

493 **A Preliminaries**

Definition 1 (Differential Privacy). [DKM⁺06, DMNS06] A randomized algorithm \mathcal{M} achieves (ε, δ) -DP if for all $\mathcal{S} \subseteq \text{Range}(\mathcal{M})$ and for any two database instances $D, D' \in \mathcal{D}$ that differ only in one tuple:

$$\Pr[\mathcal{M}(D) \in \mathcal{S}] \leq e^\varepsilon \Pr[\mathcal{M}(D') \in \mathcal{S}] + \delta.$$

494 The privacy cost is measured by the parameters (ε, δ) also referred to as the privacy budget. Smaller
 495 values of ε correspond to stricter privacy guarantees, and it is standard in literature to set $\delta \ll \frac{1}{n}$,
 496 where n is the size of the database. We set the δ_f in our work to $\frac{1}{n}$ scaled down to the nearest
 497 power of 10. Complex DP algorithms can be built from the basic algorithms following two important
 498 properties of differential privacy: 1) Post-processing states that for any function g defined over the
 499 output of the mechanism \mathcal{M} , if \mathcal{M} satisfies (ε, δ) -DP, so does $g(\mathcal{M})$; 2) Basic composition states
 500 that if for each $i \in [k]$, mechanism \mathcal{M}_i satisfies $(\varepsilon_i, \delta_i)$ -DP, then a mechanism sequentially applying
 501 $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_k$ satisfies $(\sum_{i=1}^k \varepsilon_i, \sum_{i=1}^k \delta_i)$ -DP.

502 Given a function $f : \mathcal{D} \rightarrow \mathbb{R}^d$, the *Gaussian mechanism* adds noise drawn from a normal distribution
 503 $\mathcal{N}(0, S_f^2 \sigma^2)$ to each dimension of the output, where S_f is the ℓ_2 -sensitivity of f , defined as $S_f =$
 504 $\max_{D, D' \text{ differ in a row}} \|f(D) - f(D')\|_2$. For $\varepsilon \in (0, 1)$, if $\sigma \geq \sqrt{2 \ln(1.25/\delta)}/\varepsilon$, then the Gaussian
 505 mechanism satisfies (ε, δ) -DP.

506 The Gaussian mechanism is used to privatize optimization algorithms. In contrast to non-private
 507 optimizers where batches are sliced from the training dataset, DP optimizers at each iteration work
 508 by sampling “lots” from the training with probability L/n , where L is the (expected) lot size and n is
 509 the total data size. A set of queries are computed over those samples. These queries include gradient
 510 computation, updates to batch normalization or accuracy metric calculations. As there is not any a
 511 priori bound on these query outputs, the sensitivity S_f is set by clipping the maximum ℓ_2 norm of the
 512 gradient to a user-defined parameter C . The gradient of each point is then noised and published. All
 513 DP optimizers follow the same framework in which they take steps on the computed noisy gradient
 514 as in its non-private counterpart [MAE⁺18]. The privacy cost of the whole training procedure is
 515 calculated by advanced composition techniques such as the Moments accountant [ACG⁺16].

516 **A.1 DP Optimizers**

517 **DP-SGD:** The most popular private optimizer is the differentially private stochastic gradient descent
 518 (DPSGD) [WM10, BST14, SCS13, ACG⁺16]. DPSGD takes individual steps for each point in
 519 the sampled lot just like in SGD. Due to these individual steps, SGD is more locally unstable and
 520 empirically generalizes better than other optimizers [ZFM⁺20]. However, SGD requires the learning
 521 rate to be properly tuned when changing architectures or datasets, without which SGD may show
 522 subpar performance.

523 There are five main hyperparameters involved in DPSGD. We start with those also present in the
 524 non-private setting, highlighting any differences that arise due to privacy.

- 525 • Training iterations (T) - In the private setting, more iterations results in a larger privacy cost.
- 526 • Lot size (L) - Lot size factors into the privacy calculation, due to amplification by subsam-
 527 pling [BBG18].
- 528 • Learning rate (α) - Learning rate has an important interplay with the clipping threshold C ,
 529 discussed in Section 4.1.

530 The following hyperparameters are new in the private setting.

- 531 • Clipping threshold (C) - To limit sensitivity, per-example gradients are clipped to have
 532 ℓ_2 -norm bounded by C .
- 533 • Noise scale (σ) - Scale of the noise added, as a multiple of C . A larger value gives higher
 534 privacy but (typically) lower accuracy.

