Optimal OLAT Alignment for Image Based Relighting with Color-Multiplexed OLAT Sequence

Arvin Lin
Imperial College London
London, United Kingdom
arvin.lin19@imperial.ac.uk

Abhijeet Ghosh
Imperial College London
London, United Kingdom
abhijeet.ghosh@imperial.ac.uk

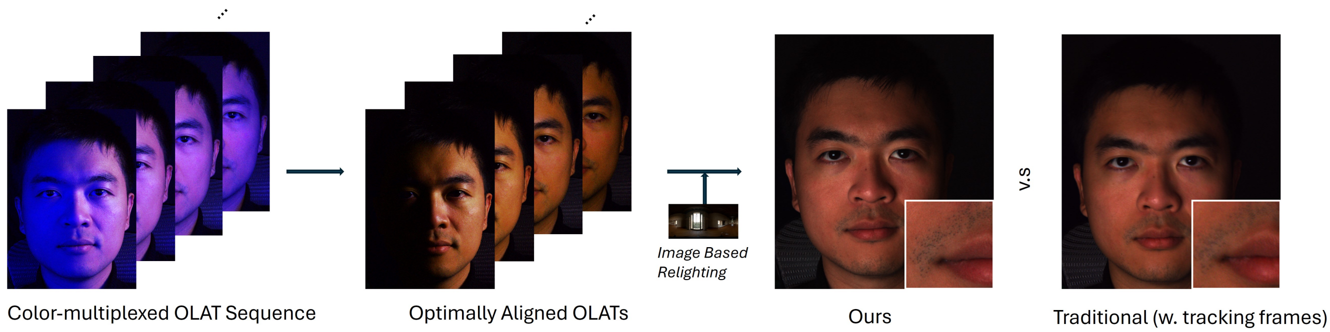

Figure 1: Our proposed color-multiplexing strategy for OLAT captures provides optimal alignment and allows sharper rendering results compared to traditional OLAT capture methods with tracking frames.

ABSTRACT

We present two color-multiplexed illumination sequences for optimally-aligned one-light-at-a-time (OLAT) captures. We leverage color-multiplexing strategies to embed tracking frames within the OLAT photographs to correct for subject motion. Our method allows better motion estimation via optical flow than traditional methods, which interleaves tracking frames between OLATs. Comparison between rendered results and user study on comfortability both demonstrate that color-multiplexed sequences give better-aligned OLATs and are more comfortable for the subject during data capture. Our proposed sequences can replace traditional OLAT sequences for better data acquisition, which would benefit both light-stage rendering results and any state-of-the-art relighting methods that are trained on OLAT-generated data.

CCS CONCEPTS

• **Computing methodologies → Image-based rendering**; **Computational photography**.

KEYWORDS

Appearance Capture, Image-based Rendering, Image Alignment

ACM Reference Format:
Arvin Lin and Abhijeet Ghosh. 2024. Optimal OLAT Alignment for Image Based Relighting with Color-Multiplexed OLAT Sequence. In *Proceedings of the 21st ACM SIGGRAPH European Conference on Visual Media Production (CVMP 2024).* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

CVMP 2024, Nov. 18–19, London, UK
2024. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

One-light-at-a-time (OLAT) images are the most common image-based reflectance representation for high-quality subject rendering. More and more advanced light stages have been built to either efficiently capture a large dataset of OLAT sequences or predict OLAT images with similar illumination sequences [Bi et al. 2021; Meka et al. 2019; Sun et al. 2019]. The ability to capture high-quality OLATs is now essential for subject relighting research and the dataset of OLAT images also drives most of the state-of-the-art photo-realistic subject relighting work.

However, even though current high end light stages can capture up to 90 fps, photographing the subject under hundreds of OLAT patterns in a couple of seconds, subjection motion during the capture still cannot be avoided. Failure to address subject motion between the OLAT images results in a sub-optimal reflectance model and will hinder the quality of both the directly rendered images and any subsequent deep-learning based relighting model that is based on the OLAT images. This issue is even more severe with lower-end capture setups where cameras capture at a lower frequency [Lin et al. 2024; Schreiber et al. 2022] where subject motion is more prominent.

