# Large Scale Private Learning via Low-rank Reparametrization

Da Yu, Huishuai Zhang, Wei Chen, Jian Yin, Tie-Yan Liu
*{yuda3@mail2,issjyin@mail}.sysu.edu.cn*
*{huishuai.zhang,wche,tyliu}@microsoft.com*

# Background

Differential privacy:

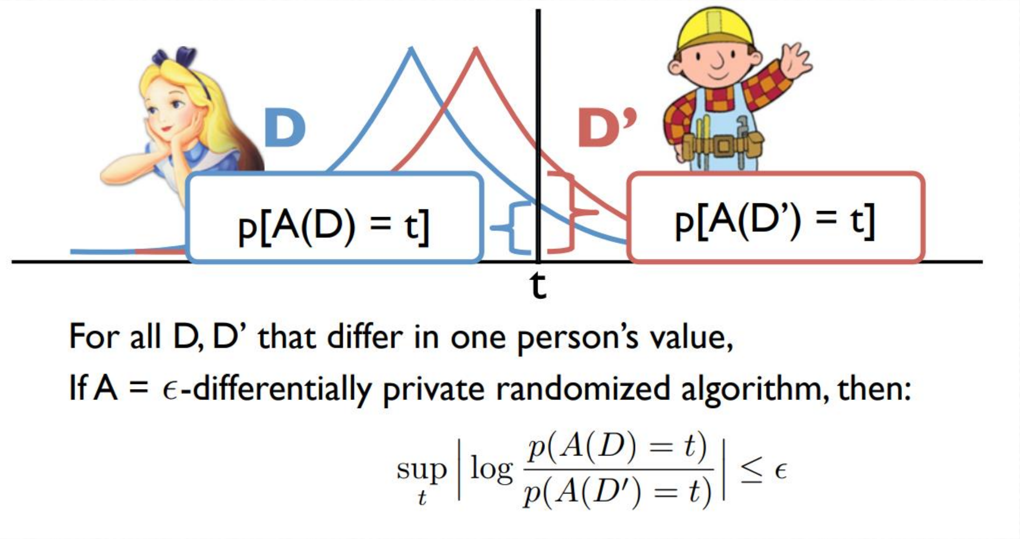


*The image is from* https://www.ece.rutgers.edu/~asarwate/nips2017/NIPS17_DPML_Tutorial.pdf

Deep learning with differential privacy: DP-SGD

1. Clip the gradients of individual samples for sensitivity control.
2. Add noise coordinate-wisely to the gradient.

# Challenges of Applying DP-SGD in Large Models

- **High computational cost**

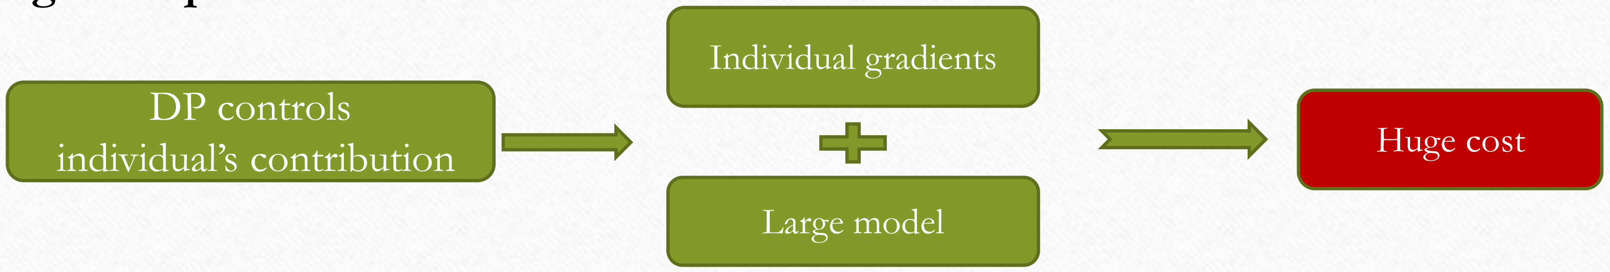


- **Bad dimensional dependence**

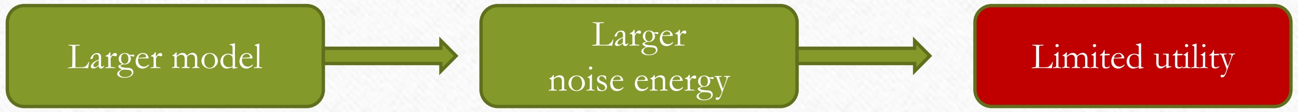

# A Reparametrization approach

**Reparametrization:**

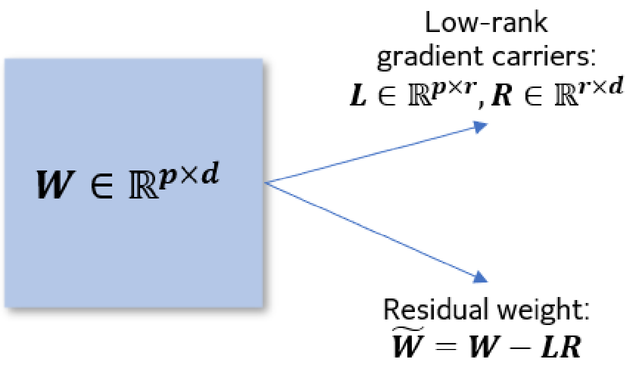


**Forward:**

Input: $\boldsymbol{x} \in \mathbb{R}^{d}$

Normal forward: $\boldsymbol{h} = \boldsymbol{Wx}$

Reparametrized forward:

$$\boldsymbol{h} = \boldsymbol{LRx} + \widetilde{\boldsymbol{W}}\boldsymbol{x}$$

**Backward:**

We show $\partial \boldsymbol{L}$ and $\partial \boldsymbol{R}$ naturally satisfy:

$$\partial \boldsymbol{L} = (\partial \boldsymbol{W})\boldsymbol{R}^{T}$$
$$\partial \boldsymbol{R} = \boldsymbol{L}^{T}(\partial \boldsymbol{W})$$

