Real-Time Global Illumination Decomposition of Videos: Supplementary Document

ABHIMITRA MEKA*, Max Planck Institute for Informatics, Saarland Informatics Campus and Google
MOHAMMAD SHAFIEI*, Max Planck Institute for Informatics, Saarland Informatics Campus
MICHAEL ZOLLHÖFER, Stanford University
CHRISTIAN RICHARDT, University of Bath
CHRISTIAN THEOBALT, Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

CCS Concepts: • **Computing methodologies** → **Computational photography**; **Image processing**.

Additional Key Words and Phrases: illumination decomposition, direct and indirect illumination, real-time sparse-dense optimization

ACM Reference Format:
Abhimitra Meka, Mohammad Shafiei, Michael Zollhöfer, Christian Richardt, and Christian Theobalt. 2021. Real-Time Global Illumination Decomposition of Videos: Supplementary Document. *ACM Trans. Graph.* 1, 1, Article 1 (January 2021), 4 pages. https://doi.org/10.1145/3374753

1 COMPARISON WITH CARROLL ET AL. [2011]

We provide a comparison with the method of Carroll et al. [2011] based on the input images and results provided in their supplementary material. The method of Carroll et al. [2011] takes as input the decomposed intrinsic layers of reflectance and illumination from a state-of-the-art intrinsic image decomposition method (in this case, the method of Bousseau et al. [2009]), and further decomposes them into the global indirect and direct illumination layers. Our method performs a joint decomposition of the input image into the reflectance, illumination and global illumination layers through optimization of a single cost function, without the need for an external intrinsic image decomposition technique. Please note that the method of Carroll et al. [2011] focuses on intrinsic decomposition of single images in an offline process, and hence is able to use extensive user interaction in the form of scribbles. Our method aims to decompose videos at real-time frame rates, which makes it infeasible to use user scribbles as additional input. Instead, we only use a limited amount of user interaction through mouse clicks in the first frame of the video.

While our technique achieves state-of-the-art quality for global illumination decomposition in videos, it still has limitations for a certain type of scene. Since our method relies on user clicks and a flood-fill algorithm to decompose regions of strong inter-reflections, it fails when such the reflectance of such a region and bleeds into a surrounding region of similar color. For example, in Figure 5, the additional constraint provided by a user click on the under arm region with strong inter-reflections from the shirt is not sufficient for our method to disambiguate the reflectance from the global illumination.

Another limitation of our method can be seen in scenes with objects of very similar reflectance. For example, in Figure 6, our method fails to decompose the illumination layer around the head of the baby due to the proximity of its reflectance to shirt's reflectance in chromaticity space. Overall, this leads to an inaccurate global illumination decomposition. Nonetheless, the reflectance and direct illumination layers are still plausibly estimated.

REFERENCES

Adrien Bousseau, Sylvain Paris, and Frédo Durand. 2009. User-Assisted Intrinsic Images. *ACM Trans. Graph.* 28, 5 (December 2009), 130:1–10. https://doi.org/10.1145/1618452.1618476

Robert Carroll, Ravi Ramamoorthi, and Maneesh Agrawala. 2011. Illumination decomposition for material recoloring with consistent interreflections. *ACM Trans. Graph.* 30, 4 (July 2011), 43:1–10. https://doi.org/10.1145/2010324.1964938

*Authors contributed equally to this work.

Authors' addresses: Abhimitra Meka, Max Planck Institute for Informatics, Saarland Informatics Campus and Google, abhim@google.com; Mohammad Shafiei, Max Planck Institute for Informatics, Saarland Informatics Campus, moshafie@eng.ucsd.edu; Michael Zollhöfer, Stanford University, zollhoefer@stanford.edu; Christian Richardt, University of Bath, christian@richardt.name; Christian Theobalt, Max Planck Institute for Informatics, Saarland Informatics Campus, Campus E1.4, Saarbrücken, 66123, Germany, theobalt@mpi-inf.mpg.de.

