NEURAL INFORMATION PROCESSING SYSTEMS

DP-Mix: Mixup-based Data Augmentation for Differentially Private Learning

Wenxuan Bao[1], Francesco Pittaluga[2], Vijay Kumar B G[2], Vincent Bindschaedler[1]

[1]University of Florida, [2]NEC Labs America

UF Herbert Wertheim College of Engineering
Department of Computer & Information Science & Engineering
UNIVERSITY *of* FLORIDA

NEC

NEC Laboratories **America**

Background | Mixup Data Augmentation

$$
\begin{aligned}
\hat{x} &= \lambda x_i + (1 - \lambda) x_j, \\
\hat{y} &= \lambda y_i + (1 - \lambda) y_j,
\end{aligned}
$$

where $\lambda \in [0, 1]$ is a random number

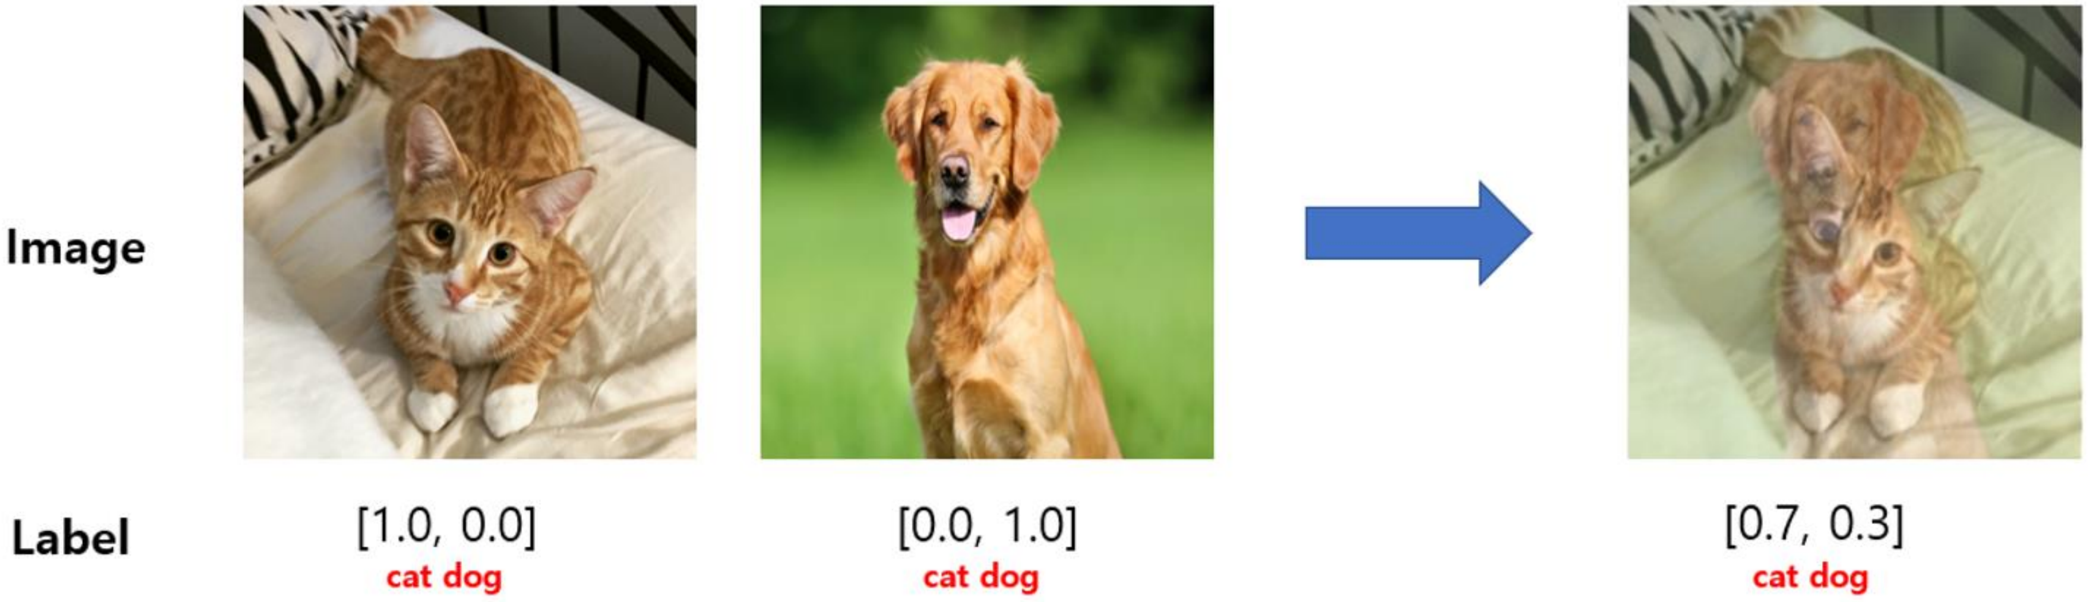

Image source: https://hoya012.github.io/blog/Bag-of-Tricks-for-Image-Classification-with-Convolutional-Neural-Networks-Review/

