

## 339 A Relation to Koh and Liang [2017]

### 340 A.1 Brief Review

341 As we mentioned in Section 6, Koh and Liang [2017] proposed to estimate the influence by (5), which  
342 is

$$\hat{\theta}_{-j} - \hat{\theta} \approx \frac{1}{N} \hat{H}^{-1} \nabla_{\theta} \ell(z_j; \hat{\theta}),$$

343 where  $\hat{H} = \frac{1}{N} \sum_{z \in D} \nabla^2 \ell(z; \hat{\theta})$  is the Hessian of the problem (1) for the optimal model  $\hat{\theta}$ .

344 Note that,  $\hat{H}^{-1} \nabla_{\theta} \ell(z_j; \hat{\theta})$  is equivalent to the solution to the following optimization problem:

$$\operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2} \langle \beta, \hat{H} \beta \rangle - \langle \nabla_{\theta} \ell(z_j; \hat{\theta}), \beta \rangle. \quad (6)$$

345 Koh and Liang [2017] proposed computing  $\hat{H}^{-1} \nabla_{\theta} \ell(z_j; \hat{\theta})$  by solving this optimization problem using  
346 conjugate gradient descent or its improved version. In the optimization, they also proposed to use  
347 the mini-batch approximation of the Hessian matrix: they proposed to use  $\hat{H}_S = \frac{1}{|S|} \sum_{z \in S} \nabla^2 \ell(z; \hat{\theta})$   
348 on the mini-batch  $S \subset D$  instead of the Hessian matrix  $\hat{H}$  computed on the all training instances  $D$ .

### 349 A.2 Relation to the Proposed Method

350 Here, we show the relationship between the proposed method and the method of Koh and Liang  
351 [2017]. Suppose that we solve the problem (6) using SGD. In the  $t$ -th step of SGD, we update  $\beta$  by

$$\beta^{[t+1]} = \beta^{[t]} - \gamma_t (\hat{H}_{S_t} \beta^{[t]} - \nabla_{\theta} \ell(z_j; \hat{\theta})) = (I - \gamma_t \hat{H}_{S_t}) \beta^{[t]} + \gamma_t \nabla_{\theta} \ell(z_j; \hat{\theta}),$$

352 where  $S_t$  is the mini-batch and  $\gamma_t > 0$  is a learning rate. Suppose that we initialized  $\beta^{[1]} = \nabla_{\theta} \ell(z_j; \hat{\theta})$   
353 and  $\gamma := \max_t \gamma_t$ . Then, the SGD for the problem (6) can be expressed as

$$\begin{aligned} \beta^{[2]} &= (I - \gamma_1 \hat{H}_{S_1}) \nabla_{\theta} \ell(z_j; \hat{\theta}) + \gamma_1 \nabla_{\theta} \ell(z_j; \hat{\theta}), \\ \beta^{[3]} &= (I - \gamma_2 \hat{H}_{S_2}) \beta^{[2]} + \gamma_2 \nabla_{\theta} \ell(z_j; \hat{\theta}) = (I - \gamma_2 \hat{H}_{S_2}) (I - \gamma_1 \hat{H}_{S_1}) \nabla_{\theta} \ell(z_j; \hat{\theta}) + O(\gamma), \\ &\vdots \\ \beta^{[T]} &= (I - \gamma_{T-1} \hat{H}_{S_{T-1}}) (I - \gamma_{T-2} \hat{H}_{S_{T-2}}) \dots (I - \gamma_1 \hat{H}_{S_1}) \nabla_{\theta} \ell(z_j; \hat{\theta}) + O(\gamma). \end{aligned}$$

354 Here, let  $\hat{Z}_t := I - \gamma_t \hat{H}_{S_t}$ , and we obtain

$$\hat{\theta}_{-j} - \hat{\theta} \approx \frac{1}{N} \hat{H}^{-1} \nabla_{\theta} \ell(z_j; \hat{\theta}) \approx \beta^{[T]} = \frac{1}{N} \hat{Z}_{T-1} \hat{Z}_{T-2} \dots \hat{Z}_1 \nabla_{\theta} \ell(z_j; \hat{\theta}) + O\left(\frac{\gamma}{N}\right).$$

355 When  $\gamma$  is small and the last term is ignorable, this result resembles to the proposed estimator  $\Delta \theta_{-j}$   
356 in Section 4. Instead of  $\hat{Z}_t := I - \gamma_t \hat{H}_{S_t}$  computed at the optimal model  $\hat{\theta}$ , the proposed estimator  
357 uses  $Z_t = I - \eta_t H^{[t]}$  based on the model  $\theta^{[t]}$  in the  $t$ -th SGD step in the training.

## 358 B Proof of Theorems

359 Before proving Theorems 5 and 6, we first prove the next lemma.

360 **Lemma 7.** Assume that  $\ell(z; \theta)$  is twice differentiable with respect to the parameter  $\theta$ , and assume  
361 that there exist  $\lambda, \Lambda > 0$  such that  $\lambda I \prec \nabla_{\theta}^2 \ell(z; \theta) \prec \Lambda I$  for all  $z, \theta$ . If  $\eta_s \leq 1/\Lambda$ , then we get

$$h_j(\Lambda) \leq \|\Delta \theta_{-j}\| \leq h_j(\lambda), \quad (7)$$

$$h_j(\Lambda) \leq \|\theta_{-j}^{[T]} - \theta^{[T]}\| \leq h_j(\lambda), \quad (8)$$

362 where  $h_j(a) := \frac{\eta_{\pi(j)}}{|\mathcal{S}_{\pi(j)}|} \prod_{s=\pi(j)+1}^{T-1} (1 - \eta_s a) \|g(z_j; \theta^{[\pi(j)]})\|$ .

