

Federated Linear Bandits with Finite Adversarial Actions

Li Fan¹, Ruida Zhou², Chao Tian², Cong Shen¹

¹Electrical & Computer Engineering, University of Virginia

²Electrical & Computer Engineering, Texas A&M University

November 14, 2023

Agenda

- Problem Formulation
- Preliminaries
- FedSupLinUCB Framework
- Algorithms
- Experiments
- Conclusion

Problem Formulation

- **Star-shape system:** M clients communicate with the central server.
- At round t , some clients $I_t \in [M]$ activate and receive K arms associated with contexts $\{x_{t,a}^i\}_{a \in [K]}$.
- Client $i \in I_t$ pulls arm a_t^i and receives reward $r_{t,a_t^i} = \theta^\top x_{t,a_t^i}^i + \epsilon_t$, where $\theta \in \mathbb{R}^d$ is unknown parameter, and ϵ_t is sub-Gaussian noise.
- **Two arrival patterns:** 1) Synchronous: all M clients are active at each round. 2) Asynchronous: one client is active at each round.
- **Total regret:** $R_T = \sum_{i=1}^M R_T^i = \sum_{i=1}^M \mathbb{E} \left[\sum_{t \in P_T^i} r_{t,a_t^i}^i - r_{t,a_t^i}^i \right]$, where P_T^i is the set of time indices at which client i is active.
- **Communication cost:** total number of communication rounds between clients and the server.

Preliminaries

Algorithm 1 Sync(s , server, client 1, ... client n)

- 1: **for** $i = 1, 2, \dots, n$ **do** ▷ Client-side local information upload
- 2: Client i sends the local new layer s information $(\Delta A_s^i, \Delta b_s^i)$ to the server
- 3: **end for**
- 4: Update server's layer s information: ▷ Server-side information aggregation and distribution

$$A_s^{ser} \leftarrow A_s^{ser} + \sum_{i=1}^n \Delta A_s^i, \quad b_s^{ser} \leftarrow b_s^{ser} + \sum_{i=1}^n \Delta b_s^i$$

- 5: Send server information A_s^{ser}, b_s^{ser} back to all clients
 - 6: **for** $i = 1, 2, \dots, n$ **do**
 - 7: $A_s^i \leftarrow A_s^{ser}, b_s^i \leftarrow b_s^{ser}, \Delta A_s^i \leftarrow 0, \Delta b_s^i \leftarrow 0$ ▷ Client i updates the local information
 - 8: **end for**
-

- **Information encoding**: matrix $A_n = \sum_{t=1}^n x_t x_t^\top$ and vector $b_n = \sum_{t=1}^n r_t x_t$, encoding function: $A_n, b_n \leftarrow \text{Encode}(\{x_t, r_t\}_{t=1}^n)$.
- **Communication criterion**: 'doubling trick' introduced in [Abbasi-Yadkori et al., 2011] and the protocol by [Wang et al., 2019, He et al., 2022].
- **Synchronization procedure**: clients upload the local information, and the server aggregates and distributes updated information.

FedSupLinUCB Framework

Algorithm 2 S-LUCB

1: **Initialization:** $S = \lceil \log d \rceil$, $\bar{w}_0 = d^{1.5} / \sqrt{T}$, $\bar{w}_s \leftarrow 2^{-s} \bar{w}_0, \forall s \in [1 : S]$.
2: $\alpha_0 = 1 + \sqrt{d \ln(2M^2T/\delta)}$, $\alpha_s \leftarrow 1 + \sqrt{2 \ln(2KMT \ln d/\delta)}$, $\forall s \in [1 : S]$
3: **Input:** Client i (with local information $A^i, b^i, \Delta A^i, \Delta b^i$), contexts set $\{x_{t,1}^i, \dots, x_{t,K}^i\}$
4: $A_{t,s}^i \leftarrow A_s^i + \Delta A_s^i, b_{t,s}^i \leftarrow b_s^i + \Delta b_s^i$ or $A_{t,s}^i \leftarrow A_s^i, b_{t,s}^i \leftarrow b_s^i$ for lazy update
5: $\hat{\theta}_s \leftarrow (A_{t,s}^i)^{-1} b_{t,s}^i, \hat{r}_{t,s,a}^i = \hat{\theta}_s^\top x_{t,s,a}^i, w_{t,s,a}^i \leftarrow \alpha_s \|x_{t,s,a}^i\|_{(A_{t,s}^i)^{-1}}, \forall s \in [0 : S], \forall a \in [K]$
6: $s \leftarrow 0; \mathcal{A}_0 \leftarrow \{a \in [K] \mid \hat{r}_{t,0,a}^i + w_{t,0,a}^i \geq \max_{a' \in [K]} (\hat{r}_{t,0,a'}^i - w_{t,0,a'}^i)\}$ \triangleright Initial screening
7: **repeat** \triangleright Layered successive screening
8: **if** $s = S$ **then**
9: Choose action a_t^i arbitrarily from \mathcal{A}_S
10: **else if** $w_{t,s,a}^i \leq \bar{w}_s$ for all $a \in \mathcal{A}_s$ **then**
11: $\mathcal{A}_{s+1} \leftarrow \{a \in \mathcal{A}_s \mid \hat{r}_{t,s,a}^i \geq \max_{a' \in \mathcal{A}_s} (\hat{r}_{t,s,a'}^i) - 2\bar{w}_s\}; s \leftarrow s + 1$
12: **else**
13: $a_t^i \leftarrow \arg \max_{\{a \in \mathcal{A}_s, w_{t,s,a}^i > \bar{w}_s\}} w_{t,s,a}^i$
14: **end if**
15: **until** action a_t^i is found
16: Take action a_t^i and receive reward $r_{t,a_t^i}^i$
17: $\Delta A_s^i \leftarrow \Delta A_s^i + x_{t,a_t^i}^i x_{t,a_t^i}^{i\top}, \Delta b_s^i \leftarrow \Delta b_s^i + r_{t,a_t^i}^i x_{t,a_t^i}^i$ \triangleright Update local information
18: **Return** layer index s

- Combination of the principles of SupLinUCB [Chu et al., 2011] and OFUL [Abbasi-Yadkori et al., 2011, Ruan et al., 2021].
- Maintain $S = \lceil \log d \rceil$ information layers.
- Accuracy strengthens as the layer index increases.