While optical-flow techniques have been used to address subject motion by attempting to align OLAT images with the estimated optical-flow field, the optical-flow field cannot be reliably estimated between images under different illumination (i.e. between OLAT images). As a result, most OLAT captures rely on inserting tracking frames regularly between OLAT images to compute the optical flow. The optical-flow field for the OLAT images is then estimated by interpolating the optical-flow field between the tracking frames. This, however, introduces errors as the subject motion between the tracking frames is rarely linear. This issue is often addressed by optimizing a deep neural architecture with appropriate loss functions that deal with misalignment between images [Lin et al. 2024; Meka et al. 2019; Schreiber et al. 2022].

Instead of relying on optimizing a deep learning problem, this work aims to resolve the issue of subject motion and misalignment by exploring a different approach to OLAT illumination. Specifically, we explore the use of color-multiplexing to embed tracking frames within the OLAT illumination for optical flow computation. This strategy allows direct estimation of optical flow between OLAT images without requiring any interpolation of the optical-flow field. We present two designs ('additive color-plex' and 'negative color-plex') of color-multiplexed OLAT capture sequences and perform comparisons between the two designs and traditional sequences with interleaving tracking frames. Our results show that color-multiplexed OLAT sequences outperform traditional OLAT sequences, while also being more comfortable to the subject under high frequencies.

To summarize, this work presents:

- Two color-multiplexed OLAT capture sequence that allows optimal alignment between OLAT images.
- A comparison between color-multiplexed sequences and traditional sequences with tracking frames showing visible differences between the rendered quality.
- A user study showing preference towards color-multiplexed sequences in terms of comfortability for captures under a high frame rate.

Our study shows that color-multiplexed OLAT sequences outperform traditional tracking-frame OLAT sequences in general, and can be used to replace traditional tracking-frame sequences to capture better-quality OLAT data under most illumination setups.

2 RELATED WORK

Most state-of-the-art subject reflectance models use one-light-at-a-time (OLAT) images as an image-based representation of reflectance. OLAT images are most commonly captured using a light stage [Debevec et al. 2000], by capturing a series of images of a subject lit by each single light source of the light stage independently. One problem with capturing multiple images of a subject is that the subject will inevitably move during the capture, affecting the quality of the OLAT reflectance representation. This issue is mainly solved by inserting fully-lit (FULLON) tracking frames between the OLAT captures to account for movement by using either face tracking [Wu et al. 2018] or optical-flow [Anderson et al. 2016]. Researchers have reported inserting tracking frames every 10 [Meka et al. 2020] or 11 OLAT images [Meka et al. 2019; Sun et al. 2019] with 60fps high-speed cameras, and even having tracking frames inserted every 3 frames with 90fps high-speed cameras [Bi et al. 2021; Saito et al. 2024].

While estimating movement with the interleaved tracking frames accounts for the subject movement to some extent, movement between the OLAT images still is not directly measured. This affects the quality of the attainable reflectance model. While joint optical-flow methods have been proposed to align complementary illuminated photographs [Wilson et al. 2010], most work in OLAT captures addresses this issue by designing deep neural architectures with appropriate loss functions that are robust to misalignments. Meka et al [2019] proposed a slide-pooling loss to account for slight translations within patches of two images, while Lin et al [2024] used a misalignment-aware perceptual loss to address misalignments restricted in a local area. These works, however, adopt the standard capturing sequence of interleaving FULLON tracking frames between OLAT captures, without investigating alternative illumination sequences that offer better alignment between OLATs. Yang et al [2023] added ambient constant background lighting to their lighting patterns to eliminate the need for tracking frames, however, their method still requires proper design of a neural network that learns to disentangle the ambient lighting from the OLAT images. To the best of our knowledge, no work has attempted to design an optimal OLAT capture sequence that directly leads to optimal alignment between OLAT images.

On the other hand, other works in structured illumination and photometric stereo have investigated color-multiplexing strategies. Hernandez et al [2007] used 3 different colored LEDs and extracted the 3 color channels of the captured photograph to compute the normal and track a moving un-textured surface. Similarly, Gotardo et al [2015] used 9 colored lights to attain high-frequency dynamic 3D reconstruction. They eliminated the need for tracking frames by relighting the input images based on known surface normal to create frames lit under similar illumination for estimating optical flow. Other work has also used color-multiplexing methods to reduce the required illumination patterns to capture SVBRDF [Fyffe 2009; Kampouris et al. 2018].

Inspired by these works, we investigate color-multiplexing methods that allow us to efficiently capture OLAT images while negating the need for interleaving FULLON tracking frames. By aligning OLAT images without interpolating movement between tracking frames we also obtain better quality alignment that was otherwise not achievable.

