

Lecture 13 Deep Learning with DP

Yu-Xiang Wang



COMPUTER SCIENCE

UC SANTA BARBARA

Computing. ReInvented.

Recap: Last lecture

- Utility of NoisyGD mechanism
 - Main tools: the convergence of SGD.
- Privacy Amplification by Sampling and Noisy SGD
- Computational consideration

Recap: NoisyGD summary

	Lipschitz + convex	Lipschitz + Smooth + convex	Smooth + Lipschitz + convex + GLM
Output Pert	$\frac{d^{1/4}L\ \theta^*\ \log(\frac{1}{\delta})^{1/4}}{n^{1/2}\epsilon^{1/2}}$	$\frac{d^{1/3}\beta^{1/3}L^{2/3}\ \theta^*\ ^{4/3}\log(\frac{1}{\delta})^{1/3}}{n^{2/3}\epsilon^{2/3}}$	Same as left
ObjPert	Not applicable	$\frac{dL\ \theta^*\ \sqrt{\log(\frac{1}{\delta})}}{n\epsilon}$ Lower order terms and dependence on β hidden.	$\frac{\sqrt{d}L\ \theta^*\ \sqrt{\log(\frac{1}{\delta})}}{n\epsilon}$
NoisyGD	$\frac{\sqrt{d}L\ \theta^*\ \sqrt{\log(\frac{1}{\delta})}}{n\epsilon}$	$\frac{\sqrt{d}L\ \theta^*\ \sqrt{\log(\frac{1}{\delta})}}{n\epsilon}$	$\frac{\sqrt{d}L\ \theta^*\ \sqrt{\log(\frac{1}{\delta})}}{n\epsilon}$

	Lipschitz + Strongly convex	Lipschitz + Smooth + Nonconvex
NoisyGD	$\frac{dL^2\log(1/\delta)}{n\lambda\epsilon^2}$	$\frac{\sqrt{n\beta d}L^2(f(\theta_1) - f^*)\log(1/\delta)}{n\epsilon}$ Stationary point convergence

Recap: Computational Complexity of NoisyGD

- General pattern: First term goes to 0, second term dominates with large T. (How large does T need to be?)

- Convex + Lipschitz

$$T \geq \frac{n^2 \rho}{d} \asymp \frac{n^2 \epsilon^2}{d}$$

- Convex + Smooth

$$T \geq \frac{2n\beta \|x_1 - x^*\| \sqrt{\rho}}{\sqrt{d}L} \asymp n\epsilon / \sqrt{d}$$

- Strongly convex + Lipschitz

$$T \geq \frac{n^2 \rho}{\lambda} \asymp \frac{n^2 \epsilon^2}{\lambda}$$

$O\left(\frac{1}{T}\right) \sim O\left(\frac{1}{\sqrt{T}}\right) + O\left(\frac{\sigma^2}{T}\right)$

$P\text{-CDP} \Rightarrow (P + \sqrt{2P \log \frac{1}{\delta}} - \delta) \cdot DP$

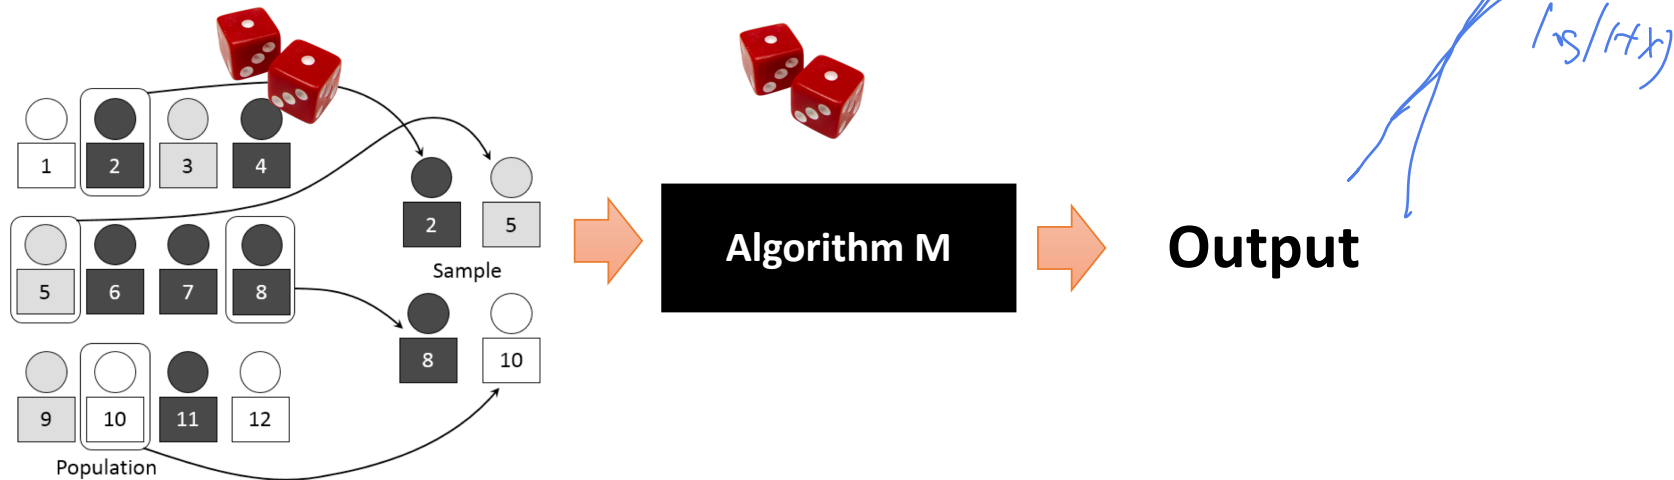
of incremental gradient call

$O(n^3 \epsilon^2)$

$O(n^3 \epsilon^2)$

Each iteration requires n incremental gradient evaluation!