363 *Proof.* Since  $(1 - \eta_s \Lambda)I \prec Z_s \prec (1 - \eta_s \lambda)I$  we immediately obtain (7) from the definition (2) of  
 364  $\Delta\theta_{-j}$ .

365 We next show the inequality (8). There exists  $r \in [0, 1]$  such that for  $\theta_*^{[s]} := r\theta_{-j}^{[s]} + (1 - r)\theta^{[s]}$ ,

$$\frac{1}{|S_s|} \sum_{i \in S_s} \left( \nabla_{\theta} \ell(z_i; \theta_{-j}^{[s]}) - \nabla_{\theta} \ell(z_i; \theta^{[s]}) \right) = H_*^{[s]} (\theta_{-j}^{[s]} - \theta^{[s]}),$$

366 where  $H_*^{[s]} := \frac{1}{|S_s|} \sum_{i \in S_s} \nabla_{\theta}^2 \ell(z_i; \theta_*^{[s]})$ . Therefore, by setting  $Z_s^* := (I - \eta_s H_*^{[s]})$ , we can show the  
 367 inequality (8) in a similar way to the proof of (7).  $\square$

### 368 B.1 Proof of Theorem 5

369 *Proof.* From Lemma 7

$$\begin{aligned} \|(\theta_{-j}^{[T]} - \theta^{[T]}) - \Delta\theta_{-j}\|^2 &= \|\theta_{-j}^{[T]} - \theta^{[T]}\|^2 + \|\Delta\theta_{-j}\|^2 - 2\langle \theta_{-j}^{[T]} - \theta^{[T]}, \Delta\theta_{-j} \rangle \\ &\leq h_j(\lambda)^2 + h_j(\lambda)^2 + 2h_j(\Lambda)^2 = 2(h_j(\lambda)^2 + h_j(\Lambda)^2). \end{aligned}$$

370  $\square$

### 371 B.2 Proof of Theorem 6

*Proof.*

$$\theta_{-j}^{[s+1]} - \theta^{[s+1]} = Z_s(\theta_{-j}^{[s]} - \theta^{[s]}) + \eta(H^{[s]} - H_*^{[s]})(\theta_{-j}^{[s]} - \theta^{[s]}),$$

372 where  $H_*^{[s]}$  is the same as that in the proof of Lemma 7. We set  $D_s := \eta(H^{[s]} - H_*^{[s]})(\theta_{-j}^{[s]} - \theta^{[s]})$ .

373 Applying this equalities recursively over  $s \in \{\pi(j), \dots, T-1\}$ , we get

$$\theta_{-j}^{[T]} - \theta^{[T]} = \Delta\theta_{-j} + \sum_{s=\pi(j)}^{T-1} \prod_{k=s+1}^{T-1} Z_k D_s.$$

374 Hence, a remaining problem is to bound the norm of the second term in the right hand side of this  
 375 equality, which corresponds to a gap we want to evaluate. Since  $\|Z_k\| \leq 1 + \eta\Lambda$ ,  $\|\theta_{-j}^{[s]} - \theta^{[s]}\| \leq 2\eta GT$   
 376 and  $\|H^{[s]} - H_*^{[s]}\| \leq L\|\theta_{-j}^{[s]} - \theta^{[s]}\|$ ,

$$\begin{aligned} \left\| \sum_{s=\pi(j)}^{T-1} \prod_{k=s+1}^{T-1} Z_k D_s \right\| &\leq \sum_{s=1}^{T-1} \prod_{k=s+1}^{T-1} \|Z_k\| \|D_s\| \leq \sum_{s=1}^{T-1} (1 + \eta\Lambda)^{T-s-1} \eta L \|\theta_{-j}^{[s]} - \theta^{[s]}\|^2 \\ &= 4 \frac{(1 + \eta\Lambda)^{T-1} - 1}{(1 + \eta\Lambda) - 1} \eta^3 T^2 G^2 L \leq 4 \frac{(1 + O(\gamma\Lambda/\sqrt{T}))^T}{\Lambda} \gamma^2 T G^2 L. \end{aligned}$$

377  $\square$

## 378 C Details and Results of Experiments

### 379 C.1 Setups in Section 7.1

380 **Datasets** We used three datasets: Adult [Dua and Karra Taniskidou, 2017], 20Newsgroups<sup>6</sup> and  
 381 MNIST [LeCun *et al.*, 1998]. These are common benchmarks in tabular data analysis, natural  
 382 language processing, and image recognition, respectively. We adopted these three datasets to  
 383 demonstrate the validity of the proposed algorithm across different data domains.

384 We preprocessed each dataset as follows. In Adult, we transformed categorical features to numerical  
 385 attributes [7]. In 20Newsgroups, we selected the two document categories `ibm.pc.hardware` and  
 386 `mac.hardware`. As a preprocessing, we transformed the documents into numerical vectors using  
 387 tf-idf, while removing frequent and scarce words. In MNIST, we selected the images from the two  
 388 categories one and seven, so that the problem to be binary classification.

