Analyzing Camera Motion in Immersive Dome Movies using Optic Flow

Emma Broman* Karljohan Lundin Palmerius † Anders Ynnerman ‡ Alexander Bock§

Linköping University

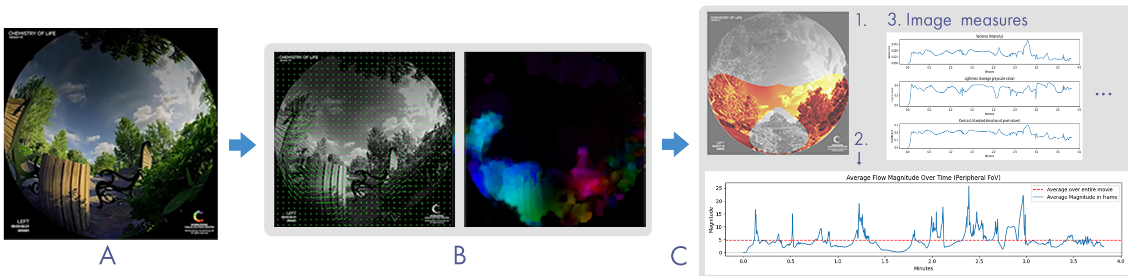

Figure 1: Overview of the optic flow analysis pipeline for dome movies. (A) Input image frame. (B) Compute dense 2D optic flow, here shown in arrow representation (left) and flow direction and magnitude mapped to color hue and brightness (right). The flow vectors are mapped into 3D spherical coordinates. (C) Peripheral flow (1) is then extracted and averaged for each frame, resulting in plots with peaks corresponding to moments of motion (2). Additional frame-wise image metrics (e.g., brightness, contrast, variance) are also calculated per fisheye frame to provide contextual information (3).

Abstract

This work presents early results for an ongoing project on analyzing optic flow to quantify motion effects in movies for immersive digital dome theaters and planetariums. Measuring optic flow and mapping flow peaks to moments of high motion indicates the overall sense of flow in a movie, thus supporting filmmakers in both designing and evaluating dome-specific content. Initial feedback from producers suggests that the measurements have strong practical value and a key next step is to design a visualization tool to facilitate a detailed analysis of the collected data.

1 Introduction

Since their introduction in the 2000s, digital planetarium domes have been used to inspire a sense of awe and wonder in visitors who get to see a fully immersive, and sometimes interactive, dome experience. However, compared to traditional movie making, the authoring tools are still in their infancy, leaving producers to rely on instinct when designing a new experience. This lack of documented guidelines for camera motion is a challenge in creating content for dome theaters. It will likely become more important as digital domes are featured in a growing range of venues and through the democratization of tools for content production.

As with virtual reality (VR) applications, camera movements in dome productions must be carefully designed to minimize audience discomfort caused by visually induced motion sickness (VIMS), also known as cybersickness or VR sickness. At the same time, motion is essential to maintain depth cues and support an engaging immersive experience, sometimes even involving fast-paced camera moves to "wake up the audience" or create excitement. Balancing these goals puts unique constraints on camera design. In VR, VIMS is often mitigated by reducing the sense of vection, for example, by limiting optic flow in the peripheral field of vision using methods such as adaptive field of view restriction, vignetting, blurring, or blocking parts of the periphery [2, 4, 1]. However, these approaches have the disadvantage of reducing immersion and are typically not suitable for domes, where the focus and periphery differ for each audience member. Instead, dome producers rely on techniques like slowing down motion and guiding audience gaze; approaches that depend heavily on instinct and professional experience regarding efficacy and what will risk making the audience uncomfortable. To date, much of the practical knowledge for dome filmmaking resides informally among a small community of experienced producers and is often shared verbally at professional gatherings, rather than being formally published. This leads to a lack of documented guidelines. An exception is a two-part article by Yu et al. [5] in *The Planetarian* in 2016 and 2017, which offers some practical recommendations for dome content creators. However, even in this work, camera motion guidelines are limited to loose advice such as avoiding cuts and slowing down movements.

