

Streamlining IBL workflows with computer vision and USD

Stefano Cieri
DNEG
sci@dneg.com

Alexander Schwank
DNEG
axs@dneg.com



Figure 1: Clarisse scene with extracted lights and overlaid HDR. ©2018 Sony Pictures. All rights reserved.

ABSTRACT

DNEG is constantly improving its global tools and processes to make them more efficient and artist-friendly, while leveraging state-of-the-art technologies and trends in the industry. The special requirements for Sony Pictures' *Venom* movie were the perfect opportunity for us to improve our Image Based Lighting (IBL) workflows. In this paper we present the *iblManager*: a semi-automated system to allow fast and artist-friendly extraction of numerous lights from HDR images (HDRIs), using computer vision. The system uses *LightCache*, a USD-based implementation of light descriptions, to allow cross-DCC usage and pipeline integration.

CCS CONCEPTS

• **Computing methodologies** → **Computer graphics**; *Computer vision*; *Image processing*; • **Applied computing** → **Media arts**.

KEYWORDS

lighting, HDRI, computer vision, workflow, automation, USD, pipeline

ACM Reference Format:

Stefano Cieri and Alexander Schwank. 2019. Streamlining IBL workflows with computer vision and USD. In *Proceedings of SIGGRAPH '19 Talks*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3306307.3328170>

1 RETHINKING LIGHT EXTRACTION

As part of DNEG's pipeline evolution process a working group was set up to make the creation of quality light rigs easier and faster while retaining artistic control and freedom. A Nuke template

was created to define standard procedures for artists to process HDRIs which also contained nodes to extract lights. However, the extraction relied heavily on manual work, causing slow iterations and leaving space for human error. Furthermore, the system was not integrated into the pipeline which made dependency tracking between light rigs and HDRIs very hard. We realized soon that fixing such a key process was crucial.

1.1 Related Work

IBL has been a well-established technique for years now but, although much literature exists about capturing and generating HDRIs, little can be found on extracting textured lights from them.

Originally, since global illumination was computationally expensive, light extraction was aimed at approximating environment maps. An interesting example is presented in [Debevec 2006]. Although partly superseded by other techniques, we think that its programmatic approach to HDRIs analysis is an effective strategy worth recovering. Nowadays, the use of area lights with HDR textures to complement environment maps is a common practice. Solutions for extracting such textures were first presented by Digital Domain [Mihashi et al. 2009] and later by ILM [Snow 2010]. Since then practices have settled mostly around this approach and the tools in use across the industry seem to be similar.

A recent innovation is [Suter et al. 2016] where a computer-aided solution is presented to determine the size and location of lights via pairs of HDRIs. The method was not viable for us at the time as it didn't scale well on many lights and it had the limitation of doubling the time of onset surveys.

1.2 Employing Computer Vision

When dealing with IBLs most of the automation applies to the capture and stitching of HDR panoramas or the creation of light rigs from extracted textured lights, while the extraction process itself still relies on a certain amount of manual work by the artist. This makes sense as the choice of lights to extract requires an

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGGRAPH '19 Talks, July 28 - August 01, 2019, Los Angeles, CA, USA

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6317-4/19/07.

<https://doi.org/10.1145/3306307.3328170>

understanding of the technical and artistic requirements of each scenario. Still, our objective was to significantly ease this phase.

We decided to leverage computer vision to easily detect large numbers of lights while providing high level controls for artists to steer the algorithm's selection criteria. Setting a luminance threshold and a minimum and maximum area determines what the algorithm will consider a light source. A dilate factor helps to cluster close, small emissive sources and use them as a single area light.

2 IMPLEMENTATION

Since Nuke is the standard package for HDRI processing and it is familiar to artists, we decided to use it as our front-end. This allowed for a high degree of customization and left space for a creative use of the tool. Python was chosen to allow fast iteration while leveraging OpenCV features and the speed provided by its C/C++ implementation.