Recap: Privacy Amplification by Sampling



$\mathcal{M} \circ \text{Sample} : \text{Data} \rightarrow \text{Output}$

Subsampling Lemma: If \mathcal{M} obeys (ϵ, δ) -DP, then $\mathcal{M} \circ \text{Subsample}$ obeys that (ϵ', δ') -DP with $\delta' = \gamma\delta$

$$\epsilon' = \log(1 + \gamma(e^\epsilon - 1)) = O(\gamma\epsilon)$$

$\log(1+x) \leq x$

Recap: The Noisy **Stochastic** Gradient Descent Mechanism (NoisySGD)

- Privacy analysis:

- A composition of T **subsampled** gaussian mechanism
- Satisfying that:

ρ_0 -CDP

$$\gamma^2 \rho_0\text{-truncated CDP} \Rightarrow \mathbb{E}_{\text{NoisySGD}}(\alpha) \leq O(\gamma^2 \rho_0 \alpha)$$

- Utility analysis:

- Same convergence theorem of SGD, but different parameters due to sampling.

$$\mathbb{E}_{\text{NoisySGD}}(\alpha) \leq O(T \gamma^2 \rho_0 \alpha)$$

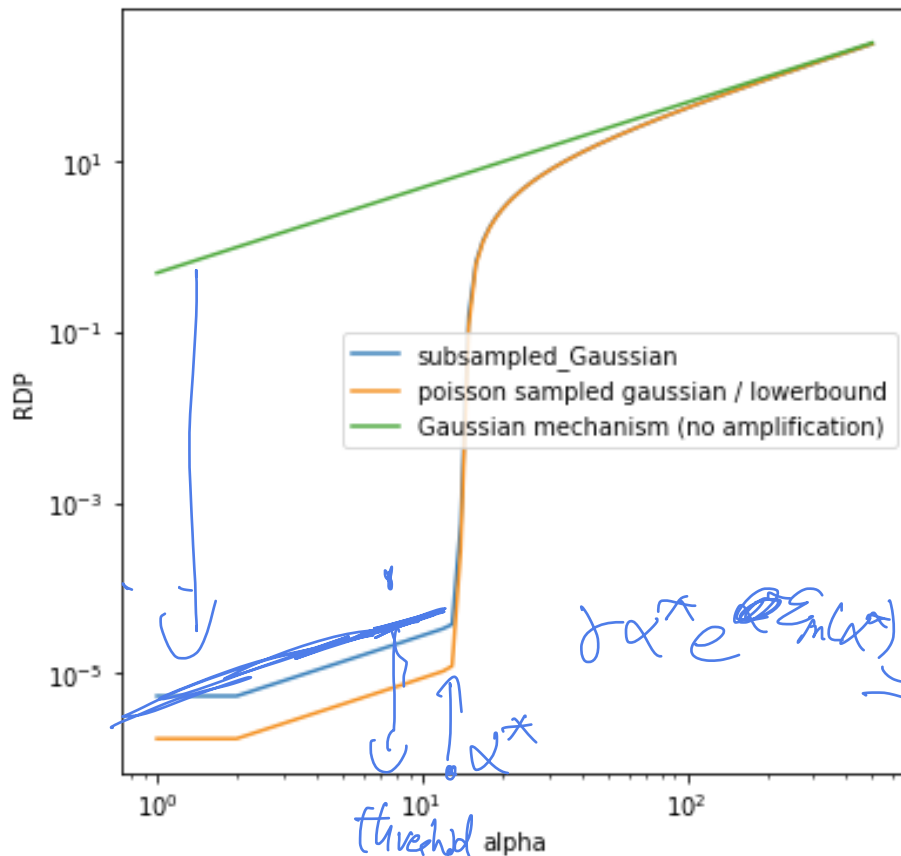
for $\alpha \leq \text{threshold}$

$$\hat{g}_t = \frac{1}{\gamma} \left(\sum_{i \in I} \nabla \ell_i(\theta_t) + \mathcal{N}(0, \sigma^2 I_d) \right);$$

\uparrow
Random

$$\mathbb{E}[\|\hat{g}_t - \mathbb{E}\hat{g}_t\|^2] \leq \frac{d\sigma^2}{\gamma^2} + \frac{\gamma n L^2}{\gamma^2}$$

Recap: The Renyi DP of sampled mechanisms



A concise asymptotic bound of the RDP for sampled Gaussian mechanism via truncated CDP.