In this work, we aim to extend this body of knowledge and identify guidelines for camera motion in dome movies by quantifying and analyzing the visual effect of motion. We analyze movies deemed state-of-the-art by domain experts, with the intent of deriving metrics that represent current best practices. The goal is to provide useful insight for the production process and help identify motion-related issues that could induce VIMS. Specifically, we analyze the optic flow from the point of view of the audience and let the average flow quantify the motion for each image frame. This represents a first step toward connecting empirical analysis with production insights, which has already been appreciated in informal interviews with three of the leading dome movie production teams. Our contributions thus far include the initial design of a pipeline to extract optic flow data from full-dome fisheye images, perform automatic magnitude and region-of-interest analysis, and provide the data in an interactive toolkit for dome-movie producers.

2 Optic Flow Analysis

The optic flow analysis is inspired by Bala et al. [1], where peripheral optic flow was used to dynamically restrict the viewer's field of view during 360-video viewing and consequently reduce VIMS. However, rather than limiting flow exposure, we let the resulting flow quantify the amount of peripheral motion.

Our proposed approach for flow estimation is summarized in Figure 1. First, we compute the dense optic flow field for pair-wise image frames in the movie using the Farnebäck method [3], resulting in one (2D) flow vector per pixel in the fisheye image frame.

*e-mail: emma.broman@liu.se
†e-mail: karljohan.lundin.palmerius@liu.se
‡e-mail: anders.ynnerman@liu.se
§e-mail: alexander.bock@liu.se

To better represent the optic flow experienced by the audience, the flow vectors are transformed into spherical coordinates (3D). We then extract the flow in the central and peripheral fields of view, and compute the average magnitude of the flow vectors in these regions over time. This gives a single numerical value per frame that quantifies the optic flow in the respective region, which can be visualized over time as shown in Figure 1 (C, 2.). For simplicity, we assume a fixed viewer focus on the dome surface and front-facing seating when estimating the focus point for the central region. This is a current limitation of the work, since in practice, the focus point varies with content, dome design (such as the tilt of the dome), and producer techniques [5]. Figure 1 (C, 1.) shows a representation of the location of the peripheral field in the fisheye frame. Note that the values for the computed magnitudes depend on the frame rate, frame resolution, and window size used for the computation, so these parameters must match for valid cross-comparison of magnitudes between movies. Since optic flow estimation depends on the homogeneity of the signal [3], we also compute frame-wise image metrics such as brightness, contrast, and local variance to provide contextual information. These metrics help determine whether low flow values reflect actual stillness or are due to insufficient visual features for motion detection. They also give insight into scene characteristics, such as lighting conditions, level of detail, and visual density, which can also be incorporated as part of the work on developing best-practice guidelines.

Finally, we also segment the movies into individual sequences and compute the average flow and image metrics per section. This allows for a higher-level analysis of entire movies as well as a more detailed segment-level comparison. For example, it allows us to visualize segment similarities using dimensionality reduction techniques such as UMAP and PCA. Figure 2 shows some initial similarity results for 5 movies, separated into 53 segments. As dome movies often rely on continuous camera motion and gradual transitions, sometimes with no cuts at all, we manually define the segments at noticeable visual changes, such as transitions in environments, visual key elements, or types of camera motion.

3 Initial Results and Implications

The preliminary results show promising patterns in the optic flow, with peaks aligning with moments that may require producer attention, both in terms of flow magnitude and the frequency of flow peaks. Figure 1 (C) shows an example of the flow magnitude plots for a scene of the dome film Chemistry of Life. In this example, the camera follows a rollercoaster-like trajectory through a park, with a bumblebee in focus for roughly half of the shot. The average flow is about 5, which is a relatively high value in the context of this analysis; all the analyzed movies have an average flow below 3 (see Figure 2), and the peaks tend to lie around a value of about 20[1]. The plot also informs about the visual pacing; in the example, peaks are spaced out over time, aligning with a common design principle of alternating faster-paced motions with slower ones, to “wake up” the audience at the peaks. There is also a period of very low flow at about 1.6 minutes, which correlates to a moment where motion halts due to a simulated stop in time. Soft peaks indicate smooth acceleration, while sharp peaks seemingly indicate fast camera motion or objects moving close to the camera.

