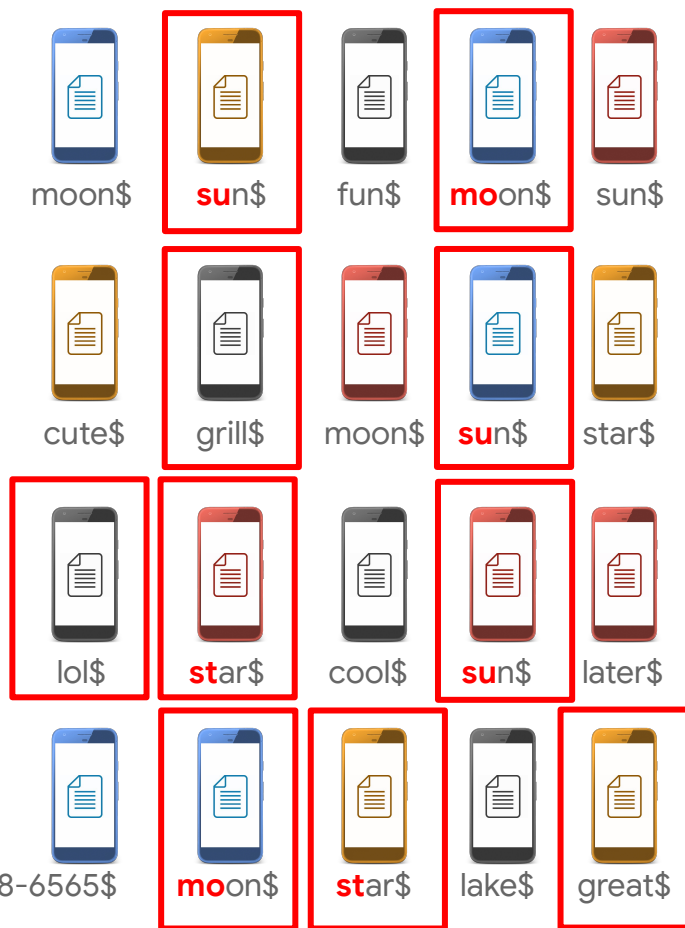
# Round 2: voting stage



'su': 3 votes


'st': 2 votes

'mo': 2 votes



Devices with words having a prefix in trie vote





















# Round 2: vote thresholding stage




'su': 3 votes





















'st': 2 votes

'mo': 2 votes

 moon\$	 sun\$	 fun\$	 moon\$	 sun\$
 cute\$	 grill\$	 moon\$	 sun\$	 star\$
 lol\$	 star\$	 cool\$	 sun\$	 later\$
 650-918-6565\$	 moon\$	 star\$	 lake\$	 great\$





