View Reviews

Paper ID
664

Paper Title
Towards a Deeper Understanding of Adversarial Losses

Reviewer #1

Questions

1. Please enter a detailed review describing the strengths and weaknesses of the submission.
This paper proposed the method to quantitatively compare the performance of adversarial losses and regularization approaches. In addition, this paper provided theoretical analyses for identifying what type of component functions are valid as adversarial loss functions. This paper conducted many experiments to compare a large variety of adversarial losses, and gave insights about how to build a GAN model.

The experimental results look promising and comprehensive. However, this paper is little difficult to follow from several reasons:
- The connection between the theoretical analysis about favorable properties for adversarial losses and the DANTest at first sight.
- DANTest does not use the original setting (generate images from hidden state vector), and it's the conditional GAN training setting, however, it is not explicitly explained until the Section 4.
- The experimental results focus on DAN, not the original GAN. It is unclear how much the results of this paper can be applied to the original GAN.
The organization of this paper can be improved.

2. Please provide an overall score for the submission.
Weak Reject: Borderline, tending to reject

3. Please enter a 2-3 sentence summary of your review explaining your overall score.
The experimental results look promising and comprehensive. However, this paper is little difficult to follow from several reasons:
- The connection between the theoretical analysis about favorable properties for adversarial losses and the DANTest at first sight.
- DANTest does not use the original setting (generate images from hidden state vector), and it's the conditional GAN training setting, however, it is not explicitly explained until the Section 4.
- The experimental results focus on DAN, not the original GAN. It is unclear how much the results of this paper can be applied to the original GAN.
The organization of this paper can be improved.

5. Please rate your confidence in the score assigned.
Low: Reviewer is making an educated guess.

8. I agree to keep the paper and code submissions confidential, and delete any submitted code at the end of the review cycle to comply with the confidentiality requirement.
Agreement accepted

Reviewer #2

Questions

1. Please enter a detailed review describing the strengths and weaknesses of the submission.

Summary: The paper considers recently-proposed loss functions and regularization strategies for GANs. The authors present a theoretical view on the loss functions, deriving necessary and sufficient conditions for the losses to be either pseudo-divergences ($p=q \Rightarrow D=0$) or divergences ($p=q$ iff $D=0$). The authors develop a benchmark for GAN losses and regularization strategies called DANTest, and use their benchmark to evaluate different combinations of the two.

Review: This paper presents a clean discussion of GAN loss functions and certain motivating properties behind them. The theoretical contributions are welcome, although their significance isn't obvious to me. The proposed evaluation benchmark is interesting and a good contribution, although the practical recommendations made on its basis aren't well-justified in my opinion, since the evaluation setup and criteria are very different from how GANs are used in practice. Overall I consider this paper a borderline accept.

Detailed comments:

This paper gets high marks for clarity in my opinion. Beautiful writing, tables, and figures throughout.

"Interestingly, while such a theoretical analysis has not been done before, it happens that all the adversarial loss functions listed in Table 1 have such favorable properties." This doesn't strike me as particularly surprising since most of these losses were explicitly motivated through their equivalence to various divergences in the optimal-discriminator limit.

The unique-global-minimum condition behind Prop 2 and Thms 2 and 4 seems irrelevant to me since in practice, the maximum in L_G is always constrained to a parametric family (a neural network), which makes all the generator losses only pseudo-divergences. Indeed, this is what makes GANs special in many ways (see e.g. Huang et al 2017 https://arxiv.org/abs/1708.02511).

The proposed DANTest is interesting. It's straightforward to define and it's a good sanity-check for GAN methods, since a good method should certainly do well on it. However it somewhat avoids much of the key difficulty behind GAN training, namely that the generator's output lives on a low-dimensional manifold embedded in a high-dimensional space. In a low-dimensional setting, plenty of non-GAN methods (e.g. plain old max-likelihood) should work very well, so the benchmark is "trivial" in a certain sense.

