

# Fast forwarding Egocentric Videos by Listening and Watching

Vinicius S. Furlan<sup>1</sup>, Ruzena Bajcsy<sup>2</sup>, Erickson R. Nascimento<sup>1,2</sup>  
Universidade Federal de Minas Gerais (UFMG), Brazil<sup>1</sup>  
University of California Berkeley, USA<sup>2</sup>

{viniciusfurlan, erickson}@dcc.ufmg.br

bajcsy@eecs.berkeley.edu

## Abstract

*The remarkable technological advance in well-equipped wearable devices is pushing an increasing production of long first-person videos. However, since most of these videos have long and tedious parts, they are forgotten or never seen. Despite a large number of techniques proposed to fast-forward these videos by highlighting relevant moments, most of them are image based only. Most of these techniques disregard other relevant sensors present in the current devices such as high-definition microphones. In this work, we propose a new approach to fast-forward videos using psychoacoustic metrics extracted from the soundtrack. These metrics can be used to estimate the annoyance of a segment allowing our method to emphasize moments of sound pleasantness. The efficiency of our method is demonstrated through qualitative results and quantitative results as far as of speed-up and instability are concerned.*

## 1. Introduction

Thanks to the recent technology advances, a flood of low-cost wearable cameras equipped with high-quality sensors is reaching the consumers. The low-cost of these cameras and the large number of sharing and storing websites are popularizing the use of these well-equipped devices. People are increasingly logging their daily routines, generating massive amounts of egocentric videos rich in visual and sound information. However, since most parts of the egocentric videos are tedious to watch, long egocentric videos are doomed to be forgotten.

Video Summarization and Semantic-aware Hyperlapse are two popular approaches for reducing the size of egocentric videos. Although Video Summarization techniques can find the meaningful moments of a video, they return only disconnected fragments of the whole video [12, 2, 6]. Semantic-aware Hyperlapse works [4, 8, 10, 11, 9], on the other hand, can identify relevant moments while preserving the timeline of the video. Despite the remarkable advances

on video summarization and hyperlapse works, they are still restricted to visual information. Virtually all hyperlapse and summarization methods disregard an important piece of information provided by additional sensors like microphones, the sound. The sound is informative and may provide important clues of the context of a scene and be used to assign importance to segments of a video based on metrics like loudness and annoyance.

In this work, we propose a novel methodology that combines the psychoacoustic metrics with visual features to fast-forward first-person videos. By also considering the sound, we can measure the relevance of a frame using the psychoacoustic annoyance metrics and avoid selecting unpleasure moments like a segment with a crowd or a noisy street. We present quantitative and qualitative results that show the ability of our method in fast-forwarding a video using the sound and images.

## 2. Related Work

We can roughly divide the methods of selecting meaningful moments in long-duration videos into two categories: Video Summarization and Semantic-Aware Hyperlapse.

Video Summarization techniques try to identify the most informative segments of a video and create a compact summary of these moments. There are many ways to classify a segment as informative, e.g., user's preference [12], abnormal behaviors [2], and story-telling [6]. For instance, Lee *et al.* [5] propose a method to summarize a video using both face detection and speech recognition. The authors start by identifying regions in the images containing human faces. Then, they generated coefficients based on the face regions and used these coefficients as features in a classification step. After extracting the audio from the speaker in the video, the face detection and speaker verification results are used to summarize the video.

