Computational Photography Lab.

Find the project web page here:
https://yaksoy.github.io/2DGraphicsComp/

Physically-Based Compositing of 2D Graphics

Tyrus Tracey
Simon Fraser University
Canada

Stefan Diaconu
Simon Fraser University
Canada

Sebastian Dille
Simon Fraser University
Canada

S. Mahdi H. Miangoleh
Simon Fraser University
Canada

Yağız Aksoy
Simon Fraser University
Canada

Figure 1: We propose an interactive pipeline that enables the seamless integration of a 2D graphic into a target image, adapting to the surface geometry and lighting conditions of the scene to ensure realistic appearance. Image credits: @Toby Parsons, @Clker

ACM Reference Format:
Tyrus Tracey, Stefan Diaconu, Sebastian Dille, S. Mahdi H. Miangoleh, and Yağız Aksoy. 2025. Physically-Based Compositing of 2D Graphics. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH Posters '25), August 10-14, 2025*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3721250.3743006

1 INTRODUCTION

Inserting a 2D graphic (e.g., a logo, poster, icon, etc.) into an image so that it looks natural is no trivial task. It can potentially take several hours of manual adjustment and retouching to create a realistic result. To perform a natural-looking insertion, two key properties of the scene must be respected:

First, the source image must align with the geometry of the target scene. Failing to do so will result in a "floating" 2D graphic disconnected from the rest. Projective methods may solve this problem for simple cases, but they critically fail to handle situations found in everyday images. For example, projecting onto a side of a cube such that the source image wraps around an edge will result in strong warping.

Second, the source's final color intensities must respect the lighting of the target scene. While human vision is surprisingly weak at detecting image manipulations, they can be easily spotted based on illumination inconsistencies [Nightingale et al. 2022]. We propose a solution to both problems, provided as an efficient, yet intuitive tool for users to download and experiment with.

2 PHYSICALLY-BASED 2D GRAPHIC COMPOSITING

Our pipeline addresses these aspects step by step. We provide two Python scripts to retrieve a geometry and a lighting representation for the scene, as well as a Blender [Blender 2025] plugin to perform the insertion. We show an overview of the overall workflow in Figure 2 and describe the details of our scene-aware integration in the following paragraphs:

Monocular Geometry Estimation. We first recover a 3D representation of the scene to accurately align our 2D graphic with the geometry. For this task, we run the target image through the monocular depth-estimation method MoGe [Wang et al. 2025], which estimates a 3D point map of the image as well as camera intrinsics for rendering the scene. We then use utils3d [Jain 2024]

SIGGRAPH Posters '25, August 10-14, 2025, Vancouver, BC, Canada

ACM ISBN 979-8-4007-1549-5/2025/08
https://doi.org/10.1145/3721250.3743006

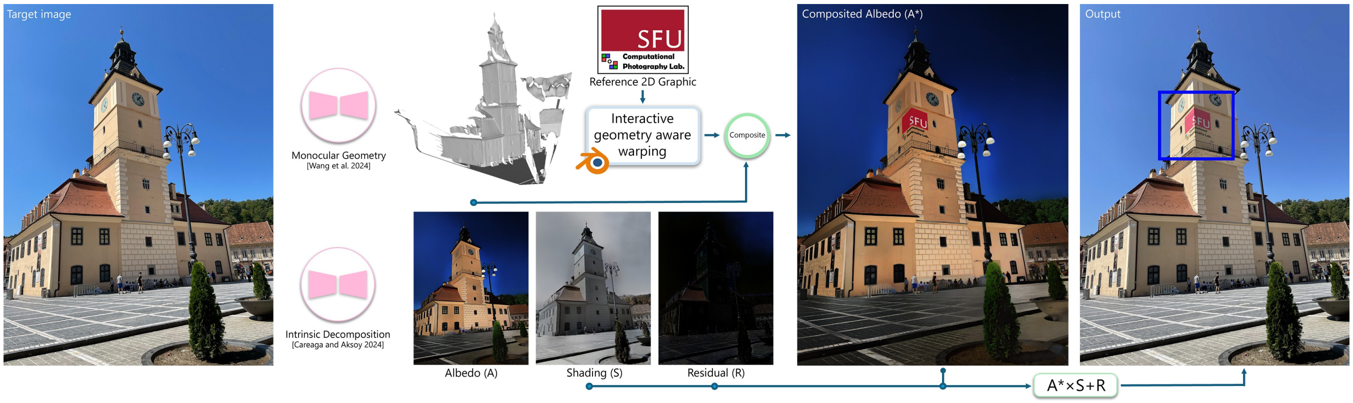

Figure 2: Overview of our pipeline. We use the intrinsic decomposition method of Careaga and Aksoy [2024] to extract intrinsic components and MoGe [Wang et al. 2025] for geometry estimation. The user selects a target region via our interactive Blender add-on over the estimated mesh. The reference logo is then warped and composited into the albedo. The final output is generated by recombining the intrinsic layers, producing a realistic compositing consistent with the scene's illumination and geometry.