3 COLOR MULTIPLEXED OLAT SEQUENCE

Consider an OLAT illumination sequence for a lighting apparatus with N unit lightings, with the n^{th} light producing a Dirac delta illumination centered around $\vec{\omega_n}$. The OLAT sequence would require a minimum amount of N captures:

$$OLAT(n) = \delta(\vec{\omega_n}), \quad 1 \leq n \leq N \tag{1}$$

where $OLAT(n)$ denotes the environmental illumination of the n^{th} OLAT pattern. Additionally, tracking frames are often interleaved between OLAT images for alignment, a typical illumination for tracking frames is the FULLON illumination:

$$FULLON = \sum_{n=1}^{N} \delta(\vec{\omega_n}) \tag{2}$$

.

This popular method, however, does not directly track the movement between the OLAT images since no available tracking frames are present in the OLAT images. Movement in the OLAT images has to be estimated by interpolating the computed optical flow of the tracking frames.

On the other hand, color-multiplexed illumination can be used to encode different illumination patterns into separate color channels. This allows us to insert a tracking frame directly into an OLAT image. For a lighting device capable of producing the three primary colors (R, G, B), assuming the RGB camera observes each of the three primary colors separately without cross-talks, the captured RGB photograph can observe up to 3 different illumination patterns in a single photograph.

In practice, however, RGB cameras do not observe the primary colors without cross-talks. The specific spectral characteristics of the capturing apparatus should be taken into consideration. For example, we observe non-trivial cross-talk between the red and green channels, with less cross-talk to the blue channel given our cameras and illumination setup. Thus, Our color-multiplexed sequences have color channels chosen accordingly to minimize cross-talks for our capturing apparatus.

3.1 Additive Color-Plexed OLATs

Designing a color multiplexed OLAT sequence can be quite straightforward. An implementation of adding the tracking frame in the blue channel will look like this:

$$OLAT_{RG}(n) = \begin{cases} R: & \delta(\vec{\omega}_n) \\ G: & \delta(\vec{\omega}_n) \\ B: & FULLON \end{cases}, \quad 1 \leq n \leq N \tag{3}$$

This allows us to obtain a set of directly aligned OLAT images under the red and green channels, a similar pattern can then be used to obtain aligned OLAT images under the blue channel:

$$OLAT_{B}(n) = \begin{cases} R: & FULLON \\ G: & FULLON \\ B: & \delta(\vec{\omega}_n) \end{cases}, \quad 1 \leq n \leq N \tag{4}$$

We term the sequence that is made up of $OLAT_{RG}$ and $OLAT_B$ the additive color-plexed OLAT sequence, the difference between this sequence and a standard OLAT sequence with tracking frames is illustrated in Figure 2.

Figure 2: Illustrating the color-multiplexed OLAT capture sequence. Top row: traditional OLAT sequence with multiple tracking frames. Bottom row: color-multiplexed OLAT sequence using different color channels (switching between red+green/blue for every image) for tracking.

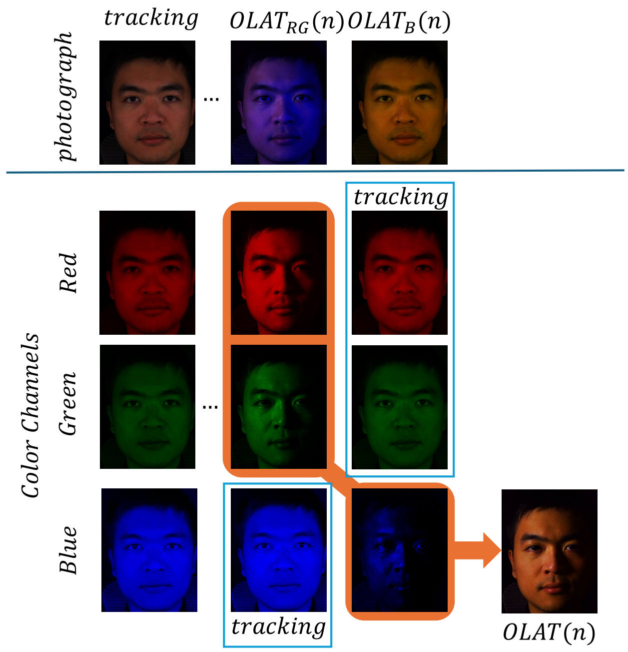

Figure 3: Illustration of the tracking and OLAT generation with additive color-multiplexed OLAT sequence. Only one tracking frame is required up front, in constrast to traditional FULLON tracking methods requiring multiple tracking frames.

