Workshop on Federated Learning for Natural Language Processing (FL4NLP 2022)

Bill Yuchen Lin[1] Chaoyang He[1] Tian Li[2] Ninareh Mehrabi[1]
Fatemehsadat Mireshghallah[3] Chulin Xie[4] Mahdi Soltanolkotabi [1] Xiang Ren [1]
[1] University of Southern California [2] Carnegie Mellon University
[3] University of California, San Diego [4] University of Illinois at Urbana-Champaign
{yuchen.lin, chaoyang.he, ninarehm, soltanol, xiangren}@usc.edu
tianli@cmu.edu, fmireshg@eng.ucsd.edu, chulinx2@illinois.edu

1 Introduction

Background. Due to increasing concerns and regulations about data privacy (e.g., General Data Protection Regulation), coupled with the growing computational power of edge devices, emerging data from realistic users have become much more fragmented, forming distributed private datasets across different clients (i.e., organizations or personal devices). Respecting users' privacy and restricted by these regulations, we have to assume that users' data in a client are not allowed to transfer to a centralized server or other clients. For example, a hospital does not want to share its private data (e.g., conversations, questions asked on its website/app) with other hospitals. This is despite the fact that models trained by a centralized dataset (i.e., combining data from all clients) usually enjoy better performance on downstream tasks (e.g., dialogue, question answering). Therefore, it is of vital importance to study NLP problems in such a scenario, where data are *distributed* across different isolated organizations or remote devices and cannot be shared for privacy concerns.

The field of **federated learning** (FL) aims to enable many individual clients to jointly train their models, while keeping their local data *decentralized* and completely *private* from other users or a centralized server. A common training schema of FL methods is that each client sends its model parameters to the server, which updates and sends back the global model to all clients in each round. Since the raw data of one client has never been exposed to others, FL is promising to be an effective way to address the above challenges, particularly in the NLP domain where many user-generated text data contain sensitive, personal information.

Recent advances in FL for NLP. There are an increasing number of papers starting to apply FL methods in NLP models. For example, federated learning has been applied to many keyboard-related applications (Hard et al., 2018; Stremmel and Singh, 2020; Leroy et al., 2019; Ramaswamy et al., 2019; Yang et al., 2018; Thakkar et al., 2021), sentence-level text intent classification using Text-CNN (Zhu et al., 2020), and pretraining and fine tuning of BERT using medical data from multiple silos without gathering all the data to the same place (Liu and Miller, 2020). FL methods also have been proposed to train high quality language models that can outperform the the models trained without federated learning (Ji et al., 2019; Chen et al., 2019). Besides these applications, some work has been done in medical relation extractions (Ge et al., 2020), medical name entity recognition (Sui et al., 2020; Jana and Biemann, 2021) and spoken language understanding (Huang et al., 2020). These methods use federated learning to preserve privacy of sensitive medical data and learn data in different platforms excluding the need of exchanging data between different platforms.

Topics of Interest. We organize this workshop to encourage discussion of current progress on federated learning research in the context of NLP tasks and models. We aim to bring together researchers from different areas (e.g., federated learning, privacy-preserving ML, distributed computing, etc.) to communicate and provide promising working directions in FL for NLP. Topics of interest include, but are not limited to:

- Federated learning methods for NLP tasks and models (e.g., Transformer-based LMs, dialog systems, etc).
- New learning frameworks to tackle data heterogeneity, label deficiency, data shift, generalization ability related issues in FL for NLP, including continual learning, online learning, multi-task learning, self/semi/un-supervised learning, etc.
- Efficient training methods for resource-constrained on-device NLP, including training-time compression, communication/computation/memory-efficient methods.
- Security and privacy for FL for NLP, including new attack methods (e.g., data and model poisoning) and defense methods (e.g., empirical and certifiable defenses), robust aggregation methods, differential privacy (DP), HE (Homomorphic Encryption), etc.
- Fair FL for NLP, including introducing different fairness notions in FL, mitigating different types of biases in different NLP applications in FL settings, introducing benchmark datasets and tasks for fair FL in NLP applications along with auditing different NLP applications in FL settings.
- Interpretability of FL for NLP, especially understanding how NLP models work in data heterogeneity.
- Scalability of FL4NLP: e.g., client sampling algorithms.
- Network typologies beyond parameter server: e.g., vertical, decentralized, and hierarchical FL for NLP.
- Benchmarking datasets (with realistic non-I.I.D. partitions) and new applications in NLP and beyond.

