Improving Multiplane Images with Deformable Layers

Lachezar Topalov
ltopalov@tudelft.nl
Responsible Professor: Petr Kellnhofer
EEMCS, Delft University of Technology, The Netherlands

Background

Novel view synthesis aims to generate images of a scene from new camera viewpoints. In the single-image setting, this is difficult because the method must infer both 3D structure and content hidden behind visible objects.

Multiplane Images (MPIs) represent a scene as a stack of semi-transparent RGBA planes placed at fixed depths. Each plane is projected into the novel viewpoint, and the projected planes are blended using alpha compositing. This makes MPIs efficient and useful for modelling layered visibility.

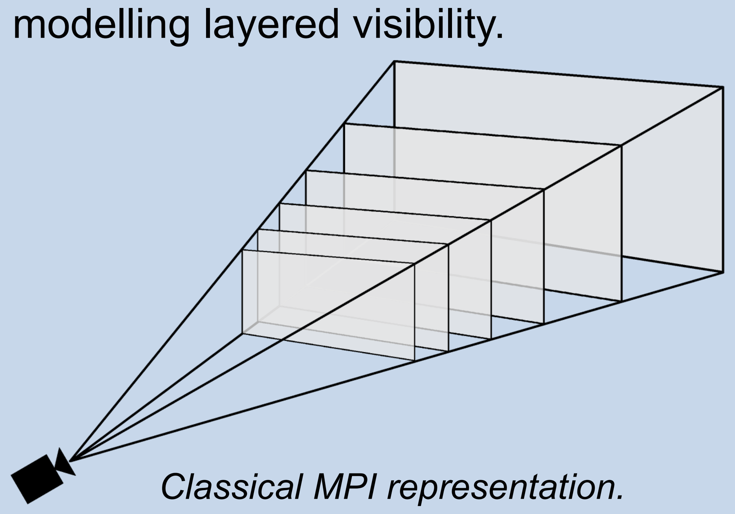

Classical MPI representation.

However, classical MPI planes are flat and fronto-parallel. This is a limited approximation for slanted, curved, or irregular surfaces, which may be spread across several depth planes. As a result, synthesized views can show inaccurate parallax, ghosting, or cardboard-like geometry.

Key limitation: MPIs model visibility well, but their flat layers struggle with surfaces that are not well described by fronto-parallel planes.

Research Question

Can deforming flat MPI planes into alpha-weighted mesh layers improve how MPIs represent scene geometry?

This project investigates whether local mesh deformation can help MPI layers better align with scene geometry while preserving MPI-style layered visibility and alpha compositing.

The hypothesis is that alpha-weighted deformation can improve parallax on slanted or curved surfaces, especially when the input layered geometry is reliable.

Methodology

1. **Predict a flat MPI**
   A pretrained single-view MPI model based on Tucker and Snavely predicts RGBA layers and plane depths from one input image.
2. **Compute deformation signals**
   The MPI alpha values are used to estimate visible-scene disparity and determine which layers are responsible for each image region.
3. **Convert planes into meshes**
   Each flat fronto-parallel MPI plane is subdivided into a regular triangular mesh so that its geometry can be locally adjusted.
4. **Deform mesh layers**
   Mesh vertices are displaced according to the computed disparity signal, weighted by layer ownership and constrained to avoid layer collapse.
5. **Render and compare**
   The deformed layers are rendered using MPI-style alpha compositing and compared against the original flat MPI baseline using **PSNR, SSIM, and LPIPS**.