0730-0301/2021/1-ART1 $15.00
https://doi.org/10.1145/3374753

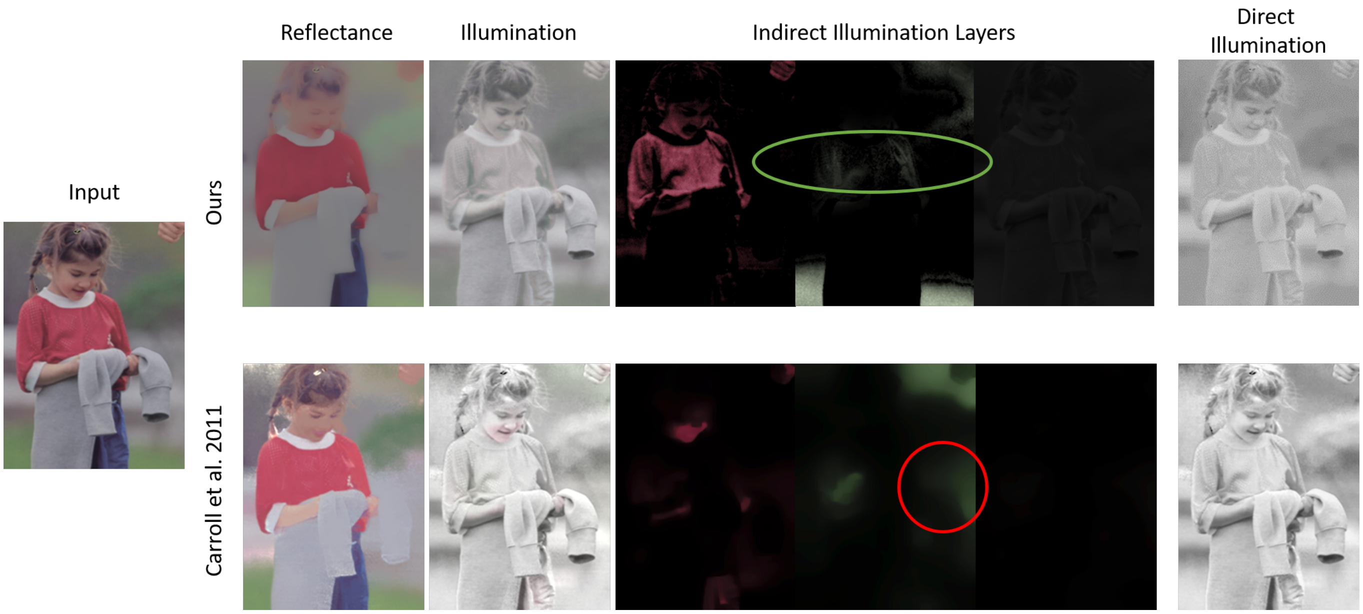

Fig. 1. Comparison to Carroll et al. [2011] on the Girl sequence. Note that our method is able to decompose the strong inter-reflection within the shirt, on the face and the neck of the girl. Our method also correctly identifies the green spill from the grass on the shirt. The method of Carroll et al. [2011] incorrectly identifies green spill in the background (circled in red).