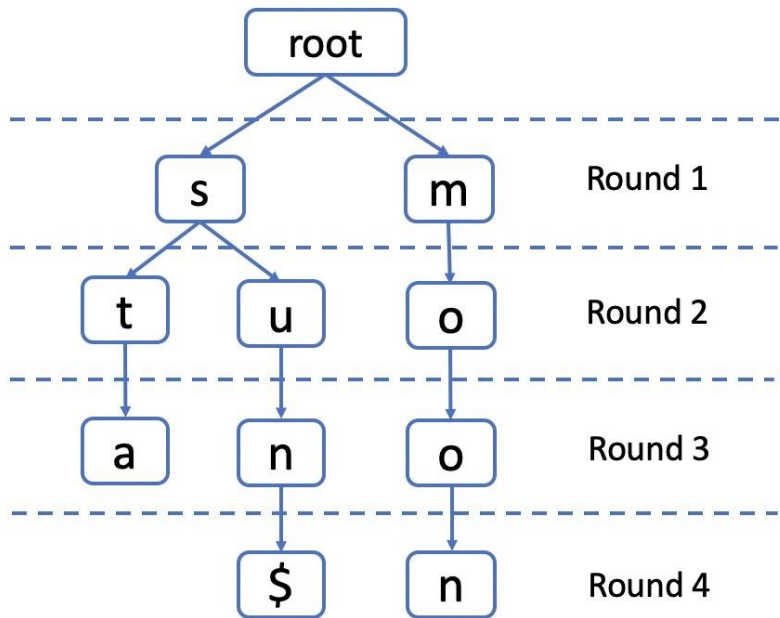
# At the end of round 3



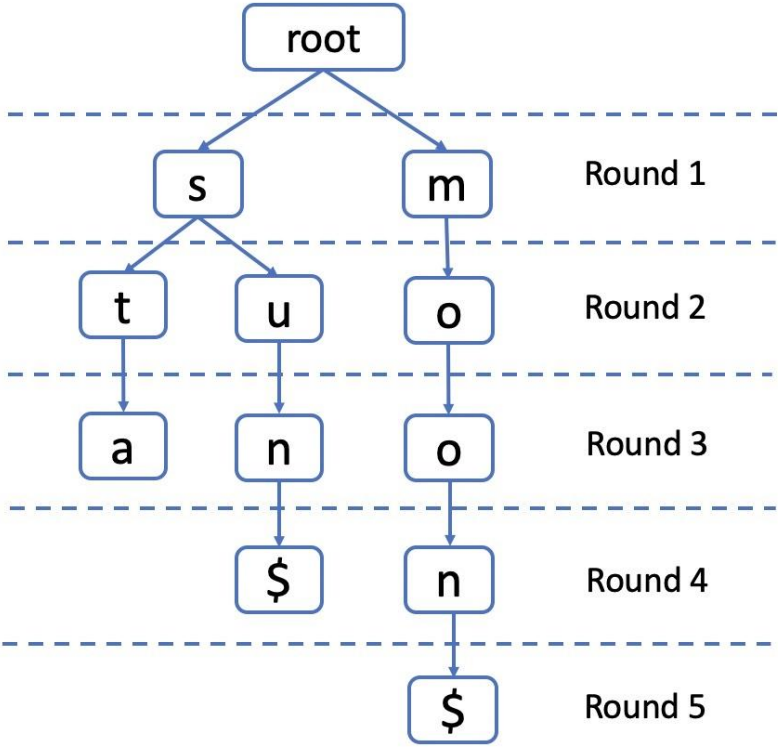
				
moon\$	sun\$	fun\$	moon\$	sun\$
				
cute\$	grill\$	moon\$	sun\$	star\$
				
lol\$	star\$	cool\$	sun\$	later\$
				
650-918-6565\$	moon\$	star\$	lake\$	great\$























# At the end of round 4








# At the end of round 5



				
moon\$	sun\$	fun\$	moon\$	sun\$
				
cute\$	grill\$	moon\$	sun\$	star\$
				
lol\$	star\$	cool\$	sun\$	later\$
				
650-918-6565\$	moon\$	star\$	lake\$	great\$

# TrieHH



## Algorithm

*Input:* List of strings

*Output:* A trie containing popular subsequences

## Key Idea

Interactively build a trie data structure that keeps track of popular prefixes. Aggregate votes on single character extensions to existing prefixes in the trie. Threshold the counts to ensure that you are only keeping track of popular prefixes.

**Paper:** [Federated Heavy Hitters Discovery with Differential Privacy](#) (TPDP19, AISTATS2020)

# Inherent strong privacy guarantees

TrieHH algorithm is differentially private!

- Structural k-anonymity
- User-level (epsilon, delta) central DP
- Great privacy-utility trade-offs
- Limited information exposed to server

Table 1: Choices of  $\theta$  and  $\gamma$  to achieve  $\epsilon = 2$  in two cases:  $\delta \leq \frac{1}{300n}$  and  $\delta \leq \frac{1}{n^2}$ .

$n$	$L = 10$			
	$\delta \leq \frac{1}{300n}$		$\delta \leq \frac{1}{n^2}$	
	$\theta$	$\gamma$	$\theta$	$\gamma$
$10^4$	10	1.81	12	1.51
$10^5$	11	5.21	14	4.09
$10^6$	12	15.10	15	12.08
$10^7$	13	44.09	17	33.71

Zhu, Kairouz *et al.* **Federated Heavy Hitter Discovery with Differential Privacy**. *TPDP19, AISTATS2020*.

