

SPHERICAL CLUSTERING OF USERS NAVIGATING 360° CONTENT

Silvia Rossi*

Francesca De Simone†

Pascal Frossard‡

Laura Toni*

* Department of Electronic & Electrical Engineering, UCL, London (UK)

† DIS, Centrum Wiskunde & Informatica, The Netherlands

‡ LTS4, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

E-mails: {s.rossi, l.toni}@ucl.ac.uk, F.De.Simone@cw.nl, pascal.frossard@epfl.ch

ABSTRACT

In Virtual Reality (VR) applications, understanding how users explore the visual content is important in order to optimize content creation and distribution, develop user-centric services, or even to detect disorders in medical applications. In this paper, we propose a graph-based method to identify clusters of users who are attending the same portion of spherical content, within one frame or a series of frames. With respect to other clustering methods, the proposed solution takes into account the spherical geometry of the content and correctly identifies clusters that group viewers who actually display the same portion of spherical content. Results, carried out by using a set of publicly available VR user navigation patterns, show that the proposed method identifies more meaningful clusters, i.e., clusters of users who are consistently attending the same portion of spherical content, with respect to other methods.

Index Terms— Virtual Reality, 360° video, user behaviour analysis, data clustering

1. INTRODUCTION

Virtual Reality (VR) systems are expected to become wide spread in a near future, with applications in a variety of fields, ranging from entertainment to e-healthcare. These systems involve omnidirectional (i.e., 360°) videos, which are visual signals defined on a virtual sphere, depicting the 360° surrounding scene. The viewer, virtually positioned at the centre of the sphere, can navigate the scene with three Degrees-Of-Freedoms (3-DOF), i.e., yaw, pitch and roll, by rotating his head and changing his viewing direction. This interactive navigation is typically enabled by a head-mounted display (HMD), which renders at each instant in time only the portion of the spherical content attended by the user, i.e., the *viewport*.

Understanding how users explore the VR content is important in order to optimize content creation [1] and distribution [2–6], develop user-centric services [7,8], and even for medical applications that use VR to study psychiatric disorders [9]. In the last few years, many studies have appeared collecting and analysing the navigation patterns of users watching VR content [6,8,10–16]. Most studies build content-dependent *saliency maps* as main outcome of their analysis, which compute the most probable region of the sphere attended by the viewers, based on their head or eye movements [6,10,17–19]. Some studies also provide additional quantitative analysis based on metrics, such as the average angular velocity, frequency of fixation, and mean exploration angles [8,13]. Models to predict future

saliency maps have also been proposed [20–22]. Nevertheless, none of these studies performs clustering of the navigation patterns, i.e., none provides quantitative data indicating how many groups of users consistently share the same behaviour over time, by attending a significantly overlapping portion of the 360° content. This information can be useful in order to improve the accuracy and robustness of algorithms predicting users navigation patterns. A proper clustering could also be useful to refine user-centric distribution strategies, where for example different groups of users might be served with high quality content in the different portions of the sphere that will be more likely attended by the viewers.

To the best of our knowledge, studies identifying clusters for omnidirectional content delivery have appeared only recently [23,24]. User clustering is employed to identify the number of Region of Interests (RoIs) over time and to perform long-term prediction, associating to each user the future trajectory of the cluster that user belong to. In [23], the viewport center, i.e., the viewing direction of each user at each instant in time, is considered as a point on the equirectangular planar representation of the spherical content. These points on the plane are then clustered based on their Euclidean distance, which unfortunately ignores the actual spherical geometry of the navigation domain. Conversely, in [24] each user navigation pattern is modelled as independent trajectories in roll, pitch, and yaw angles, and spectral clustering is then applied. While it is efficient in discovering general trends of users' navigation, this clustering methodology might fail to identify clusters that are consistent in terms of actual overlap between viewports displayed by different users. It means that users in the same cluster do not necessarily consume the same portion of content. At the same time, this consistency needs to be guaranteed for clustering methods to be used for prediction purposes or for implementing accurate user-based delivery strategies.