The color-plexed OLAT photographs can all be aligned to a white FULLON tracking frame, and then by extracting the OLAT color channels we obtain the desired RGB OLAT images:

$$OLAT(n) = \begin{cases} R: & OLAT_{RG}(n)[R] \\ G: & OLAT_{RG}(n)[G] \\ B: & OLAT_{B}(n)[B] \end{cases}, \quad 1 \leq n \leq N \tag{5}$$

This process is further illustrated in Figure 3.

3.2 Negative Color-Plexed OLATs

While additive color-plexed OLAT sequence provides tracking frames with constant illumination, it significantly increases the number of captures. To obtain an OLAT sequence of N images, an additive color-plexed OLAT sequence requires $2N + 1$ photographs in total. The longer capture sequence is undesirable and may cause the subject to move more during the capture. To resolve this issue, we can design another color-plexed strategy by leveraging the fact that optical-flow algorithms still work sufficiently well under slight changes in illumination. By allowing the tracking frame to have slight illumination variations we can design a negative color-plexed OLAT sequence as follows:

$$OLAT_{NEG}(n) = \begin{cases} R: & \delta(\vec{\omega}_n) \\ G: & \delta(\vec{\omega}_n) \\ B: & FULLON - \delta(\vec{\omega}_n) \end{cases}, \quad 1 \leq n \leq N \tag{6}$$

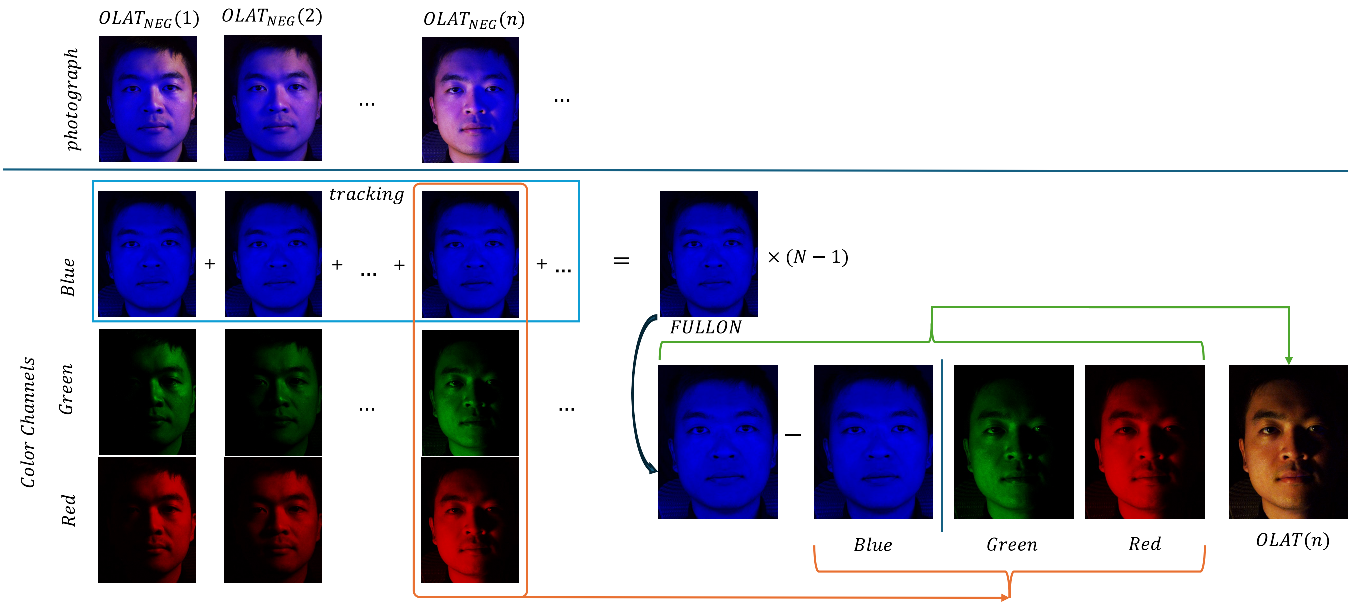

Figure 4: Illustration of the tracking and OLAT generation with negative color-multiplexed OLAT sequence. No explicit white tracking frame is required for this strategy. A FULLON blue image is computed by summing up all the blue channels of the OLAT sequence, then dividing by $N-1$ (N is the total number of the OLATs). Each OLAT image's blue channel can then be computed by subtracting the blue channel of the negative color-plexed OLAT from the FULLON blue image.

