View Reviews

Paper ID
3555

Paper Title
TensoRF: Tensorial Radiance Fields

Reviewer #1

Questions

2. Summarize the paper's claimed primary contributions: In 5-7 sentences, describe the key ideas, results, findings, and significance as claimed by the paper's authors.
The paper considers the 3D scene representation problem in the context of NeRF-type methods which have shown remarkable performance for modeling and reconstructing 3D scenes. The core idea of this paper is to represent the scene with per-voxel multi-channel features and factorize the scene tensor into multiple compact low-rank tensor components. In addition to the commonly used CP decomposition, the authors also propose a vector-matrix decomposition of tensors which shows better performance. Extensive experiments on multiple datasets demonstrate that the proposed method can improve the rendering quality while greatly reducing the training time and memory requirements.

3. What do you see as the main strengths of this work? Consider, among others, the significance of critical ideas, validation, writing quality, and data contribution. Explain clearly why these aspects of the paper are valuable. ACs are instructed to ignore unsupported responses.
The strengths of this work are multi-fold as I see it.
First, the motivation of this paper is clear and intuitively effective. The idea of using tensor decomposition for more efficient scene representation is smart, and one can immediately understand the logic. The presentation is well-organized, self-contained, and easy to follow.
Second, I believe the paper is inspiring to the researchers in this field and can potentially make a big impact. The practicability is good but more importantly, the presented results tell more about scene representation itself, that the low-rank property can and shall be exploited in other methods. The proposed method utilizes this in a simple yet principled manner which is informative.
In addition, the experimental parts are extensive and solid, comparing multiple very recent competitors and showing advantages in both rendering quality and efficiency, which well supports the claim. The authors have done a good job presenting their results with the video.

4. What do you see as the main weaknesses of this work? Clearly explain why these are weak aspects of the paper, e.g., why a specific prior work has already demonstrated the key contributions, why the experiments are insufficient to validate the claims, etc. ACs are instructed to ignore unsupported responses.
The paper is well-written in my opinion and I did not find many weak points. But I am curious that since DVGO [35] and Plenoxels [44] optimize voxel grids of features directly, why do they perform worse in terms of rendering quality compared to the proposed method which optimize over low-rank approximations.

Typo: In first sentence of caption of Fig.3, TensorRF -> TensoRF

5. Reproducibility: Could the work be reproduced by a talented graduate student from the information in the paper?
Agreement accepted

6. [Rate the paper as it stands now (pre-rebuttal). Borderline will not be an option for your final post-rebuttal recommendation, and so it should only be used rarely now.]

Accept

7. Justify your rating. Be specific: What are the most critical factors in your rating? What points should the authors cover in their rebuttal? Your reply should clearly explain to the authors what you need to see in order to increase your rating.
As commented in Q3, I think the paper is very informative and practically important. Thus I think the paper should be accepted.

11. Justify your post-rebuttal assessment. Acknowledge any rebuttal and be specific about the final factors for and against acceptance that matter to you. (Will be visible to authors after author notification)
The paper has a smart idea which solves an important problem in novel view synthesis, and the experiments are sufficient to demonstrate the effectiveness. Thus I think this is a good paper that should be accepted.

12. Give your final rating for this paper. Don't worry about poster vs oral. Consider the input from all reviewers, the authors' feedback, and any discussion. (Will be visible to authors after author notification)
Accept. I vote and argue for accepting this submission.

Reviewer #2

Questions

2. Summarize the paper's claimed primary contributions: In 5-7 sentences, describe the key ideas, results, findings, and significance as claimed by the paper's authors.
This paper proposed an improved version of NeRF. Instead of the coordinate-based MLPs used by NeRF, the authors proposed to represent radiance fields as an explicit voxel grid of features.
It achieves high efficiency in training time and high compactness in memory footprint. Moreover, it provides state-of-the-art rendering quality, especially for instance detail rendering.