Although video summarization techniques achieved remarkable results creating summaries, they do not generate a pleasant experience for the user, once they miss the temporal continuity. Semantic-aware Hyperlapse techniques,

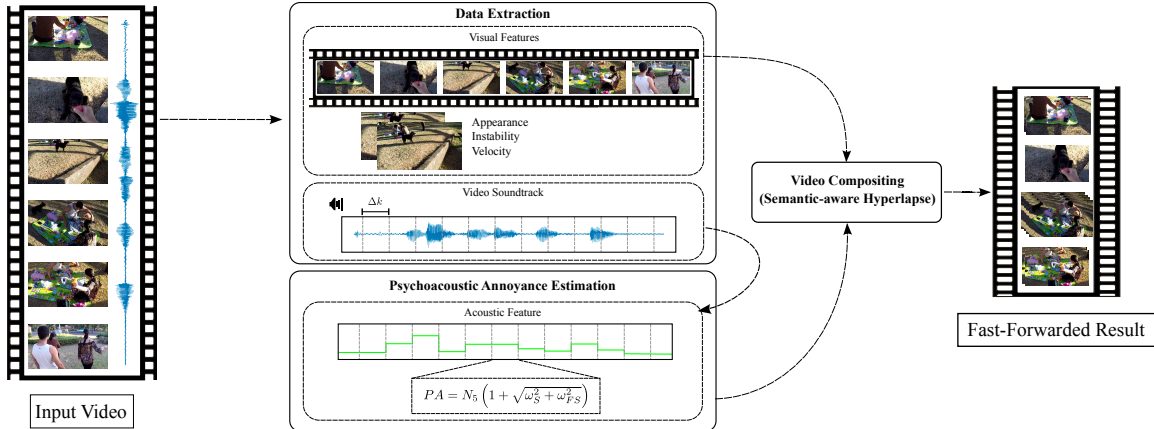


Figure 1. Diagram illustrating the three main steps of methodology. After segmenting the soundtrack of the input video into slices of size  $\Delta k$ , we compute the PA metric (green curve). The PA value is a semantic score assigned to each segment. The semantic score is used to create a relevant profile of the video. This profile is used in the video compositing step to select the relevant frames.

like the works of Silva *et al.* [10, 9, 11], Ramos *et al.* [8] and Lai *et al.* [4], preserve the temporal continuity and the smoothness of a video while emphasizing relevant segments. However, these Semantic-aware Hyperlapse techniques rely only on visual information, they overlook the information provided by other sensors such as the sound.

Several recent studies have been trying to combine sound and sight. A recent and representative approach is the work of Owens *et al.* [7]. The authors propose a self-supervised way to learn a multi-sensory representation that jointly models audio and visual information from a video. The learned representation is then used to predict sound source localization, audio-visual action recognition, and on/off-screen audio separation. Arandjelovic *et al.* [1] present a system that can learn the semantic information of a scene by looking and listening to unlabelled videos.

Similar to the works of Owens *et al.* and Arandjelovic *et al.*, in this work we combine sound and visual data. We present an approach that uses psychoacoustic metrics extracted from the sound of a scene in conjunction with visual information to fast-forward first-person videos.

### 3. Methodology

Our fast-forward method consists of three primary steps, outlined in Figure 1: data extraction, the psychoacoustic annoyance estimation and video compositing.

**Data Extraction.** In this step, we extract the soundtrack from a video of length  $n$  and segment it into slices of lower length  $\Delta k$  resulting in the following set  $\mathcal{S} = \{s_1, s_2, s_3, \dots, s_m\}$ , where  $s_i$  is the  $i$ -th slice of the soundtrack and  $m = \frac{n}{\Delta k}$ .

As stated, our methodology combines acoustical and visual information. Thus, after collecting the sound data,

we compute the instability, appearance, and velocity between frames. The instability metric is estimated by the average distance Focus of Expansion (FOE) to the center of the frame. To estimate appearance, we compute the Earth Mover’s Distance between the color histogram of the frames. At last, the velocity factor is given by the difference between the average magnitude of the optical flows of the frames along with the optical flows of the whole video. All these metrics and the psychoacoustic annoyance metric (computed from the sound) are joined in a linear combination that is used in the video compositing step.