2.1 User Interface and Nuke Integration

We developed a set of Nuke Gizmos for the main process steps: *iblLightFind*, *iblLightExtract* and *iblRigPublish*.

The *iblLightFind* gizmo is the process entry point where artists set the high level parameters which drive the algorithm. Since we are dealing with HDR images it is easy to separate emissive areas from non emissive ones: based on the artist input, the HDRI is pre-processed to produce a LDR image where only the values above the energy threshold are kept. This minimizes errors in identifying the lights' contours, making the process trivial, fast and reliable. Once the contours are found, the minimal bounding box is computed to detect each light's roll axis rotation.

A custom UI in the node properties previews what the algorithm sees as the input image as well as the contours and the bounding boxes identified. This way, the artist can interactively correct and find the parameters which best fit his needs.

When happy with the parameters, the artists can confirm their choice and an *iblLightExtract* node per light is created to fine-tune the lights position, size and rotation before publishing.

At this point the *iblRigPublish* will take care of exporting all the required textures and generating a USD *LightCache* which can be either previewed in Clarisse or published to the database.




Figure 2: Light extraction node set up in Nuke.

2.2 Pipeline

Prior to work on the *iblManager*, DNEG made efforts to restructure and consolidate its light rig transfer across departments and DCCs.

The main goal was to have a single flexible light representation that can be imported and exported from any of the main DCCs Maya, Clarisse, Houdini and Nuke.

Rather than implementing a new custom framework we decided to use USD as the foundation, as it offers a schema for most of our required light types, a powerful API and is modifiable and extensible for our specific needs.

Although the existing light schemas provide common light attributes, for a full representation of our lights we had to extend them. Furthermore, DCC specific conversions had to be added on import and export to achieve portable light rigs. However, manual tweaks to lights can still be required for close visual matching because some DCC light types are incompatible or missing.

3 CONCLUSION




Figure 3: A close shot of Venom using the extracted light rig. ©2018 Sony Pictures. All rights reserved.

Given the simplicity and the openness of the system, artists were able to produce light rigs very quickly and naturally began to develop personalized ways to use it (see result in Figure 3). Both the *iblManager* and *LightCache* are now part of DNEG's IBL workflow; they concurrently provide a solid framework which is at once flexible and transparently integrated with the pipeline, upon which new tools can be built. A variety of performance improvements and additional features are now being considered to leverage the potential of the current paradigm.

ACKNOWLEDGMENTS

We thank Bernardo Spadafora for his key contribution to the tool development; Ignacio Thomas, Sachin Narendra for their precious input. A big thank you goes to Matthew Smith, Bernhard Kerschbaumer, Mark Pascoe and to our colleagues, especially Keanan Cantrell, Alison Wortman, Peter Seager, Tom Bracht.

REFERENCES

- Paul Debevec. 2006. A Median Cut Algorithm for Light Probe Sampling. In *ACM SIGGRAPH 2006 Courses (SIGGRAPH '06)*. ACM, New York, NY, USA, Article 6. <https://doi.org/10.1145/1185657.1185688>
- Tadao Mihashi, Dan Abrams, Paul Lambert, Marco Maldonado, and Jesse James Chisholm. 2009. The Light Kit: HDRI-Based Area Light System for "The Curious Case of Benjamin Button". In *SIGGRAPH 2009 Talks: Two Bolts and a Button*.
- Ben Snow. 2010. Terminators and Iron Men: Image-Based Lighting and Physical Shading at ILM. In *SIGGRAPH 2010 Course: Physically Based Shading Models in Film and Game Production*.
- Alex Suter, Dan Lobl, Victor Schutz, and Brian Gee. 2016. LightCraft: Extract Light Position and Information from On-set Photography. In *ACM SIGGRAPH 2016 Talks (SIGGRAPH '16)*. ACM, New York, NY, USA, Article 15, 2 pages. <https://doi.org/10.1145/2897839.2927464>