The concrete recommendations don't really make sense to me. "In sum, according to the overall performance and the robustness to different settings, for the component functions, we recommend the hinge, the asymmetric and the two relativistic losses." / "we recommend to use hinge loss for the discriminator and minimax loss for the generator" It's not clear why robustness to the regularization strategy is a desirable property. Further, the DANTest setting is too different from the normal settings GANs are applied in for anything but very large differences in performance to suggest anything, in my opinion.

Minor: "Note that L_G is not a divergence since $L_G \geq 0$ does not always hold" – this is true but somewhat irrelevant, since L_G^* is a constant term irrelevant to optimization.

Minor: Your cite for DANs in the intro is wrong – should be dos Santos et al.

2. Please provide an overall score for the submission.

Weak Accept: Borderline, tending to accept

3. Please enter a 2-3 sentence summary of your review explaining your overall score.

This paper presents a clean discussion of GAN loss functions and certain motivating properties behind them. The theoretical contributions are welcome, although their significance isn't obvious to me. The proposed evaluation benchmark is interesting and a good contribution, although the practical recommendations made on its basis aren't well-justified in my opinion, since the evaluation setup and criteria are very different from how GANs are used in practice. Overall I consider this paper a borderline accept.

5. Please rate your confidence in the score assigned.

Medium: Reviewer has understood the main points in the paper, but skipped the proofs and technical details.

8. I agree to keep the paper and code submissions confidential, and delete any submitted code at the end of the review cycle to comply with the confidentiality requirement.
Agreement accepted

Reviewer #5

Questions

1. Please enter a detailed review describing the strengths and weaknesses of the submission.
The authors try to theoretically and empirically understand what component functions perform better or worse than others in GANs, and which ones have certain desired properties (such as having the true data distribution as the unique optimum of the minimax procedure).

In section 3 the authors provide their theoretical results. The desired properties they discuss are wether the true distribution is the (unique or not) minimum of the loss for an optimal discriminator. They summarize the requirements for a certain component function based on two functions derived from the component (\psi and \Psi). They show that for the components presented in the background section 2, this properties are satisfied. Note that the fact that the unique minimum was already satisfied for these components was already known [1], [2], [3]. Perhaps more importantly, the fact that there's a unique optimum doesn't mean that that optimum is achievable (for example, the divergence as a function of the generator's parameters might be discontinuous [4]).

The experiments in sections 4-5 are quite interesting in my opinion, but suffer from two severe flaws. The first one being that the DAN test is quite different from the way people employ GANs in practice. Namely, the generator is simply generating labels, which is a sample (a) discrete and (b) low dimensional sample (c) most likely coming from a unimodal distribution; whilst most people are interested in generating from continuous, high dimensional and multimodal distribution. Thus, the test doesn't quite reflect the setup for which people want to test the capabilities of their models. The second concern, is that essentially all the experiments are in MNIST, which lacks sufficient scale for modern generative modeling approaches to be properly tested. I encourage the authors to take this feedback into account and keep pushing on this interesting avenue of research.

--------- Minor things -----------------
Regarding question R1, the authors should describe at least a bit what it means for them to be a 'valid' adversarial loss function?

In 3.2, the authors should avoid using \Psi and \psi to represent the two functions, since when orally discussing the paper this will lead to confusion :). (Which 'Psi' are we referring to?)

017R: when empirically compare -> when we empirically compare

[1]: https://arxiv.org/abs/1712.07822
[2]: https://arxiv.org/abs/1406.2661
[3]: https://arxiv.org/abs/1609.03126
[4]: https://arxiv.org/abs/1701.07875

2. Please provide an overall score for the submission.
Reject: Clearly below the acceptance threshold

3. Please enter a 2-3 sentence summary of your review explaining your overall score.
The first reason for my score is simple: the theoretical results don't yield any new understanding or conclusions of known component functions (which is the main object of study of the paper), furthermore, the fact that there's a unique optimum (the property considered) doesn't mean that that optimum is achievable (for example, the divergence as a function of the generator's parameters might be discontinuous). The second reason is that the experimental setup lacks a strong resemblance to the way people employ these models in practice due to the

generator creating samples that come from a unimodal distribution, are discrete and low dimensional. Furthermore, the dataset considered (MNIST) lacks sufficient difficulty to test these models properly.