Microbatch

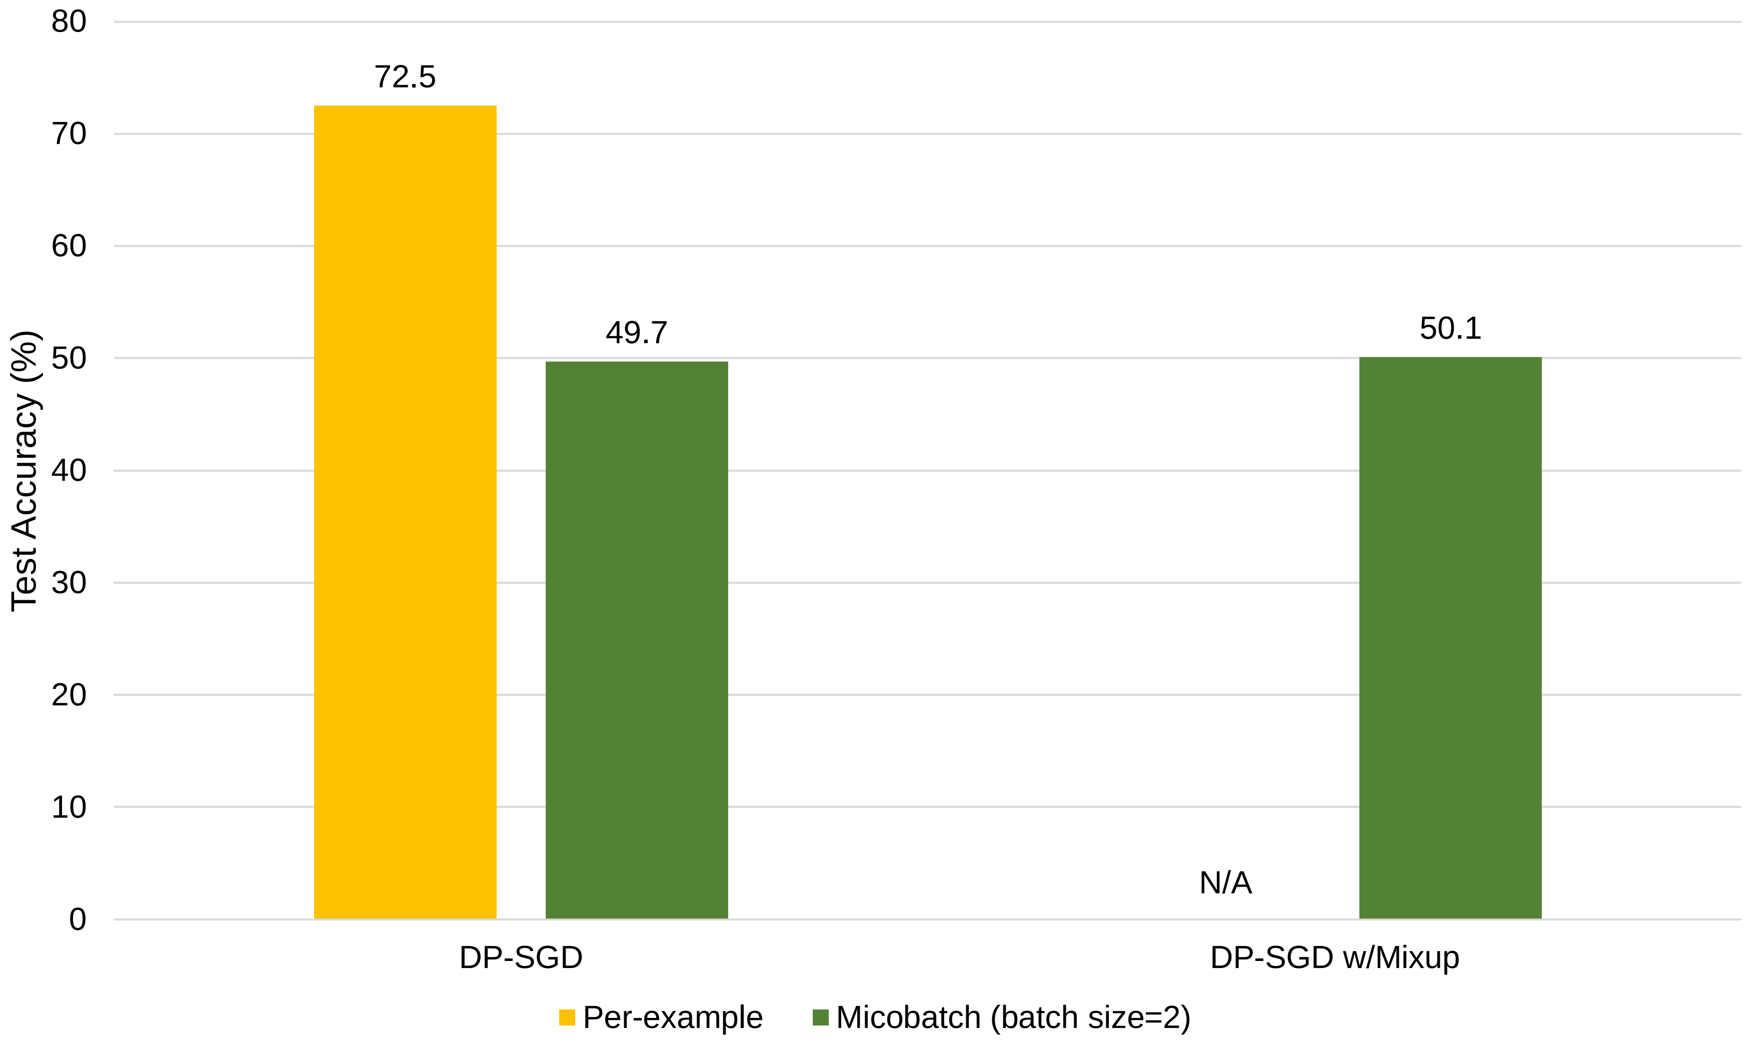
80
70
60
50
40
30
20
10
0
Test Accuracy (%)
72.5
49.7
50.1
N/A
DP-SGD
DP-SGD w/Mixup
Per-example
Micobatch (batch size=2)