- Systems, including user-friendly, efficient, scalable, reproducible, flexible distributed training systems, on-device training engine tailored for NLP tasks and models.

2 Workshop Description

2.1 Keynote Speakers

We invited and **confirmed** 6 keynote speakers:

Salman Avestimehr, Professor, USC
Salman Avestimehr is a Dean's Professor of ECE and CS at the University of Southern California (USC) and the inaugural director of the USC-Amazon Center on Secure and Trusted Machine Learning (Trusted AI). He is also an IEEE Fellow and an Amazon Scholar at Alexa AI. His group is particularly active in studying security/privacy and scalability aspects of FL, and more recently in applications of FL to NLP.

Virginia Smith, Assistant Professor, CMU
Her research interests span machine learning, optimization, and distributed systems. Her work has been recognized via a Facebook Faculty Award, Google Faculty Awards, MLconf Industry Impact Award, and MIT TR35 Innovator Award. Virginia co-founded MLSys, a conference that brings together machine learning and systems researchers.

Bo Li, Assistant Professor, UIUC
Recipient of the Symantec Research Labs Fellowship, Rising Stars, MIT Technology Review TR-35 award, Intel Rising Star award, Amazon Research Award. Her research focuses on both theoretical and practical aspects of security, machine learning, privacy, game theory, and adversarial machine learning.

Tong Zhang, Professor, HKUST
Tong Zhang is a professor of Computer Science and Mathematics at the Hong Kong University of Science and Technology (HKUST). Previously, he was a professor at Rutgers university, and worked at IBM, Yahoo, Baidu, and Tencent. He is a fellow of ASA, IEEE, and IMS, and he has been in the editorial boards of ICML, NeurIPS, JMLR, etc. Tong has also published intensively in NLP-related conferences such as ACL/EMNLP since 2001.

Manzil Zaheer, Research Scientist, Google
Ph.D. at CMU, under the able guidance of Prof Barnabas Poczos, Prof Ruslan Salakhutdinov, and Prof Alexander Smola. He is the winner of Oracle Fellowship in 2015.

Rahul Gupta, Applied Science Manager, Alexa AI
He focuses on Trustworthy Alexa technologies. His research is geared towards making Alexa more private and equitable for the customer. His recent research has been published in venues such as ACL, EMNLP, Interspeech and ICASSP.

2.2 Panel Discussion (pending confirmations)

We plan to have a panel discussion with the invited speakers as well as the following invited panelists: **Anna Rumshisky** (Associate Prof. at Umass), **Lin Xiao** (Facebook AI Research), **Gauri Joshi** (Assistant Prof. at CMU), **Tom Diethe** (Amazon Research), and **Xu Zheng** (Google FL team).

2.3 Lightning Talks

Accepted papers will be presented as lightning talks at three different points in the workshop schedule. There will also be poster sessions (virtually or in-person).

2.4 Shared Tasks

To encourage the development of this area, we also host a shared task based on an existing open-source platform, FedNLP[1], where there are unified interfaces for developing novel methods easily. The shared task will focus on benchmarking novel FL methods for a wide range of NLP task formulations, including text classification, sequence tagging, reading comprehension, seq2seq generation, etc. We will host an online tested for assessing all submissions in multiple evaluation metrics.