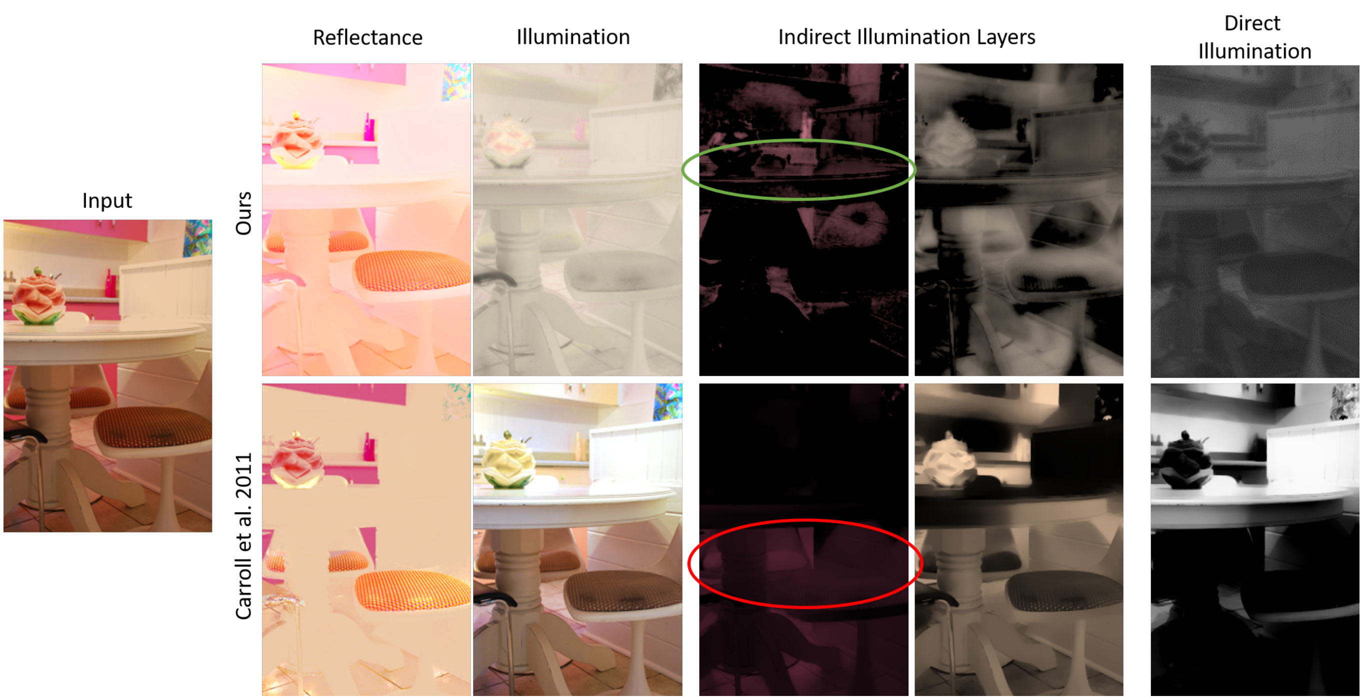

Fig. 2. Comparison to Carroll et al. [2011] on the Kitchen sequence. Our method correctly identifies the pink inter-reflections on the glass tabletop, while their method misses it. Their method also incorrectly identifies pink inter-reflections on the surface of the chairs instead of the yellow inter-reflections from the background light. Our method also better decomposes the illumination and reflectance on the chair and leaves less color in the illumination layer.

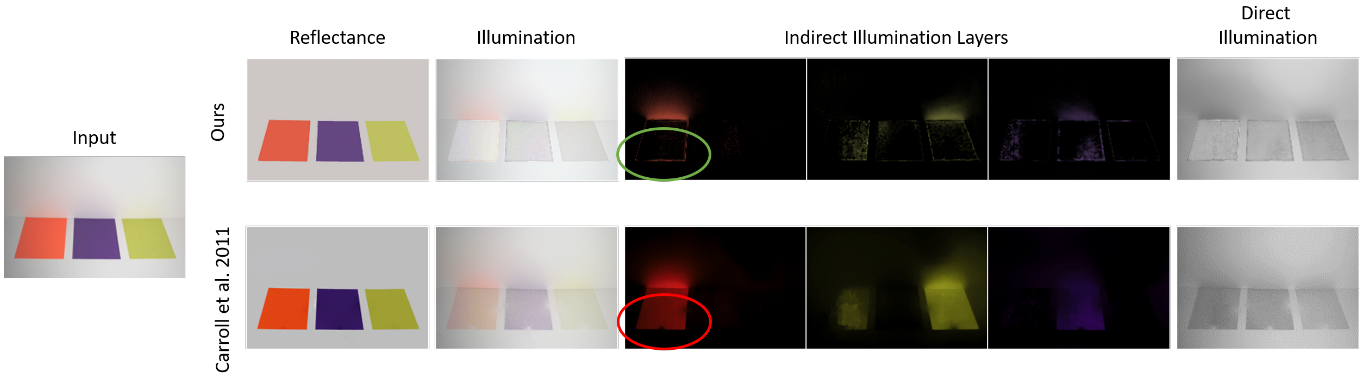

Fig. 3. Comparison to Carroll et al. [2011] on the PAPER sequence. Our method is able to accurately decompose the orange colored paper into reflectance and illumination, leading to a more accurate global illumination decomposition.