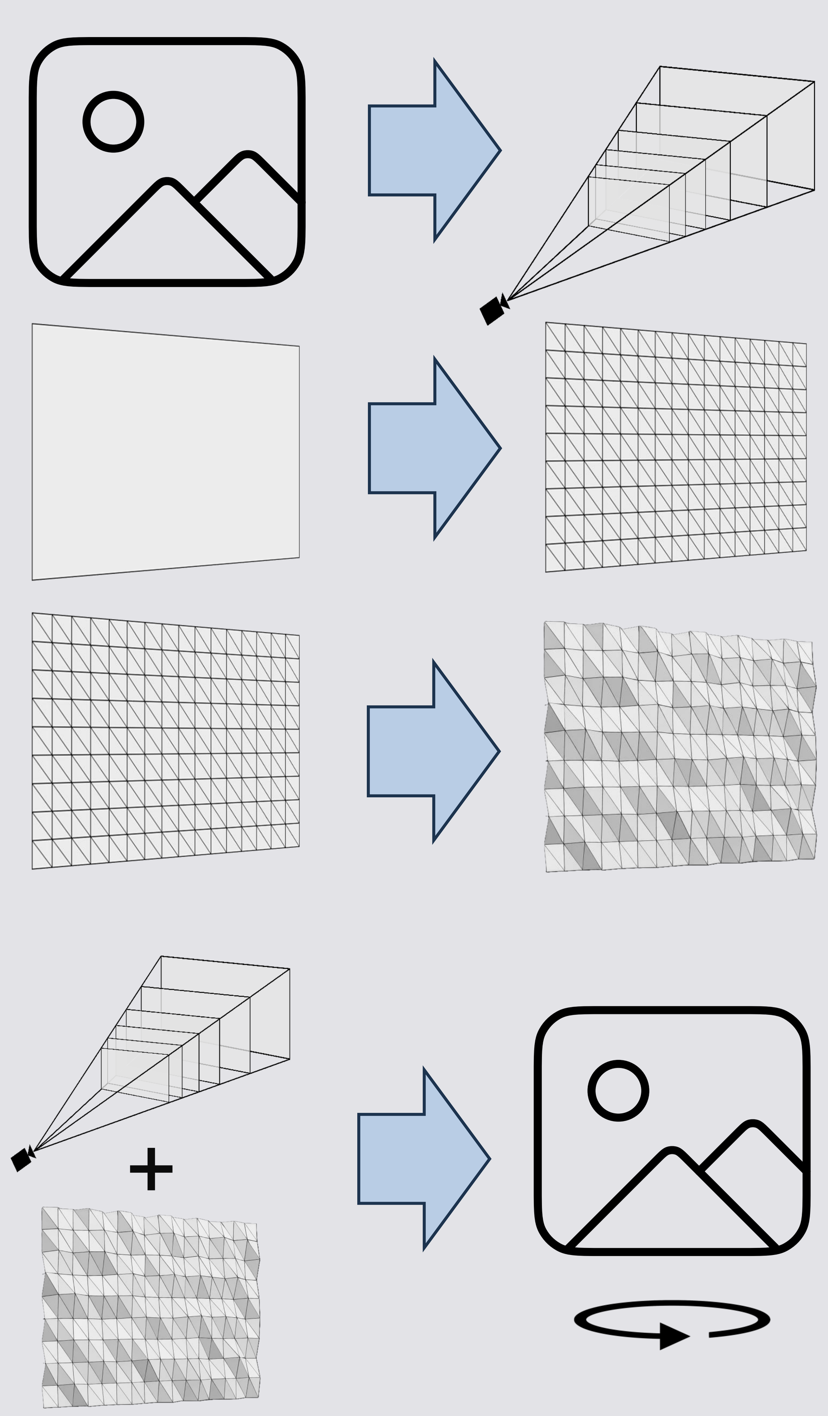

Results

Predicted MPI benchmarks: marginal improvement

Deformation produces almost no measurable change on RealEstate10K and only small positive changes on LLFF.