2.5 Submission Formats

We welcome submissions of the following types:

- Long papers (8 pages + references) or short papers (4 pages + references), describing new, previously unpublished research in this field.
- Extended abstracts (posters) describing preliminary work which would benefit from exposure but is not ready for publication. (non-archival)
- Previously published papers at other conferences (e.g., ICLR, NeurIPS, and *CL Findings) on topics relevant to the workshop theme (non-archival).
- System reports of the submissions on our shared task (8 pages + references).

2.6 Logistics

Attendees We expect the workshop to attract researchers working on federated learning, privacy, trustworthy AI for NLP applications. Based on the success of similar workshops, such as SpicyFL@NeurIPS, DPML@ICLR, and FL@IJCAI, we expect around 70∼100 attendees.

Virtual Option We aim to develop a program that is accessible to both in-person and virtual audiences. Consequently, our workshop program will consist primarily of invited talks and lightning talks for accepted papers. Offline, we maintain a community Slack workspace that all attendees are invited to join where they can converse with paper authors in dedicated channels.

Preferred Venues: ACL, NAACL, EMNLP.

Diversity & Inclusion We have organized our program (invited talks; lightning talks) to be widely accessibly to both on-site and virtual attendees. Contents will be recorded to be available to those who could not attend. As for the composition of our organizing teams and program, nearly 50% of our organizing committee, speakers, and panelists are women. Our speakers are from both universities and industry. Finally, our invited speakers work in diverse areas to bring in novel viewpoints to the problems we tackle in FL for NLP.

[1] https://github.com/FedML-AI/FedNLP/

3 Organizers

(Bill) Yuchen Lin (USC) is a Ph.D. candidate in Computer Science at the University of Southern California, working with Prof. Xiang Ren at the Intelligence and Knowledge Discovery Research Lab (USC-INK). His research aims to integrate information extraction, knowledge graphs, logical reasoning, graph neural networks, explanations, robustness, etc. Apart from that, he is also interested in cross-task generalization (i.e., meta-learning, continual learning) and federated learning in the NLP domain. He mainly publishes (as Bill Yuchen Lin) and serves as Program Committee members for ACL, EMNLP, NAACL, NeurIPS, ICLR, AAAI, etc. He got Best Paper Award at TrustNLP 2021, Best Paper Runner-Up at WWW 2020, and he was selected as one of the AI Rising Stars in Chinese Students by Baidu Research. He was a co-leader of the first open-sourced FedNLP platform.

Chaoyang He (USC) is a Ph.D. Candidate in the CS department at the University of Southern California, Los Angeles, USA. He is advised by Salman Avestimehr (USC), Professor Mahdi Soltanolkotabi (USC), Professor Murali Annavaram (USC), and Professor Tong Zhang (HKUST). He also works closely with researchers/engineers at Google, Facebook, Amazon, and Tencent. Previously, He was an RD Team Manager and Staff Software Engineer at Tencent (2014-2018), a Team Leader and Senior Software Engineer at Baidu (2012-2014), and a Software Engineer at Huawei (2011-2012). His research focuses on distributed/federated machine learning algorithms, systems, and applications. He was a co-leader of the first open-sourced FedNLP platform.

Tian Li (CMU) is a Ph.D. student in the Computer Science Department at Carnegie Mellon University working with Virginia Smith. Her research interests are in distributed optimization, federated learning, and data-intensive systems. Prior to CMU, she received her undergraduate degrees in Computer Science and Economics from Peking University.

Fatemehsadat Mireshghallah (UCSD) Fatemehsadat Mireshghallah is a Ph.D. student at the CSE department of UC San Diego. Her research interests are Trustworthy Machine Learning and Natural Language Processing. She received her B.S. from Sharif university of technology in Iran. She is a recipient of the National Center for Women IT (NCWIT) Collegiate award in 2020 for her work on privacy-preserving inference, and a finalist of the Qualcomm Innovation Fellowship in 2021. She has interned twice at Microsoft Research's Language and Intelligent Assistance group, where she worked on private training of large language models. She is also serving as a NAACL 2022 D&I co-chair and WiNLP committee member.