535 **DPMomentum:** The private counterpart of SGD-Momentum [RHW86, Qia99], which adds the
 536 momentum parameter to the update rule of DPSGD [GAYB17]. This optimizer adds an extra
 537 hyperparameter to tune as no default value for momentum is known.

538 **DP-Adam:** Adam [KB14] is an adaptive optimizer that combines the advantages from Ada-
539 Grad [DHS11] and RMSProp [HSS12]. At the core of Adam, exponentially averaged first and
540 second moment estimates of the gradients are used to take a step. Converting Adam to its differ-
541 entially private counterpart DPAdam can be done trivially by replacing the standard gradients with their
542 clipped and noised counterparts. Adam adds two extra hyperparameters (β_1, β_2) to tune in the DP
543 setting. However, default values of these parameters are known in the non-private setting. We will
544 tune these parameters to the private setting in Section 4.2. The adaptivity of these optimizers imply
545 they need not be tuned across learning rates, hence reducing a hyperparameter to tune.

546 **ADADP:** This DP adaptive optimizer finds the best learning rate at every alternate iteration [KH20].
547 It does so by leveraging the ℓ_2 error of taking a full step and taking two half steps. If the error
548 computed is greater than a threshold τ , the learning rate is updated using a closed form expression.
549 As suggested by the authors, for all our experiments using ADADP, we use the threshold $\tau = \sqrt{\frac{d}{2T}}$,
550 where d is the model dimension and T is the total number of iterations.

551 B Dataset details

Table 2: Datasets used in experiment

Dataset	Type	#Samples	#Dims	#Classes
MNIST	Image	70000	784	10
Gisette	Image	6000	5000	2
Adult	Structured	45222	202	2
ENRON	Structured	5172	5512	2

552 C Parameter grid for DPSGD and DPAdam comparison

Table 3: Parameter grid for comparing DPSGD and DPAdam

Optimizer	Parameter	Values
DPSGD	α	0.001, 0.002, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1
	C	0.1, 0.2, 0.5, 1
DPMomentum	α	0.001, 0.002, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1
	C	0.1, 0.2, 0.5, 1
	m	0.5, 0.6, 0.7, 0.8, 0.9, 0.95, 0.99
DPAdam	C	0.1, 0.2, 0.5, 1

553 D Proof of Theorem

554 **Theorem 4.** Let f be a convex and β -smooth function, and let $x^* = \arg \min_{x \in \mathcal{S}} f(x)$. Let x_0 be
555 an arbitrary point in \mathcal{S} , and $x_{t+1} = \Pi_{\mathcal{S}}(x_t - \alpha(g_t + z_t))$, where $g_t = \min(1, \frac{C}{\|\nabla f(x)\|^2}) \nabla f(x)$
556 and $z_t \sim \mathcal{N}(0, \sigma^2 C^2)$ is the noise due to privacy. After T iterations, the optimal learning rate is
557 $\alpha_{opt} = \frac{R}{CT\sqrt{1+\sigma^2}}$, where $\mathbb{E}[f(\frac{1}{T} \sum_i x_t) - f(x^*)] \leq \frac{RC\sqrt{1+\sigma^2}}{\sqrt{T}}$ and $R = \mathbb{E}[\|x_0 - x^*\|]$.

Proof.

$$\begin{aligned}
\mathbb{E}[\|x_{t+1} - x^*\|^2 \mid x_t] &= \mathbb{E}[\|x_t - \alpha(g_t + z_t) - x^*\|^2 \mid x_t] \\
&= \mathbb{E}[\|x_t - x^*\|^2 - 2\alpha(g_t + z_t)(x_t - x^*) + \alpha^2\|(g_t + z_t)\|^2 \mid x_t] \\
&= \|x_t - x^*\|^2 - 2\alpha \mathbb{E}[(g_t + z_t) \mid x_t]^T (x_t - x^*) + \alpha^2 \mathbb{E}[\|(g_t + z_t)\|^2 \mid x_t] \\
&\leq \|x_t - x^*\|^2 - 2\alpha[f(x_t) - f(x^*)] + \alpha^2 \mathbb{E}[\|(g_t + z_t)\|^2 \mid x_t]
\end{aligned}$$