**Psychoacoustic Annoyance estimation.** In the context of first-person videos, auditory stimulus is a relevant factor for perceiving the environment. Zwicker’s metric [13] is a popular method for estimating the sound impression for a human listener. It is based on four psychoacoustic measurements:

- *Fluctuation* and *Roughness*: A complex environment has multiple frequencies sounds that constructively and destructively interfere with each other creating modulation. Fluctuation and roughness are two measurements of the modulation of a signal over the time. The fluctuation was designed to work with 20 up to modulations per second; roughness describes sounds with modulations range from 20 to 300 times per second. A modulated signal is considerably more unpleasant when having a higher roughness and fluctuation. In this work, we used the Roughness metric proposed by Daniel and Weber [3];
- *Loudness* and *Sharpness*: While loudness takes into account the distributions of critical bands in the human hearing, sharpness is a function of the spectral composition. The sharpness metric is estimated by a

weighted sum of specific loudness levels in different bands. A sound with high sharpness is more annoying. The loudness is a psychological phenomenon. Different from a sound level that is a physical measurement, the loudness was developed based on human subject studies in persons with normal hearing. People in these studies listened to a tone at frequency  $f$  Hz and a particular dB level, a second tone was then played at a different frequency. The level of this second tone would be altered until it sounded equally as loud as the  $f$  Hz tone.

Zwicker proposes to compute the psychoacoustic annoyance (PA) as a function of sharpness (S), loudness (N), fluctuation (F), and roughness (R) as:

$$PA = N_5 \left( 1 + \sqrt{\omega_S^2 + \omega_{FS}^2} \right), \quad (1)$$

$$\omega_S = \mathbb{1}[S > 0] \times (1.75 - S) \log(N_5 + 10), \quad (2)$$

$$\omega_{FS} = 2.78 \times N_5^{-0.4} \times (0.4F + 0.6R), \quad (3)$$

where  $N_5$  is the 95th percentile of loudness and  $\mathbb{1}[X]$  is the indicator function, having the value 1 if  $X$  is true and the value 0 otherwise.

After computing the median values of Roughness, Fluctuation, Sharpness and Loudness for each slice  $s_i$ , we estimate the PA value. We used the shape function  $f(PA) = \lambda e^{-\lambda PA}$  to compute the semantic score since the video compositing step relates high values of the score with semantic information. In our case, high PA values represent irrelevant information. The parameter  $\lambda$  controls the decay of the semantic score when the PA increases. In our experiments, we used  $\lambda = 5$ .

**Video compositing.** In this step, our methodology removes frames from the original video using the semantic curve. Firstly, for each frame of the video, we assigned a semantic score concerning the sound segments. Then, according to the scores of each frame, the video is segmented into semantic and non-semantic parts, where a semantic segment has low PA values. After the segmentation, we compute for each segment a different speed-up proportional to the semantic level of the segment. Thus, segments with higher semantic levels receive a lower speed-up whilst respecting the overall required speed-up.

To select which frames to remove, we use the Multi-Importance Fast-Forward (MIFF) method, the state-of-the-art semantic-aware hyperlapse method [9]. The MIFF method is based on a directed graph, where each node represent the frames of the video and edges representing the transition between pair frames. A weight is attributed to each edge of the graph according to metrics as shakiness, visual appearance, the speed of motion and semantics (in our case, the PA of the segment that the frame is part of).

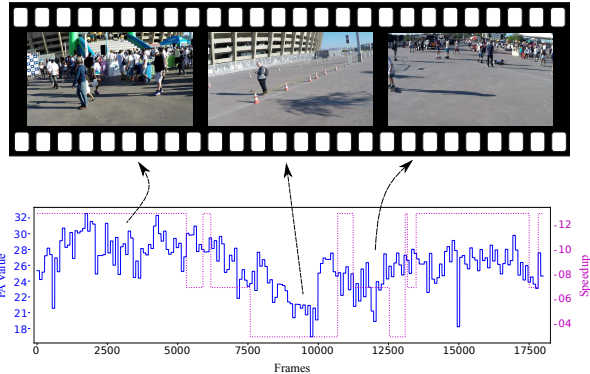


Figure 2. A qualitative analysis of the speed-up rates and PA values. On the left side, a frame from a segment with high PA value. In the middle a frame from a segment with a low PA and on the right a frame from a segment with medium PA value.