<sup>6</sup><http://qwone.com/~jason/20Newsgroups/>

<sup>7</sup>We used the implementation available at <https://www.kaggle.com/kost13/us-income-logistic-regression/notebook>

389 **Models** To see the validity of the proposed  
 390 method beyond convexity, we adopted two mod-  
 391 els, which are linear logistic regression and deep  
 392 neural networks. For linear logistic regression,  
 393 we adopted the  $\ell_2$ -regularized loss  $\ell(z; \theta) =$   
 394  $\log(\exp(-y\langle \theta, x \rangle) + 1) + \frac{\alpha}{2} \|\theta\|^2$  where  $y \in$   
 395  $\{-1, 1\}$ . In the experiments, we determined the  
 396 regularization parameter  $\alpha$  using cross validation.  
 397 For deep neural networks, we used a network with  
 398 two fully connected layers each of which has eight  
 399 units with ReLU as an activation function. We used the sigmoid function at the output layer, and  
 400 adopted the cross entropy as the loss function. To run SGD, we used the parameters shown in Table 2.  
 401 We note that the loss function for the linear logistic regression is convex, while that for the deep  
 402 neural networks is non-convex.

Table 2: Parameters used in SGD:  $K$  denotes the number of epochs.  $|S_t|$  denotes the batch size.

	LogReg			DNN		
	$K$	$ S_t $	$\eta_t$	$K$	$ S_t $	$\eta_t$
Adult	20	5	$\frac{0.1}{\sqrt{t}}$	10	20	0.1
20News	10	5	$\frac{0.01}{\sqrt{t}}$	10	20	0.1
MNIST	5	5	$\frac{0.1}{\sqrt{t}}$	10	20	0.1

403 **Target Linear Influence** In the experiments, we randomly subsampled 200 instances for the  
 404 training set  $D$  and the validation set  $D'$ . We then estimated the linear influence for the validation loss  
 405 using Algorithm 2. Here, we set the query vector  $u$  as  $u = \frac{1}{|D'|} \sum_{z' \in D'} \nabla_{\theta} \ell(z'; \theta^{[T]})$ . Estimation of  
 406 the linear influence thus amounts to estimating the change in the validation loss

$$\langle u, \theta_{-j}^{[T]} - \theta^{[T]} \rangle \approx \frac{1}{|D'|} \sum_{z' \in D'} \left( \ell(z'; \theta_{-j}^{[T]}) - \ell(z'; \theta^{[T]}) \right).$$

407 We note that the instances with large negative linear influences are deemed to be negatively affecting  
 408 the resulting models. Removing such instances can improve the validation loss, and thus the users  
 409 can prioritize the inspection of such instances.

410 **Baseline Method** We adopted the method of Koh and Liang [2017] as the baseline, which we  
 411 abbreviated as K&L. In K&L, we estimate the influence by (5), which is

$$\hat{\theta}_{-j} - \hat{\theta} \approx \frac{1}{N} \hat{H}^{-1} \nabla_{\theta} \ell(z_j; \hat{\theta}),$$

412 where  $\hat{H} = \frac{1}{N} \sum_{z \in D} \nabla^2 \ell(z; \hat{\theta})$  is the Hessian of the problem (1) for the optimal model  $\hat{\theta}$ . For a  
 413 query vector  $u \in \mathbb{R}^p$ , the linear influence can be estimated as

$$\langle \hat{\theta}_{-j} - \hat{\theta}, u \rangle \approx \frac{1}{N} \langle \hat{H}^{-1} \nabla_{\theta} \ell(z_j; \hat{\theta}), u \rangle = \frac{1}{N} \langle \nabla_{\theta} \ell(z_j; \hat{\theta}), \hat{H}^{-1} u \rangle.$$

414 Here, the last equality follows from the symmetricity of the Hessian matrix. Thus, for estimating the  
 415 linear influence for all the training instances, we first compute  $\hat{H}^{-1} u$ , and then take an inner product  
 416 with  $\nabla_{\theta} \ell(z_j; \hat{\theta})$  for each training instance  $z_j \in D$ .

417 Note that,  $\hat{H}^{-1} u$  is equivalent to the solution to the following optimization problem:

$$\operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2} \langle \beta, \hat{H} \beta \rangle - \langle u, \beta \rangle. \quad (9)$$

418 Koh and Liang [2017] proposed computing  $\hat{H}^{-1} u$  by solving this optimization problem using  
 419 conjugate gradient descent or its improved version. In the optimization, they also proposed to use the  
 420 mini-batch approximation of the Hessian matrix: they proposed to use  $\hat{H}_S = \frac{1}{|S|} \sum_{z \in S} \nabla^2 \ell(z; \hat{\theta})$   
 421 on the mini-batch  $S \subset D$  instead of the Hessian matrix  $\hat{H}$  computed on the all training instances  
 422  $D$ . In the experiment, we used the implementation available at [https://github.com/kohpangwei/](https://github.com/kohpangwei/influence-release)  
 423 [influence-release](https://github.com/kohpangwei/influence-release).

424 **Evaluation Metrics** In the experiments, we ran the counterfactual SGD for all  $z_j \in D$ , and  
 425 computed the true linear influence. We then used this ground truth to evaluate the goodness of the  
 426 estimated linear influences. For evaluation, we adopted the following two metrics. The first metric  
 427 is Kendall’s tau. Kendall’s tau is a typical metric for measuring ordinal associations between two  
 428 observations. Kendall’s tau takes the value between plus and minus one, where the value one indicates  
 429 that the orders of the two observations are identical.