We term this a negative color-plexed OLAT sequence since the tracking frame (in our case the blue channel) contains a FULLON illumination subtracted by the target OLAT illumination. The slight variation in illumination will not affect the quality of image alignment, while at the same time providing information on the OLAT images' blue channel. To compute the OLAT blue channels we first compute the FULLON blue image:

$$FULLON[B] = \frac{1}{N-1}\sum_{n=1}^{N} OLAT_{NEG}(n)[B] \tag{7}$$

then we can estimate $OLAT(n)$ as follows:

$$OLAT(n) = \begin{cases} R: & OLAT_{NEG}(n)[R] \\ G: & OLAT_{NEG}(n)[G] \\ B: & FULLON[B] - OLAT_{NEG}(n)[B] \end{cases}, \quad 1 \le n \le N \tag{8}$$

The required photographs to capture a negative color-plexed OLAT sequence with N OLATs are N. An illustration of this sequence as well as the method of extracting the desired OLAT image is shown in Figure 4.

4 EXPERIMENTS

To evaluate the proposed color-plex patterns, we capture 7 subjects under the 32 OLAT basis used by [Lin et al. 2024] with color-plexed sequences and FULLON tracking sequences. We find that using the capture setup in [Lin et al. 2024] for color-multiplexing introduces a color shift in the synthesized OLAT images, which we correct by computing a color correction matrix per color-multiplexed sequence with a colorchart. For the additive color-plex sequence, we capture 2 sequences with the same patterns but arranged in a different order, the first interleaving $OLAT_{RG}$ and $OLAT_B$ patterns (add. color-plex (A)), and the second placing all $OLAT_{RG}$ patterns before the $OLAT_B$ patterns (add. color-plex (B)). We also capture a white FULLON image at the beginning of each sequence to serve as a reference for our evaluation.

We capture four traditional OLAT sequences with FULLON tracking frames for comparison, the tracking frame is inserted between every 1, 2, 4, and 11 OLATs respectively for these sequences. Table 1 summarizes the number of required photographs and whether OLAT images can be directly aligned for each capture sequence. Note that increasing the frequency of the tracking frame to one every OLAT capture (the highest frequency possible) would require the same amount of photographs as the additive color-plex method, while still not being able to provide direct alignment between OLATs.

4.1 OLAT Alignment

To evaluate the quality of OLAT image alignment, we render a FULLON image by summing up the aligned OLAT basis images for each capture sequence, then compare the rendered image to the reference FULLON image that the OLAT images are aligned to. Figure 5 shows an example of the rendered FULLON image and the per-pixel difference image of a subject captured under different OLAT sequences. The 3 color-plexed patterns create the sharpest detailed rendering as the OLAT images are well aligned. In contrast, FULLON images rendered with OLATs aligned with interleaved

add. color-plex (A) **add. color-plex (B)** **neg. color-plex** **tracking(1)** **tracking(2)** **tracking(4)** **tracking(11)**

Figure 5: Top row: Rendered FULLON image using the aligned patterns. Color-plexed methods provide better alignment which results in sharper rendered images. Bottom row: per-pixel difference image compared to the reference FULLON tracking frame.

Table 1: Comparison of the OLAT sequences with different tracking strategies. Method tracking(n) denotes adding a FULLON tracking frame for every n OLAT images.

Method	No. images	direct alignment
add. color-plex	$2N + 1$	✓
neg. color-plex	N	✓
tracking (1)	$2N + 1$	✗
tracking (2)	$\lceil \frac{3N}{2} \rceil + 1$	✗
tracking (4)	$\lceil \frac{5N}{4} \rceil + 1$	✗
tracking (11)	$\lceil \frac{12N}{11} \rceil + 1$	✗