Our Method | DP-Mix$_{\text{self}}$

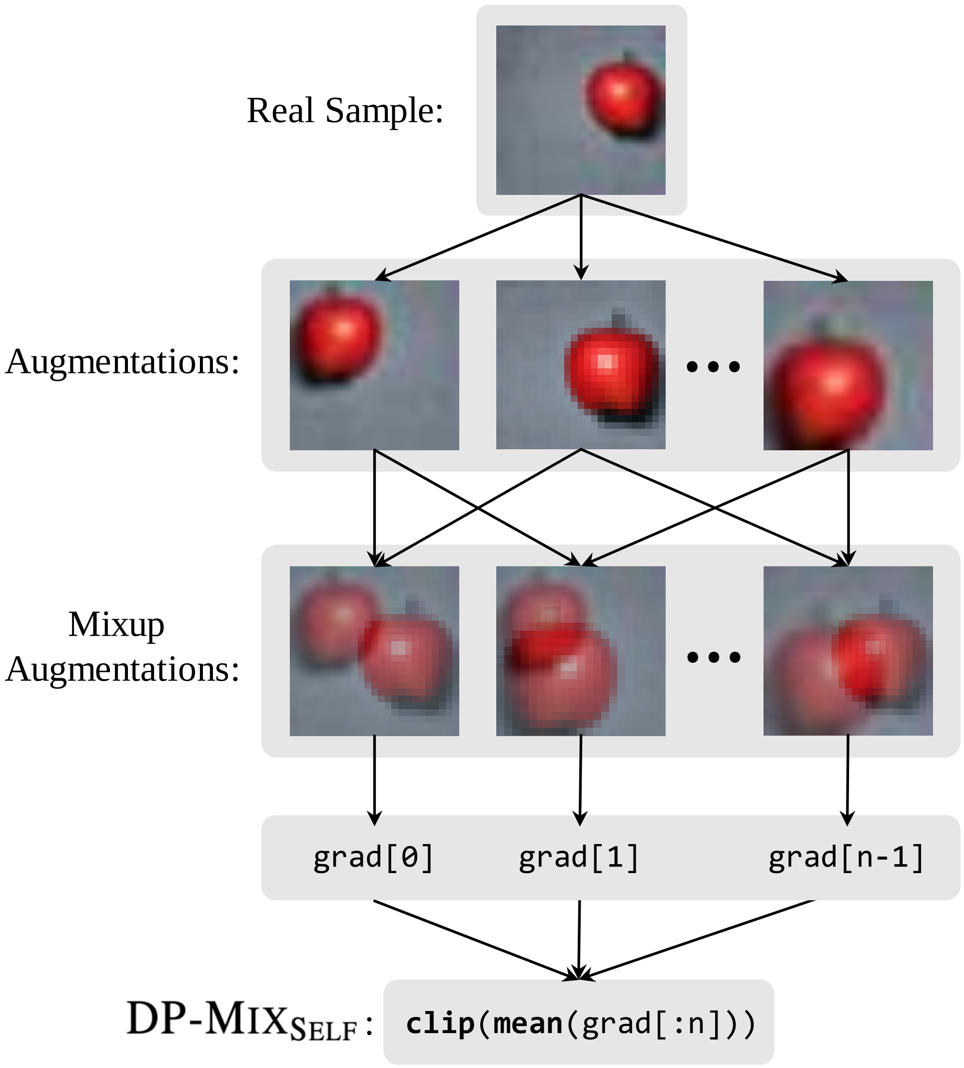

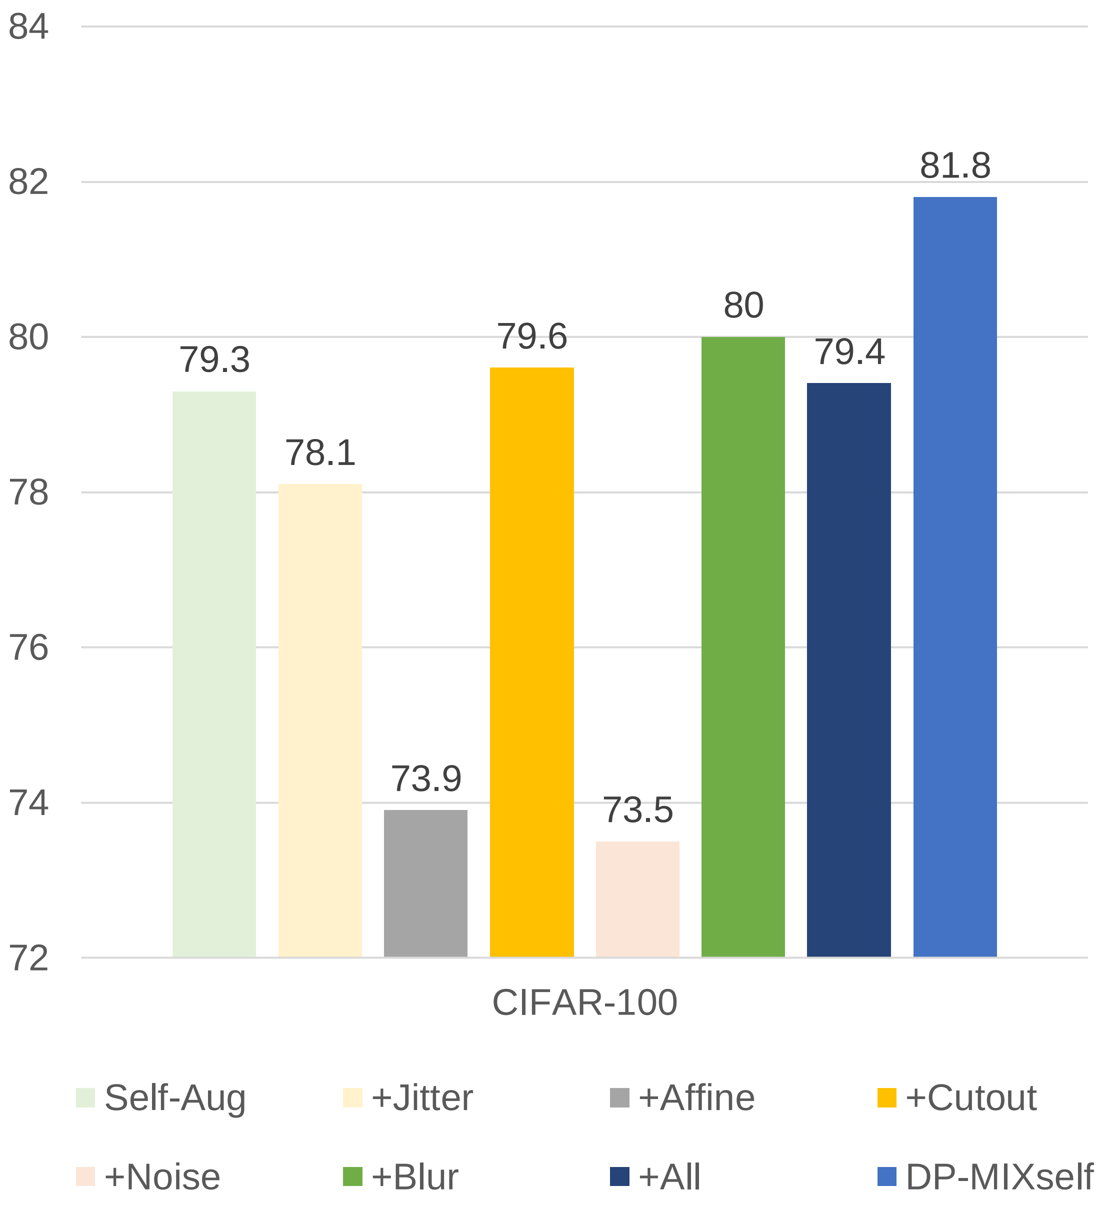

Background | Stable Diffusion

```
{
  "prompt:" "Futuristic architectures
             with planets in the
             background."
}
```

```
{
  "prompt:" "View of a cyberpunk
             city."
}
```

```
{
  "prompt:" "cyberpunk city at night
             with transparent neon
             billboards."
}
```