430 The second metric is Jaccard index. For data cleansing, the users are interested in instances with  
 431 large positive or negative influences. We measured how accurately those important instances could be  
 432 identified using the estimated influences. To this end, we selected 10 instances with largest positive  
 433 and negative true influences, and constructed a set of 20 important instances. We compared this  
 434 true important instances with the estimated important instances using Jaccard index. Jaccard index  
 435 measures the similarity of the two sets. Jaccard index takes the value between zero and one, where  
 436 the value one indicates that the sets are identical.

## 437 C.2 Setups in Section 7.2

438 **Datasets** We used MNIST [LeCun et al., 1998] and CIFAR10 [Krizhevsky and Hinton, 2009].  
 439 The MNIST dataset contains 60,000 training instances, while the CIFAR10 dataset contains 50,000  
 440 training instances. Both datasets also contain 10,000 test instances. From the original training  
 441 instances, we held out randomly selected 10,000 instances for the validation set, and used the  
 442 remaining instances as the training set. Thus, in the experiment, we used 50,000 instances in MNIST  
 443 and 40,000 instances in CIFAR10 for training, and the held out 10,000 instances for validation.

444 **Models** We used convolutional neural networks (CNNs) in the experiment. The network structures  
 445 can be found in Figure 4. In SGD, we set the epochs  $K = 20$ , batch size  $|S_t| = 64$ , and learning rate  
 446  $\eta_t = 0.05$ . In the training, we used a simple data augmentation. For MNIST, we applied horizontal  
 447 and vertical shifts in  $\pm 2$  pixels. For CIFAR10, we applied horizontal and vertical shifts in  $\pm 4$  pixels  
 448 and horizontal flipping.

449 **Target Linear Influence** We set the query vector  $u$  as  $u = \frac{1}{|D'|} \sum_{z' \in D'} \nabla_{\theta} \ell(z'; \theta^{[T]})$ . Estimation  
 450 of the linear influence thus amounts to estimating the change in the validation loss

$$\langle u, \theta_{-j}^{[T]} - \theta^{[T]} \rangle \approx \frac{1}{|D'|} \sum_{z' \in D'} \left( \ell(z'; \theta_{-j}^{[T]}) - \ell(z'; \theta^{[T]}) \right).$$

451 We note that the instances with large negative linear influences are deemed to be negatively affecting  
 452 the resulting models. Removing such instances can improve the validation loss, and thus the users  
 453 can prioritize the inspection of such instances.

454 **Baseline Methods** For K&L [Koh and Liang, 2017], to solve the problem (9), we ran momentum-  
 455 SGD for two epochs, where we set the learning rate to be 0.005, the size of momentum to be 0.9,  
 456 and the batch size to be 1000. As baselines for data cleansing, in addition to K&L [Koh and Liang,  
 457 2017], we also adopted two outlier detection methods, Autoencoder [Aggarwal, 2016] and Isolation  
 458 Forest [Liu et al., 2008]. In outlier detection, we treated the validation set as a healthy dataset. We  
 459 then computed outlierness of each training instance using outlier detection methods, as follows.

- 460 • **Autoencoder:** We trained an autoencoder using the validation set. See Figure 5 for the  
 461 structures of autoencoders used. We adopted the squared loss as the training objective  
 462 function. For training, we used Adam with the learning rate set to 0.001 and the batch  
 463 size set to 128. We used the same data augmentation as the training of CNNs. After the  
 464 autoencoder is trained, we fed each training input  $x$  into the autoencoder and obtained an  
 465 reconstructed input  $\hat{x}$ . We measured the outlierness of the input  $x$  by  $a = \|x - \hat{x}\|^2$ .
- 466 • **Isolation Forest:** We first fed each validation input  $x'$  into the trained CNN, and obtained  
 467 its latent representation  $r'$  from the flatten layer in Figure 4. We trained an isolation forest  
 468 using the latent representations of the validation set. In the experiment, we used the `fit`  
 469 method of `sklearn.ensemble.IsolationForest` with default configurations. After the  
 470 isolation forest is trained, we fed each training input  $x$  into the isolation forest and obtained  
 471 its outlierness score  $a$  using the `score_samples` method.

472 We also adopted random data removal as the baseline.

473 **Proposed Method** For the proposed method, we introduced an approximate version in this experi-  
 474 ment. In Algorithm 2, the proposed method retraces the entire SGD steps. In the approximate version,  
 475 we retrace only one epoch, which requires less computation than the original algorithm. Moreover, it  
 476 is also storage friendly because we need to store intermediate information only in the last epoch of  
 477 SGD.

478 **Procedure** We proceeded the experiment as follows. First, we trained the model with SGD using  
479 the training set. We then computed the influence of each training instance using the proposed method  
480 as well as the other baseline methods. Finally, we removed the top- $m$  influential instances from the  
481 training set and retrained the model. For the model retraining, we considered the two settings.

- 482 • Retrain All: In this setting, we ran counterfactual SGD for all the 20 epochs with influential  
483 instances omitted.
- 484 • Retrain Last: In this setting, we ran normal SGD for 19 epochs and switched to counterfactual  
485 SGD in the last epoch with influential instances omitted.

486 If the misclassification rate of the retrained model decreases, we can conclude that the training set  
487 was effectively cleansed.