# Out-of-vocab<sup>(1)</sup> experiments on Sentiment140

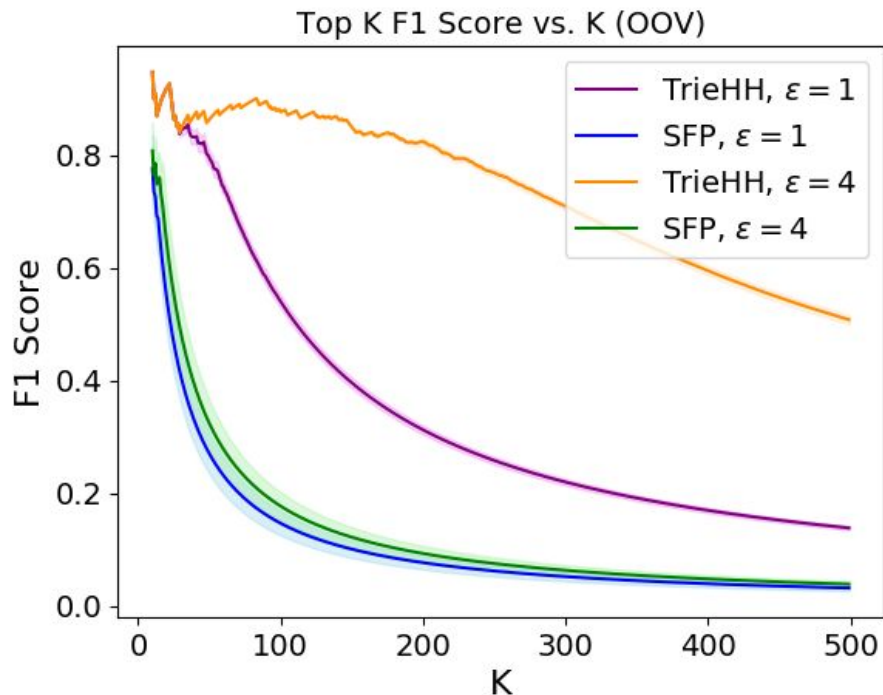


Table 2. Comparison between SFP and TrieHH of recall at  $K = 50$  and  $\delta = 1/n^2$ .

	<sup>(3)</sup> $\epsilon = 1$		$\epsilon = 4$	
	Recall	Prec	Recall	Prec
TrieHH	0.65	1	0.76	1
SFP <sup>(2)</sup> (20)	0.17	0.853	0.19	0.867
SFP (80)	0.25	0.494	0.325	0.456

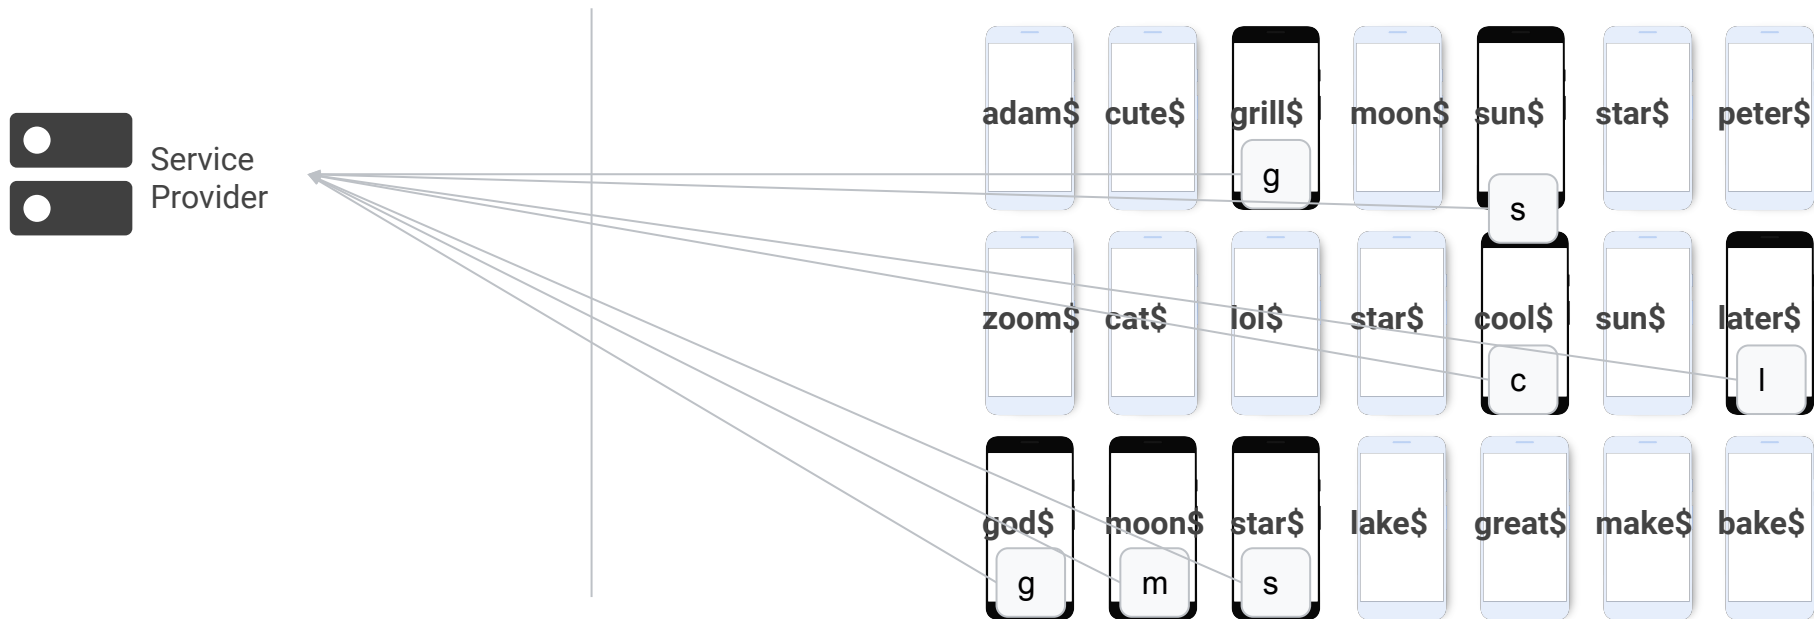
(1) Dictionary contains over 260k words. After removing dictionary words, we lose over 160k users