Theorem: Assume \mathcal{M} satisfies ρ -CDP.
then

$$\epsilon_{\mathcal{M} \circ \text{Sample}_\gamma}(\alpha) \leq 13\alpha\gamma^2\rho \text{ for } \alpha \leq \frac{\log(1/\gamma)}{4\rho}$$

$$\rho_0 = \frac{\sigma^2}{2\sigma^2}$$

More precise bound: <https://arxiv.org/abs/1808.00087>

Theorem 11 of Bun et al. "Composable and Versatile Privacy via Truncated CDP"
https://projects.iq.harvard.edu/files/privacytools/files/bun_mark_composable.pdf

Working out the constraint due to this phases transition in amplification by sampling

$$\delta = \text{~~100~~} e^{-8}$$

(ϵ, δ)-DP, $\epsilon = P + \frac{\sqrt{P \log \frac{1}{\delta}}}{n}$ for P-COP

$$\epsilon = \min_{\alpha} \epsilon(\alpha) + \frac{\log \frac{1}{\delta}}{\alpha - 1}$$

$$\geq \min_{\alpha} \alpha P + \frac{\log \frac{1}{\delta}}{\alpha - 1} = P + (\alpha - 1)P + \frac{\log \frac{1}{\delta}}{\alpha - 1}$$

$$\alpha^* = \sqrt{\frac{\log \frac{1}{\delta}}{P}} + 1$$

$$\frac{13\gamma^2 L^2 \alpha}{2G^2} = 13\gamma^2 P \cdot \alpha = P$$

$$1 + \sqrt{\frac{T \cdot \frac{1}{\delta}}{P}} \leq \frac{\log \frac{1}{\delta} G^2}{4 L^2} \Rightarrow$$

Constraint on T being large

$$\frac{L^2}{2G^2} \leq \frac{\log \frac{1}{\delta}}{4 \left(1 + \sqrt{\frac{\log \frac{1}{\delta}}{P}}\right)}$$

Today

- Computational complexity of NoisySGD
- Application to Deep Learning with Differential Privacy
- Other models in Deep Learning with DP

Readings

- NoisySGD: Bassily et al. <https://arxiv.org/abs/1405.7085>
- Deep Learning with DP: Abadi et al. <https://arxiv.org/abs/1607.00133>
- Amplification by sampling in RDP: <https://arxiv.org/abs/1808.00087>

NoisySGD's Convergence Bound under the smooth convex Case

$$\frac{G^2}{T\eta^2} = \frac{\beta L^2}{2\rho}$$

$$\mathbb{E} \left[\frac{1}{T} \sum_t (f(x_t) - f^*) \right] \leq \frac{\|x_1 - x^*\|^2}{T\eta} + \eta d \sigma^2 \left[\frac{dG^2}{\delta^2} + \frac{nL^2}{\delta} \right]$$

- Larger variance

Constraint 1: $\frac{L^2}{\delta G^2} \leq \frac{\log \frac{1}{\delta}}{4 \left(\frac{\eta \sqrt{10\eta\delta}}{\rho} \right)}$

Constraint 2: $\frac{\beta T \delta^2 L^2}{\geq G^2} = \rho$

Constraint 3: (optimal rate of privacy utility)

$$\left\{ \frac{\max \eta^2 \rho}{\delta}, \sqrt{\frac{n^3 \rho \beta \rho \rho \|x_1 - x^*\|}{\delta L}} \right\}$$

$$\begin{aligned} & \frac{\eta \frac{1}{\rho}}{= \max \eta \rho} \left\{ \frac{\|x_1 - x^*\|^2}{T} \left(\frac{dG^2}{\delta^2} + \frac{nL^2}{\delta} \right) \right\} \\ & \approx \frac{n\beta \|x_1 - x^*\|^2}{T} + \sqrt{\frac{d \|x_1 - x^*\|^2 L^2}{P}} + \sqrt{\frac{\|x_1 - x^*\|^2 n L^2}{T\delta}} \\ & \text{# of grad} \geq \frac{\delta n^2 \beta \rho \|x_1 - x^*\|}{\sqrt{\delta} L} \end{aligned}$$

Constraint 1 and 2: $T \geq \frac{\rho \log \frac{1}{\delta}}{\sqrt{\log \frac{1}{\delta}} \delta}$

of grad: $\approx \frac{\delta n \rho}{\delta} = \frac{n \rho}{\delta}$

NoisySGD's Convergence Bound under the non-smooth convex Case

$$\mathbb{E} \left[\frac{1}{T} \sum_t (f(x_t) - f^*) \right] \leq \frac{\|x_1 - x^*\|^2}{T\eta} + \eta \left(\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \|\partial f(x_t)\|^2 \right] \right) + \boxed{d\sigma^2}$$

$n^2 L^2$

$$\leq \sqrt{\|x_1 - x^*\|^2 \left(n^2 L^2 + \frac{dG^2}{\delta^2} + \frac{nL^2}{\delta} \right)} \quad \frac{dG^2}{\delta^2} + \frac{nL^2}{\delta}$$

$$\leq \sqrt{\|x_1 - x^*\|^2 \left(\sqrt{\frac{n^2}{\delta T}} + \sqrt{\frac{d}{p}} + \sqrt{\frac{n}{T\delta}} \right)}$$

of gradient oracles

$$\max \left\{ \frac{n \sqrt{p} L^{1.5}}{\delta \sqrt{1.5} T}, \frac{n^2 p}{d} \right\}$$

$$\delta^2 = \frac{n^4 p^{1.5}}{d} \left(\frac{1}{T} \right)$$

$$\max \left\{ \frac{n^2 p}{d}, \sqrt{\frac{n^4 p^{1.5}}{d}} \right\}$$

NoisySGD's convergence under the strongly convex case

$$\mathbb{E} \left[\underline{f\left(\frac{2}{T(T+1)} \sum_{t=1}^T tx_t\right)} \right] - f(x^*) \leq \frac{4(n^2 L^2 + \boxed{d\sigma^2})}{\lambda(T+1)}$$