558 The inequality is due to convexity of the loss function and $E[(g_t + z_t)] = g_t$ due to 0-mean noise.
559 Taking expectation on both sides and reordering,

$$\begin{aligned}
2\alpha[f(x_t) - f(x^*)] &\leq \mathbb{E}[\|x_{t+1} - x^*\|^2] - \mathbb{E}[\|x_t - x^*\|^2] + \alpha^2 \mathbb{E}[\|(g_t + z_t)\|^2] \\
&\leq \mathbb{E}[\|x_{t+1} - x^*\|^2] - \mathbb{E}[\|x_t - x^*\|^2] + \alpha^2(C^2 + C^2\sigma^2)
\end{aligned}$$

560 Summing for T steps and dividing both sides by $2\alpha T$,

$$\mathbb{E}[f(\frac{1}{T} \sum_i x_t) - f(x^*)] \leq \frac{R^2}{2\alpha T} + \frac{\alpha C^2(1 + \sigma^2)}{2} \tag{1}$$

561 Taking derivative and finding best value of α ,

$$\alpha_{opt} = \frac{R}{C\sqrt{1 + \sigma^2}T}$$

562 Plugging α_{opt} to Eq. 1,

$$\mathbb{E}[f(\frac{1}{T} \sum_i x_t) - f(x^*)] \leq \frac{RC\sqrt{1 + \sigma^2}}{\sqrt{T}}$$

563 □

564 E LT Algorithm

565

Algorithm 1 Hard stopping private selection algorithm for (ϵ, δ) -DP input algorithms

Require: $\gamma \leq 1$, $\delta_2 > 0$, and sampling access to $Q(D)$

- 1: Initialize the list $S = \emptyset$
 - 2: Initialize $\Upsilon = \frac{1}{\gamma} \log \frac{1}{\delta_2}$
 - 3: **for** $j \in [1, \Upsilon]$ **do**
 - 4: Draw $(x, q) \sim Q(D)$
 - 5: $S \leftarrow S \cup (x, q)$
 - 6: Flip a γ -biased coin, output highest scored candidate from S and halt;
 - 7: **end for**
 - 8: Output highest scored candidate from S
-

566 F LT vs MA with varying candidate size

567 Continuing from Section 3.1, in this section we show an additional experiment in which we compare
568 the LT (Liu and Talwar) and MA (Moments Accountant) algorithms with varying number of hyper-
569 parameter candidates. In Figure 6, we run the LT and MA algorithms for $T = 10000$ with $\sigma = 4$
570 and $L = 250$ with varying candidate size and compare the final privacy costs. The γ value for the
571 LT algorithm is set to $1/k$, where the k is the number of candidates. It can be seen that the privacy
572 cost of LT (blue) remains almost constant for with increasing number of candidates. Figure 6 also
573 demonstrates the exact number of candidates when the cost of MA (orange) remains below the LT
574 cost. This insight is valuable in practice to a practitioner to decide the which algorithm to choose for
575 hyperparameter tuning with respect to the number of candidates.

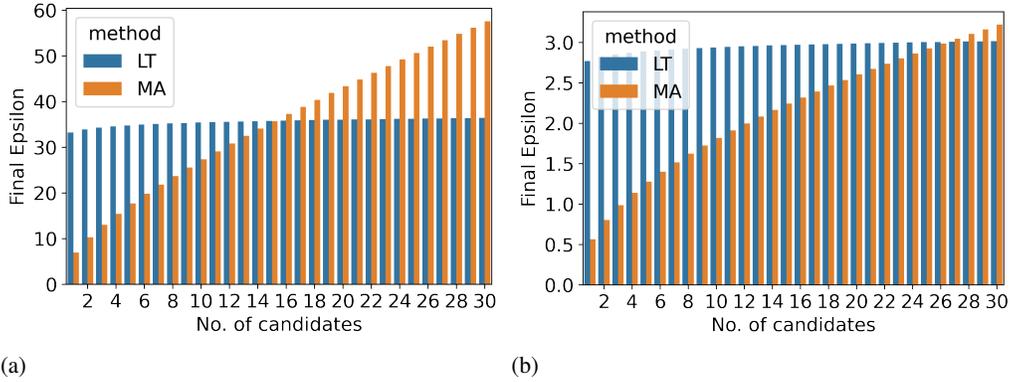


Figure 6: Comparing LT vs MA with varying number of candidates at setting $\sigma = 4$, $L = 250$, $T = 10000$. MA can compose up to 14 and 26 candidates for the same cost of LT for dataset sizes 5k (left) and 60k (right) respectively.