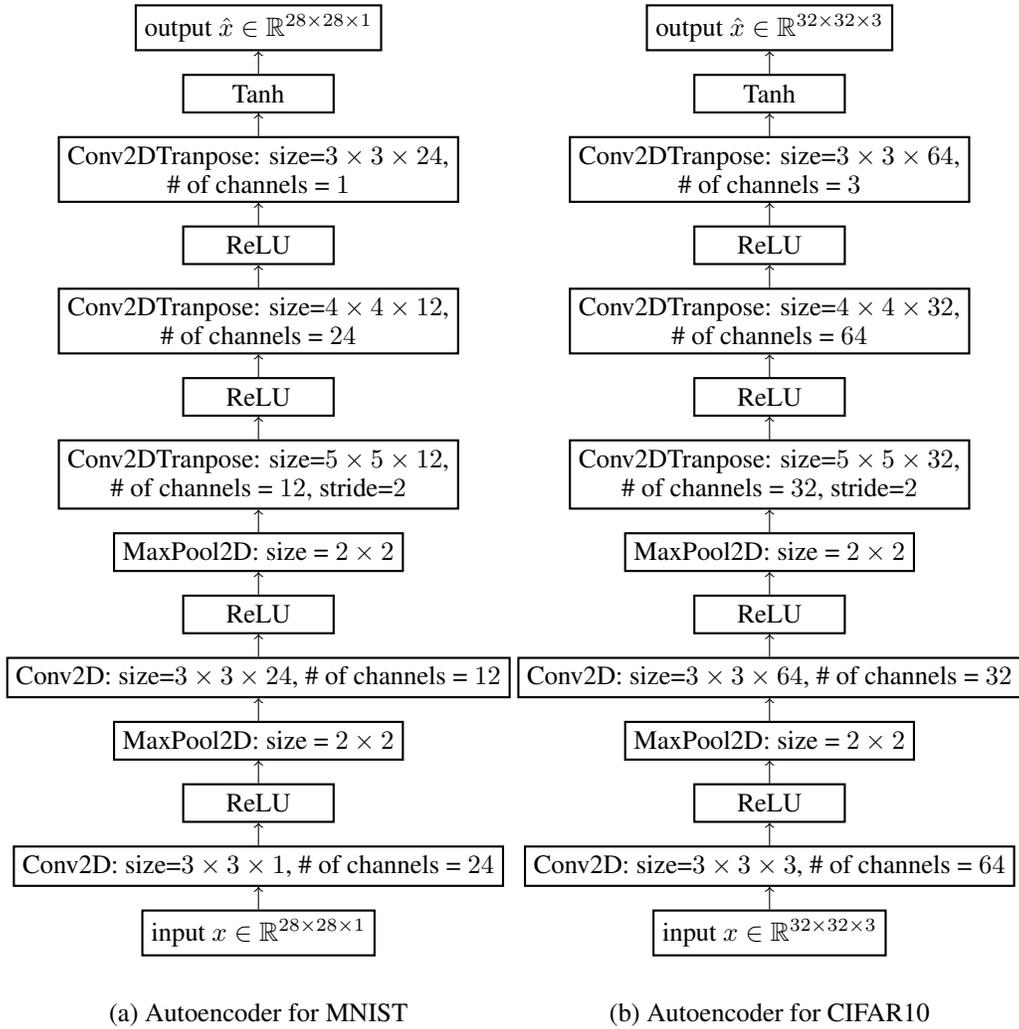


Figure 5: Structures of Autoencoders

488 **C.3 Full Results in Section 7.2**

489 The full results for MNIST are shown in Figures 6 and 7. The full results for CIFAR10 are shown in  
490 Figures 8 and 9. In the figures, it is evident that the misclassification rates have decreased after data  
491 cleansing with the proposed method and its approximate version. We compared the misclassification  
492 rates before and after the data cleansing using t-test with the significance level set to 0.05. We observed  
493 that non of the baseline methods have attained statistically significant improvements. By contrast,  
494 the proposed method and its approximate version attained statistically significant improvements:  
495 in MNIST, their improvements were statistically significant for the number of removed instances  
496 between 10 and 100, and in CIFAR10, their improvements were statistically significant for the number  
497 of removed instances between 100 and 1000. Figure 10 also confirms the effectiveness of the data  
498 cleansing with the proposed method. Out of 30 experiments, the misclassification rates decreased  
499 with the proposed method for 25 cases in MNIST, and for 26 cases in CIFAR10. These results  
500 confirm that the proposed method can effectively suggest influential instances for data cleansing.  
501 We also note that the proposed method and its approximation version performed comparably well.  
502 This observation suggests that, in practice, we only need to trace back only one epoch for inferring  
503 influential instances, which requires less computation and storing intermediate information only in  
504 the last epoch of SGD.

505 Figures 11 and 12 show the examples of found influential instances. An interesting observation is that  
506 Autoencoder tended to find images with noisy or vivid backgrounds. Visually, it seems reasonable to  
507 select them as outliers. However, as we have seen in Figure 2, removing these outliers didoes not help  
508 improving the models. On the other hand, the proposed method found images with confusing shapes  
509 or backgrounds. Although they are not strongly visually appealing as outliers, Figure 2 confirms  
510 that these instances have high impacts to the models. These observations indicate that the proposed  
511 method could find influential instances, which can be missed even by users with domain knowledge.