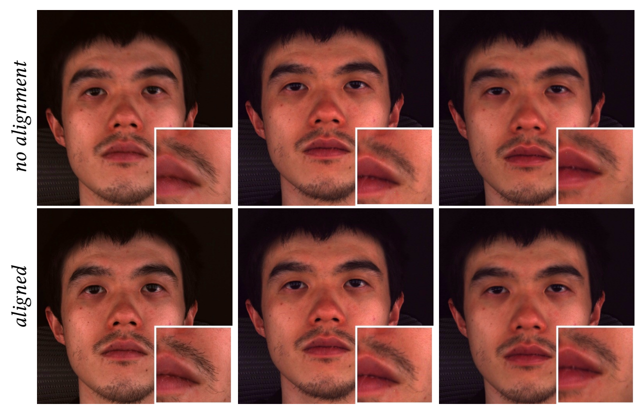

(a) neg. color-plex **(b) tracking(2)** **(c) tracking(11)**

Figure 6: Comparing the effect of optical-flow alignment between different OLAT sequences. The top row shows the rendered FULLON with unaligned OLAT images, and the bottom row shows the rendered FULLON with aligned OLAT images.

tracking frames are much blurrier, indicating that there are still misalignments between the OLAT images.

Figure 6 further shows the efficacy of using optical flow for alignment between color-multiplexed sequences and traditional tracking sequences. The top row shows the rendered FULLON with OLAT images that are not aligned, while the bottom row shows the rendered result using aligned OLATs. We can see the rendered FULLONs on the top row are all blurry, indicating movement during all captures. After optical-flow alignment, the negative color-plex sequence produces the sharpest FULLON rendering compared to either 'tracking(2)' or 'tracking(11)'.

It is worth noting that with 'tracking(11)' the optical flow does not significantly improve image alignment, however, the overall movement of the subject is less than 'tracking(2)' due to the faster capture. This results in similar rendering quality even though the movement cannot be accurately measured. This result suggests that the total number of photographs required is an important factor that needs to be considered when designing an OLAT capture sequence.

For quantitative analysis, we compute the average mean squared error (MSE) and structural similarity scores (SSIM) for each sequence across 7 subjects. We masked out the background when computing these metrics to eliminate the effect of ambient lighting changes. Table 2 shows the average MSE and SSIM across 7 subjects. We find that the negative color-plex sequence consistently outperforms other sequences with its shortest capture time and direct optical-flow alignment. The two additive color-plex sequences outperform the traditional FULLON tracking sequences.

4.2 Environment Map Relighting

The accuracy of OLAT alignment directly affects the quality of the renderings. The effect is shown in Figure 7 where tracking methods for OLAT alignment produce blurry image-based relighting results. Note that both color-multiplexed methods produce some color shift even after white-balancing the rendered images. This is caused by using computer monitors as illumination devices, as monitors do not provide a linear response when illuminating different patterns. The color shift should not be an issue if the subject is captured under

Table 2: Comparison of the average mean square error (MSE) and structural similarity metric (SSIM) of different capturing strategies over the 7 subjects captured.

Method	MSE (10^{-4})	SSIM
add. color-plex (A)	4.110	0.787
add. color-plex (B)	4.557	0.782
neg. color-plex	**3.381**	**0.805**
tracking (1)	5.702	0.747
tracking (2)	5.152	0.759
tracking (4)	5.530	0.757
tracking (11)	4.888	0.768

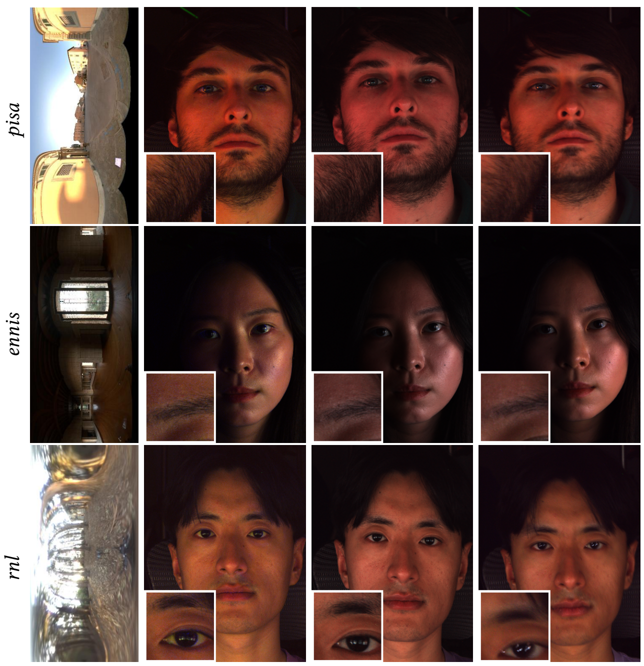

(a) env. (b) neg. plex (c) add. plex (d) tracking

Figure 7: Results of relighting under various HDR environments. Color-plexed methods (a, b) produce sharper alignments while the negative color-plex method has more color-shift (even after white balance) due to our illumination device.

more high-end setups that provide more direct control over lighting intensity (e.g. a light stage). Even with more practical setups, a slight color shift is a worthy trade-off for sharper relighting results.