576 G Pruning hyperparameter grid for SGD

577 Figure 7 demonstrates a heat map plot of the candidate hyperparameter pairs for DPSGD. Each point
 578 on this heatmap is assigned a score (totalling 2400) that reflects how many times that (α, C) pair has
 579 performed the best among all the candidates, and we score across all iterations (at a granularity of
 580 every 100 iterations) of training.

581 We justify this as a fair metric of ‘goodness’, for candidates as one could in practice stop training at any
 582 iteration. Furthermore this metric is quite critical of quality, in that it only awards a hyperparameter
 583 set a point, if it appeared as the best candidate at one of the intervals. Hence we deem this to be a
 584 generous pruning of the search space, which will imbue the best possible advantage to DPSGD with
 585 regards to a pruned hyperparameter search space.

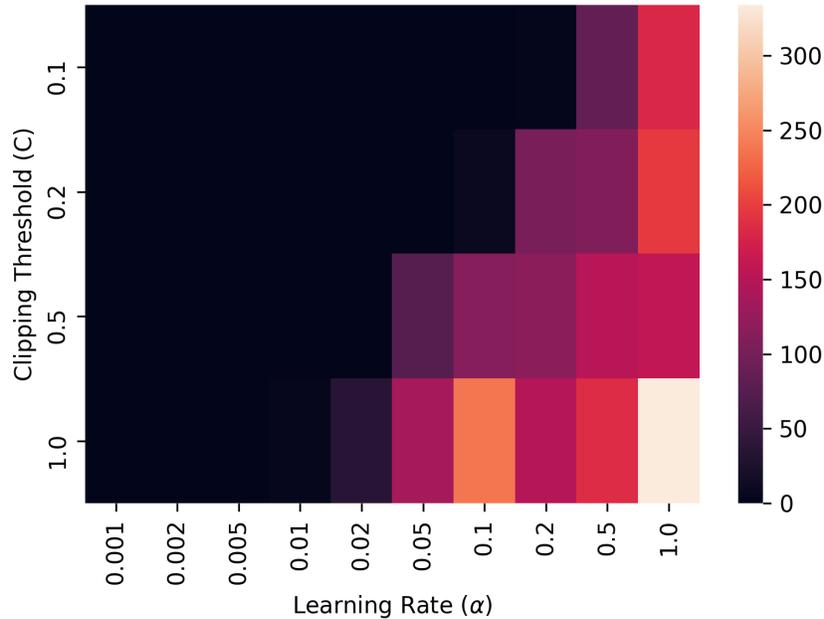


Figure 7: Pruning for DPSGD. Each (α, C) point on the heatmap shows how many times it has performed best among all candidates

586 H Implementation Details

587 The code for our paper is written in Python3.6 using the PyTorch library. The implementation of all
588 private optimizers are done using the Pylvacy library². We run our code on ComputeCanada servers.
589 Each allocation of the server includes 2 CPU cores, 8 GB RAM and 1 GPU from the list – P100,
590 V100, K80. We report results from all our experiments after averaging over 3 runs. The code is
591 attached with our supplementary material submission.

592 All datasets used in our experiments are publicly available. We split all datasets into 80% training
593 and 20% validation sets. For our experiments, we assume that all our datasets start in its preprocessed
594 state, i.e. the numerical features are scaled to the range [0,1], as is standard practice in machine
595 learning. However when considering an end-to-end private algorithm, this preprocessing itself may
596 need to be performed in a privacy-preserving fashion. In this work, we do not account for privacy in
597 this step. Note that for our work this only effects the ENRON and Adult datasets, where scaling the
598 values does require computing the maximum possible values of features in a differentially-private
599 fashion, whereas the max values for image datasets (Gisette and MNIST) are known a priori due to
600 max pixel value and does not involve any privacy cost.

601 I Additional experiment results for Section 5 and Section 6

602 In Figures 8 and 9, we display our results for the same experiments described in Section 5, with
603 $\sigma = 2$, and $\sigma = 8$ respectively. Similarly Figure 10 and 11 displays our results of the experiments
604 detailed in Section 6 with $\sigma = 2$, and $\sigma = 8$.