3. What do you see as the main strengths of this work? Consider, among others, the significance of critical ideas, validation, writing quality, and data contribution. Explain clearly why these aspects of the paper are valuable. ACs are instructed to ignore unsupported responses.
+ The paper is well-written. The structure is clear and easy to follow. Figures 1-3 are very helpful for understanding the main idea and the tensor factorization techniques.

+ The proposed method significantly reduces the memory footprint through decomposition strategies.
The space complexity is reduced from O(n^3) to O(n) by CP, or O(n^2) by VM.

+ The idea of tensorizing radiance fields seems general and can be extended to the related tasks.

4. What do you see as the main weaknesses of this work? Clearly explain why these are weak aspects of the paper, e.g., why a specific prior work has already demonstrated the key contributions, why the experiments are insufficient to validate the claims, etc. ACs are instructed to ignore unsupported responses.
- It is relatively easy to understand that the proposed tensor decomposition techniques lead to higher compactness and efficiency. However, I do not fully understand why the proposed techniques can also improve rendering quality. If possible, please provide a more theoretical analysis regarding which element is the main reason for quality improvement.

- In the abstract, the authors do not indicate the meaning of "CP" (see line 12). I suggest using "CANDECOMP/PARAFAC (CP) decomposition" to facilitate understanding.

- It seems that the font size of figure/table captions is not consistent with that in the template (other papers I reviewed do not have this problem). Please solve this problem to save space.

- Table 3 is too compact. Please consider re-organizing this table.

5. Reproducibility: Could the work be reproduced by a talented graduate student from the information in the paper?
Agreement accepted

6. [Rate the paper as it stands now (pre-rebuttal). Borderline will not be an option for your final post-rebuttal recommendation, and so it should only be used rarely now.]
Weak Accept

7. Justify your rating. Be specific: What are the most critical factors in your rating? What points should the authors cover in their rebuttal? Your reply should clearly explain to the authors what you need to see in order to increase your rating.
The tensor decomposition techniques used in this paper seem sound. Extensive experiments have demonstrated the superiority of the proposed method. As I mentioned above, in the rebuttal file, I would like to see a more theoretical analysis of render quality improvement.

11. Justify your post-rebuttal assessment. Acknowledge any rebuttal and be specific about the final factors for and against acceptance that matter to you. (Will be visible to authors after author notification)
I am satisfied with the authors' response/discussion regarding low-rank optimization. If possible, please add the core part of this discussion to the main paper.

12. Give your final rating for this paper. Don't worry about poster vs oral. Consider the input from all reviewers, the authors' feedback, and any discussion. (Will be visible to authors after author notification)
Accept. I vote and argue for accepting this submission.

Reviewer #3

Questions

2. Summarize the paper’s claimed primary contributions: In 5-7 sentences, describe the key ideas, results, findings, and significance as claimed by the paper’s authors.
The authors propose to leverage a tensor-based radiance field representation (that maps input 3D location and viewing direction to volume density and view-dependent appearance) for scene modeling in the context of novel view synthesis to reduce the memory footprint in comparison to recent comparable approaches such as PlenOctrees or DVGO. In particular, the scene tensor (three spatial dimensions used for representing the underlying voxel grid and one dimension for representing the per-voxel features) is factorized into compact low-rank tensor components, i.e. a set of components that describe the scene along the corresponding XYZ axes (used for the geometry grid and to estimate the local density) and additional appearance characteristics (encoded in a 4D tensor used with the view direction and a encoding function to predict color). For this purpose, the authors use (1) CANDECOMP/PARAFAC decomposition that factorizes a tensor into the sum of outer products and (2) an apparently novel vector-matrix decomposition that factorizes a tensor into a set of vectors and matrices. For both options, trilinear interpolation allows a continuous field representation.

In the evaluation, the authors show that their approach performs better than several recent voxel-based NeRF extensions in terms of reconstruction time, model size and rendering quality, and also provide an evaluation of the effect of different design choices.

