Statement of Purpose **UC Berkeley**
Ken Ziyu Liu (kenziyuliu.github.io, kenziyuliu@gmail.com) PhD Program in Computer Science

Research Interests & Motivation

I am an MS student at Carnegie Mellon University working with Prof. **Virginia Smith**, **Artur Dubrawski**, and **Steven Wu**. My research focuses on **privacy-preserving machine learning** with an emphasis on **federated learning** (FL) and **differential privacy** (DP); more broadly, I'm interested in both developing and understanding ML methods and systems that can be easily and responsibly deployed in the real world. These interests have drawn me to two broad directions I'd like to explore at Berkeley:

1. **Evolving notions of ML privacy**: What notions of (differential) privacy are suitable for what scenarios, and when can they be relaxed? More generally, what do we mean for ML to be "private" when training examples do not align well with the sensitive information we want to protect?
2. **Practical privacy**: Can we build methods that drastically improve the privacy-utility tradeoff so as to afford privacy as a default for ML? How do these methods interact with other desirable properties of ML, such as efficiency and generalization?

My decision to pursue a PhD has been a natural outcome of my past exploration. I've had the opportunities to work on a spectrum of problems from the applied side [1, 2, 3] to those that involve careful analysis [4, 5, 6, 7] both in the industry and as a research student, and I've also engineered and deployed some of my published methods in practice [8, 9]. Through these experiences, I gained interest in ideas and methods from both theoretical and empirical research, and I believe ML privacy is one of the rare areas situated at this intersection where better privacy mechanisms often directly translate to tangible improvements in real-world applications. I believe a PhD is an exciting next step because making progress in open directions, such as the above, would require me to both deepen my understanding of this field and explore into related areas; a PhD is the unique opportunity that gives me the time, freedom, and curiosity to do both. In the following, I outline my initial work, learnings, and motivation in my intended directions as well as other interests that I hope to explore.

Evolving notions of ML privacy

As ML methods, systems, and problem settings evolve, so should how we define and implement privacy. My first exposure in this direction was through my work on **distributed differential privacy** at Google, published at ICML'21 [4] and NeurIPS'21 [5] and **deployed at scale for applications like Gboard**. Distributed DP asks the following: can we achieve *central* DP-like guarantees under a *local* DP setup, where n clients each noise their vector *locally* to minimize their trust in the central aggregator? Prior work bridges the $O(\sqrt{n})$ local/central utility gap assuming trusted shufflers or execution environments, but it remained open whether we could do so leveraging cryptography, specifically *secure aggregation* (SecAgg). With my collaborators, I *put together the first working system* for distributed DP with SecAgg using **distributed discrete Gaussians** (DDG)—client vectors are appropriately clipped, scaled, rotated/reflected, rounded, locally noised, and securely summed—and I showed that it matches the performance of the *central, continuous* Gaussian even under low bitwidths [4]. However, discrete Gaussian noise is tricky to sample and not closed under summation, and tuning the discretization factor γ that simultaneously touches on privacy, communication, and utility can be challenging. To address these issues, we also *analyzed the* **Skellam noise** *as a practical alternative* [5]: it is closed under summation, easy to sample, and has an RDP overhead over Gaussians shrinking quickly at $1/\text{poly}(\text{variance})$—a small compromise from DDG's $1/\exp(\text{variance})$ in practice—and it thus has a matching performance in production. Moreover, I proposed a simple method leveraging private quantiles of client vector ℓ_2-norms to **automatically tune** γ, substantially simplifying the deployment of our distributed DP algorithms. These experiences meddling with cryptography, compression, empirical heuristics, and engineering issues showed me that there is meaningful headroom for better privacy tools if we see privacy as a multi-faceted problem rather than a standalone statistical discipline.

In another work along this direction at CMU, I studied the application of DP for **cross-silo federated learning**. Past work on differentially private FL, including my distributed DP efforts, has centered on *client-level* DP, but in cross-silo FL there are much fewer clients with large local datasets and varying privacy targets, and more importantly the entities that enjoy protection (e.g. institutions) may not align with those we want to protect (e.g. people). In my paper at NeurIPS'22 [6], I found that the need for granular protection and *model personalization* in cross-silo FL naturally points to a relaxation which we call **silo-specific sample-level** privacy, where each client caters to the privacy of its own samples as it participates in FL protocols. I thoroughly analyzed many of its interesting implications, such as the removal of trust on the server, the breakdown of local fine-tuning for (provably) effective personalization, and the emergence of a **privacy-heterogeneity cost tradeoff** and the role of personalization in navigating it. I showed that *mean-regularized multi-task learning* forms a simple but surprisingly strong baseline under these conditions, but also investigated its practical complications in terms of hyperparameter tuning. This project reinforced my understanding that privacy is interwoven with many aspects of learning and good privacy algorithms should not be developed in isolation.

Practical privacy

A major roadblock in the wide adoption of differential privacy in ML is its practicality, primarily in terms of the stark privacy-utility tradeoff. Currently, I'm leading a project at CMU exploring the mitigation of **per-example clipping bias** of DP training through various *generalization methods*. I noticed that the recent empirical success of [10] suggests an outsized role of *data augmentation* in boosting DP-SGD performance that correlates to smaller augmentation-averaged gradient norms. Drawing connections to the body of work on sharpness-aware optimization which (debatably) links generalization to flatter minina, I'm investigating whether various generalization methods would remain equally effective across private and non-private optimization; one curious observation is that while both augmentation and gradient norm penalty reduces the *magnitude* part of clipping bias, only the latter sees to reduce the *direction* part which manifests as smaller utility drop from clipping. Understanding and mitigating the clipping bias would be a key step towards making DP practical, and I hope to explore this further with both empirical and theoretical analysis.