Construct an update on $W$ from $\partial L, \partial R$ via

$$(\partial L)R + L(\partial R) - LL^{T}(\partial L)R,$$

**Corollary**: The above update is equivalent to **projecting $\partial \boldsymbol{W}$ into the subspace spanned by $\boldsymbol{L}$ and $\boldsymbol{R}$.**

# Why Does Our Approach Work?

- Gradients are of low-stable rank.
  - Low-rank gradient carriers.

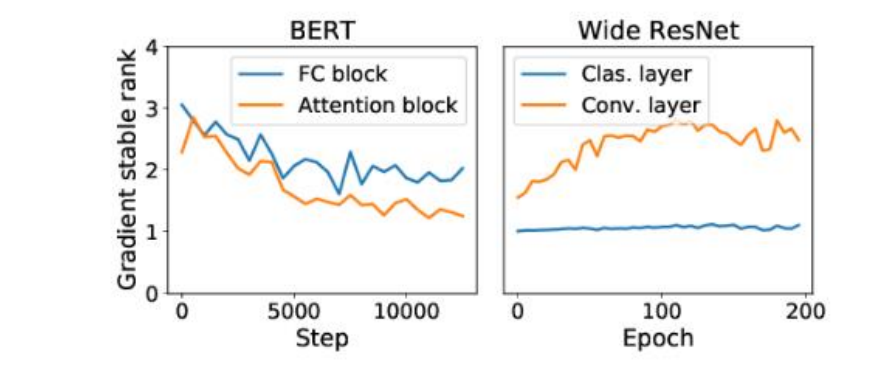


- How to generate proper $L, R$?
  - The principal components of the historical updates.
  - For linear regression, the current gradient stays exactly in the subspace spanned by historical updates.
  - For deep models, most energy of the current gradient also lives in such subspace.

# RGP on Downstream Tasks of BERT

- Experiment architecture: the $\text{BERT}_{\text{base}}$ model, 110M parameters

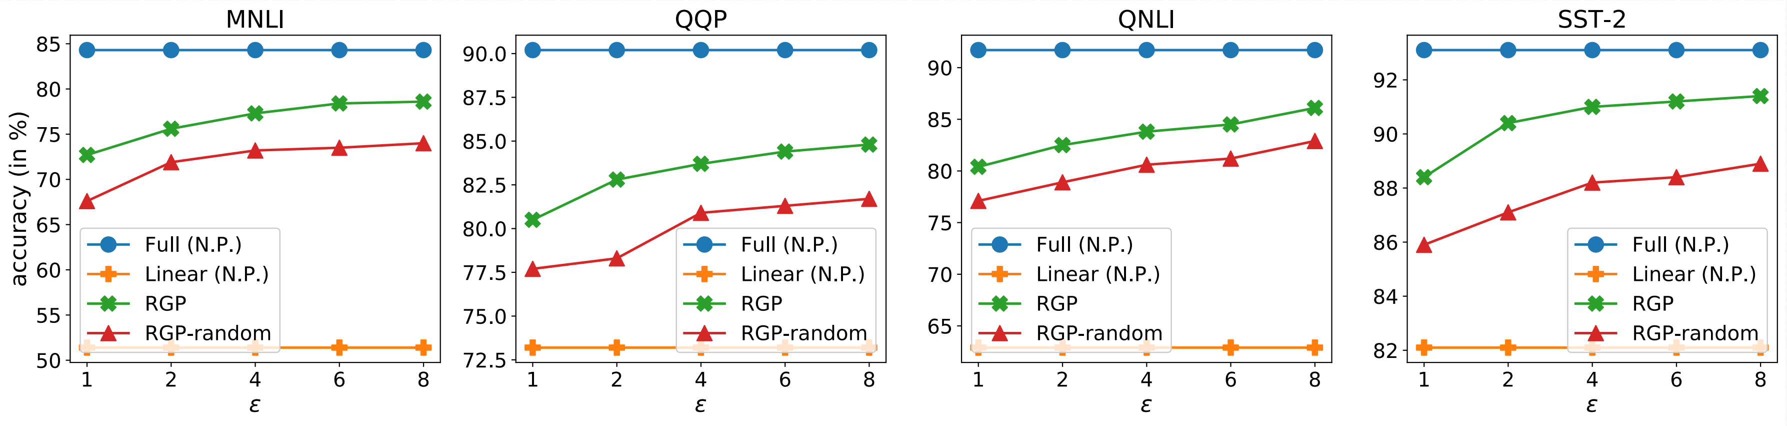

# Thanks!

- Our source code is available at

  https://github.com/dayu11/Differentially-Private-Deep-Learning

- If you have any question, please feel free to contact us.

*{yuda3@mail2,issjyin@mail}.sysu.edu.cn*
*{huishuai.zhang,wche,tyliu}@microsoft.com*