3. What do you see as the main strengths of this work? Consider, among others, the significance of critical ideas, validation, writing quality, and data contribution. Explain clearly why these aspects of the paper are valuable. ACs are instructed to ignore unsupported responses.
Technical soundness
+ The approach is reasonable as also indicated by the achieved results.

+ Though left for future work, the combination with orthogonal approaches such as InstantNGP offers interesting potential.

Exposition:
+ The paper is well-structured and easy to follow. Text quality is mostly good (a few things should be handled in a proofreading process).
+ In the abstract and introduction, the relevance and the contribution is clearly stated.
+ Figures/tables and the respective captions are informative.

Evaluation
+ The authors provide quantitative and qualitative evaluations on several scenes with comparisons to competing techniques, where the proposed approach seems to be beneficial with respect to rendering quality, model size and reconstruction time.
+ Ablation studies show the effect of the used number of optimization steps on the rendering quality and reconstruction time as well as the effect of using spherical harmonics or MLP-based representations on rendering quality, model size and reconstruction time.

Reproducibility
+ The descriptions seem adequate to allow reproducibility. Code is available via github.

4. What do you see as the main weaknesses of this work? Clearly explain why these are weak aspects of the paper, e.g., why a specific prior work has already demonstrated the key contributions, why the experiments are insufficient to validate the claims, etc. ACs are instructed to ignore unsupported responses.
Exposition
- There are some misleading things.E.g. there should be more citations on NeRF improvements than [47,48] in the brackets in the introduction.
- (minor) On page 12, where the authors mention that except the largest VM model, all models would finish within less than 30 min, the reference seems to rather go for Table 2 instead of Table 1.

Evaluation
- Limitations and failure cases have not been discussed in detail.
- The effect of choosing different values for the feature numbers than the proposed value of 27 would be interesting, but is missing in the current version.

References
- The discussion of related work regarding NeRF approaches is lacking. I would recommend to also include:
Trevithick & Yang, GRF: Learning a General Radiance Field for 3D Representation and Rendering
--> The authors use higher-dimensional features than RGB in a voxel-based representation with trilinear interpolation.
Garbin et al., FastNeRF: High-Fidelity Neural Rendering at 200FPS
--> The authors use a graphics-inspired factorization which enables (i) compactly caching a deep radiance map at each position in space, (ii) efficiently querying that map using ray directions to estimate the pixel values in the rendered image. This leads to huge speedup by a factor of about 3000 in comparison to the original NeRF approach.
- The authors could cite the following survey on tensor representations/decomposition.
Ji et al., A survey on tensor techniques and applications in machine learning

5. Reproducibility: Could the work be reproduced by a talented graduate student from the information in the paper?
Agreement accepted

6. [Rate the paper as it stands now (pre-rebuttal). Borderline will not be an option for your final post-rebuttal recommendation, and so it should only be used rarely now.]

Weak Accept

7. Justify your rating. Be specific: What are the most critical factors in your rating? What points should the authors cover in their rebuttal? Your reply should clearly explain to the authors what you need to see in order to increase your rating.
I would strongly encourage the authors to add discussions of further related work and failure cases/limitations, but the performance of the presented approach makes me lean towards acceptance.

11. Justify your post-rebuttal assessment. Acknowledge any rebuttal and be specific about the final factors for and against acceptance that matter to you. (Will be visible to authors after author notification)
After reading the rebuttal and the other reviewers' comments, I support accepting the paper. The contribution is adequate and proven to perform well. However, I encourage the authors to incorporate the insights from the rebuttal into their paper and supplemental to strengthen their submission.

12. Give your final rating for this paper. Don't worry about poster vs oral. Consider the input from all reviewers, the authors' feedback, and any discussion. (Will be visible to authors after author notification)
Accept. I vote and argue for accepting this submission.