

Practical Secure Aggregation for Privacy-Preserving Machine Learning (Ch4~)

이종진

Seoul National University

ga0408@snu.ac.kr

August 25, 2020

- ▶ Multiparty computation, Federated learning setting.
- ▶ Considering training neural network to predict the next word.
 - Mobile devices, central server.
 - Communication is extremely expensive and user dropouts are common.
 - How to securely aggregate the data from mobile devices.

Cryptographic Primitives

- ▶ Secret Sharing(Shamir's t -out-of- n)
- ▶ Key Agreement(g, g^a, g^b, g^{ab})
- ▶ Authenticated Encryption
- ▶ Pseudorandom Generator(seed, $[0, R)^m$)
- ▶ Signature Scheme (for active adversary.)
- ▶ Public Key Infrastructure.

How to aggregate the inputs

► Notation

- A single server \mathcal{S}
 - n client parties \mathcal{U}
 - A private vector for user $u \in \mathcal{U}$, \mathbf{x}_u , dimension m , in \mathbb{Z}_R ($u \in \{1, \dots, n\}$)
- The protocol can guarantee that the server only learns a sum of the clients' inputs.

Masking with One-Time Pads

- ▶ Mask x_u in a particular way.
- ▶ Suppose each pair of users (u, v) , $u < v$ agree on some random vector $r_{u,v}$

$$y_u = x_u + \sum_{v \in \mathcal{U}: u < v} r_{u,v} - \sum_{v \in \mathcal{U}: u > v} r_{v,u} \pmod{R}$$

- ▶ then,

$$z = \sum y_u = \sum x_u \pmod{R}$$

Masking with One-Time Pads

- ▶ Two shortcomings.
- ▶ The users must exchange the random vectors $\mathbf{r}_{u,v}$
 - (Requires quadratic communication overhead($|\mathcal{U}| \times |\mathbf{x}|$))
- ▶ No tolerance for a party failing to complete the protocol.

Efficient Communication and Handling Dropped Users.

- ▶ Reduce the communication by handling the parties agree on common seed($s_{u,v}$) for PRG rather than on the entire mask $r_{u,v}$
 - Seed may have comparatively low dimension
- ▶ Notify the surviving users of the drop-out, and to have them each reply with the common seed they computed with the dropped users.
 - Additional users may drop out in the recovery phase before replying with the seed.
 - Leading the number of rounds up to at most the number of users.
 - They use a threshold secret sharing scheme.

Double-Masking to Protect Security

- ▶ Each user u distributes shares of $s_{u,v}$ to each of the other users.
- ▶ Secret sharing scheme allows dropped user's seed to be recovered
- ▶ This approach can solve the problem of unbounded recovery rounds, but still has issue.
 - If user u 's device is too slow in sending her y_u to the server.
 - Adversarial server in the active model can learn x_u by lying about whether user u has dropped out.
- ▶ Double masking.

Double-Masking to Protect Security

- ▶ To resolve this new security problem, they introduce doubly masking structure.
- ▶ Each user u samples an additional random seed \mathbf{b}_u and distributes shares of \mathbf{b}_u to each of the other users.

$$\mathbf{y}_u = \mathbf{x}_u + \mathbf{PRG}(\mathbf{b}_u) + \sum_{v \in \mathcal{U}: u < v} \mathbf{PRG}(s_{u,v}) - \sum_{v \in \mathcal{U}: u > v} \mathbf{PRG}(s_{v,u}) \pmod{R}$$

- ▶ Dropped user's $s_{u,v}$ and surviving user's \mathbf{b}_u are needed.

Secure Aggregation Protocol

- ▶ They present a protocol which has a constant number of rounds, low communication overhead, robustness to failures.
- ▶ The protocol consists of 4 rounds.
- ▶ They present two variants of the protocol
 - Secure against honest but curious adversaries.
 - Secure against active adversaries.

Secure Aggregation Protocol, Setup

► Setup

- The protocol is run between a server(S) and a set of n users.
- $\mathbf{x}_u \in \mathbb{Z}_R^m$ is a input vector for user u
- The security parameter k , a threshold value t , and honestly generated $pp \leftarrow \mathbf{KA.gen}(k)$
- The server can communicate with the users with secure channels. (public keys, encoded secret keys)
- All users u receive their signing key d_u^{SK} from the trusted third party, together with verification keys d_v^{PK} bound to each user identity v

Secure Aggregation Protocol, Round 0 (AdvertiseKeys)

► User u :

- Generate key pairs

$$(c_u^{PK}, c_u^{SK}) \leftarrow \mathbf{KA.gen}(pp), (s_u^{PK}, s_u^{SK}) \leftarrow \mathbf{KA.gen}(pp),$$

$$\sigma_u \leftarrow \mathbf{SIG.sign}(d_u^{SK}, c_u^{PK} \| s_u^{PK})$$

- Send $(c_u^{PK} \| s_u^{PK} \| \underline{\sigma_u})$ to the server.