Our Method | DP-Mix$_{\text{diff}}$

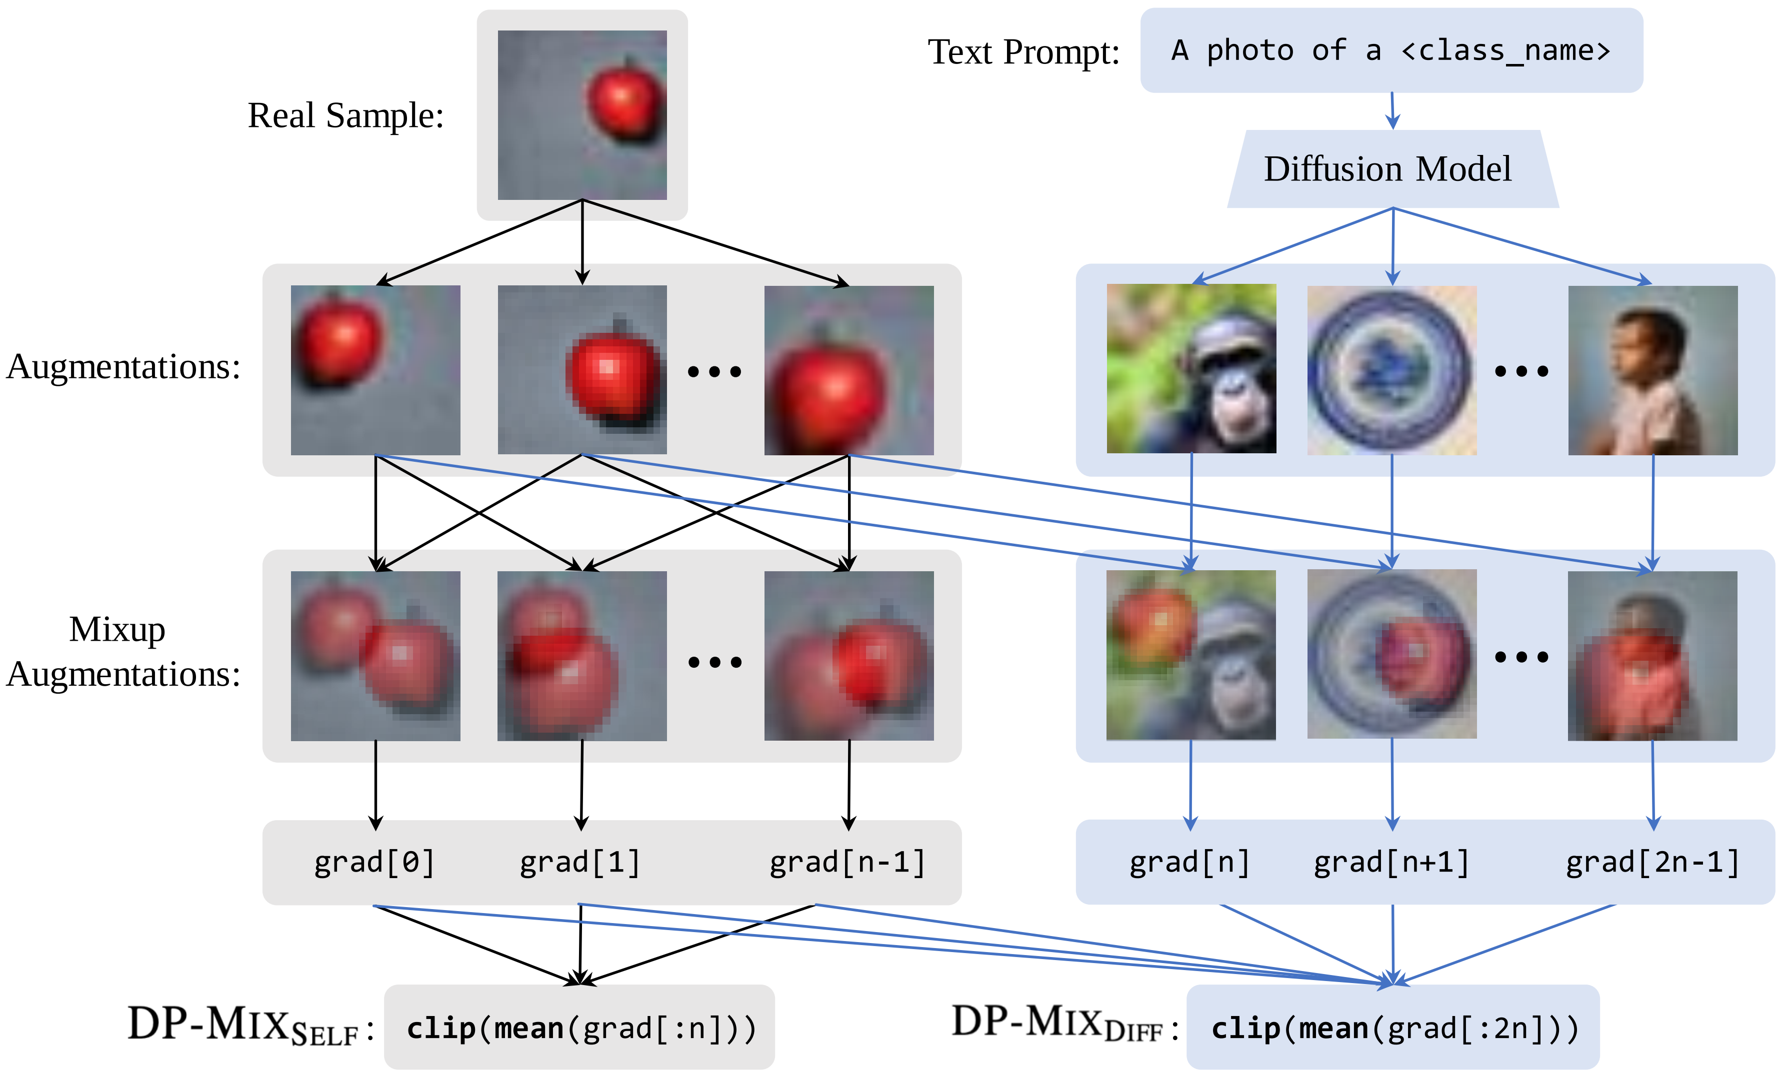

DP-Mix$_{diff}$ vs Pretraining with Diffusion Data

- Pre-training on diffusion samples does **not** improve performance; mixing up training samples with them does