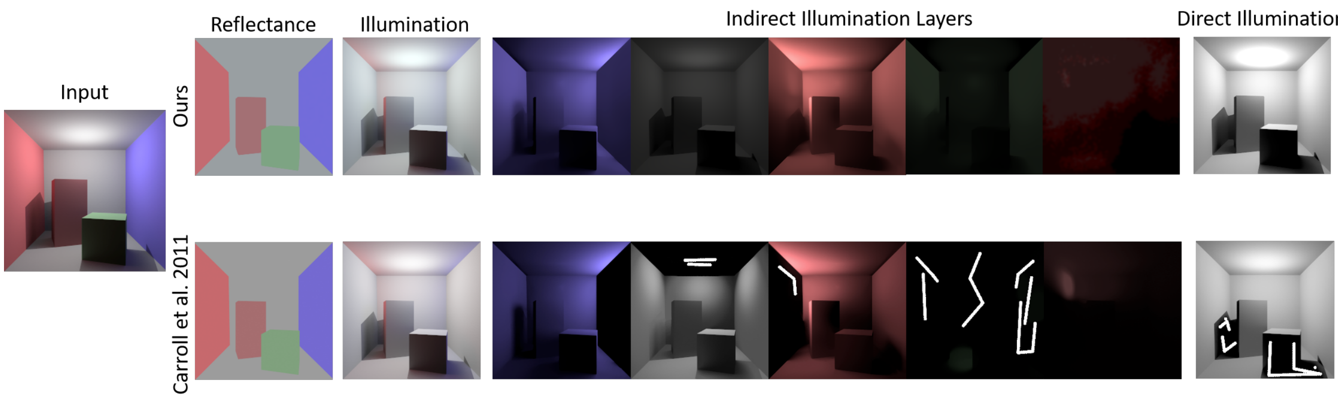

Fig. 4. Comparison to Carroll et al. [2011] on the CORNELL Box sequence. We compare our method on the intrinsic decomposition of the empty Cornell box image. Our method produces an equally plausible decomposition for the simple input. They are able to additionally correct the decomposition using multiple user interaction strokes.

Fig. 5. Comparison to Carroll et al. [2011] on the TEASER sequence. Our method fails to perform an accurate decomposition of the strong inter-reflection in the under arm region, which cannot be easily corrected due to the neighbouring dark region with similar chromaticity. The method of Carroll et al. [2011] uses user scribbles to correct it.

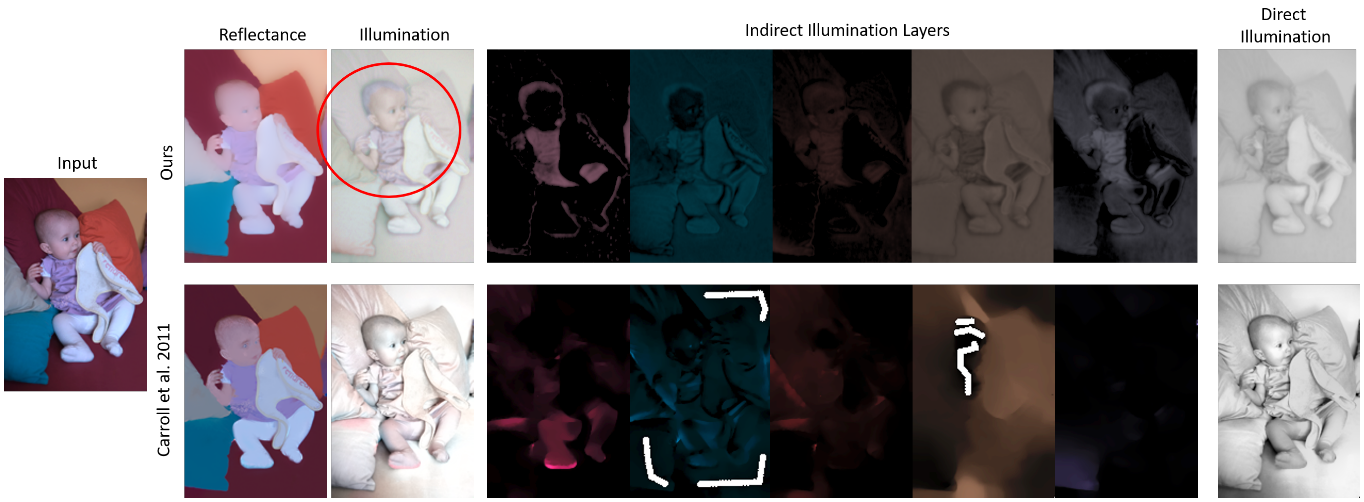

Fig. 6. Comparison to Carroll et al. [2011] on the Baby sequence. Our method fails to correctly decompose the image due to the regions neighbouring the baby's head and the shirt having similar underlying chromaticities nearby. The method of Carroll et al. [2011] is able to improve the decomposition through extensive user interaction. Please note that their method also does not perform the decomposition into the reflectance and illumination, but assumed it to be provided, while our method estimates all layers jointly.