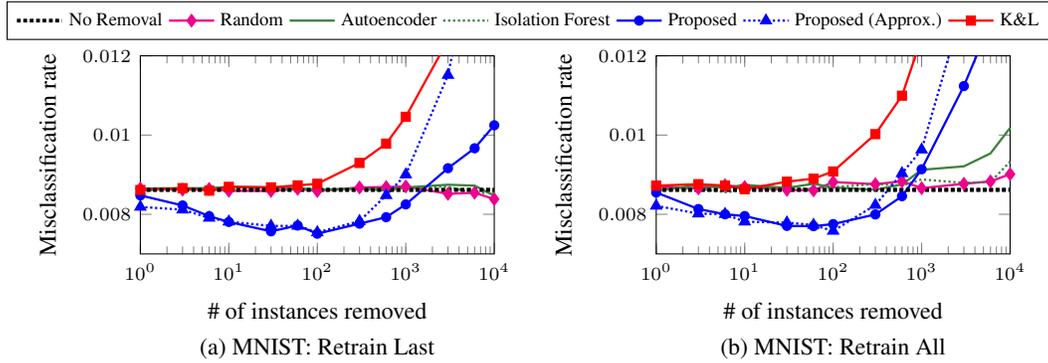


Figure 6: MNIST: Average misclassification rates on the test set after data cleansing over 30 experiments.

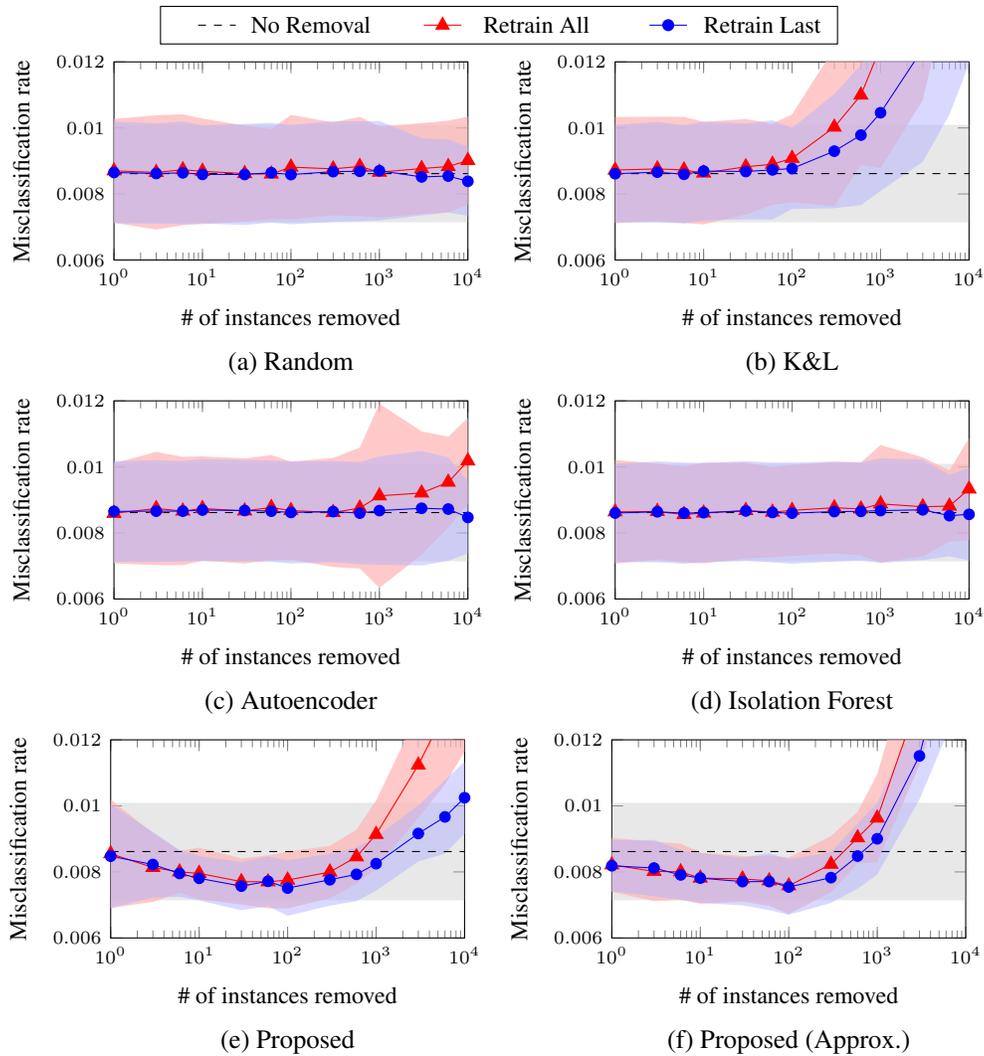


Figure 7: Exhaustive results on MNIST: [Thick lines] Average misclassification rates on the test set after data cleansing over 30 experiments. [Shaded Regions] Average  $\pm$  standard deviation.