The goal of this paper is to propose a novel clustering strategy able to detect meaningful clusters in the spherical domain. We consider as meaningful cluster a set of users *attending the same portion of spherical content* at a given time instant or over a series of frames. This implies that the overlap between the viewports of all users in a cluster must be substantial. With this goal in mind, first we define a metric to quantify the geometric overlap between two viewports on the sphere (Section II). Then, we use this metric to build a graph whose nodes are the centers of the viewports associated to different users. Two nodes are connected only if the two corresponding viewports have a significant overlap (Section III). Finally, we propose a clustering method based on the Bron-Kerbosch (BK) algorithm [25] to identify clusters that are cliques, i.e., sub-graphs of inter-connected nodes (Section III). Results demonstrate the con-

sistency of the proposed clustering method in identifying clusters where the overlap between the portions of the spherical surface corresponding to different viewports is higher than in state-of-the-art clustering (Section IV). In summary, the main contribution of this paper is to propose a clustering algorithm that *i*) considers the spherical geometry of the data, *ii*) identifies clusters in which there is a consistent and significant geometric overlap between the portions of spherical surface corresponding to viewports attended by different users (by imposing that clusters are cliques), *iii*) can be applied to a single frame or to a series of frames. This is a useful new tool to improve the accuracy of user's navigation prediction algorithms and user-dependent VR content delivery strategies, such as those proposed in [23, 24].

2. GEODESIC DISTANCE AS PROXY OF VIEWPORT OVERLAP

Our goal is to identify clusters of users who are displaying the same portion of spherical content within a frame or over a series of consecutive frames. We derive a similarity metric that reliably quantifies how similar the portions attended by two users are. More specifically, each user attends a portion of the spherical surface. This is the projection on the spherical surface of a plane tangent to the sphere (i.e., *viewport*) in the point that identifies the user's viewing direction (*center of the viewport*)¹. The overlap between the viewports attended by two users at an instant in time is a clear indicator of how similar users are with respect to their displayed viewports. For example, an overlap equal to the area of the viewport corresponds to two users attending exactly the same portion of visual content. The geometric overlap could be analytically computed, knowing the rotation associated to each user head's position (i.e., roll, pitch, and yaw) and the horizontal and vertical fields of view that define the viewport. However, this is non trivial. Thus, we propose the simple and straightforward solution of using the *geodesic distance* between two viewport centres as a proxy for the viewport overlap.

By geodesic distance we denote the length of the shortest arc connecting the viewport centers on the sphere. Such distance is clearly an approximation of the actual area overlap: it does not account for the three degrees of freedom of the user's head rotation, which define the exact viewport. As a result, viewports whose centers have the same geodesic distance could correspond to a different viewport overlap (example in Figure 1). Nevertheless, the smaller the distance between viewport centers, the smaller the approximation error with geodesic distance. As an example, Figure 2 shows the pairwise geodesic distance (in blue) and the pairwise area overlap (in red) between the viewport attended by one user and those of 58 other users, for a frame of a video sequence, extracted from the public dataset proposed in [13]. The correlation between the two metrics is evident: if the overlap is high, the geodesic distance between the two viewport centres is low. Particularly, a viewport area overlap larger than 75% of the viewport area corresponds to a geodesic distance smaller than $3\pi/4$. We are therefore interested in identifying a threshold value below which the geodesic distance is a robust proxy of the viewports overlap.

To empirically define this threshold, we built the Receiver Operating Characteristic (ROC) curve as follows. We assume that two users are attending the same portion of content if their viewports overlap by at least O_{th} of the total viewport area. We then define a threshold value for the geodesic distance G_{th} such that users are

¹Without loss of generalization, we consider a scenario in which the viewports of all users have the same horizontal and vertical field of view.




Fig. 1. Viewports (in green and blue) with $\pi/10$ centre distance. (a) viewports are aligned with an overlap of 87%, (b) one viewport is rotated by $\pi/2$ resulting an overlap of 58%.




Fig. 2. Comparison between pairwise geodesic distance and viewport overlap in one frame of video Rollercoaster from [13].

neighbours if their geodesic distance is below threshold. Anytime users are neighbors but their overlap is less than O_{th} , we experience a false positive. Conversely, a true positive is experienced if users that are neighbors also experience an overlap equal or higher than O_{th} . Equipped with these definitions, we can compute the ROC by considering all the videos and user's navigation patterns included in the dataset described in [13]. Figure 3 shows the curve obtained in our scenario with $O_{th} = 80\%$. On the x axis of the ROC curve there is the False Positive Rate (FPR) i.e., probability to have a wrong classification over the number of actual negative events. This rate should be as small as possible. On the contrary, the True Positive Rate (TPR) on the y axis represents the probability to correctly classify an event. The best value of geodesic distance is $\pi/10$ since it corresponds to a TPR value equal to 1, which in our application means a sure identification of viewports with an overlap of at least 80% based on the geodesic distance between their centers. Therefore, in the following we assume $G_{th} = \pi/10$ as a suitable threshold to robustly approximate the area overlap between two viewports by means of the geodesic distance between their centers.