4.3 Comfortability Study

Comfortability of subject during capture is an important factor influencing the likelihood of subject movement, especially blinking during the capture sequence. To study the comfortability between each capture sequence we ask the captured subjects to rate the comfortability of each illumination sequence after each capture. The subjects are asked to rate each capture on a scale of 1 to 10,

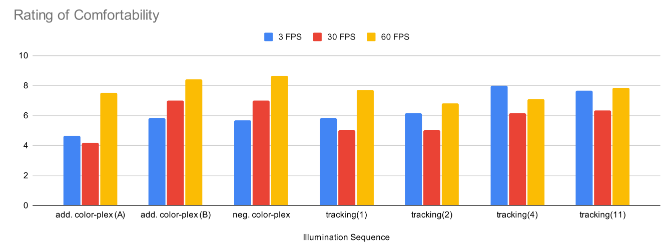

Figure 8: Average rating of comfortability across 6 subjects for each illumination sequence under different frame rates. While traditional tracking sequences are more comfortable at lower frame rates, at high frame rates subjects report better comfortability under color-multiplexed sequences.

10 being the most comfortable, and without the urge to move and blink in between captures, and 1 being the least comfortable, and the subject cannot hold the eyes open without blinking during the capture sequence.

We collected the ratings from 6 out of the 7 subjects captured (excluding the author), since our capture setup works at a much slower frame rate (3fps) compared to light stages, we also have the subject go through the illumination patterns at 30fps and 60fps without taking any photographs. The subjects then are asked to rate the comfortability of the sequences at different frame rates. We do note that illuminating some of the sequences (especially 'add. color-plex (A)' and 'tracking(1)') at high frequencies may have the risk of triggering epilepsy. As a result, we only perform the capture after confirming the subject does not have a history of epilepsy.

Figure 8 shows the result of our comfortability study. With slower captures around 3 fps, we find that traditional tracking-frame sequences are preferable, especially with longer intervals between tracking frames. This trend, however, starts to diminish with higher frame rates. At 60 fps, 'additive color-plex (B)' and 'negative color-plex' are preferred over tracking-frame sequences, possibly due to the lack of flashing FULLONs which changes the environmental brightness rapidly. We also find that for additive color-plex sequences, interleaving $OLAT_{RG}$ and $OLAT_B$ is less comfortable than placing the patterns sequentially. This again indicates the constant change and flashing of lights is the main source of discomfortability.

Overall, we find that the negative color-plex sequence provides the best alignment and relighting results while requiring the least amount of photographs. The negative color-plex sequence also provides the most comfortability (in high-speed captures) and does not require rapid switching between drastically different illuminations. We conclude that the negative color-plex sequence is the best choice for efficient and high-quality OLAT captures.

5 CONCLUSION

In this paper, we present two color-multiplexed sequences for OLAT capture. We demonstrate that these sequences provide better alignment between OLAT images, thus giving better rendering qualities compared to traditional tracking-frame sequences. Both additive color-plex and negative color-plex sequences create optimally aligned OLATs which gives better rendering results compared to traditional tracking-frame OLAT sequences. Additionally, the negative color-plex sequence requires the least amount of photographs. While we observe color-shifting in the rendered images for both color-multiplexed sequences, this issue can be resolved to some extent by pre-computing color correction matrices and white balancing the images. We also note that captures with high-end illumination devices may not suffer from this color-shifting. We also studied the comfortability of each capture sequence under various frame rates. Under high frame rates, color-multiplexed sequences can be more favorable compared to tracking-frame sequences which is possibly due to the lack of interleaving FULLON frames. Overall, the negative color-plex design offers the best relighting results with the least amount of required photographs and offers slightly more comfort to the subject being captured. We believe such lighting designs would be preferable over tracking-frame sequences for any OLAT capture sequences under most lighting apparatus.