Then, the method generates the shortest path of each graph and concatenates these paths. In the last step, the MIFF method stabilizes the concatenate path generating a more visually pleasant fast-forward video.

## 4. Experiments

**Experimental setup.** To evaluate our method we used the Dataset of Multimodal Semantic Egocentric Videos (DoMSEV) proposed by Silva *et al.* [11]. This dataset is composed of 80 hours egocentric videos. It provides videos, sound, IMU measurements, GPS, and depth. Each video has annotations specifying the ambient where the video was recorded (indoor, outdoor), the activity performed, and recorder attention/interaction during the video.

We manually selected 11 segments from videos in the DoMSEV dataset with a rich diversity of sound and visual events (e.g., crowded places and quiet streets and parks). For each segment, we split its soundtrack into segments of 3 seconds. Regarding the semantic hyperlapse creation, all parameters were set with the same values used by the authors in their experiments [9].

**Evaluation Metrics and baseline.** We used two metrics to evaluate the results of the experiments: the speed-up rate and the instability index [10]. The Speed-up rate describes the ratio of frames between the input and the output video. The instability index measures the amount of instability generated by the method in the resultant video.

We pit our method against the MIFF method *et al.* [9], where each frame receives a semantic score according to the number of detected faces and pedestrians on the image.

**Results.** Figure 2 shows the visual results of our method. The blue curve represents the semantic score, and the pink curve shows the speed-up value. We highlighted three

Table 1. Speed-up and instability results for the 11 videos segments used. We used a 10× speed-up factor.

Videos	Speed-up <sup>1</sup>		Instability <sup>2</sup>		Mean PA <sup>2</sup>	
	Ours	MIFF	Ours	MIFF	Ours	MIFF
Academic_Life_03	10.0	10.0	39.4	34.6	20.5	21.5
Attraction_01	10.0	10.0	37.5	37.5	20.9	21.8
Daily_Life_01	10.0	15.8	49.0	50.0	21.4	22.0
Party_01	10.0	10.0	29.6	29.0	28.1	28.3
Recreation_01	10.0	10.0	41.0	37.8	13.8	15.1
Recreation_02	10.0	10.0	45.0	40.1	23.5	24.7
Recreation_04	10.0	10.0	38.9	33.3	25.8	27.4
Recreation_06	10.0	12.3	38.6	42.4	22.0	22.9
Recreation_08	10.0	10.2	27.9	27.8	24.2	24.8
Shopping_01	10.0	10.2	41.5	37.1	22.0	23.3
Tourism_02	10.0	12.4	39.8	40.0	22.3	23.2
<i>Mean</i>	<b>10.0</b>	<i>11.0</i>	<b>38.9</b>	<b>37.2</b>	<b>22.2</b>	<i>23.2</i>

<sup>1</sup>Better closer to 10.      <sup>2</sup>Lower is better.

frames on the figure representing different segments of the video with different PA values. The first frame shows a segment with a high PA that received a high speed-up. We can see several sources of acoustic annoyance in the image as a soundbox near to the recorder and a crowd. The second frame is related to a low value of psychoacoustic annoyance and happens when the recorder goes away from the crowd and the music playing in the background. This segment received the lowest speed-up in the video. In the third part, the record decides to go back to the crowd. Thus, we have a medium value of PA, since loudness and sharpness metrics increase when close to the crowd and music.