3. CLIQUE-BASED CLUSTERING ALGORITHM

We now describe the proposed clustering algorithm, aimed at identifying clusters of users having a common viewport overlap. We model the evolution of users' viewports over a time-window T , i.e., a series of consecutive frames, as a set of graphs $\{\mathcal{G}_t\}_{t=1}^T$. Each unweighted and undirected graph $\mathcal{G}_t = \{\mathcal{V}, \mathcal{E}_t, \mathcal{W}_t\}$ represents the set of users' viewports at a particular instant t , where \mathcal{V} and \mathcal{E}_t de-

²Without loss of generality, we assume that the set of users does not change over time. This covers also cases in which users' devices are not synchronized in the acquisition time, as users' positions are usually interpolated to create a synchronized dataset.




Fig. 3. ROC curve to evaluate optimal G_{th} considering all video in database [13] and $O_{th} = 80\%$.

note the node and edge sets of \mathcal{G}_t . Each node in \mathcal{V} corresponds to a user interacting with the 360° content. Each edge in \mathcal{E}_t connects neighbouring nodes, where two nodes are neighbours if the geodesic distance between the viewport centers associated to the users represented by the nodes is lower than G_t , as defined in Section II. The binary matrix W_t is the adjacency matrix of \mathcal{G}_t , with $w_t(i, j) = 1$ if the geodesic distance between the two viewport centres of users i and j at time t is below a threshold. More formally:

$$w_t(i, j) = \begin{cases} 0, & \text{if } g(i, j) \leq G_{th} \\ 1, & \text{otherwise} \end{cases} \quad (1)$$

where $g(i, j)$ is the geodesic distance between the viewport centres of users i and j and G_{th} is thresholding value, discussed in Section II. Note that the clique-based clustering algorithm that we present in the following gets in input binary adjacency matrices. Hence, W_t is binary.

Looking at the graphs over time $\{\mathcal{G}_t\}_{t=1}^T$, we are interested in clustering users based on their trajectories within a time window T . Similarly to other clusters of trajectories [26], we derive an affinity matrix A that will be the input to our clustering algorithm, with

$$a(i, j) = \mathcal{I}_D \left(\sum_{t=1}^T w_t(i, j) \right) \quad (2)$$

where $\mathcal{I}_D(x) = 1$ if $x \geq \tau$ and 0 otherwise. This means that in the final graph two nodes, representing two users, are neighbours, i.e., connected by an edge, only if the corresponding viewports have a significant overlap in D instants over T . In the case of $\tau = T$, we obtain $a(i, j) = \mathcal{I}_T \left(\sum_{t=1}^T w_t(i, j) \right) = \prod_{t=1}^T W_t$, and users' viewport centers need to be *always* at a distance below threshold G_t . This condition is however too constraining, therefore we introduce the threshold value τ .

The goal of our clustering algorithm is to identify groups of users that are consistently attending the same portion of the spherical surface. To ensure that all users belonging to a cluster are attending the same area, they all need to be neighbors (i.e., $a(i, j) = 1$ for all pairs of users i and j in the cluster). Therefore, we propose a *clique-based clustering*. In graph theory, a set of nodes all connected to each other is called a *clique*. A clique perfectly matches with the definition of cluster needed in our application, which identifies a set of users all having significant pairwise viewport overlaps, thus attending a common portion of video. We consider the *Bron-Kerbosch (BK) algorithm* [25] to find all *maximal cliques* present in our graph (i.e., the most populated sub-graphs forming cliques). However, maximal cliques identified by the BK algorithm can intersect, i.e., one user can belong to more than one clique. Conversely,




Fig. 4. Graphical example of the proposed clique clustering.

Algorithm 1 Clique-Based Clustering

Input: $\{\mathcal{G}_t\}_{t=1}^T, D$
Output: $K, \mathcal{Q} = [Q_1, \dots, Q_K]$
Init: $i = 1, A^{(1)} = \mathcal{I}_D(\sum_t W_t), \mathcal{Q} = [\{\emptyset\}, \dots, \{\emptyset\}]$
repeat
 $\mathcal{C} = [C_1, \dots, C_L] \leftarrow KB(A^{(i)})$
 $l^* = \arg \max_l |C_l|$
 $Q_i = C_{l^*}$
 $A^{(i+1)} = A^{(i)}(\mathcal{C} \setminus C_{l^*})$
 $i \leftarrow i + 1$
until $A^{(i)}$ is not empty;
 $K = i - 1$

we are interested in identifying disjoint sets³. Hence, our clustering method consists of iterations of BK instances, as depicted in Figure 4. We initialize the clustering method by evaluating the affinity matrix from Eq. (2). Then, we perform the following steps (Algorithm 1):

1. Maximal cliques in the graph are detected by the BK algorithm.
2. Among the resulting cliques, only the most populated one (i.e., the one with the highest cardinality) is kept as a cluster.
3. A new affinity matrix is built, eliminating the entries corresponding to the elements of the cluster identified in Step 2.

