Reflectance reconstruction by HDR Inverse Rendering for SfM/MVS workflow

Introduction

Assessing a surface's Reflectance is a longstanding problem in the digitisation process since the colours we see, determined during SfM/MVS workflows, are the result of the combined contribution of the reflectance measured for each R,G,B wavelengths. However, while reflectance is a physical property of the surface, the colours we see result from the interaction of incoming light with the surface itself and are hence just apparent, while the combined reflectance in RGB channels can be visualised as a colour which is lighting-independent. This is the intrinsic colour, or albedo, which for Lambertian materials accounts for the whole amount of light reflected by the surface.

The pioneering work of Yu et al. (1998,1999) and further developments by Debevec et al. (2004) posed a first solution via Inverse Global Illumination, an analytical approach which determined the unknown radiance by applying the known lighting setup as an HDR map to a 3D model of the scene until convergence of the outcome with a sparse set of photograph. The research moved from these experiments mainly with the purpose of relighting pictures and frames. In the late 20ies, this was accomplished by abandoning the analytical approach for a Neural Networks-based one (Yu et al., 2020; Srinivasan et al., 2020). The flaw of these methods in respect to the SfM/MVS requirements is that they do not provide the albedo for the whole surface of objects, which in turn is the main purpose of digitisation and hence do not solve the aforementioned problem.

An alternative approach is Shadow Removal, which directly works on images but is specifically aimed at removing sharp direct shadows while is less effective for objects' self-shadowing and diffused environmental light sources. Cipriani, Fantini, and Bertacchi (2015) proposed a method for Shadow Removal based on Inverse Rendering where the shadowed and lit portions were separately colour-calibrated and then merged, but the method is not suitable for the more complex and diffused lighting schemes which are typical of small-object's digitisation cases.

The method we are proposing reconnects with the pioneering work of Yu and its Inverse Global Illumination approach, merged with Debevec's work about HDR Image Based Lighting and aware of Cipriani's work, but overcomes the difficulties of a dedicated analytical algorithm by directly acquiring an environmental HDR image via multibracketing with a 360° camera and performing the Inverse Rendering with any render engine. By doing so, the digitised 3D assets would become more suitable for RTR applications and would hence acquire improved accessibility and dissemination potential.

Material/Data

The method proposed derives from an extensive SfM/MVS digitisation of museum collections aimed at their virtualisation. Most of the assets could not be moved and hence had to be acquired under conditions which do not allow for a complete removal of site lighting contribution. The digitised objects thus presented extensive self-shadowing which limited the possibility of changing the virtual lighting setup to avoid the double shadow effect.

- Well-established workflow for outdoor scenes under direct sunlight is not applicable (blending of differently colour-calibrated subsets via B/W inverse rendering) because of coloured complex bounced light and lack of sharp direct shadows.

Methodology

The methodology comprises two phases: a preliminary synthetic scene case study for the validity assessment of the intrinsic colour derivation from the apparent colour that uses a purely digital scene where intrinsic colour is known and a real-case scenario case study which includes SfM/MVS acquisition of a test object and its intrinsic colour determination via the technique validated in the first case.

1_Preliminar synthetic scene case study. Scope: assessing the solidity of the technique

- 3D asset modelling with known UV-mapped intrinsic colour, pure Lambertian
- Scene lighting setup and HDR spherical map calculation
- Baking of coloured 3D asset rendered shading on the same UV set used for the colour
- 3D asset colour is set to pure white (R:255, G:255, B:255)
- Baking of white 3D asset rendered shading on the same UV set
- Export both versions as HDR to avoid data loss and ambiguities in gamma correction
- The two baked maps are blended in Photoshop using Divide mode
- The resulting map is the reconstructed intrinsic colour
- Mean Squared Error between the original known intrinsic colour and reconstructed one

2_Real scenario case study. Scope: testing the technique

- Digitisation of object via SfM/MVS
- Import in 3D modelling software
- UV parametrisation (unwrapping)
- Baking of unlit self-shadowed digitised asset
- Changing mat to pure white Lambertian
- HDR extraction of environmental lighting via 360 camera multi-bracketing
- IBL lighting setup using HDR environment spherical map
- Baking of pure white version
- Export of baked rendered colour textures as HDR into Photoshop
- Blending of the two maps in Photoshop using Divide mode
- Save reconstructed real colour as Albedo map for the 3D asset.

Results & Conclusions

The method proposed results to be technically valid and produces consistent results in comparing the synthetic reconstructed intrinsic colour and the previously known one. The method loses efficacy with self-illumination by bounced light because, in the pure-white calculations, the bounced light is not tinted by the surface directly hit by incoming light, hence is the less accurate the more complex the geometry of the digitised object and the more diverse is the reflectance of real object amongst RGB channels. This could be accounted for in an iterative process which might be the aim of future work.

Discussion

The method is easily implementable in any SfM/MVS workflow because it mostly relies on tools that are already implied in it, with the sole exception of HDR IBL reconstruction that can be assessed by using any 360° camera capable of capturing the whole dynamic range of the scene and to output digital negatives to avoid on-board uncontrollable optimisations. The adoption of this method allows the production of more ductile 3D assets by minimising self-shadowing in Albedo texture maps and thus allowing drastic relighting within the virtual scenes and consistent shading of the 3D assets.

Bibliographic References

Cipriani, L., Fantini, F., Bertacchi, S., 2015. El color en las piedras y en los mosaicos de Rávena: nuevas imágenes de los monumentos antiguos a través de la fotogrametría no convencional de última generación. EGA Expresión Gráfica Arquitectónica 190–201. https://doi.org/10.4995/ega.2015.4052

Debevec, P., Tchou, C., Gardner, A., Hawkins, T., Poullis, C., Stumpfel, J., Jones, A., Yun, N., Einarsson, P., Lundgren, T. and Fajardo, M., 2004. Estimating surface reflectance properties of a complex scene under captured natural illumination. *Conditionally Accepted to ACM Transactions on Graphics*, *19*, p.2.

Srinivasan, P.P., Deng, B., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., Barron, J.T., Barron, J.T., 2020. NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis. arXiv: Computer Vision and Pattern Recognition. https://doi.org/10.1109/cvpr46437.2021.00741

Stumpfel, J., Tchou, C., Jones, A., Hawkins, T., Wenger, A., Debevec, P., 2004. Direct HDR capture of the sun and sky, in: Proceedings of the 3rd International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa, AFRIGRAPH '04. Association for Computing Machinery, New York, NY, USA, pp. 145–149. https://doi.org/10.1145/1029949.1029977

Yu, Y., Debevec, P., Malik, J., Hawkins, T., 1999. Inverse global illumination: recovering reflectance models of real scenes from photographs, in: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '99. ACM Press/Addison-Wesley Publishing Co., USA, pp. 215–224. https://doi.org/10.1145/311535.311559

Yu, Y., Malik, J., 1998. Recovering photometric properties of architectural scenes from photographs, in: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98. Association for Computing Machinery, New York, NY, USA, pp. 207–217. https://doi.org/10.1145/280814.280874

Yu, Y, Meka, A., Elgharib, M, Seidel, H., Theobalt, C., Smith, W., 2020. Self-supervised Outdoor Scene Relighting | SpringerLink [WWW Document]. URL https://link.springer.com/chapter/10.1007/978-3-030-58542-6_6 (accessed 6.20.24).