$\eta = \frac{2}{t+1}$
 $\frac{d\sigma^2}{T^2} + \frac{nL^2}{\lambda}$

$$\frac{n^2 L^2}{\lambda T} + \frac{dL^2}{\lambda \rho} + \frac{nL^2}{\lambda \gamma T}$$

of grade class = $\max \left\{ \frac{\gamma n^3 \rho}{d}, \frac{n^2 \rho}{d}, \frac{n \sqrt{\rho}}{\gamma} \right\}$

$$\gamma = \sqrt{\frac{d}{n^2 \rho^{1.5}}} \Rightarrow \max \left\{ \frac{n^2 \rho}{d}, \frac{n^2 \rho^{0.75}}{\sqrt{d}} \right\}$$

Requires more advanced algorithm to obtain subquadratic time algorithm for **non-smooth** convex ERM

- Very interesting recent paper by
 - Kulkarni, Lee and Liu (2021): <https://arxiv.org/pdf/2103.15352.pdf>
- By a **smoothing trick** they obtained an algorithm with the following oracle calls

$$O\left(\frac{n^{3/2}}{d^{1/8}} + \frac{n^2}{d}\right)$$

Handwritten annotations: A blue circle highlights the first term $\frac{n^{3/2}}{d^{1/8}}$. A blue arrow points to the n^2 term in the second fraction. To the right, the expression $\frac{n^{3/2}}{d^{1/4}}$ is written in blue.

- i.e., slightly worse dimension dependence than the smooth + convex case that we derived.
- Also, if we are to consider minimizing population risk instead, then subquadratic-time algorithms are possible.

Checkpoint: Comparing NoisyGD and NoisySGD computationally

- Both optimal information-theoretically.
 - If we ignore computation and add very large noise, but use infinitesimal step-size
- Table to compare computation
 - in terms of **the number of incremental gradient calls** to achieve **information theoretic limit** up to a constant

	Lipschitz + Smooth + Convex	Lipschitz + Convex	Lipschitz + Strongly convex
NoisyGD	$\frac{n^2 \beta \ x_1 - x^*\ \sqrt{\rho}}{\sqrt{d} L}$	$\frac{n^3 \rho}{d}$	$\frac{n^3 \rho}{\lambda}$
NoisySGD	$\frac{n^{3/2} \beta^{1/2} \ x_1 - x^*\ \rho^{1/2}}{d^{1/4} L^{1/2}} + \frac{n^2 \rho}{d}$	$\frac{n^2 \rho^{3/4}}{d^{1/2}} + \frac{n^2 \rho}{d}$	$\frac{n^2 \rho^{3/4}}{d^{1/2}} + \frac{n^2 \rho}{d}$

Open problem: what is the optimal computational complexity?

Remainder of the lecture

- A survey of differentially private deep learning
- Reference:
 - NoisySGD: Abadi et al. (2016)
 - PATE: Papernot et al. (2018), PrivateKNN: Zhu et al. (2020)
 - Autodp tutorial on private deep learning:
https://github.com/yuxiangw/autodp/blob/master/tutorials/tutorial_private_deep_learning.ipynb

Deep Learning with Differential Privacy

- Main challenge: It is difficult to enforce global sensitivity

- Abadi et al: Per-example Gradient Clipping

$$\theta_{t+1} = \theta_t - \eta \left(\sum_{i \in \mathcal{I}_t} \nabla \ell_i(\theta_t) + \text{Noise}(\sigma, G^2 \eta) \right)$$

Handwritten note: $\sum_{i \in \mathcal{I}_t}$ ← mini-batch

$$\hat{g}_t = \sum_{i \in \mathcal{I}_t} \nabla \ell_i(\theta_t) \cdot \min \left\{ \frac{L}{\|\nabla \ell_i(\theta_t)\|_2}, 1 \right\}$$

Handwritten note: \hat{g}_t has $\Delta_2 \leq L$

- Other tricks:
 - Use small neural networks
 - Private-PCA to reduce dimension in the original data

What does Deep Learning with Differential Privacy look like in autodp?

```
from autodp.autodp_core import Mechanism
from autodp.transformer_zoo import Composition
from autodp import mechanism_zoo, transformer_zoo

class NoisySGD_mech(Mechanism):
    def __init__(self, prob, sigma, niter, name='NoisySGD'):
        Mechanism.__init__(self)
        self.name=name
        self.params={'prob':prob, 'sigma':sigma, 'niter':niter}

        # create such a mechanism as in previously
        subsample = transformer_zoo.AmplificationBySampling() # by default this is using poisson sampling
        mech = mechanism_zoo.GaussianMechanism(sigma=sigma)
        prob = prob
        # Create subsampled Gaussian mechanism
        SubsampledGaussian_mech = subsample(mech, prob, improved_bound_flag=True)


        # Now run this for niter iterations
        compose = transformer_zoo.Composition()
        mech = compose([SubsampledGaussian_mech], [niter])

        # Now we get it and let's extract the RDP function and assign it to the current mech being constructed
        rdp_total = mech.RenyiDP
        self.propagate_updates(rdp_total, type_of_update='RDP')
```


Handwritten annotations:

- A blue arrow points from the left margin to the line `subsample = transformer_zoo.AmplificationBySampling()`.
- Handwritten blue scribbles and lines are present above the `self.params` dictionary.
- Handwritten blue arrows point to the `improved_bound_flag=True` parameter in the `subsample` call.
- Handwritten blue arrows point to the `[niter]` list in the `compose` call.