I'm also interested in **private adaptive optimization** and am crucially involved in the development of **DP^2** [7], a family of adaptive optimizers leveraging *delayed preconditioners* to drastically improve privacy-utility tradeoffs. DP^2 is based on the simple idea that in many tasks the underlying loss geometry doesn't change much over training, and we can learn less noisy but *stale* preconditioners using averages of past private gradients. DP^2 is very practical and efficient as it does not require public data or 2nd-order computation; I performed the empirical analyses in our ICLR'23 submission and OPT'22 talk [7] and it was motivating to see DP^2 excelling on real-world tasks. Through this work, I also discovered a tangential interest in **efficient private ML systems**—I wrote our code from scratch with auto-vectorization and JIT compilation and massively accelerated our research with over 10× speedup over our initial prototype. I'm curious to explore other higher/lower-level efficiency methods too.

Working at Berkeley & Future Plans

At Berkeley, I'm excited to explore these directions and beyond in the broader area of trustworthy ML. For example, I'm interested in relaxations of (differential) privacy for ML and FL, such as fine-grained analyses and label-only protection. Moreover, what if a person's data are distributed across multiple clients in FL? How do we protect privacy as well as *intellectual property* in the recent surge of foundation models? Catering to such intricate scenarios that necessitate tools beyond standard DP would be interesting. I'm also keen to take alternative perspectives on practical privacy; one example is to examine what 'weak' DP can offer in practice, as large ε's may still be effective against membership and per-attribute inference, data reconstruction, and (to some extent) FL backdoors, owing to DP being a uniform and worst-case guarantee. With my experiences in ML deployment [8] and benchmarking [11], I'm also interested in collaborative and open-source efforts on private ML systems and benchmarking to facilitate future research in these areas.

Following my background and interests, particularly in the areas of federated learning and differential privacy, I would love to work with Professors **Dawn Song**, **Michael Jordan**, **Jelani Nelson**, and **Nika Haghtalab**. I'm also excited to explore how the above relates to cryptography and ML security with Prof. **Raluca Ada Popa**, as well as how it may intersect with adjecent fields such as economics and policy-making.

I believe Berkeley is a great place to pursue my PhD for its strong faculty, student body, and interdisciplinary culture, and while my past experiences inform my current interests, I'm open and eager to explore new directions too. Looking forward, I intend to become a professor after my PhD for the unique opportunities to make my impact through research and mentoring and to pass on the positive influence I've been fortunate to receive from my research advisors. I believe my skillsets and interests make me the ideal candidate and a valuable component to Berkeley's diverse intellectual community as I continue to learn and contribute in the above areas.

References

(† denotes alphabetical authorship, * denotes equal contribution)

[1] **Ziyu Liu**, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), 2020. Oral Presentation.

[2] Meng Zhou*, **Ziyu Liu***, Pengwei Sui, Yixuan Li, and Yuk Ying Chung. Learning implicit credit assignment for cooperative multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems* (NeurIPS), 2020.

[3] Shanshan Wu, Tian Li, Zachary Charles, Yu Xiao, **Ziyu Liu**, Zheng Xu, and Virginia Smith. Motley: Benchmarking Heterogeneity and Personalization in Federated Learning. In *International Workshop on Federated Learning: Recent Advances and New Challenges* (FL-NeurIPS), 2022.

[4] Peter Kairouz†, **Ziyu Liu**†, and Thomas Steinke†. The Distributed Discrete Gaussian Mechanism for Federated Learning with Secure Aggregation. In *International Conference on Machine Learning* (ICML), 2021. Also Oral Presentation at TPDP 2021.

[5] Naman Agarwal†, Peter Kairouz†, and **Ziyu Liu**†. The Skellam Mechanism for Differentially Private Federated Learning. In *Advances in Neural Information Processing Systems* (NeurIPS), 2021. Also Oral Presentation at PPML 2021.

[6] **Ziyu Liu**, Shengyuan Hu, Zhiwei Steven Wu, and Virginia Smith. On Privacy and Personalization in Cross-Silo Federated Learning. In *Advances in Neural Information Processing Systems* (NeurIPS), 2022.

[7] Tian Li, Manzil Zaheer, **Ziyu Liu**, Sashank Reddi, Brendan McMahan, and Virginia Smith. Differentially Private Adaptive Optimization with Delayed Preconditioners. In *International Conference on Learning Representations* (ICLR), 2023. **Under Review at ICLR**. Oral Presentation at OPT 2022.

[8] Distributed DP via SecAgg Algorithms and Implementations in TensorFlow led by **Ziyu Liu**, 2021. https://www.tensorflow.org/federated/api_docs/python/tff/learning/ddp_secure_aggregator.

[9] **Ziyu Liu**, Shengyuan Hu, Zhiwei Steven Wu, and Virginia Smith. Pushing the Privacy-Utility Frontier with Heterogeneity-Aware FL, 2022. Technical report for the UK/US PETs Prize Challenge (https://petsprizechallenges.com).

[10] Soham De, Leonard Berrada, Jamie Hayes, Samuel L Smith, and Borja Balle. Unlocking high-accuracy differentially private image classification through scale. *arXiv preprint arXiv:2204.13650*, 2022.

[11] Shanshan Wu, Tian Li, Zachary Charles, Yu Xiao, **Ziyu Liu**, Zheng Xu, and Virginia Smith. Motley: Benchmarking Heterogeneity and Personalization in Federated Learning. In *International Workshop on Federated Learning* (FL-NeurIPS), 2022.