605 J Omitted Pseudocode for DPAdamWOSM

Algorithm 2 Optimization using DPAdamWOSM

Require: Training set $A : \{x_1, \dots, x_n\}$, Loss function $\mathcal{L}(\theta)$, Parameters: Lot size L , Learning rate α ,
Gradient norm bound C , Noise scale σ , Total number of iterations T , Exponential decay rate β_1

- 1: Initialize model with θ_0 randomly
 - 2: Initialize first moment vector $m_0 = 0$
 - 3: Set learning rate to ESS $\alpha = \frac{10^{-3}}{(\sigma C/L)+10^{-8}}$;
 - 4: **for** $t \in [1, T]$ **do**
 - 5: Sample a random subset $L_t \subseteq A$, by independently including each element of A with
probability L/n
 - 6: Compute gradient $\forall x_i \in L_t$
 $g_t(x_i) = \nabla_{\theta} \mathcal{L}(\theta_t, x_i)$
 - 7: Clip each gradient in ℓ_2 norm to C $\bar{g}_t(x_i) = g_t(x_i) / \max(1, \frac{\|g_t(x_i)\|_2}{C})$
 - 8: Add noise $\tilde{g}_t = \frac{1}{|L_t|} (\sum_i \bar{g}_t(x_i) + \mathcal{N}(0, \sigma^2 C^2 I))$
 - 9: Exponentially average the first moment
 $m_t = \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot \tilde{g}_t$
 - 10: Perform bias correction
 $\hat{m}_t = \frac{m_t}{1 - \beta_1^t}$
 - 11: Update model $\theta_t = \theta_{t-1} - \alpha \cdot \hat{m}_t$
 - 12: **end for**
 - 13: Compute privacy cost using Moments Accountant.
-

606 Broader Impacts

607 Our work points out a false sense of security afforded by prior work in the space of differentially
608 private machine learning, as true privacy losses are much larger than what is typically reported in
609 papers. That said, regardless, differential privacy is a very difficult topic to properly deploy and
610 genuinely provide its theoretical guarantees, rather than just a mirage of privacy. These issues can

²<https://github.com/ChrisWaites/pyvacy>

611 be avoided via training and/or consultation with data privacy experts, although this may be more
 612 challenging for smaller, resource-constrained organizations.

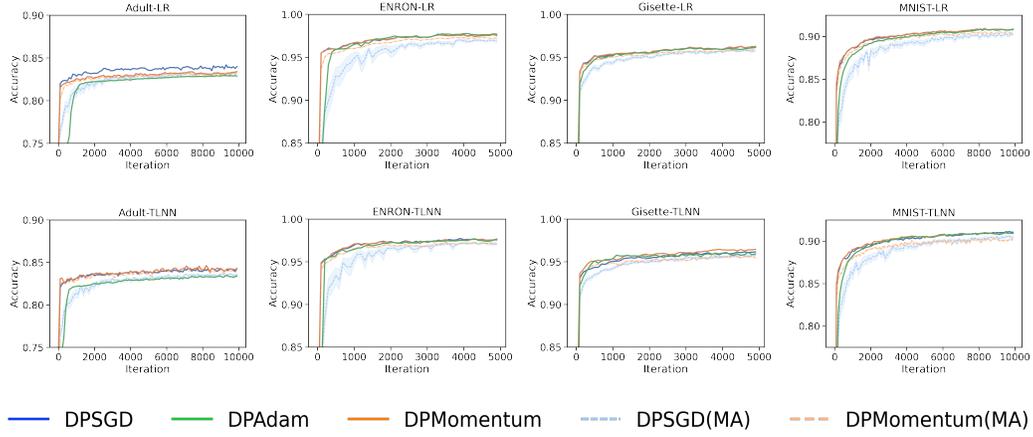


Figure 8: Comparing the testing accuracy curves of DPAdam and DPSGD models across hyperparameter tuning grid from Table 3 with $\sigma = 2$.

