We are grateful to the anonymous reviewers for the insightful feedback on our work. Please see our responses below.

Question 1 What is the motivation of using GAN (adversarial learning) for learning from label proportions?

Response There are mainly three reasons for using GAN to solve LLP problem. Firstly, as described in paragraph 3 of the Introduction, GAN is an elegant recipe for solving WeLL problems, especially semi-supervised learning [1]. From this viewpoint, our approach is in line with the idea of applying GAN to incomplete label scenarios. More important, the success of generative models for WeLL stems from the explicit or implicit representation learning, which has been an essential method for unsupervised learning for a long time. In LLP-GAN, the conv layers in discriminator can perform as a feature extractor for downstream tasks, which is proved to be efficient [2]. Hence, our work can be regarded as **solving LLP based on representation learning with GAN**. In this scheme, generated fake samples encourage the discriminator to not only detect the difference between the real and the fake instances but also distinguish *true K classes* for real samples (*K+1 classifier*). Thirdly, most LLP methods assume that the bags are i.i.d., which cannot sufficiently explore the underlying distribution in the data and may contradict in some applications. Instead, the generator in LLP-GAN is designated to learn the data distribution through the adversarial scheme without this assumption.

Question 2 What is the performance of the baseline of using entropy regularization for DLLP?

Response Firstly, this straightforward improvement for DLLP is a side contribution of our work. We consider not to include it as a baseline because the experimental results suggest that the original DLLP has already converged to the solution with fairly low instance-level entropy, which makes the regularization term redundant. Please see Figure 1.

Question 3 What is the advantage of performing gradient method on the lower bound instead of original objective?

Response The most important advantage of this trick is it allows to *perform SGD upon every instance* instead of every bag. It can greatly accelerate the convergence. However, in DLLP, we follow the original setting without this trick.

Question 4 Please offer more details on the experimental results.

Response Two issues should be clarified for experiments. Firstly, as shown in Figure 3 of our paper, the results demonstrate oscillation as bag size soaring. This phenomenon indicates a common drawback of deep models: For more complex objective surfaces (more possible label candidates), normally the convergence will be dramatically getting worse, due to more chances to attain local minima or saddle points of the objective. Secondly, because our results are based on original datasets without data augmentation, the reported DLLP performance is worse than that in [3].

Question 5 How much influence of the bag construction on the final results?

Response Indeed, the distribution of proportions has an huge impact on LLP algorithm performance. Hence, fixing bag size, we randomly construct bags for multiple times and present the accuracy performance in Table 1. The result shows the stability of our method. Currently, we can only artificially build LLP datasets from supervised ones. However, the gap between the importance of LLP in real-life and lack of specific LLP datasets exactly suggests the meaning of our work: It is worthy of devoting efforts to further study in order to draw more attention from the community.

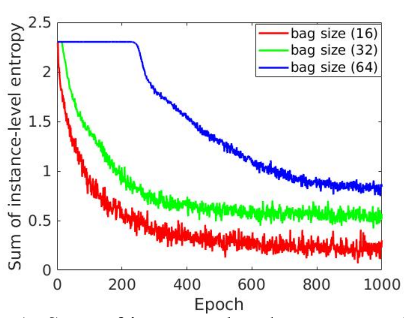

Figure 1: Sum of instance-level entropy on MNIST.

Bag Size (# of Random)	# of Errors	Accuracy (%) (Deviation)	Baseline (CNN)
16 (7)	106	98.94 (0.0285)	99.64
32 (22)	124	98.76 (0.0542)	
64 (45)	147	98.53 (0.11)	
128 (85)	335	96.65 (0.4)	

Table 1: The performance on accuracy with deviation under multiple random bag generations on MNIST. (Due to the time limitation of response, # of random are differently chosen.)

References

[1] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, et al. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*, pages 2234–2242, 2016.

[2] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

[3] Gabriel Dulac-Arnold, Neil Zeghidour, Marco Cuturi, Lucas Beyer, and Jean-Philippe Vert. Deep multi-class learning from label proportions. *arXiv preprint arXiv:1905.12909*, 2019.