Empirical results on MNIST



(1) Large noise



(2) Medium noise



(3) Small noise

Figure 3: Results on the accuracy for different noise levels on the MNIST dataset. In all the experiments, the network uses 60 dimension PCA projection, 1,000 hidden units, and is trained using lot size 600 and clipping threshold 4. The noise levels (σ, σ_p) for training the neural network and for PCA projection are set at $(8, 16)$, $(4, 7)$, and $(2, 4)$, respectively, for the three experiments.

Empirical results on CIFAR10

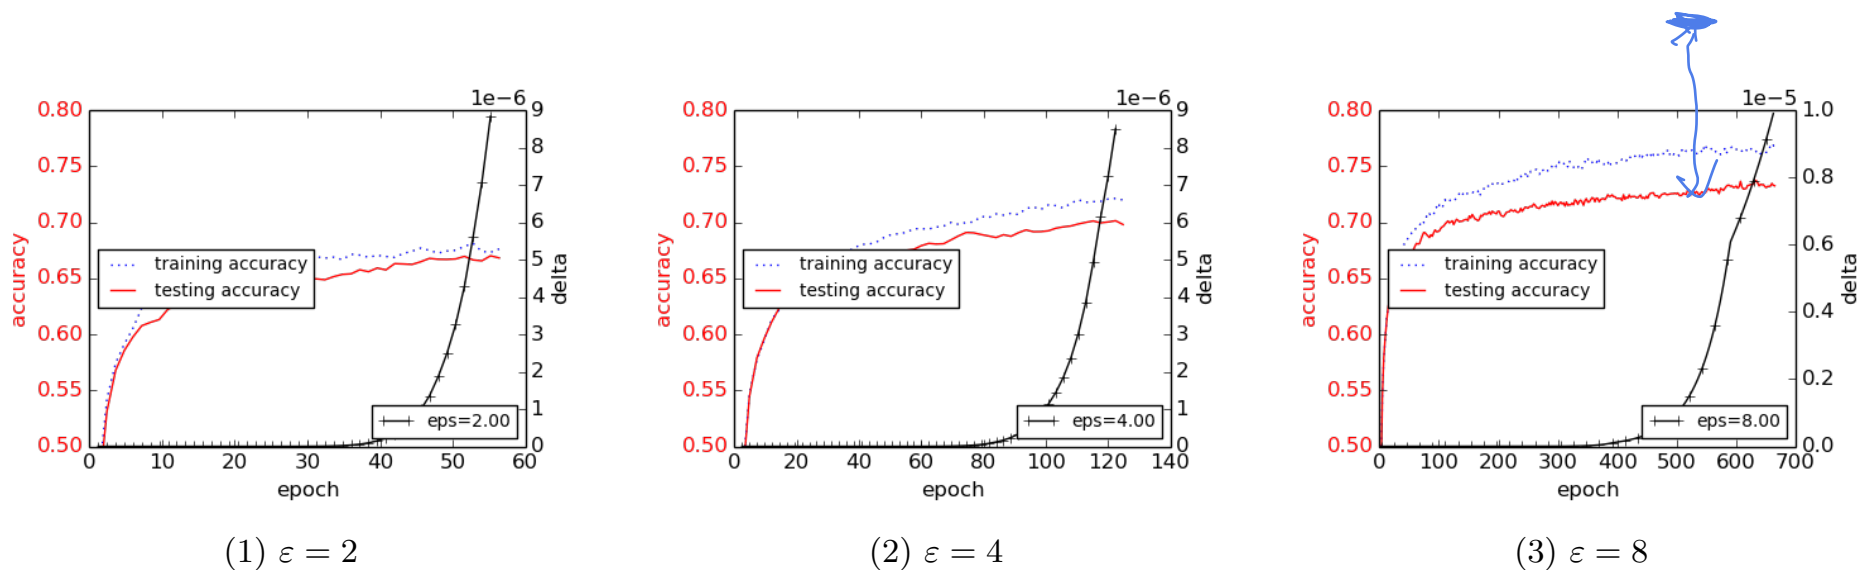


Figure 6: Results on accuracy for different noise levels on CIFAR-10. With δ set to 10^{-5} , we achieve accuracy 67%, 70%, and 73%, with ϵ being 2, 4, and 8, respectively. The first graph uses a lot size of 2,000, (2) and (3) use a lot size of 4,000. In all cases, σ is set to 6, and clipping is set to 3.