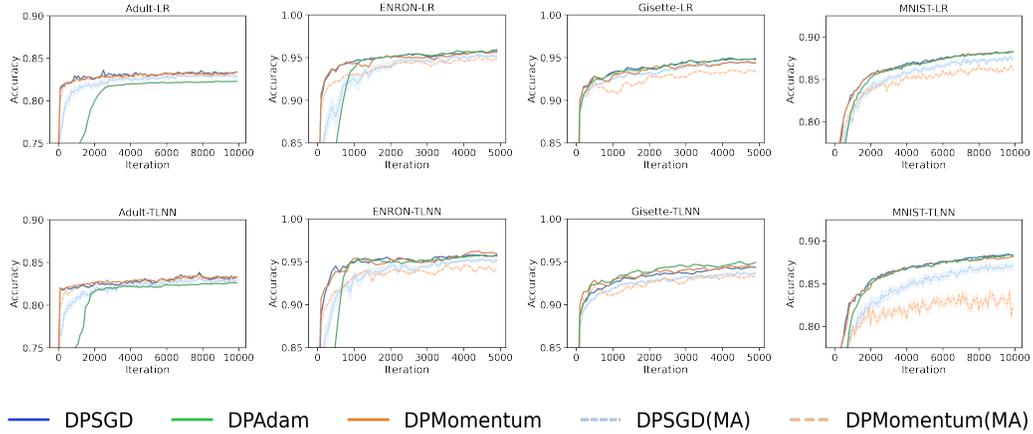


Figure 9: Comparing the testing accuracy curves of DPAdam and DPSGD models across hyperparameter tuning grid from Table 3 with $\sigma = 8$.

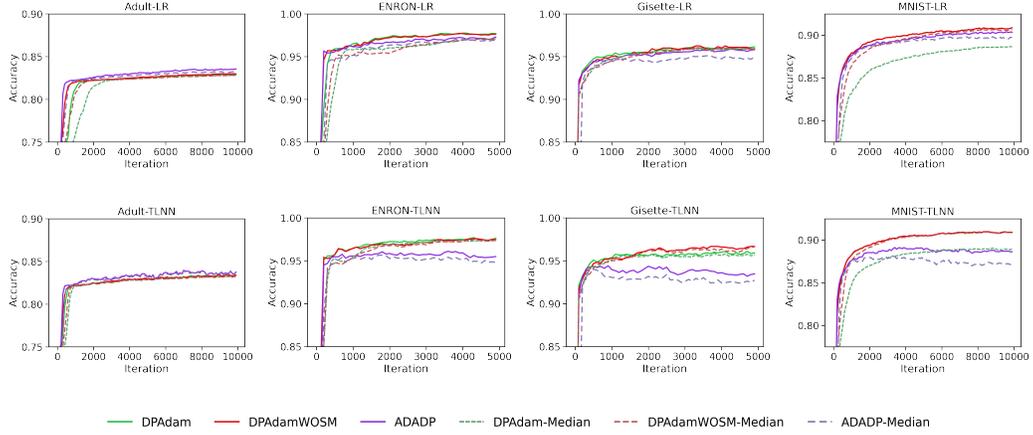


Figure 10: Comparing the testing accuracy curves of DPAdam, ADADP and DPAdamWOSM models across hyperparameter tuning grid from Table 3 with $\sigma = 2$. The limits for the y-axes are adjusted based on the dataset while maintaining a 15% range for all.

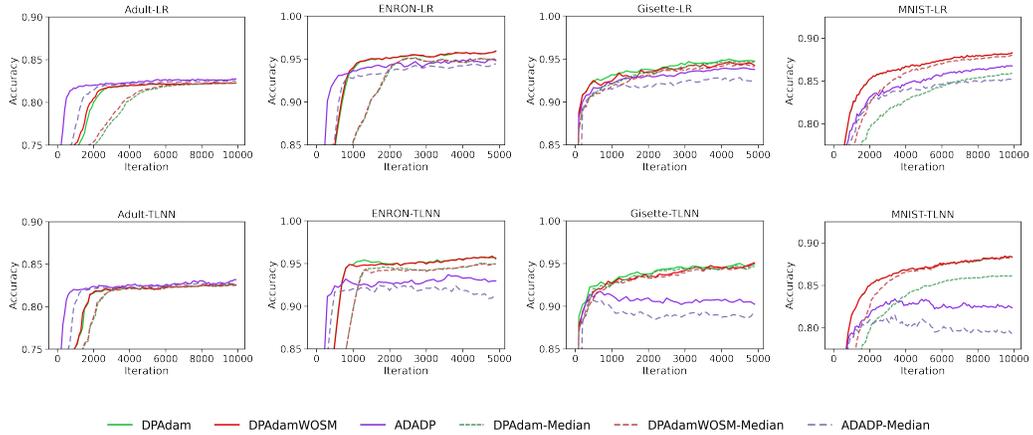


Figure 11: Comparing the testing accuracy curves of DPAdam, ADADP and DPAdamWOSM models across hyperparameter tuning grid from Table 3 with $\sigma = 8$. The limits for the y-axes are adjusted based on the dataset while maintaining a 15% range for all.