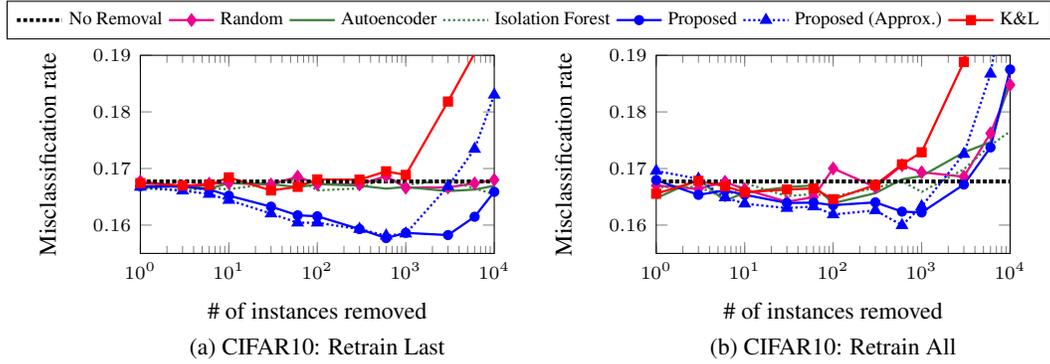


Figure 8: CIFAR10: Average misclassification rates on the test set after data cleansing over 30 experiments.

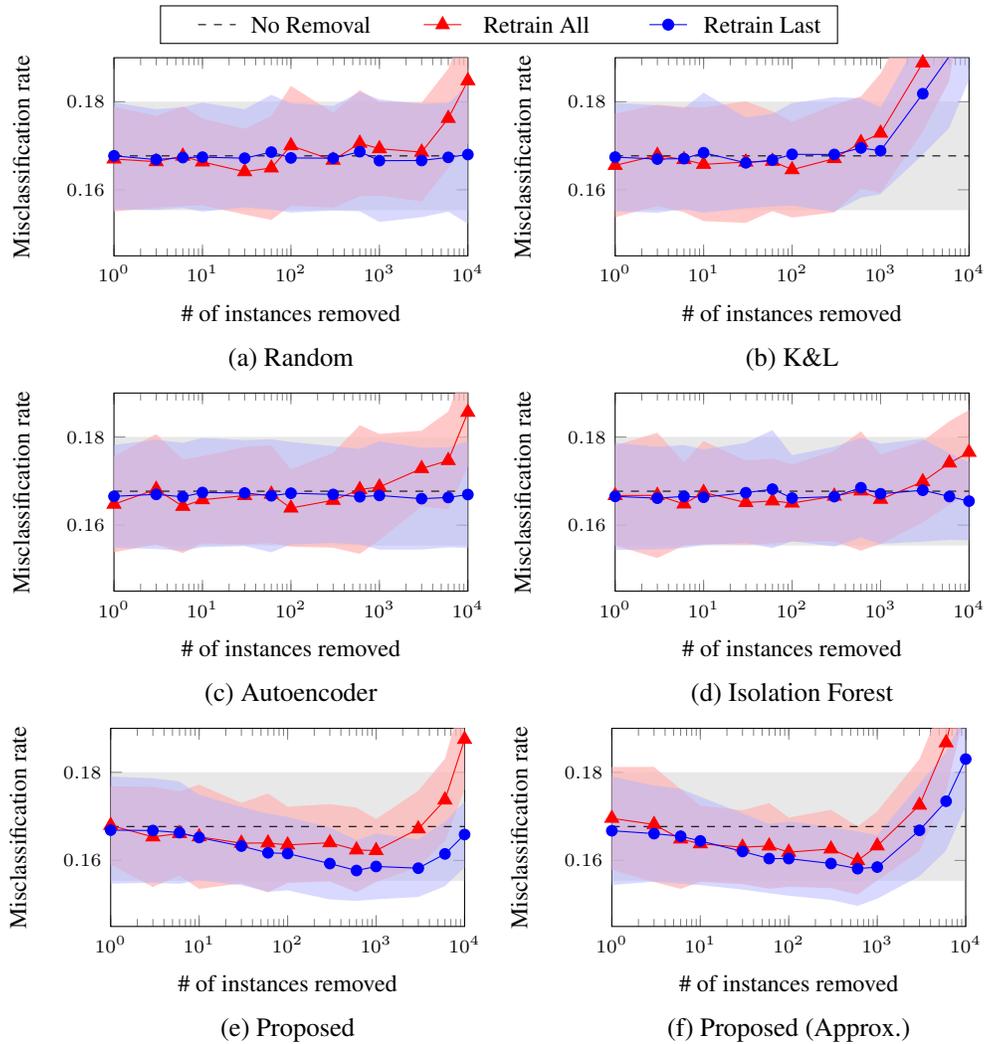
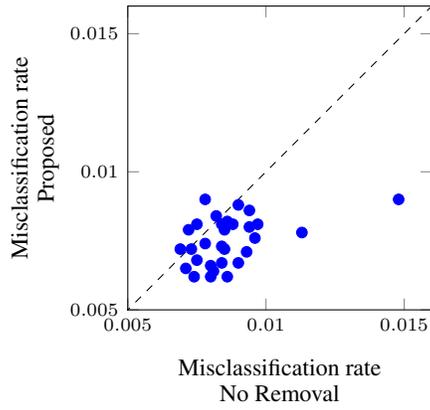
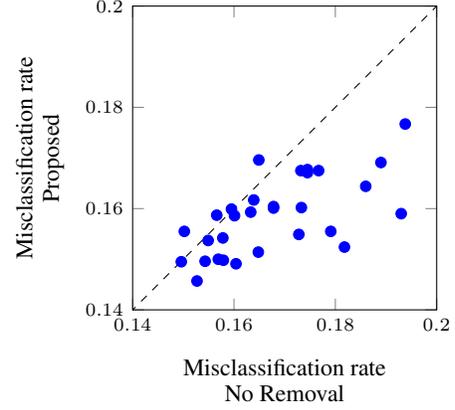


Figure 9: Exhaustive results on CIFAR10: [Thick lines] Average misclassification rates on the test set after data cleansing over 30 experiments. [Shaded Regions] Average  $\pm$  standard deviation.