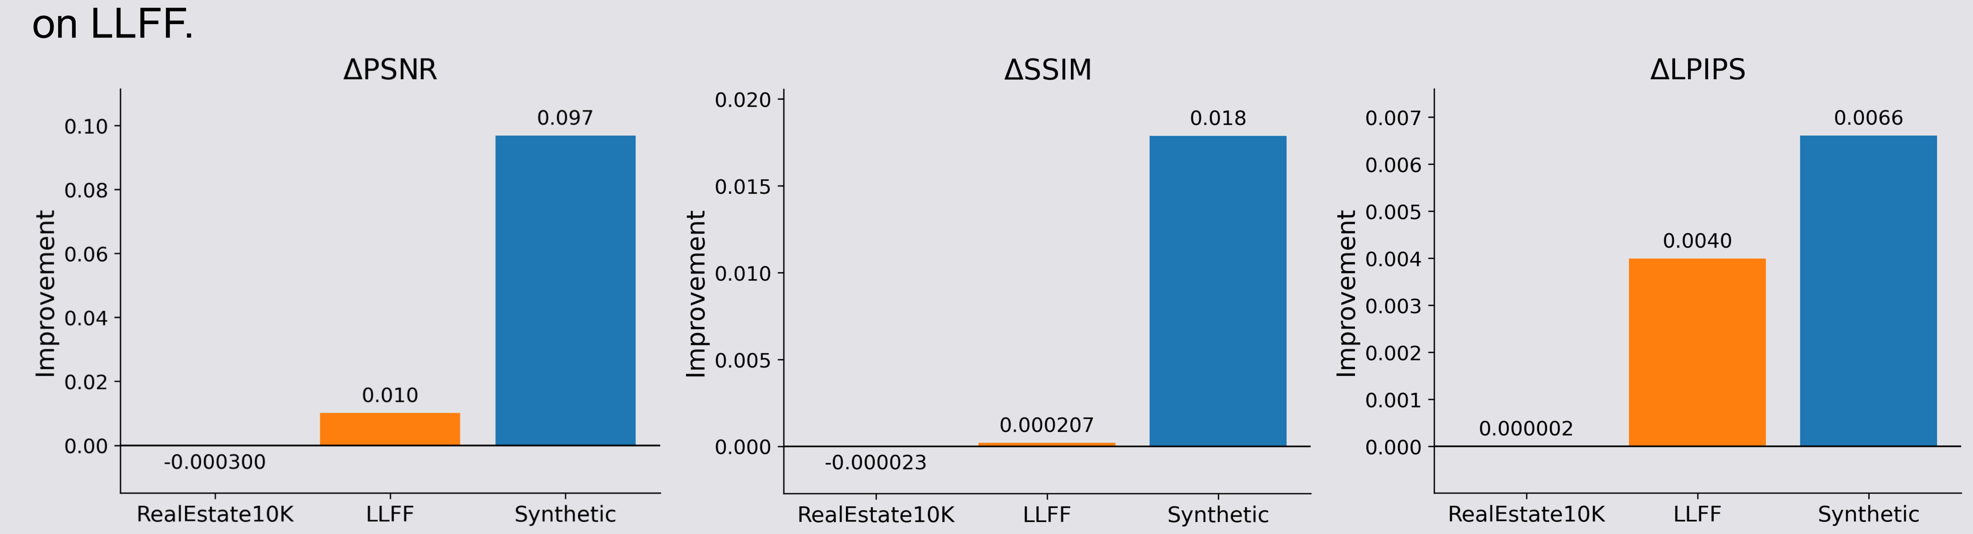

Positive Δ means improvement over the flat MPI baseline.

Controlled synthetic scene

To isolate the effect of deformation, a Blender wall scene was converted into an RGB-D-derived MPI with accurate depth. In this cleaner setting, deformable layers better follow the slanted wall and produce clear improvements across all metrics.