Ninareh Mehrabi (USC-ISI) is a Ph.D. candidate at University of Southern California's Information Sciences Institute. Her research is on developing trustworthy AI systems with an emphasis on algorithmic fairness in Machine Learning and Natural Language Processing. She received her B.Sc. degree in Computer Science and Engineering from University of Southern California.

Chulin Xie (UIUC) is a PhD student in the CS Department at University of Illinois at Urbana-Champaign, advised by Bo Li. Her research interests include machine learning, adversarial robustness, privacy, and federated learning. She received her Bachelor degree in the CS Department from Zhejiang University.

Xiang Ren (USC) is an assistant professor at the USC Computer Science Department, a Research Team Leader at USC ISI, and the PI of the Intelligence and Knowledge Discovery (INK) Lab at USC. Priorly, he spent time as a research scholar at the Stanford NLP group and received his Ph.D. in Computer Science from the University of Illinois Urbana-Champaign. Dr. Ren works on knowledge acquisition and reasoning in natural language processing, with focuses on developing human-centered and label-efficient computational methods for building trustworthy NLP systems. He received NSF CAREER Award, The Web Conference Best Paper runner-up, ACM SIGKDD Doctoral Dissertation Award, and several research awards from Google, Amazon, JP Morgan, Sony, and Snapchat. He was named Forbes' Asia 30 Under 30 in 2019.

Mahdi Soltanolkotabi (USC) is an associate professor in the Ming Hsieh Department of Electrical and Computer Engineering and Computer Science at the University of Southern California where he holds an Andrew and Erna Viterbi Early Career Chair. Prior to joining USC he spent a year as a postdoc in the AMPLAB at UC Berkeley mentored by Ben Recht and Martin Wainwright. He obtained his Ph.D. in Electrical Engineering from Stanford in 2014 advised by Emmanuel Candes. His research focuses on developing the mathematical foundations of learning from signals and data spanning optimization, machine learning, signal processing, high dimensional probability/statistics, computational imaging and artificial intelligence. Recently, he has also contributed extensively to federated/distributed machine learning for NLP from both a theorectical and application perspective. His team also helped in the development of https://FedML.ai and https://DistML.ai

4 Tentative Schedule

Start	End	Event
8:00	8:10	Opening remarks
8:10	8:55	Invited Speaker # 1
8:55	9:40	Invited Speaker # 2
9:40	10:00	Break
10:00	10:45	Invited Speaker # 3
10:45	11:00	Lightning talks # 1
11:00	12:00	Panel Discussion
12:00	13:30	Lunch Break
13:30	14:15	Invited Speaker # 4
14:15	14:30	Lightning talks # 2
14:30	15:15	Invited Speaker # 5
15:15	15:45	Break
15:45	16:00	Invited Speaker # 6
16:00	16:45	Shared Task Results
16:45	17:00	Closing Statements

5 Tentative PC (pending confirmations)

- Peter Kairouz, Google
- Zheng Xu, Google FL team
- Jakub Konečný, Google FL team
- Kshitiz Malik, Facebook
- Hongyuan Zhan, Facebook
- Shen Li, Facebook
- Anit Kumar, Amazon
- Anit Kumar Sahu, Alexa AI
- Kevin Hsieh, Microsoft
- Ali Anwar, IBM