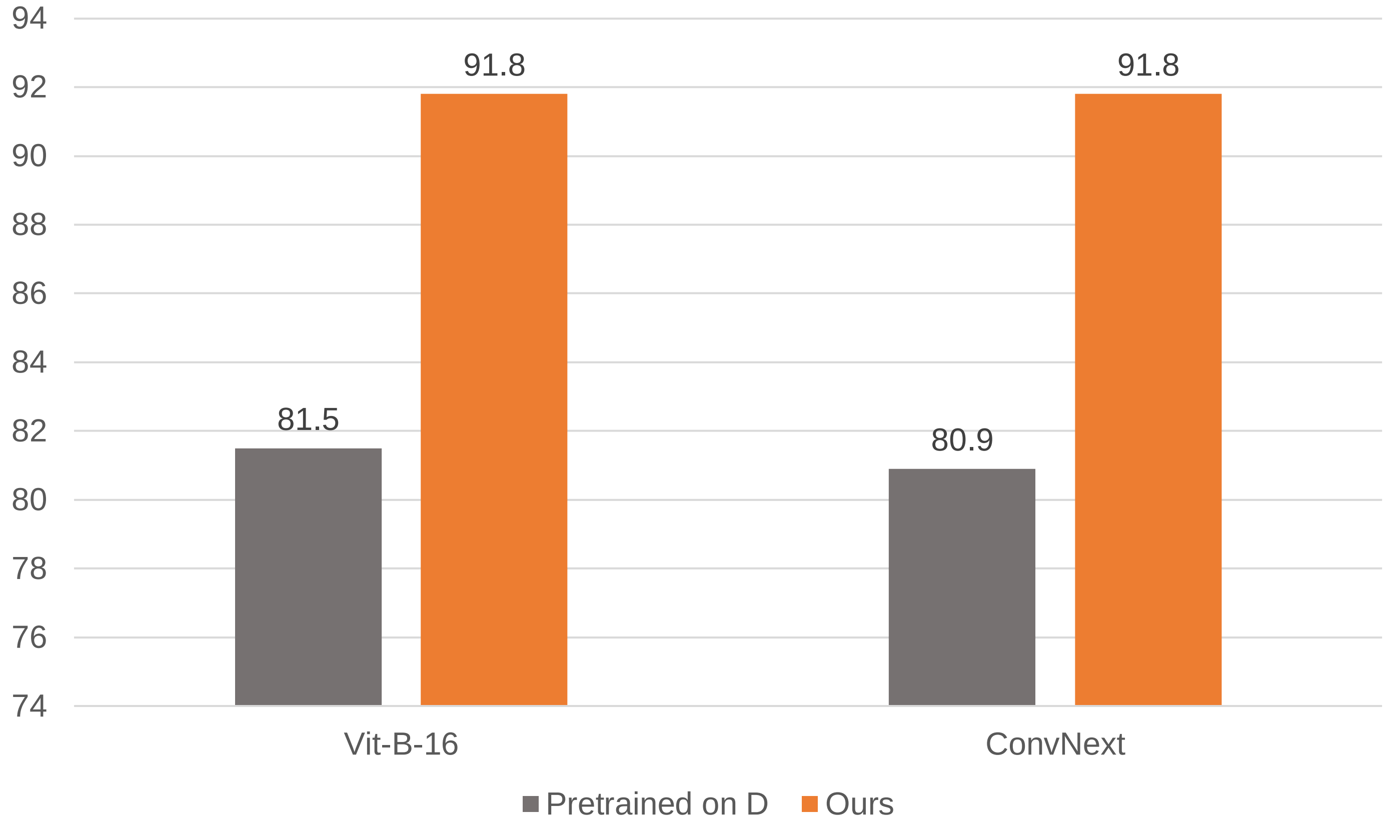

Caltech 256

Main Results

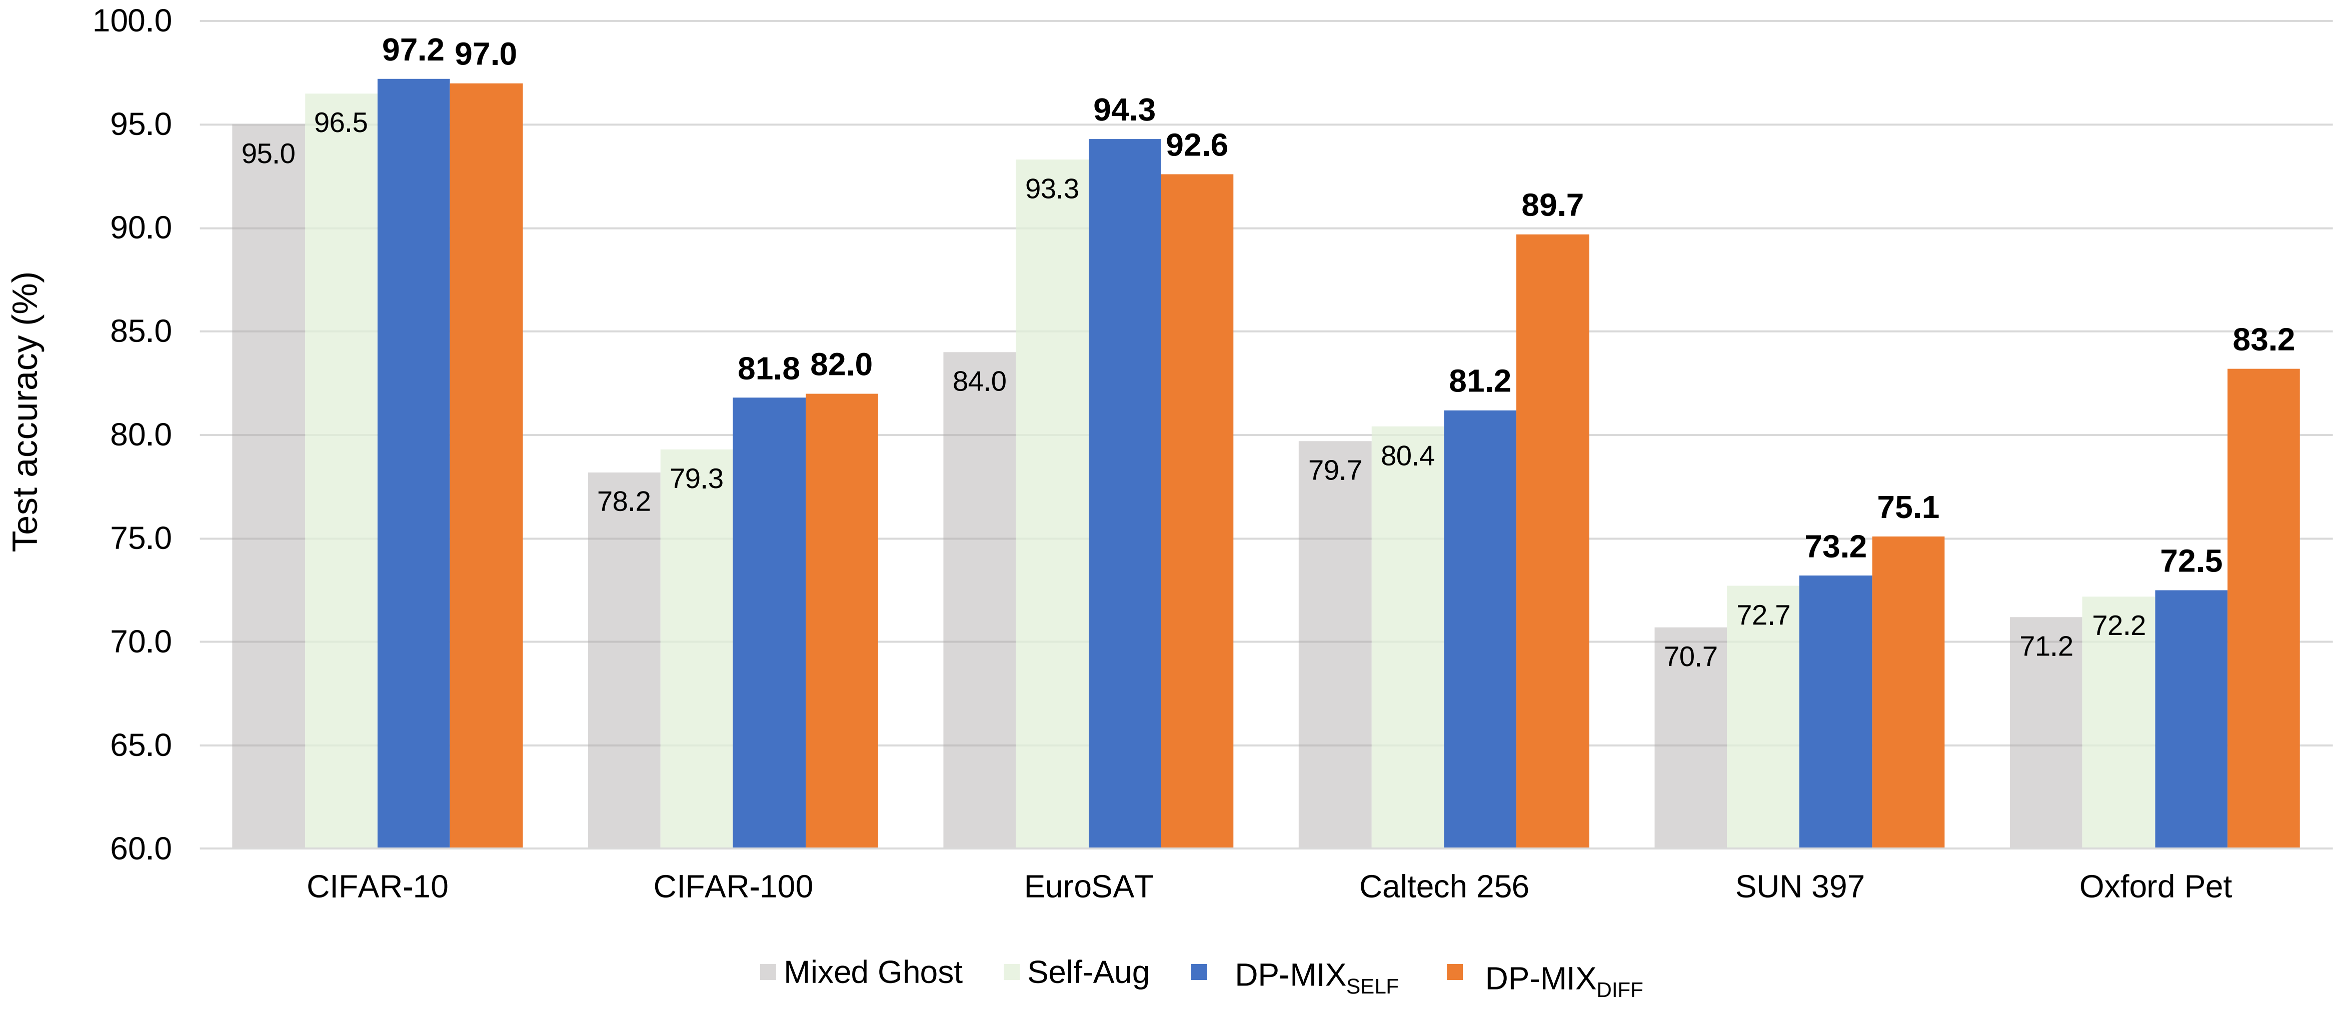