(2) SFP is an algorithm by Apple for heavy hitter discovery with local DP

(3) SFP's [local epsilon is amplified to a central \(epsilon, delta\)](#) for a fair comparison

# Weaknesses of TrieHH


# Linking devices to per-round character extensions



The server can link contributions (character extensions) to devices


# Learning votes on unpruned edges

**Trie seen by analyst**



versus

**Trie seen by server**



DP & k-anonymity properties hold with respect to the analyst (not the server!)



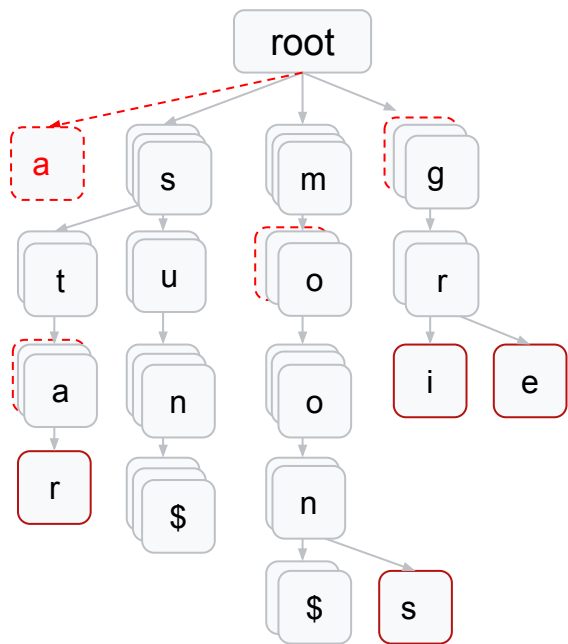
# Uniform random device selection



DP holds only when server can sample uniformly at random from the entire population

# Hardened TrieHH

# Hardened TrieHH



--- represents spurious votes

Google



fun\$



moon\$



sun\$



moon\$



sun\$



star\$



cool\$



sun\$



later\$



650-918-6565\$



star\$



great\$

## Algorithm

Input: List of strings

Output: A trie containing popular subsequences

## Key Idea

Interactively build a trie data structure that keeps track of popular prefixes. Use **SecAgg** to aggregate **noisy votes** on single character extensions to existing prefixes in the trie. Threshold the counts to ensure that you are only keeping track of “popular” prefixes.

## Privacy

Per-round securely aggregated noisy votes are **automatically differentially private**.