► Server:

- Denote users set \mathcal{U}_1 in this round
- (Assert that $|\mathcal{U}_1| \geq t$)
- Broadcast to all users in \mathcal{U}_1 the list $\{(v, c_v^{PK}, s_v^{PK}, \underline{\sigma_v})\}$

Secure Aggregation Protocol, Round 1 (ShareKeys)

► User u :

- Received the list $\{(v, c_v^{PK}, s_v^{PK}, \sigma_v)\}$
Verify that $\forall v \in \mathcal{U}_1$, $\text{SIG.ver}(d_v^{PK}, c_v^{PK} || s_v^{PK}, \sigma_v) = 1$
- Sample a random element $b_u \leftarrow \mathbb{F}$
- Generate t -out-of- $|\mathcal{U}_1|$ shares of s_u^{SK} and b_u :
 $\{(v, s_{u,v}^{SK})\}_{v \in \mathcal{U}_1} \leftarrow \text{SS.share}(s_u^{SK}, t, \mathcal{U}_1)$,
 $\{(v, b_{u,v})\}_{v \in \mathcal{U}_1} \leftarrow \text{SS.share}(b_u, t, \mathcal{U}_1)$
- Encode secret keys for each other user $v \in \mathcal{U}_1 \setminus \{u\}$, Compute
 $e_{u,v} \leftarrow \text{AE.enc}(\text{KA.agree}(c_u^{SK}, c_v^{PK}), u || v || s_{u,v}^{SK} || b_{u,v})$
- Send all the ciphertexts $\{e_{u,v}\}_{v \in \mathcal{U}_2}$ to the server

► Server:

- Collect lists of ciphertexts from at least t users (denote $\mathcal{U}_2 \subset \mathcal{U}_1$)
- (Assert that $|\mathcal{U}_2| \geq t$)
- Sent to each user $u \in \mathcal{U}_2$ all ciphertexts $\{e_{u,v}\}_{v \in \mathcal{U}_2}$

Secure Aggregation Protocol, Round 2 (MaskedInputCollection)

► User u :

- For each other user $v \in \mathcal{U}_2 \setminus \{u\}$ compute $s_{u,v} \leftarrow \mathbf{KA.agree}(s_u^{SK}, s_v^{PK})$
- Compute $\mathbf{p}_{u,v}$

$$\mathbf{p}_{u,v} = \begin{cases} \mathbf{PRG}(s_{u,v}), & \text{when, } u > v \\ \mathbf{PRG}(s_{u,v}), & \text{when, } u < v \\ \mathbf{PRG}(s_{u,v}), & \text{when, } u = v \end{cases}$$

- Compute $\mathbf{p}_u = \mathbf{PRG}(b_u)$
- Compute $\mathbf{y}_u = \mathbf{x}_u + \mathbf{p}_u + \sum_{v \in \mathcal{U}_2} \mathbf{p}_{u,v}$

► Server:

- Collect \mathbf{y}_u from at least t users. (denote $\mathcal{U}_3 \subset \mathcal{U}_2$)

Secure Aggregation Protocol, Round 3 (ConsistencyCheck)

▶ **User u :**

- Send to the server $\sigma'_u \leftarrow \mathbf{SIG.sign}(d_u^{SK}, \mathcal{U}_3)$

▶ **Server:**

- Collect σ'_u from at least t users (denote $\mathcal{U}_4 \subset \mathcal{U}_3$),
Send to each user in \mathcal{U}_4 the set $\{v, \sigma'_v\}_{v \in \mathcal{U}_4}$

Secure Aggregation Protocol, Round 4 (Unmasking)

► User u :

- $SIG.ver(d^{PK}, \mathcal{U}_3, \sigma'_v) = 1$ for all $v \in \mathcal{U}_4$
- Assert that $u = u'$ and $v = v'$
- Decrypt the ciphertext
 $v' || u' || s_{u,v}^{SK} || b_{v,u} \leftarrow AE.dec(KA.agree(c_u^{PK}, c_v^{PK}), e_{v,u})$
- Send a list of shares to the server, which consists of $s_{u,v}^{SK}$ for users $v \in \mathcal{U}_2 \setminus \mathcal{U}_3$ and $b_{v,u}$ for users in $v \in \mathcal{U}_3$

► Server:

- Collect responses from at least t users (denote with \mathcal{U}_5 this set of users).
- For $u \in \mathcal{U}_2 \setminus \mathcal{U}_3$, reconstruct $s_{v,u} \leftarrow SS.recon(\{s_{u,v}^{SK}\}_{v \in \mathcal{U}_5}, t)$ for all $v \in \mathcal{U}_3$.
- Compute $p_{v,u}$ for all $v \in \mathcal{U}_2 \setminus \mathcal{U}_3$.
- For $u \in \mathcal{U}_3$, reconstruct $b_u \leftarrow SS.recon(\{b_{u,v}\}_{v \in \mathcal{U}_5}, t)$ and then recompute p_u using the PRG.
- Compute and output $z = \sum_{u \in \mathcal{U}_3} x_u$ as
$$\sum_{u \in \mathcal{U}_3} x_u = \sum_{u \in \mathcal{U}_3} y_u - \sum_{u \in \mathcal{U}_3} p_u + \sum_{u \in \mathcal{U}_3, v \in \mathcal{U}_2 \setminus \mathcal{U}_3} p_{v,u}$$

The end.