Per-Example Clipped Gradient Descent may not converge

- Example: Two data points / linear losses

$$x_1, x_2 \in [-1, 1] \quad \theta \in [-1, 1]$$

$$\text{loss: } l_1(\theta) = \langle x_1, \theta \rangle = x_1 \cdot \theta \quad x_1 = 1.0$$

$$l_2(\theta) = \langle x_2, \theta \rangle = x_2 \cdot \theta \quad x_2 = -0.5$$

clip at 0.5

$$\nabla l_1(\theta) = 1.0$$

$$\nabla l_2(\theta) = -0.5$$



$$x_1 = 1 \quad x_2 = 0.3, \quad \theta = 0.3 \quad \text{clip at } 0.5$$

$$\nabla \sum_{i=1}^2 l_i = 1 - 0.3 - 0.3 = 0.4$$

$$\sum \nabla \text{clip}(\nabla l_i) = 0.5 - 0.3 - 0.3 = -0.1$$


Effect of gradient clipping in Generalized Linear Model

$$\ell(\theta, (x, y)) = f_y(\langle x, \theta \rangle)$$

- Fixing clipping threshold B , there exists another convex GLM that the algorithm is optimizing.

$$g_{\mathbf{x}}(y) = \begin{cases} -\frac{B}{\|\mathbf{x}\|_2}(y - y_1) + f(y_1) & \text{for } y \in (-\infty, y_1) \\ f(y) & \text{for } y \in [y_1, y_2] \cap \mathbf{R} \\ \frac{B}{\|\mathbf{x}\|_2}(y - y_2) + f(y_2) & \text{for } y \in (y_2, \infty) \end{cases}$$

$$\nabla \ell(\theta) = \mathbf{x} \cdot f'_y(\langle x, \theta \rangle)$$



Knowledge transfer model via the “Private Aggregation of Teacher Ensemble” (PATE)

- Assume a large private dataset
- A small public dataset (can be unlabeled)

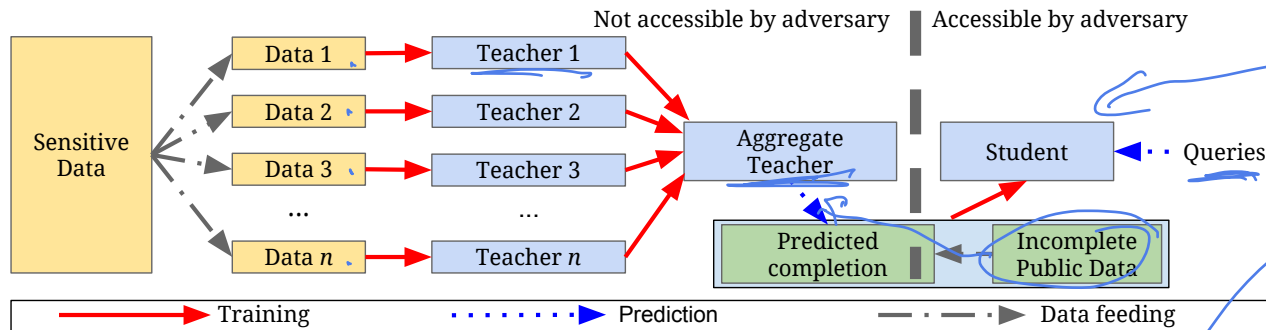



Figure 2: Overview of the approach: (1) an ensemble of teachers is trained on disjoint subsets of the sensitive data, (2) a student model is trained on public data labeled using the ensemble.



What does PATE look like from autodp?

- Just a composition of Gaussian mechanisms!

```
from autodp.mechanism_zoo import GaussianMechanism

import numpy as np

class PATE(GaussianMechanism):
    def __init__(self, sigma, m, Binary, name='PATE'):
        # sigma is the std of the Gaussian noise added to the voting scores
        if Binary:
            # This is a binary classification task
            sensitivity = 1
        else: # for the multiclass case, the L2 sensitivity is sqrt(2)
            sensitivity = np.sqrt(2)
        GaussianMechanism.__init__(self, sigma=sigma/sensitivity/np.sqrt(m), name=name)

        self.params = {'sigma':sigma}
```

- Ideas / tricks that improve to PATE
 - Noisy-Screening and semi-supervised learning (Papernot et al 2017/2018)
 - Release only the argmax rather than voting scores? (Papernot et al 2018)
 - Sparse Vector Technique (Bassily et al., 2018; Liu et al. 2020)
 - Active learning (Liu et al. 2020)

Private K-Nearest Neighbor




Figure 2. The overview of the proposed framework. Given the unlabeled public data X_{public} , we query through privacy wall for pseudo labels, where the private data and the queried public data are sent through feature extractor Φ and "Private-kNN" to assign pseudo labels. Combining the public data and the pseudo labels, the feature extractor Φ is further updated. This procedure can be iterated for rounds to achieve satisfied privacy-accuracy trade-off.

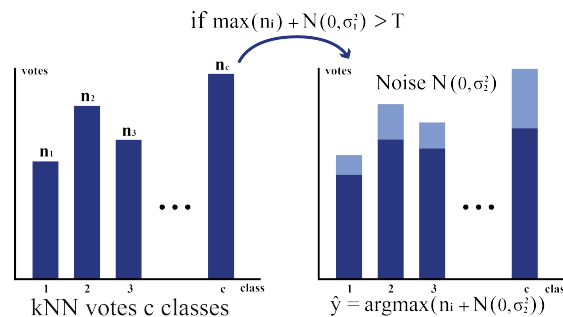


Figure 3. Illustration on the noisy screening and noisy aggregation procedure.

What does private KNN look like from autodp?

