Reviewer #1:

The paper is an evaluation of the myriad loss functions and regularizations that exist in the GAN literature,

Though the paper provides sufficient comparisons in terms of the metrics they have described but the paper lacks comparisons in terms of the widely accepted metrics of Frechet Inception Distance (FID) or Kernel Inception Distance (KID) thus necessitating major revisions which can show improvements with respect to these metrics and whether similar conclusions hold up.

The schematic diagrams of the gradient penalty and graphical representation of the different loss functions in Figure 4 would be an interesting contribution to the community.

Reviewer #2:

This paper mainly contributes two things. First, they derived in theory some necessary and sufficient conditions for 'valid' GAN losses. Second, they proposed a new evaluation metric for GANs evaluation based on the discriminative adversarial networks (DANs). The derived conditions for GAN losses are useful to check if the loss functions form a valid adversarial loss for GANs. The authors also evaluated a comprehensive set (in total 168) of different combinations of 12 component functions and 14 regularization approaches using the proposed evaluation metric.

My main concerns include the follows:

1. The authors proposed two new loss functions, absolute and asymmetric, based on the derived conditions. But the performance of the two new losses is not as good as the hinge loss. It would be better if the new losses can outperform the existing losses in some tasks.

2. The authors proposed a new evaluation metric, but they didn't compare with existing evaluation metric such as FID. The authors may explain the advantages of the proposed evaluation metric over FID (e.g., more consistent with human judgment).

3. The authors only evaluated on the MNIST dataset, which is not enough to conclude which loss performs better. The authors may evaluated on more datasets.

Reviewer #3:

The paper studies the adversarial loss. The analysis does not bring much more insight than existing f-divergence and IPM results, which are known in statistics for decades. The chosen protocol is DAN, which is not standard and not well-justified.

1. I'm sorry to say it is not clear to me how useful the studied theorems could be. The examples shown in the paper (table 3 and section 3.6) are all not beyond the scope of IPM and f-divergence. From those general results, we already know they are all valid probability discrepancies. Also, being a valid discrepancy is not enough to train a GAN, many results are studied and the paper failed to discuss those (e.g. Towards Principled Methods for Training Generative Adversarial Networks, Arjovsky et al., (2017)). The theory of adversarial loss, such as its minimax rate (e.g. Nonparametric Density Estimation with Adversarial Losses Sashank et al., (2018))., and the training dynamics are also studied a lot. Compared with all the existing works, the insights from the studied theories are not clear. Also, the proofs has to be more rigorous. For example, it does not illustrate all the technical conditions and assumptions. From (7) to (8), it is not true that we can switch integral and max. What's the condition allowing us to do so here?

2. The used protocol is not standard and the authors do not justify why it should be favored. The standard metrics and the protocol from Lucic et al., (2018) would be better.