$\varepsilon=1$

Main Results

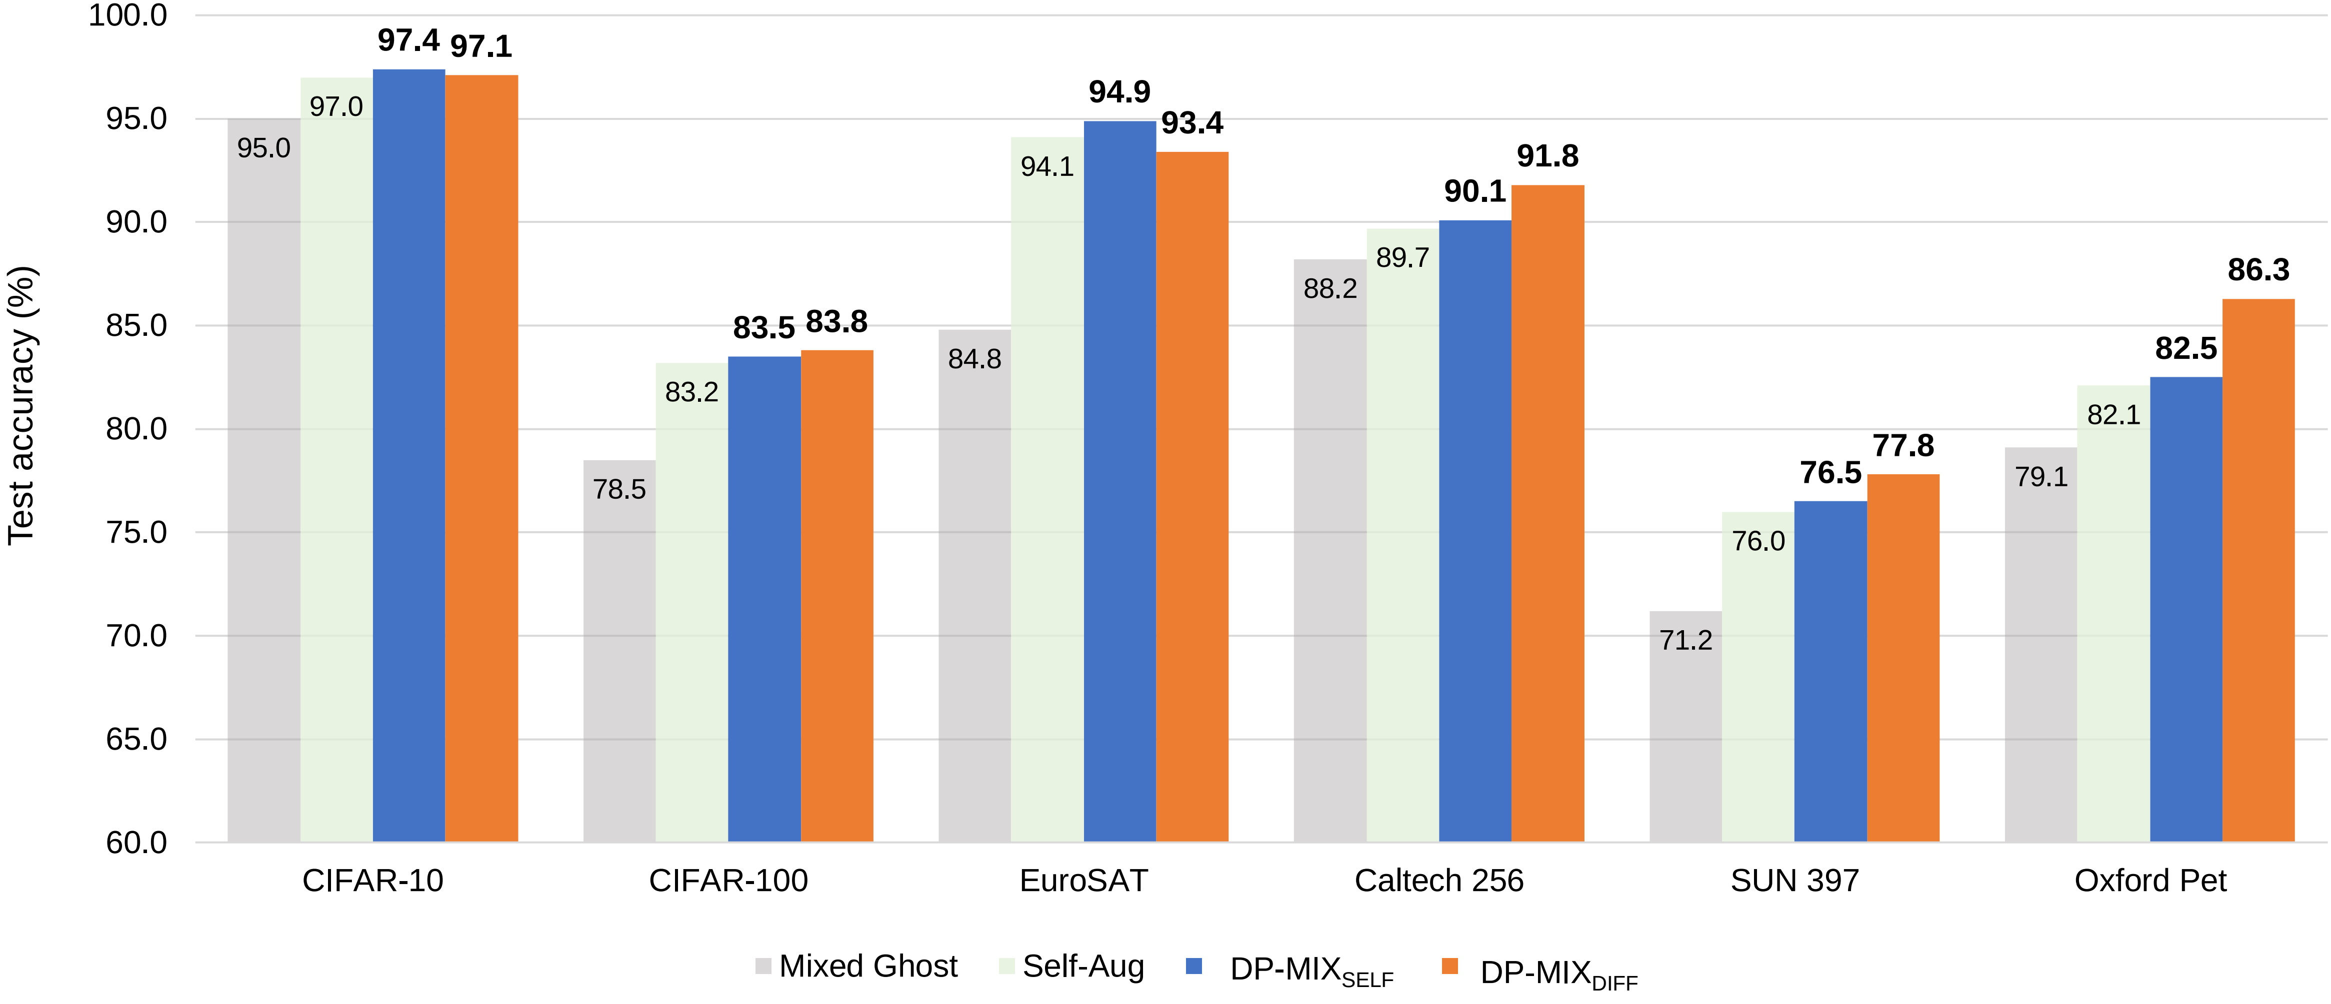

$\varepsilon=2$

Takeaway