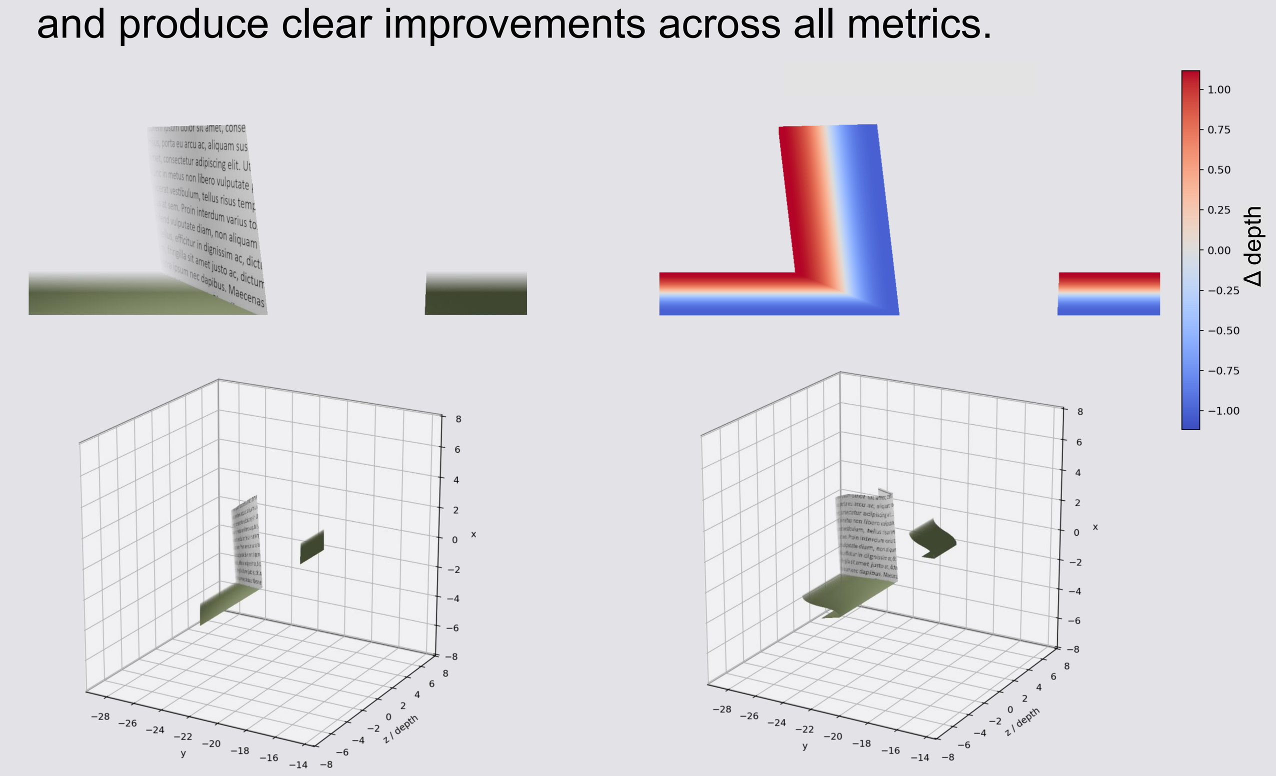

Runtime cost

Deformable mesh rendering is approximately 2.3x slower than flat MPI rendering.