5. Please rate your confidence in the score assigned.
Medium: Reviewer has understood the main points in the paper, but skipped the proofs and technical details.

8. I agree to keep the paper and code submissions confidential, and delete any submitted code at the end of the review cycle to comply with the confidentiality requirement.
Agreement accepted

Reviewer #6

Questions

1. Please enter a detailed review describing the strengths and weaknesses of the submission.
The paper analyses two questions (1) How to set the component functions in the discriminator loss to establish “valid adversarial loss functions”, and (2) what is the interplay between those component functions and the regularization approaches? To answer (1) authors suggest that F and G should be chosen such that L_G has the global minimum at p_g = p_d. They rewrite the loss L_G for a fixed discriminator as a function (psi) which takes the density ratio and the discriminators output and control the desired property directly on psi (theorems 1-5). They provide several sets of experiments on MNIST and empirically analyse (2).

Strengths:
- The proposed analysis confirms that the component functions (F, G, H) in all currently used losses (minimax, non-saturating, Wasserstein, least-squares, hinge) lead to L_G - L_G* being a divergence.
- Empirical insight on why balancing the two loss terms in L_D is useful and why two-sided GP might behave better than the 1-sided alternative.

Weaknesses:
- Very strong conclusions are derived based on a very limited empirical study (only one data set, low-capacity neural network parametrising G and D).
- The novelty is unclear as most of these results are already derived for the special cases (e.g. f-divergences [1]) and similar insights provided in other works (e.g. [2]).

[1] Nowozin, S., Cseke, B., and Tomioka, R. f-GAN: Training generative neural samplers using variational divergence minimization. NeurIPS 2016.
[2] Fedus, W., Rosca, M., Lakshminarayanan, B., Dai, A. M., Mohamed, S., and Goodfellow, I. Many paths to equilibrium: GANs do not need to decrease a divergence at every step. ICLR 2018.

2. Please provide an overall score for the submission.
Reject: Clearly below the acceptance threshold

3. Please enter a 2-3 sentence summary of your review explaining your overall score.
Not enough novel theoretical insights. Conclusions are too strong based on the limited experimental setup.

5. Please rate your confidence in the score assigned.
High: Reviewer has understood the main arguments in the paper, and has made high level checks of the proofs.

8. I agree to keep the paper and code submissions confidential, and delete any submitted code at the end of the review cycle to comply with the confidentiality requirement.
Agreement accepted

View Meta-Reviews

Paper ID
664

Paper Title
Towards a Deeper Understanding of Adversarial Losses

META-REVIEWER #1

META-REVIEW QUESTIONS

2. Please enter a detailed meta-review explaining your decision.
The paper aims to compare various component loss functions and regularization approaches for GANs. It does so in two ways: theoretically and empirically. Theoretically the authors focus on whether the loss configuration results in the objective having a unique global minimum which corresponds to the true distribution. Empirically, they propose a novel test they call DANTest.

Reviewers agree that the theoretical part of the contributions is not too exciting, since part of conclusions were already known (many of the current methods satisfy the unique global minimum requirement under optimal discriminator) and otherwise an existence of a unique global minimum is indeed attractive but does not necessarily lead to success. On a practical side, all the reviewers pointed out that the setting of proposed experiments is very different from the way GANs are used nowadays, which means it is not the best way to compare them.

7. I agree to keep the paper and code submissions confidential, and delete any submitted code at the end of the review cycle to comply with the confidentiality requirement.
Agreement accepted