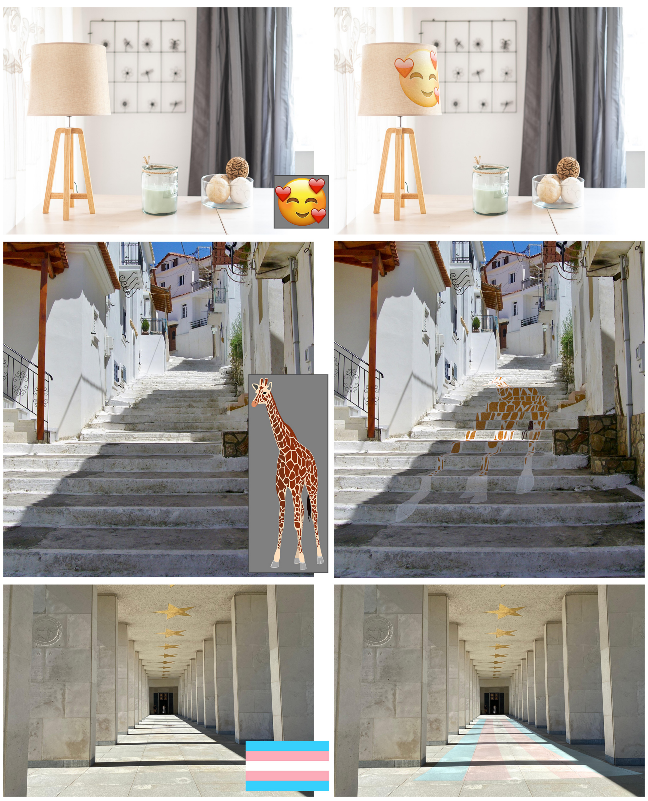

Figure 3: Example results generated using our pipeline. Logos adapt to the target surfaces, preserving consistency with the scene's geometry and varying lighting conditions. Image credits: @Joce lyne, @Maxpic, @Eisbaer, @Clker, @Frank, @Katlove

to convert the point map into a triangle mesh that can be loaded into Blender.

Lighting Representation. Secondly, we perform intrinsic decomposition for both images to retrieve the surface colors, as well as the shading as our lighting representation. The intrinsic residual model represents an image I with the surface reflectance A, the diffuse shading S, and a residual component R containing specularities and non-diffuse effects: $I = A * S + R$.

By recovering these components via a pretrained decomposition method [Careaga and Aksoy 2024], we can freely manipulate the surface colors in the albedo layer while still respecting the scene's overall lighting.

Source Image Insertion. In Blender, our plugin first performs a single raycast from the 3D viewpoint to target a face on the mesh, identifying the location of the insertion. The user can freely control the ray to select a desired point on the mesh and expand the region where the source image will be mapped onto. This separates our approach from simple projection, as the expansion follows the mesh's geometry. Our implementation enables fine-tuning of the mapping by controlling rotation, scale, and translation of the source texture within the target region of the mesh. With the source image in place, the scene is rendered back to a 2D image, using the target image's resolution along with the camera parameters estimated by MoGe [Wang et al. 2025].

We create the final result by alpha-blending the mapped, rendered source albedo with the original albedo of the target image. Next, we multiply the result by the target shading and add the residuals, effectively integrating the source colors into the target scene.

3 CONCLUSION

We demonstrate an interactive approach to physically-based 2D graphic insertion that can be completed from start to finish in a matter of minutes. Our tool provides an intuitive workflow to generate realistic composites with high fidelity.

REFERENCES

Blender. 2025. Blender - a 3D modelling and rendering package. http://www.blender.org

Chris Careaga and Yağız Aksoy. 2024. Colorful Diffuse Intrinsic Image Decomposition in the Wild. *ACM Trans. Graph.* 43, 6, Article 178 (2024), 12 pages.

Kalash Jain. 2024. utils3d. https://pypi.org/project/utils3d/

Sophie J. Nightingale, Kimberly A. Wade, and Derrick G. Watson. 2022. Investigating age-related differences in ability to distinguish between original and manipulated images. *Psychology and Aging* 37, 3 (2022), 11 pages. https://doi.org/10.1037/pag0000682

Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang. 2025. MoGe: Unlocking Accurate Monocular Geometry Estimation for Open-Domain Images with Optimal Training Supervision. In *Proc. CVPR*.