- Just a composition of Subsampled Gaussian mechanisms (same as NoisySGD!)

```
In [4]: # We will reuse the noisySGD class

# number of data points

gamma = 0.01
m=1000
sigma = 5.0

privateKNN = NoisySGD_mech(prob=gamma, sigma=sigma, niter=m)

# compute epsilon, as a function of delta
privateKNN.get_approxDP(delta=1e-6)
```

Out[4]: 0.27105623043762284

Comps (1) (2) (3) (4) (5) (6) (7) (8) (9) (10) (11) (12) (13) (14) (15) (16) (17) (18) (19) (20) (21) (22) (23) (24) (25) (26) (27) (28) (29) (30) (31) (32) (33) (34) (35) (36) (37) (38) (39) (40) (41) (42) (43) (44) (45) (46) (47) (48) (49) (50) (51) (52) (53) (54) (55) (56) (57) (58) (59) (60) (61) (62) (63) (64) (65) (66) (67) (68) (69) (70) (71) (72) (73) (74) (75) (76) (77) (78) (79) (80) (81) (82) (83) (84) (85) (86) (87) (88) (89) (90) (91) (92) (93) (94) (95) (96) (97) (98) (99) (100) (101) (102) (103) (104) (105) (106) (107) (108) (109) (110) (111) (112) (113) (114) (115) (116) (117) (118) (119) (120) (121) (122) (123) (124) (125) (126) (127) (128) (129) (130) (131) (132) (133) (134) (135) (136) (137) (138) (139) (140) (141) (142) (143) (144) (145) (146) (147) (148) (149) (150) (151) (152) (153) (154) (155) (156) (157) (158) (159) (160) (161) (162) (163) (164) (165) (166) (167) (168) (169) (170) (171) (172) (173) (174) (175) (176) (177) (178) (179) (180) (181) (182) (183) (184) (185) (186) (187) (188) (189) (190) (191) (192) (193) (194) (195) (196) (197) (198) (199) (200) (201) (202) (203) (204) (205) (206) (207) (208) (209) (210) (211) (212) (213) (214) (215) (216) (217) (218) (219) (220) (221) (222) (223) (224) (225) (226) (227) (228) (229) (230) (231) (232) (233) (234) (235) (236) (237) (238) (239) (240) (241) (242) (243) (244) (245) (246) (247) (248) (249) (250) (251) (252) (253) (254) (255) (256) (257) (258) (259) (260) (261) (262) (263) (264) (265) (266) (267) (268) (269) (270) (271) (272) (273) (274) (275) (276) (277) (278) (279) (280) (281) (282) (283) (284) (285) (286) (287) (288) (289) (290) (291) (292) (293) (294) (295) (296) (297) (298) (299) (300) (301) (302) (303) (304) (305) (306) (307) (308) (309) (310) (311) (312) (313) (314) (315) (316) (317) (318) (319) (320) (321) (322) (323) (324) (325) (326) (327) (328) (329) (330) (331) (332) (333) (334) (335) (336) (337) (338) (339) (340) (341) (342) (343) (344) (345) (346) (347) (348) (349) (350) (351) (352) (353) (354) (355) (356) (357) (358) (359) (360) (361) (362) (363) (364) (365) (366) (367) (368) (369) (370) (371) (372) (373) (374) (375) (376) (377) (378) (379) (380) (381) (382) (383) (384) (385) (386) (387) (388) (389) (390) (391) (392) (393) (394) (395) (396) (397) (398) (399) (400) (401) (402) (403) (404) (405) (406) (407) (408) (409) (410) (411) (412) (413) (414) (415) (416) (417) (418) (419) (420) (421) (422) (423) (424) (425) (426) (427) (428) (429) (430) (431) (432) (433) (434) (435) (436) (437) (438) (439) (440) (441) (442) (443) (444) (445) (446) (447) (448) (449) (450) (451) (452) (453) (454) (455) (456) (457) (458) (459) (460) (461) (462) (463) (464) (465) (466) (467) (468) (469) (470) (471) (472) (473) (474) (475) (476) (477) (478) (479) (480) (481) (482) (483) (484) (485) (486) (487) (488) (489) (490) (491) (492) (493) (494) (495) (496) (497) (498) (499) (500) (501) (502) (503) (504) (505) (506) (507) (508) (509) (510) (511) (512) (513) (514) (515) (516) (517) (518) (519) (520) (521) (522) (523) (524) (525) (526) (527) (528) (529) (530) (531) (532) (533) (534) (535) (536) (537) (538) (539) (540) (541) (542) (543) (544) (545) (546) (547) (548) (549) (550) (551) (552) (553) (554) (555) (556) (557) (558) (559) (560) (561) (562) (563) (564) (565) (566) (567) (568) (569) (570) (571) (572) (573) (574) (575) (576) (577) (578) (579) (580) (581) (582) (583) (584) (585) (586) (587) (588) (589) (590) (591) (592) (593) (594) (595) (596) (597) (598) (599) (600) (601) (602) (603) (604) (605) (606) (607) (608) (609) (610) (611) (612) (613) (614) (615) (616) (617) (618) (619) (620) (621) (622) (623) (624) (625) (626) (627) (628) (629) (630) (631) (632) (633) (634) (635) (636) (637) (638) (639) (640) (641) (642) (643) (644) (645) (646) (647) (648) (649) (650) (651) (652) (653) (654) (655) (656) (657) (658) (659) (660) (661) (662) (663) (664) (665) (666) (667) (668) (669) (670) (671) (672) (673) (674) (675) (676) (677) (678) (679) (680) (681) (682) (683) (684) (685) (686) (687) (688) (689) (690) (691) (692) (693) (694) (695) (696) (697) (698) (699) (700) (701) (702) (703) (704) (705) (706) (707) (708) (709) (710) (711) (712) (713) (714) (715) (716) (717) (718) (719) (720) (721) (722) (723) (724) (725) (726) (727) (728) (729) (730) (731) (732) (733) (734) (735) (736) (737) (738) (739) (740) (741) (742) (743) (744) (745) (746) (747) (748) (749) (750) (751) (752) (753) (754) (755) (756) (757) (758) (759) (760) (761) (762) (763) (764) (765) (766) (767) (768) (769) (770) (771) (772) (773) (774) (775) (776) (777) (778) (779) (780) (781) (782) (783) (784) (785) (786) (787) (788) (789) (790) (791) (792) (793) (794) (795) (796) (797) (798) (799) (800) (801) (802) (803) (804) (805) (806) (807) (808) (809) (810) (811) (812) (813) (814) (815) (816) (817) (818) (819) (820) (821) (822) (823) (824) (825) (826) (827) (828) (829) (830) (831) (832) (833) (834) (835) (836) (837) (838) (839) (840) (841) (842) (843) (844) (845) (846) (847) (848) (849) (850) (851) (852) (853) (854) (855) (856) (857) (858) (859) (860) (861) (862) (863) (864) (865) (866) (867) (868) (869) (870) (871) (872) (873) (874) (875) (876) (877) (878) (879) (880) (881) (882) (883) (884) (885) (886) (887) (888) (889) (890) (891) (892) (893) (894) (895) (896) (897) (898) (899) (900) (901) (902) (903) (904) (905) (906) (907) (908) (909) (910) (911) (912) (913) (914) (915) (916) (917) (918) (919) (920) (921) (922) (923) (924) (925) (926) (927) (928) (929) (930) (931) (932) (933) (934) (935) (936) (937) (938) (939) (940) (941) (942) (943) (944) (945) (946) (947) (948) (949) (950) (951) (952) (953) (954) (955) (956) (957) (958) (959) (960) (961) (962) (963) (964) (965) (966) (967) (968) (969) (970) (971) (972) (973) (974) (975) (976) (977) (978) (979) (980) (981) (982) (983) (984) (985) (986) (987) (988) (989) (990) (991) (992) (993) (994) (995) (996) (997) (998) (999) (1000)