Asynchronous FedSupLinUCB

Algorithm 3 Async-FedSupLinUCB

```
1: Initialization:  $T, C, S = \lceil \log d \rceil$ 
2:  $\{A_s^{ser} \leftarrow I_d, b_s^{ser} \leftarrow 0 \mid s \in [0 : S]\}$  ▷ Server initialization
3:  $\{A_s^i \leftarrow I_d, \Delta A_s^i, b_s^i, \Delta b_s^i \leftarrow 0 \mid s \in [0 : S], i \in [M]\}$  ▷ Clients initialization
4: for  $t = 1, 2, \dots, T$  do
5:   Client  $i_t = i$  is active, and observes  $K$  contexts  $\{x_{t,1}^i, x_{t,2}^i, \dots, x_{t,K}^i\}$ 
6:    $s \leftarrow$  S-LUCB (client  $i, \{x_{t,1}^i, x_{t,2}^i, \dots, x_{t,K}^i\}$ ) with lazy update
7:   if  $\frac{\det(A_s^i + \Delta A_s^i)}{\det(A_s^i)} > (1 + C)$  then
8:     Sync( $s$ , server, clients  $i$ ) for each  $s \in [0 : S]$ 
9:   end if
10: end for
```

- Only one client is active in each round.
- Global synchronization and coordination are not required

Synchronous FedSupLinUCB

Algorithm 4 Sync-FedSupLinUCB

```
1: Initialization:  $T_c, D, S = \lceil \log d \rceil, t_{last}^s \leftarrow 0, \forall s \in [0 : S], \text{CommLayers} \leftarrow \emptyset.$   
2:  $\{A_s^{ser} \leftarrow I_d, b_s^{ser} \leftarrow 0 \mid s \in [0 : S]\}$  ▷ Server initialization  
3:  $\{A_s^i \leftarrow I_d, \Delta A_s^i, b_s^i, \Delta b_s^i \leftarrow 0 \mid s \in [0 : S], i \in [M]\}$  ▷ Clients initialization  
4: for  $t = 1, 2, \dots, T_c$  do  
5:   for  $i = 1, 2, \dots, M$  do  
6:     Client  $i_t = i$  is active, and observes  $K$  contexts  $\{x_{t,1}^i, x_{t,2}^i, \dots, x_{t,K}^i\}$   
7:      $s \leftarrow \text{S-LUCB}(\text{client } i, \{x_{t,1}^i, x_{t,2}^i, \dots, x_{t,K}^i\})$   
8:     if  $(t - t_{last}^s) \log \frac{\det(A_s^i + \Delta A_s^i)}{\det(A_s^i)} > D$  then  
9:       Add  $s$  to CommLayers  
10:    end if  
11:  end for  
12: end for  
13: for  $s \in \text{CommLayers}$  do  
14:   Sync( $s$ , server, clients  $[M]$ );  $t_{last}^s \leftarrow t, \text{CommLayers} \leftarrow \emptyset$   
15: end for
```

- All clients are active and make decisions at each round.
- Refined communication criterion to enable time-independent communication cost

Performance

Theorem

For any $0 < \delta < 1$, if we run Async with $C = 1/M^2$, then with probability at least $1 - \delta$, the regret of the algorithm is bounded as $R_T \leq \tilde{O}\left(\sqrt{d \sum_{i=1}^M T_i}\right) = \tilde{O}\left(\sqrt{dT}\right)$.


Moreover, the corresponding communication cost is bounded by $O(dM^2 \log d \log T)$.

Theorem


For any $0 < \delta < 1$, if we run Sync with $D = \frac{T_c \log T_c}{d^2 M}$, with probability at least $1 - \delta$, the regret of the algorithm is bounded as $R_T \leq \tilde{O}(\sqrt{dMT_c})$ where T_c is the total per-client arm pulls. Moreover, the corresponding communication cost is bounded by $O(\sqrt{d^3 M^3 \log d})$.

- Both achieve **minimax optimal** regret with **sublinear** communication cost.


Experiments




(a) Regret: arrival patterns.



(b) Communication: arrival patterns.



(c) Regret: client numbers.





(d) Regret vs communications.


- Compare with different arrival patterns and amount of clients.
- The trade-off between regrets and communications.


Thank you!


References I

 Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011).
Improved algorithms for linear stochastic bandits.
Advances in neural information processing systems, 24.

 Chu, W., Li, L., Reyzin, L., and Schapire, R. (2011).
Contextual bandits with linear payoff functions.
In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, pages 208–214. JMLR Workshop and Conference Proceedings.

 He, J., Wang, T., Min, Y., and Gu, Q. (2022).
A simple and provably efficient algorithm for asynchronous federated contextual linear bandits.
arXiv preprint arXiv:2207.03106.

 Ruan, Y., Yang, J., and Zhou, Y. (2021).
Linear bandits with limited adaptivity and learning distributional optimal design.
In Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, pages 74–87.

 Wang, Y., Hu, J., Chen, X., and Wang, L. (2019).
Distributed bandit learning: Near-optimal regret with efficient communication.
In International Conference on Learning Representations.