In informal interviews, movie producers expressed potential in using this analysis to streamline production by identifying motion-related issues early, and more importantly, before even entering the dome, which is often an expensive resource. However, they emphasized the need for the analysis and visualizations to be delivered through an easy-to-use tool that can be naturally integrated into their existing workflow. This motivates our future work to incorporate the measurements in a visualization toolkit for this purpose.

[1]We used movies downsampled to 512×512 pixels at 30 FPS, computing flow for every 5th frame using an average window size of 15. The low resolution was chosen for reasons of storage and computation speed, but is also motivated by the reduced detail seen in peripheral vision.

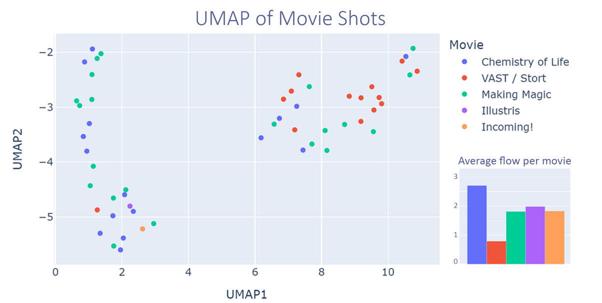

Figure 2: UMAP visualization of movie segments based on image and peripheral flow measurements. The left cluster contains more dynamic shots (higher flow), while the right shows slower-paced ones (lower flow). From this, we see that the VAST movie has more slow-paced segments overall, consistent with its lower average flow (inset). Internal subclusters seem to reflect image features; the far-right group contains dark, low-flow segments, from intros or credits.

While these initial findings are encouraging, several open challenges remain. First, more work is needed to empirically verify how the flow correlates to the audience experience. For example, too frequent and/or high peaks likely increase the risk of VIMS, but this needs to be confirmed. We plan to conduct a pilot study in domes to empirically validate the analysis and perceptual relevance of any proposed metrics. Once verified, we can apply the analysis to a larger pool of movies and begin translating our findings into actionable guidelines for dome filmmakers, which can then be incorporated in the planned visualization tool. The pipeline should also be extended to better estimate likely viewer focus and be adapted to handle hard cuts. This has not been an issue so far, as the chosen movies all use a continuous narrative style with fading transitions instead of hard cuts. However, supporting abrupt scene changes is essential for generalizing the analysis, applying it to more movies, and ultimately developing guidelines and thresholds to automatically detect potentially problematic moments.

Acknowledgments

The authors wish to thank Christian Södergren, Ka Chun Yu, and Ryan Wyatt for interesting discussions about dome movie production and the applicability of the work. Also, thanks to Charles Hansen and Ingrid Hotz for useful discussions related to the method. This work was supported by the Swedish e-Science Research Centre (SeRC) and the Knut and Alice Wallenberg Foundation through Grant KAW 2019.0024.

References

[1] P. Bala, I. Oakley, V. Nisi, and N. J. Nunes. Dynamic Field of View Restriction in 360° Video: Aligning Optical Flow and Visual SLAM to Mitigate VIMS. In *Proc. of the 2021 CHI Conf. on Human Factors in Computing Systems*, pp. 1–18. ACM, Yokohama Japan, May 2021. doi: 10.1145/3411764.3445499 1

[2] N. Biswas, A. Mukherjee, and S. Bhattacharya. “Are you feeling sick?” – A systematic literature review of cybersickness in virtual reality. *ACM Computing Surveys*, 56(11):1–38, Nov. 2024. doi: 10.1145/3670008 1

[3] G. Farnebäck. Two-frame motion estimation based on polynomial expansion. In *Proc. of the 13th Scandinavian Conf. on Image Analysis*, LNCS 2749, pp. 363–370. Gothenburg, Sweden, June 2003. 1, 2

[4] D. Page, R. W. Lindeman, and S. Lukosch. Identifying Strategies to Mitigate Cybersickness in Virtual Reality Induced by Flying with an Interactive Travel Interface. *Multimodal Technologies and Interaction*, 7(5):47, Apr. 2023. doi: 10.3390/mti7050047 1

[5] K. C. Yu, D. Neafus, and R. Wyatt. Filmmaking for Fulldome: Best Practices and Guidelines for Immersive Cinema (Part I & II). *Planetarian*. 45(4), Dec. 2016 and 46(2), Mar. 2017. 1, 2