## Advances and Open Problems in Federated Learning

Peter Kairouz<sup>7\*</sup> H. Brendan McMahan<sup>7\*</sup> Brendan Avent<sup>21</sup> Aurélien Bellet<sup>9</sup>  
Mehdi Bennis<sup>19</sup> Arjun Nitin Bhagoji<sup>13</sup> Keith Bonawitz<sup>7</sup> Zachary Charles<sup>7</sup>  
Graham Cormode<sup>23</sup> Rachel Cummings<sup>6</sup> Rafael G.L. D'Oliveira<sup>14</sup>  
Salim El Rouayheb<sup>14</sup> David Evans<sup>22</sup> Josh Gardner<sup>24</sup> Zachary Garrett<sup>7</sup>  
Adrià Gascón<sup>7</sup> Badih Ghazi<sup>7</sup> Phillip B. Gibbons<sup>2</sup> Marco Gruteser<sup>7,14</sup>  
Zaid Harchaoui<sup>24</sup> Chaoyang He<sup>21</sup> Lie He<sup>4</sup> Zhouyuan Huo<sup>20</sup>  
Ben Hutchinson<sup>7</sup> Justin Hsu<sup>25</sup> Martin Jaggi<sup>4</sup> Tara Javidi<sup>17</sup> Gauri Joshi<sup>2</sup>  
Mikhail Khodak<sup>2</sup> Jakub Konečný<sup>7</sup> Aleksandra Korolova<sup>21</sup> Farinaz Koushanfar<sup>17</sup>  
Sanmi Koyejo<sup>7,18</sup> Tancrede Lepoint<sup>7</sup> Yang Liu<sup>12</sup> Prateek Mittal<sup>13</sup>  
Mehryar Mohri<sup>7</sup> Richard Nock<sup>1</sup> Ayfer Özgür<sup>15</sup> Rasmus Pagh<sup>7,10</sup>  
Mariana Raykova<sup>7</sup> Hang Qi<sup>7</sup> Daniel Ramage<sup>7</sup> Ramesh Raskar<sup>11</sup>  
Dawn Song<sup>16</sup> Weikang Song<sup>7</sup> Sebastian U. Stich<sup>4</sup> Ziteng Sun<sup>3</sup>  
Ananda Theertha Suresh<sup>7</sup> Florian Tramèr<sup>15</sup> Praneeth Vepakomma<sup>11</sup> Jianyu Wang<sup>2</sup>  
Li Xiong<sup>5</sup> Zheng Xu<sup>7</sup> Qiang Yang<sup>8</sup> Felix X. Yu<sup>7</sup> Han Yu<sup>12</sup> Sen Zhao<sup>7</sup>

<sup>1</sup>Australian National University, <sup>2</sup>Carnegie Mellon University, <sup>3</sup>Cornell University,

<sup>4</sup>École Polytechnique Fédérale de Lausanne, <sup>5</sup>Emory University, <sup>6</sup>Georgia Institute of Technology,

<sup>7</sup>Google Research, <sup>8</sup>Hong Kong University of Science and Technology, <sup>9</sup>INRIA, <sup>10</sup>IT University of Copenhagen,

<sup>11</sup>Massachusetts Institute of Technology, <sup>12</sup>Nanyang Technological University, <sup>13</sup>Princeton University,

<sup>14</sup>Rutgers University, <sup>15</sup>Stanford University, <sup>16</sup>University of California Berkeley,

<sup>17</sup>University of California San Diego, <sup>18</sup>University of Illinois Urbana-Champaign, <sup>19</sup>University of Oulu,

<sup>20</sup>University of Pittsburgh, <sup>21</sup>University of Southern California, <sup>22</sup>University of Virginia,

<sup>23</sup>University of Warwick, <sup>24</sup>University of Washington, <sup>25</sup>University of Wisconsin–Madison

### Abstract

Federated learning (FL) is a machine learning setting where many clients (e.g. mobile devices or whole organizations) collaboratively train a model under the orchestration of a central server (e.g. service provider), while keeping the training data decentralized. FL embodies the principles of focused data collection and minimization, and can mitigate many of the systemic privacy risks and costs resulting from traditional, centralized machine learning and data science approaches. Motivated by the explosive growth in FL research, this paper discusses recent advances and presents an extensive collection of open problems and challenges.

## Advances and Open Problems in FL

58 authors from 25 institutions

[arxiv.org/abs/1912.04977](https://arxiv.org/abs/1912.04977)