Future Work

Although we showed that color-multiplexed sequences outperform tracking-frame sequences under 3 fps captures, we have yet to study the results captured under a light stage setting with more OLATs and higher frame rates. A more complete analysis done with light-stage data would be desirable. Specifically, the conclusion reached on the negative color-plex sequence may be different when hundreds of lights are present and the signal-to-noise ratio may be too small to recover OLAT images for the negative OLAT channel.

Another direction of research is to investigate the choice of color channels in which multiplexing is done. We chose to separate the blue channel from the red and green channels since with our camera and illumination setups there is some amount of crosstalk between the red and green channels. However, the blue channel generally has the weakest signal in facial captures, which is not ideal for certain computations. Different setups with different spectral characteristics would benefit from modifying the multiplexed color channels and may achieve better results.

We also acknowledge that our comfortability study is informal and would benefit from a more comprehensive user study in the future.

REFERENCES

Robert Anderson, David Gallup, Jonathan T Barron, Janne Kontkanen, Noah Snavely, Carlos Hernández, Sameer Agarwal, and Steven M Seitz. 2016. Jump: virtual reality video. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 1–13.

Sai Bi, Stephen Lombardi, Shunsuke Saito, Tomas Simon, Shih-En Wei, Kevyn Mcphail, Ravi Ramamoorthi, Yaser Sheikh, and Jason Saragih. 2021. Deep relightable appearance models for animatable faces. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–15.

Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. 2000. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques.* 145–156.

Graham Fyffe. 2009. Cosine lobe based relighting from gradient illumination photographs. In *SIGGRAPH'09: Posters.* 1–1.

Paulo FU Gotardo, Tomas Simon, Yaser Sheikh, and Iain Matthews. 2015. Photogeometric scene flow for high-detail dynamic 3d reconstruction. In *Proceedings of the IEEE international conference on computer vision.* 846–854.

Carlos Hernández, George Vogiatzis, Gabriel J Brostow, Bjorn Stenger, and Roberto Cipolla. 2007. Non-rigid photometric stereo with colored lights. In *2007 IEEE 11th International Conference on Computer Vision.* IEEE, 1–8.

Christos Kampouris, Stefanos Zafeiriou, and Abhijeet Ghosh. 2018. Diffuse-Specular Separation using Binary Spherical Gradient Illumination. *EGSR (EI&I)* 1, 10 (2018).

Arvin Lin, Yiming Lin, Xiaohui Li, and Abhijeet Ghosh. 2024. High Quality Neural Relighting using Practical Zonal Illumination. (2024).

Abhimitra Meka, Christian Haene, Rohit Pandey, Michael Zollhöfer, Sean Fanello, Graham Fyffe, Adarsh Kowdle, Xueming Yu, Jay Busch, Jason Dourgarian, et al. 2019. Deep reflectance fields: high-quality facial reflectance field inference from color gradient illumination. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.

Abhimitra Meka, Rohit Pandey, Christian Haene, Sergio Orts-Escolano, Peter Barnum, Philip David-Son, Daniel Erickson, Yinda Zhang, Jonathan Taylor, Sofien Bouaziz, et al. 2020. Deep relightable textures: volumetric performance capture with neural rendering. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–21.

Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. 2024. Relightable gaussian codec avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.* 130–141.

Frederik David Schreiber, Anna Hilsmann, and Peter Eisert. 2022. Model-based deep portrait relighting. In *Proceedings of the 19th ACM SIGGRAPH European Conference on Visual Media Production.* 1–9.

Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul Debevec, and Ravi Ramamoorthi. 2019. Single image portrait relighting. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.

Cyrus A Wilson, Abhijeet Ghosh, Pieter Peers, Jen-Yuan Chiang, Jay Busch, and Paul Debevec. 2010. Temporal upsampling of performance geometry using photometric alignment. *ACM Transactions on Graphics (TOG)* 29, 2 (2010), 1–11.

Chenglei Wu, Takaaki Shiratori, and Yaser Sheikh. 2018. Deep incremental learning for efficient high-fidelity face tracking. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–12.

Haotian Yang, Mingwu Zheng, Wanquan Feng, Haibin Huang, Yu-Kun Lai, Pengfei Wan, Zhongyuan Wang, and Chongyang Ma. 2023. Towards practical capture of high-fidelity relightable avatars. In *SIGGRAPH Asia 2023 Conference Papers.* 1–11.