- Isn't there a NoisyScreening step?

Empirical results in the Knowledge Transfer model

- Private-KNN and PATE:

Table 1. Utility and privacy of semi-supervised student model

Dataset	Methods	#Queries	ϵ	Acc.	NP Acc.
MNIST	LNMAX	1000	8.03	98.1%	
	GNMAX	286	1.97	98.5%	99.2%
	Ours	735	0.47	98.8%	
SVHN	LNMAX	1000	8.19	90.1%	
	GNMAX	3098	4.96	91.6%	92.8%
	Ours	2939	0.49	91.6%	
CIFAR-10	GNMAX			$\leq 50\%$	
	Noisy SGD		4	70%	80.5%
	Ours	3877	2.92	70.8%	

Theory behind PATE and algorithmic improvements

Table 1: Summary of our results: excess risk bounds for PATE algorithms.

Algorithm	PATE (Gaussian Mech.)	PATE (SVT-based)		PATE (Active Learning)
	Papernot et al. [2017]	Bassily et al. [2018a]	This paper	This paper
Realizable	$\tilde{O}\left(\frac{d}{(n\epsilon)^{2/3}} \vee \frac{d}{m}\right)$	$\tilde{O}\left(\frac{d}{(n\epsilon)^{2/3}} \vee \sqrt{\frac{d}{m}}\right)$	$\tilde{O}\left(\frac{d^{3/2}}{n\epsilon} \vee \frac{d}{m}\right)$	$\tilde{O}\left(\frac{d^{3/2}\theta^{1/2}}{n\epsilon} \vee \frac{d}{m}\right)$
τ -TNC	$\tilde{O}\left(\left(\frac{d^{3/2}}{n\epsilon}\right)^{\frac{2\tau}{4-\tau}} \vee \frac{d}{m}\right)$	same as agnostic	$\tilde{O}\left(\left(\frac{d^{3/2}}{n\epsilon}\right)^{\frac{\tau}{2-\tau}} \vee \frac{d}{m}\right)$	$\tilde{O}\left(\left(\frac{d^{3/2}\theta^{1/2}}{n\epsilon}\right)^{\frac{\tau}{2-\tau}} \vee \frac{d}{m}\right)$
Agnostic (vs h^*)	$\Omega(\text{Err}(h^*))$ required.	$13\text{Err}(h^*) + \tilde{O}\left(\frac{d^{3/5}}{n^{2/5}\epsilon^{2/5}} \vee \sqrt{\frac{d}{m}}\right)$	$\Omega(\text{Err}(h^*))$ required.	$\Omega(\text{Err}(h^*))$ required.
Agnostic (vs h_∞^{agg})	-	-	Consistent under weaker conditions.	-

- What do these results say:
 - Sparse-Vector-Technique-based online query release helps
 - Active learning helps
 - Label-noise models help

Liu, Zhu, Chaudhuri and W. (2020) “Revisiting model-agnostic private learning”. AISTATS and JMLR. <https://arxiv.org/pdf/2011.03186.pdf>

Next lecture

- Going beyond the worst case!
- Smoothed Sensitivity and the Median