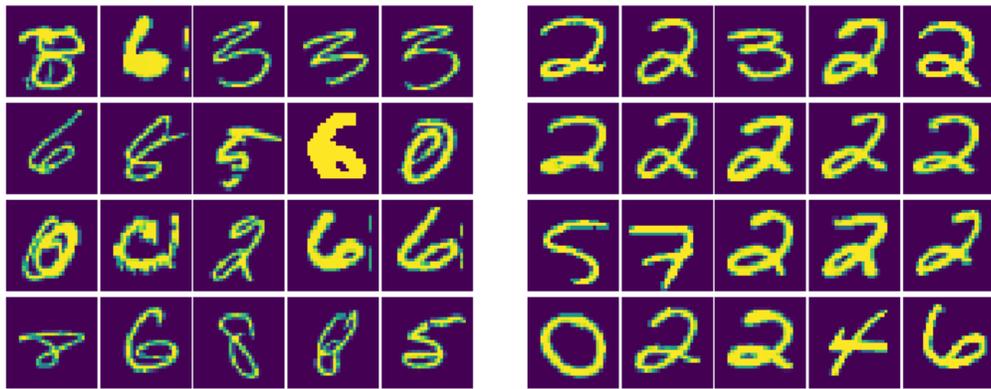


(a) MNIST



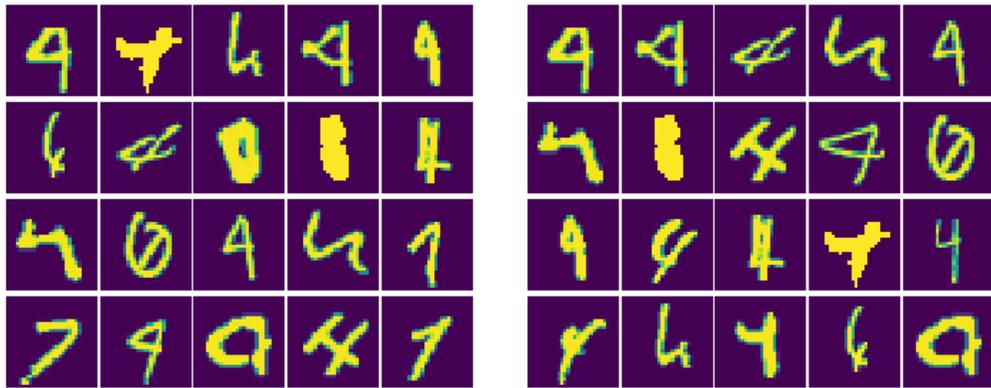
(b) CIFAR10

Figure 10: Comparison of the misclassification rates before and after the data cleansing with the proposed method. We set the number of removed instances to be 100 for MNIST and 10000 for CIFAR10.



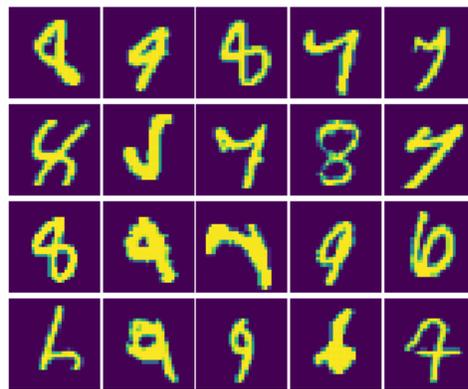
(a) Autoencoder

(b) Isolation Forest



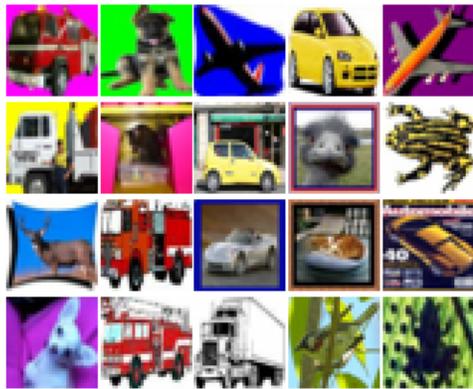
(c) Proposed

(d) Proposed (Approx.)



(e) K&L [Koh and Liang, 2017]

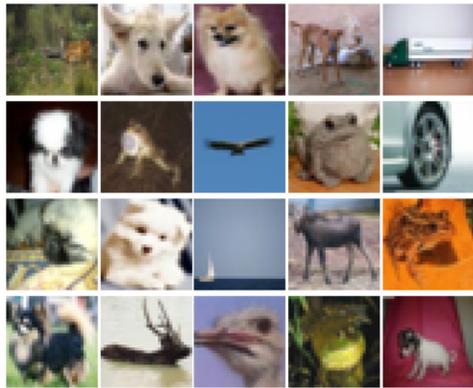
Figure 11: Examples of found top-20 influential instances in MNIST



(a) Autoencoder



(b) Isolation Forest



(c) Proposed



(d) Proposed (Approx.)



(e) K&L [Koh and Liang, 2017]

Figure 12: Examples of found top-20 influential instances in CIFAR10