Ministry of Science and ICT

IITP Institute of Information & Communications Technology Planning & Evaluation

APPLICATION FORM
(JOINT RESEARCH)
HIGH POTENTIAL INDIVIDUALS GLOBAL TRAINING PROGRAM)

AGREEMENT

As stated above, I submit this application form to IITP that conducts "High Potential Individuals Global Training Program" supported by Ministry of Science, ICT in South Korea. IITP may disclose the information below to the public for the purpose of providing information and matching a research partnership between your institute and a Korean university.

* IITP : Institute for Information & communications Technology Planning & Evaluation

Printed Name of Chief of Research	Jaewoo Lee	Date(mm-dd-yyyy)	02-16-2020
Signature of Chief of Research	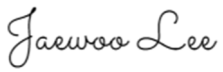		

☞ ***(Note)*** *This application is to identify the willingness to participate in this research and to find a research partnership for research institutes in Korea. Therefore, in its sole discretion, it is acceptable to contain only minimal information. (max. 3 pages)*

1. Research Title	New Algorithms for Differentially Private Deep Learning						
2. Research Area	**A.I.**	**Big Data**	**Cloud Computing**	**Block Chain**	**AR/ VR**	**ICT/SW Convergence**	**Other ICT /SW**
	X	**X**					
3. Chief of research	Title	Assistant Professor		Contact	E-mail : jaewoo.lee@uga.edu		
	Name	Jaewoo Lee			Tel : +1-706-542-8241		
4. Affiliation	Name	University of Georgia		Classification	(X) University ( ) Research Institute ( ) Industry ( ) ETC.		
5. Capacity for students (5 or less)	2		**Support for students** ***(all necessary)***		(X) Visa support (X) Research Mentoring (X) Research Space (X) Accessibility to Research equipment		

6. Research Objective	With ever-increasing applications of deep learning (DL) techniques in sensitive application domain such as healthcare, financial services, and national security, privacy concerns on individuals whose data are used for training DL models continue to grow. The needs for privacy protection in such data-intensive applications have led to burgeoning interest in the development of differentially private algorithms for training deep networks. Despite recent advances in model building techniques with differential privacy, existing algorithms often fail to provide acceptable accuracy, and hence their use in practice is severely limited. This project investigates a broad class of differentially private algorithms and will develop new tools and techniques for designing scalable, differentially private algorithms for training deep learning models.
7. Research Summary	Differential Privacy [3, 4] is widely considered to be the gold standard for protecting privacy in data sharing and machine learning [5, 6, 7, 8]. It carefully injects noise and bias into the training process in a way that ensures that no record will be accidentally memorized. Recent advances in differential privacy has shown the feasibility of applying it to deep learning tasks. Despite their promise, however, differentially private deep networks often lag far behind their non-private counterparts in accuracy, showing the need for more research in model architectures, principled methods for training, optimizers, etc. This project aims to develop new mathematical tools and techniques for building differentially private deep networks that can eliminate the efficiency gap between private and non-private deep learning. One of the barriers to this expanded research is the training time --- often orders of magnitude larger than training non-private networks. The reason is that differentially private deep learning approaches [1] modify the training algorithm and require a step called "*per-example-gradient-clipping*" whose implementations (e.g., TensorFlow Privacy [2]) slow down the processing significantly. Our own experiments in PyTorch show that the per-example-gradient-clipping technique can slow down training by a factor of 64-75x. Our research plan is to provide corresponding solutions to more complex networks (including convolutional layers, skip connections/residual blocks [9], weight sharing, recurrent networks [10, 11], and attention models [12]) as well as to arbitrary differentiable loss functions. The back-propagation (BP) algorithm combined with stochastic gradient descent (SGD) has been the workhorse for training deep neural networks. While back-propagation has shown to be effective and provides reliable performance on a variety of network architectures, its layer-to-layer error back propagation renders it difficult to extend existing tools for differential privacy to training deep learning models. In this project, we will investigate alternative learning algorithms to back-propagation algorithm to improve differentially private training processes for deep neural networks. One of major difficulties in using the gradient clipping technique with deep neural network is that gradients may grow large due successive multiplications between weight matrices. This problem stems from the fact, called *weight transport* [13, 14], that the error signals being passed on the feedback path are multiplied with weight matrices transmitted from the units on the forward path. To address this problem, we will examine and evaluate the feasibility biologically plausible [14, 15, 16, 17] alternatives to back-propagation algorithm.

8. Need for funding from Korean government	The proposed research will develop new theoretical framework, techniques, methodologies, and algorithms for differentially private algorithms for deep learning models. The successful outcomes of this research will benefit (i) researchers across disciplines by allowing them to apply deep learning models to sensitive data, (ii) companies with an established user base by enabling them to analyze data, collected from users, while protecting privacy of their customers, and (iii) Korean society at large by mitigating privacy concerns on using cutting-edge technologies.
9. Request for Korean Universities	None

References Cited

1. M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pages 308–318. ACM, 2016
2. N. Papernot, S. Chien, C. C. Choo, G. M. Andrew, and I. Mironov. TensorFlow Privacy
3. C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In Theory of Cryptography Conference, pages 265–284. Springer, 2006
4. I. Mironov. Renyi differential privacy. In Computer Security Foundations Symposium (CSF), 2017 IEEE 30th, pages 263–275. IEEE, 2017
5. K. Chaudhuri, C. Monteleoni, and A. D. Sarwate. Differentially private empirical risk minimization. Journal of Machine Learning Research, 12(Mar):1069–1109, 2011
6. D. Kifer, A. Smith, and A. Thakurta. Private convex empirical risk minimization and high-dimensional regression. In Conference on Learning Theory, pages 25–1, 2012
7. R. Bassily, A. Smith, and A. Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science, FOCS '14, pages 464–473, Washington, DC, USA, 2014. IEEE Computer Society.
8. C. Chen, J. Lee, and D. Kifer. Renyi differentially private erm for smooth objectives. In The 22nd International Conference on Artificial Intelligence and Statistics, pages 2037–2046, 2019
9. K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. CoRR, abs/1512.03385, 2015
10. S. Hochreiter and J. Schmidhuber. Long short-term memory. Neural Comput., 9(8):1735–1780, Nov. 1997.
11. K. Cho, B. van Merrienboer, Ç. Gülçehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. CoRR, abs/1406.1078, 2014
12. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I.

Polosukhin. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, pages 6000–6010, USA, 2017.

13. Stephen Grossberg. Competitive learning: From interactive activation to adaptive resonance. Cognitive science, 11(1):23–63, 1987

14. Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. Nature communications, 7:13276, 2016

15. Arild Nøkland. Direct feedback alignment provides learning in deep neural networks. In Advances in neural information processing systems, pages 1037–1045, 2016

16. Brian Crafton, Abhinav Parihar, Evan Gebhardt, and Arijit Raychowdhury. Direct feedback alignment with sparse connections for local learning. arXiv preprint arXiv:1903.02083, 2019

17. Mohamed Akrout, Collin Wilson, Peter C Humphreys, Timothy Lillicrap, and Douglas Tweed. Using weight mirrors to improve feedback alignment. arXiv preprint arXiv:1904.05391, 2019