Table 1 presents the quantitative results. It contains the Instability Index [10], speed-up, and the average PA achieved by each method in the 11 segments of video used in the experiments. Regarding speed-up, our method achieved the desired factor (10×) in most of the videos. It presented a better result when compared to MIFF method. MIFF performs slightly better than our method regarding instability. One reason may be the large number of speed-up thresholds representing different levels of semantic importance created because of granularity of the semantic information extracted by our method. For this reason, the value of the speed-up increases, generating a higher discrepancy between frames of the fast-forward version of the videos. When considering the average PA of the final video, our method produced fast-forwarded videos with lower mean PA than the MIFF method.

## 5. Conclusion

In this paper, we proposed a new approach to extract semantic from the soundtrack of egocentric videos. This semantic is combined with visual features for fast-forwarding

the input video creating a smaller and stabilized new version. We estimate the annoying sound moments from the video applying psychoacoustic metrics. In our results, we show that even though we lost some stability when compared with state-of-the-art methods, our method can fast-forward a video using the sound information and prevent annoying acoustic segments in the final video.

**Acknowledgments.** The authors would like to thank CAPES, CNPq, FAPEMIG, Petrobras, and NSF VeHICaL project (#1545126) for funding this work.

## References

- [1] R. Arandjelovic and A. Zisserman. Look, listen and learn. In *ICCV*, pages 609–617. IEEE, 2017. 2
- [2] M. Cote, F. Jean, A. B. Albu, and D. Capson. Video summarization for remote invigilation of online exams. In *WACV*, 2016. 1
- [3] P. Daniel and R. Weber. Psychoacoustical roughness: Implementation of an optimized model. *Acta Acustica united with Acustica*, 83(1):113–123, 1997. 2
- [4] W.-S. Lai, Y. Huang, N. Joshi, C. Buehler, M.-H. Yang, and S. B. Kang. Semantic-driven generation of hyperlapse from 360 video. *IEEE Trans. on Visualization and Computer Graphics*, page 99, 2017. 1, 2
- [5] Y. S. Lee, C. Y. Hsu, P. C. Lin, C. Y. Chen, and J. C. Wang. Video summarization based on face recognition and speaker verification. In *ICIEA*, pages 1821–1824, June 2015. 1
- [6] Z. Lu and K. Grauman. Story-driven summarization for egocentric video. In *CVPR*, pages 2714–2721. IEEE, 2013. 1
- [7] A. Owens and A. A. Efros. Audio-Visual Scene Analysis with Self-Supervised Multisensory Features. *ArXiv e-prints*, Apr. 2018. 2
- [8] W. L. S. Ramos, M. M. Silva, M. F. M. Campos, and E. R. Nascimento. Fast-forward video based on semantic extraction. In *ICIP*, pages 3334–3338, 2016. 1, 2
- [9] M. M. Silva, W. L. S. Ramos, F. C. Chamone, J. P. K. Ferreira, M. F. M. Campos, and E. R. Nascimento. Making a long story short: A multi-importance fast-forwarding egocentric videos with the emphasis on relevant objects. *Journal of Visual Comm. and Image Repr.*, 53:55–64, 2018. 1, 2, 3
- [10] M. M. Silva, W. L. S. Ramos, J. P. K. Ferreira, M. F. M. Campos, and E. R. Nascimento. Towards semantic fast-forward and stabilized egocentric videos. In *Int. Workshop on Egocentric Perception, Interaction and Computing at ECCV*, pages 557–571, 2016. 1, 2, 3, 4
- [11] M. M. Silva, W. L. S. Ramos, J. P. K. Ferreira, F. C. Chamone, M. F. M. Campos, and E. R. Nascimento. A weighted sparse sampling and smoothing frame transition approach for semantic fast-forward first-person videos. In *CVPR*, Jun. 2018. to appear. 1, 2, 3
- [12] T. Yao, T. Mei, and Y. Rui. Highlight detection with pairwise deep ranking for first-person video summarization. In *CVPR*, pages 982–990, 2016. 1
- [13] E. Zwicker and H. Fastl. *Psychoacoustics: Facts and models*, volume 22. Springer Science & Business Media, 2013. 2