- We show how to apply ***mixup*** for DP training of ML models and demonstrate it surpasses the prior SoTA ***at no extra privacy cost***.

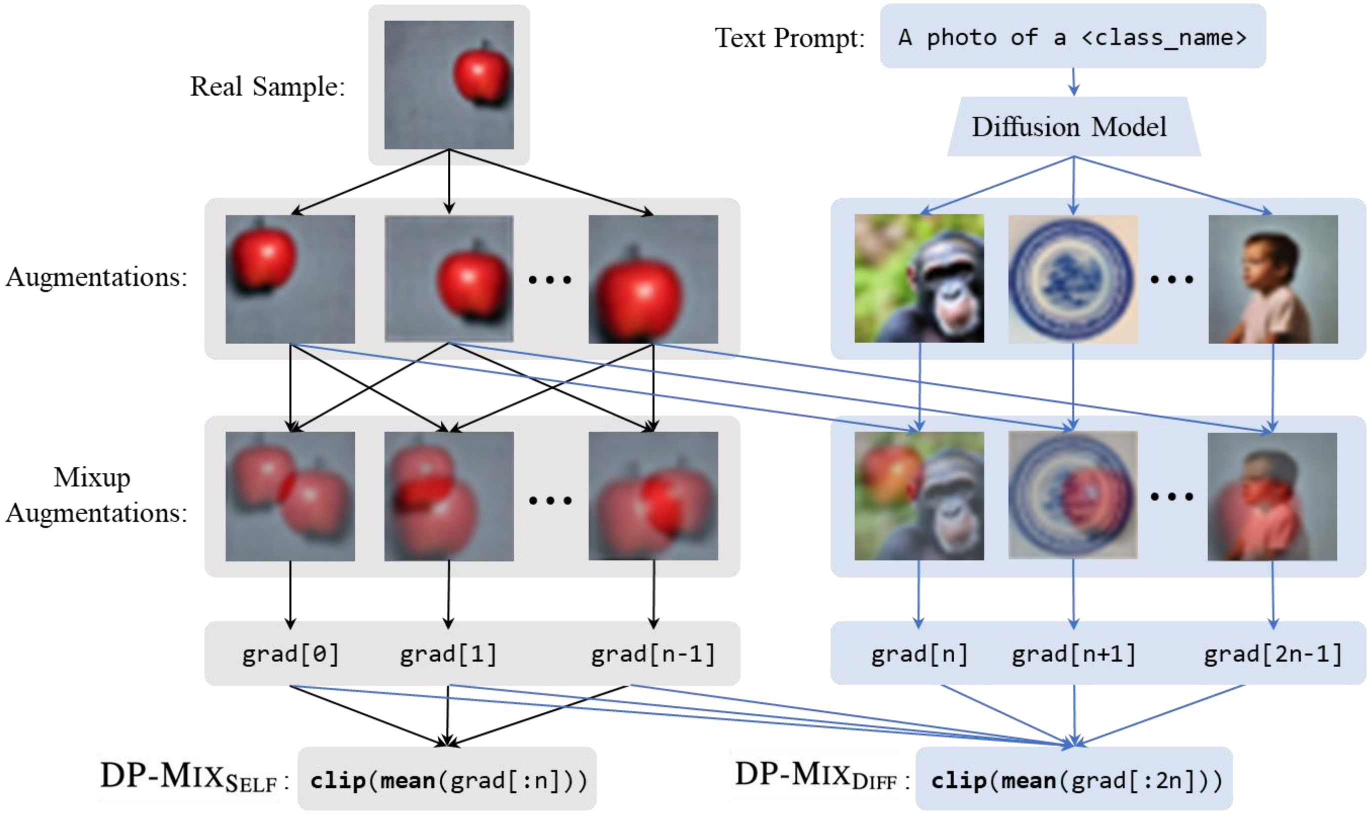

https://wenxuan-bao.github.io