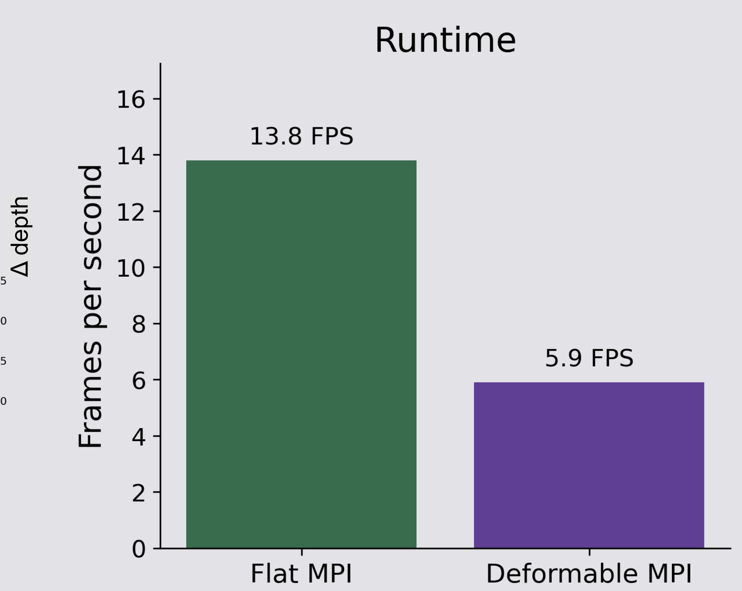

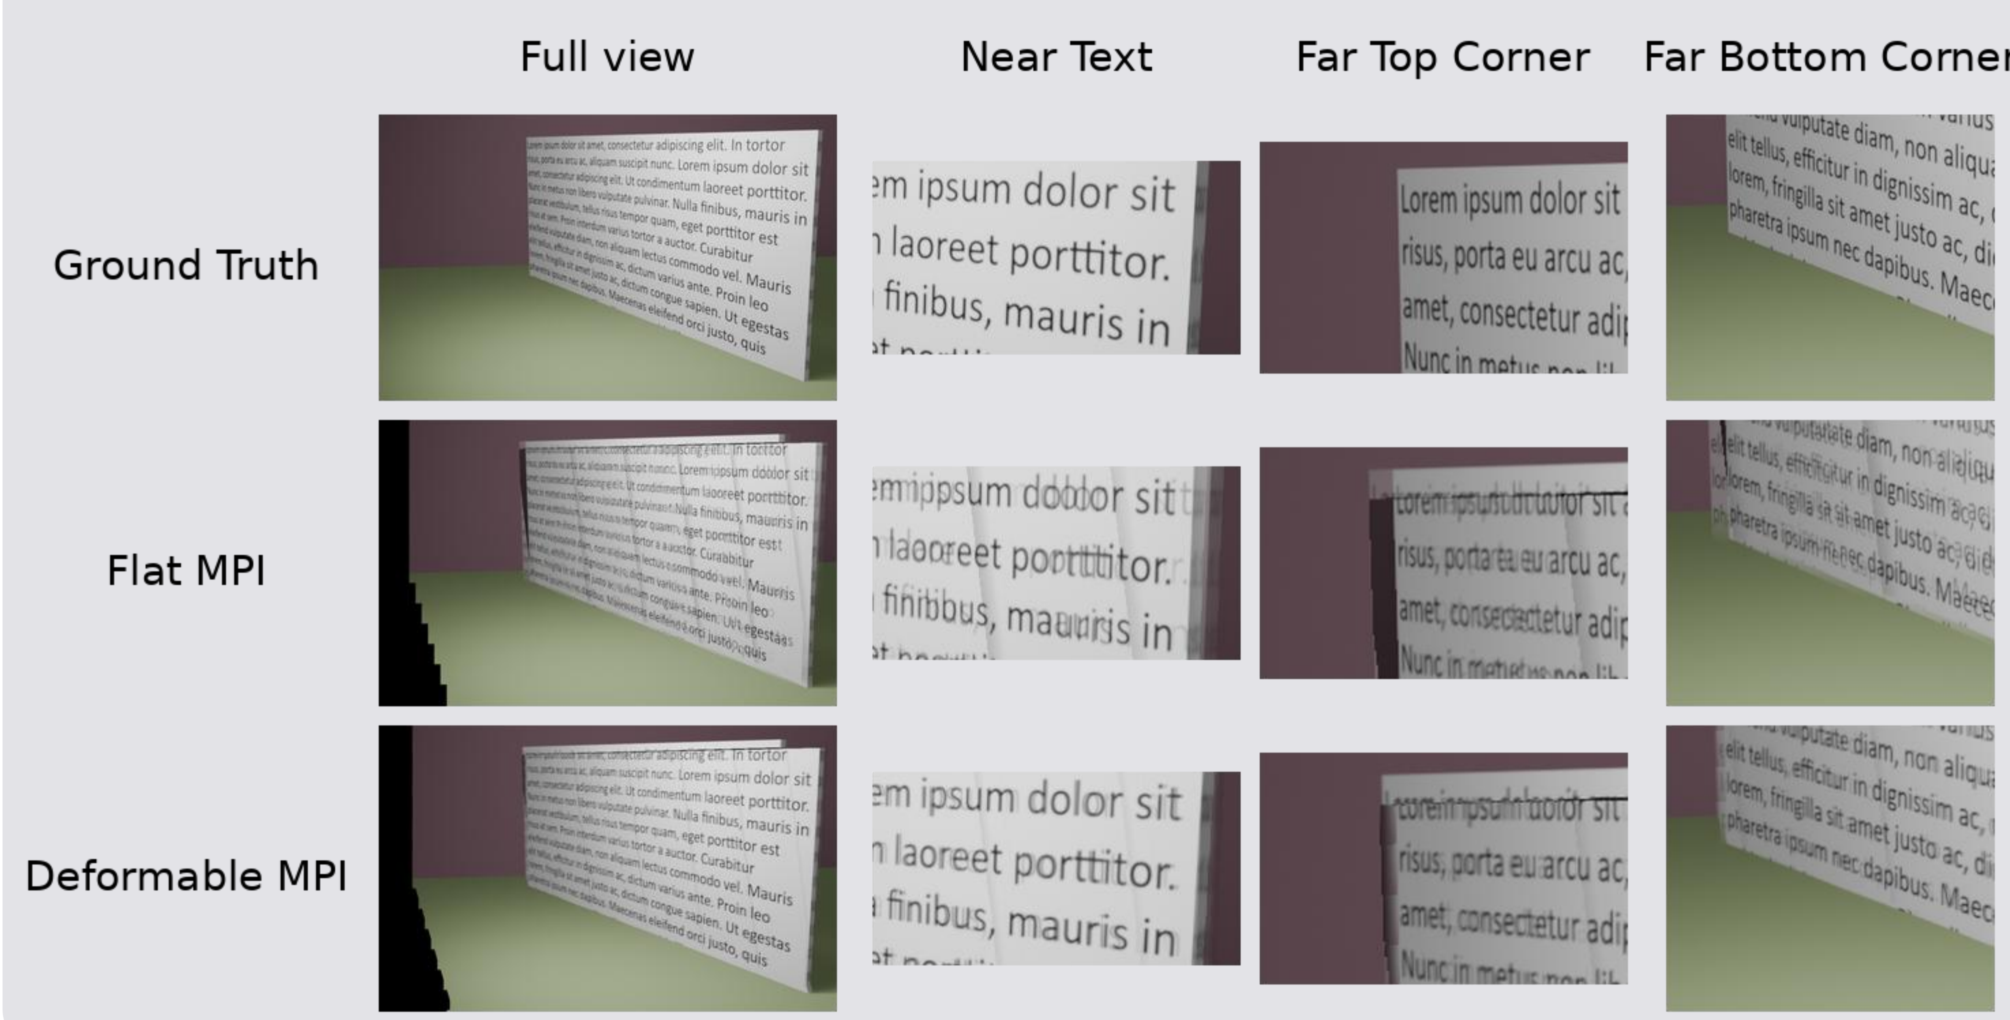

Conclusion

The proposed deformation does not meaningfully improve predicted single-view MPIs. On RealEstate10K the output is almost unchanged, and on LLFF the gains are only marginal, while rendering becomes substantially slower.

The method fails mainly because it only changes geometry after the MPI has already been predicted. It can move existing layers, but it cannot fix missing content, incorrect colors, noisy alpha maps, or wrong depth structure. Since the deformation signal is computed from the predicted MPI itself, it cannot reliably correct errors already present in that MPI. Overall, the results point to prediction quality, rather than mesh deformation alone, as the main bottleneck.

The synthetic experiment shows that deformation can help when the input MPI has reliable geometry. This suggests that deformable layers are useful only when the geometric signal is accurate and should likely be learned together with MPI appearance rather than corrected afterward. This makes the method more useful as a design direction for future MPI models than as a standalone improvement to existing predicted MPIs.

Future work

Integrate deformable geometry directly into MPI prediction, so that the model learns layer geometry and appearance jointly instead of deforming a fixed MPI afterward.

Evaluate models with fewer MPI layers, where flat planes are more widely spaced and deformation may provide a clearer benefit.

Study local improvements around slanted surfaces, object boundaries, and disoccluded regions, since full-image metrics may hide small geometric gains.

Run ablations on visibility weighting, depth-closeness weighting, layer-order constraints, and mesh resolution to better understand which parts of the deformation are useful.