- Aston Zhang, AWS AI
- Shuai Zheng, AWS AI
- Zha Sheng, AWS AI
- Bill Yuchen Lin, USC
- Jun Yan, USC
- Krishna Pillutla, UW
- Jie Ding, UMN
- Gerald Penn, UofT
- Basak Guler, UCR
- Xin Dong, Harvard
- Sijie Cheng, FDU
- Mitra Bokaei Hosseini, St. Mary's University
- Shengyuan Hu, CMU
- Tao Yu, Cornell
- Eugene Bagdasaryan, Cornell
- Arjun Nitin Bhagoji, Princeton
- Krishna Pillutla, UW
- Saurav Prakash, USC
- Jinhyun So, USC
- Ziteng Sun, Cornell
- Umang Gupta, USC-ISI
- Bahare Harandi, USC-ISI
- Pei Zhou, USC-ISI
- Yae Jee Cho, CMU
- Gauri Joshi, CMU
- Jianyu Wang, CMU
- Hongyi Wang, CMU
- Sunwoo Lee, USC
- Oscar Li, CMU
- Sai Praneeth Karimireddy, Berkeley
- Alireza Fallah, MIT
- Nic Lane, University of Cambridge
- Dimitris Stripelis, USC-ISI

References

Mingqing Chen, Ananda Theertha Suresh, Rajiv Mathews, Adeline Wong, Cyril Allauzen, Françoise Beaufays, and Michael Riley. 2019. Federated learning of n-gram language models. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*.

Suyu Ge, Fangzhao Wu, Chuhan Wu, Tao Qi, Yongfeng Huang, and X. Xie. 2020. Fedner: Privacy-preserving medical named entity recognition with federated learning. *ArXiv preprint*.

Andrew Hard, K. Rao, Rajiv Mathews, F. Beaufays, S. Augenstein, Hubert Eichner, Chloé Kiddon, and D. Ramage. 2018. Federated learning for mobile keyboard prediction. *ArXiv preprint*.

Zhiqi Huang, Fenglin Liu, and Yuexian Zou. 2020. Federated learning for spoken language understanding. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3467–3478, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Abhik Jana and Chris Biemann. 2021. An investigation towards differentially private sequence tagging in a federated framework. In *Proceedings of the Third Workshop on Privacy in Natural Language Processing*, pages 30–35, Online. Association for Computational Linguistics.

Shaoxiong Ji, Shirui Pan, Guodong Long, Xue Li, Jing Jiang, and Zi Huang. 2019. Learning private neural language modeling with attentive aggregation. *2019 International Joint Conference on Neural Networks (IJCNN)*.

David Leroy, Alice Coucke, Thibaut Lavril, Thibault Gisselbrecht, and Joseph Dureau. 2019. Federated learning for keyword spotting. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2019, Brighton, United Kingdom, May 12-17, 2019*.

D. Liu and T. Miller. 2020. Federated pretraining and fine tuning of bert using clinical notes from multiple silos. *ArXiv preprint*.

Swaroop Indra Ramaswamy, Rajiv Mathews, K. Rao, and Franccoise Beaufays. 2019. Federated learning for emoji prediction in a mobile keyboard. *ArXiv preprint*.

Joel Stremmel and Arjun Singh. 2020. Pretraining federated text models for next word prediction. *ArXiv preprint*.

Dianbo Sui, Yubo Chen, Jun Zhao, Yantao Jia, Yuantao Xie, and Weijian Sun. 2020. FedED: Federated learning via ensemble distillation for medical relation extraction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Om Dipakbhai Thakkar, Swaroop Ramaswamy, Rajiv Mathews, and Francoise Beaufays. 2021. Understanding unintended memorization in language models under federated learning. In *Proceedings of the Third Workshop on Privacy in Natural Language Processing*, pages 1–10, Online. Association for Computational Linguistics.

T. Yang, G. Andrew, Hubert Eichner, Haicheng Sun, W. Li, Nicholas Kong, D. Ramage, and F. Beaufays. 2018. Applied federated learning: Improving google keyboard query suggestions. *ArXiv preprint*.

Xinghua Zhu, Jianzong Wang, Zhenhou Hong, and Jing Xiao. 2020. Empirical studies of institutional federated learning for natural language processing. In *Findings of the Association for Computational Linguistics: EMNLP 2020*.