These three steps are repeated until all nodes are assigned to clusters. It is worth mentioning that this iterative selection does not guarantee optimal clusters (i.e., clusters with maximal joint overlap among the viewports of users belonging to a cluster). However, *i)* it imposes viewport overlap among users within a cluster, *ii)* it identifies highly populated clusters, which can be translated in reliable trajectories/behaviours shared among users.

4. EXPERIMENTAL RESULTS

The proposed clustering algorithm is compared to state-of-the-art solutions, namely the *Louvain method* [27], the *K-means clustering* [28] and the clustering of VR trajectories proposed in [24] (labelled “SC”). We use the geodesic distance between viewport centers as distance metric in all algorithms. Moreover, in the *K-means clustering*, the number of clusters K is imposed as the value achieved by the Louvain method (labelled “K-means 1”), as well as the K value obtained from our proposed clustering (labelled “K-means 2”). The proposed implementations have been made publicly available⁴. We

³Clusters should be disjoint for most content-delivery applications. For example, if clusters are used for prediction, each user must belong only to one cluster.

⁴<https://github.com/LASP-UCL/spherical-clustering-in-VR-content>.


		ROLLERCOASTER				TIMELAPSE			
		Louvain method	Clique Clustering	K-Means 1	K-Means 2	Louvain	Clique Clustering	K-Means 1	K-Means 2
Fr. - 30s	K	10	15	10	15	13	24	13	24
	Mean Overlap Cl.(% user >3)	38.90 % (84.75 %)	62.50 % (76.30 %)	53.95 % (93.20 %)	48.10 % (94.90 %)	46 % (89.70%)	72.35% (56.90%)	45.90% (96.50 %)	51.50% (50%)
	Main cl. overlap (% users)	26.70% (44.10%)	58.60% (30.50%)	48.30% (19%)	0% (20.70%)	32.90% (20.70%)	69% (12.10%)	15% (19%)	23.50% (13.80%)
	K	8	15	8	15	18	27	18	27
Fr. - 40s	Mean Overlap Cl.(% users >3)	35.60% (89.83%)	65.75% (76.30%)	44.38% (100%)	47.65% (84.75%)	47.65 % (75.90%)	72.95 % (77.60%)	60.27% (96.55%)	65.90% (84.50%)
	Main cl. overlap (% users)	24.20% (45.80%)	58.33% (35.60%)	0% (30.50%)	0% (15.25%)	51.80% (20.70%)	63.70% (17.24%)	47.50% (20.70%)	33.60% (8.60%)
	K	8	12	8	12	18	29	18	29
	Mean Overlap Cl.(% users >3)	48.20% (89.80%)	65.70% (86.45%)	43.50% (98.30%)	55.30% (96.60%)	49.12 % (77.60%)	71.40% (51.70%)	48.36 % (87.90%)	55.90 % (55.17%)
Fr. - 50s	Main cl. overlap (% users)	46.40%(30.50 %)	59.90% (57.70%)	0% (22.40%)	0% (15.25%)	30.60 (22.40%)%	70.80% (25.90%)	37% (24.15%)	62.71% (17.24%)
	K	8	12	8	12	18	29	18	29
	Mean Overlap Cl.(% users >3)	48.20% (89.80%)	65.70% (86.45%)	43.50% (98.30%)	55.30% (96.60%)	49.12 % (77.60%)	71.40% (51.70%)	48.36 % (87.90%)	55.90 % (55.17%)
	Main cl. overlap (% users)	46.40%(30.50 %)	59.90% (57.70%)	0% (22.40%)	0% (15.25%)	30.60 (22.40%)%	70.80% (25.90%)	37% (24.15%)	62.71% (17.24%)

Table 1. Clustering analysis of users in three selected frames from Rollercoaster (first half) and Timelapse (second half). In brackets, the percentage of covered population.


test these algorithms on two 1-minute long video sequences (Rollercoaster and Timelapse), which have been watched by 59 users whose navigation paths are publicly available [13]. Rollercoaster has one main RoI (i.e., the rail) while in Timelapse, there are many fast moving objects (e.g., buildings, people) along the equator line.

Frame-based Clustering. First, we consider frame-based clustering, in which users are identified by their viewport centers at one given frame. Table 1 reports results in terms of number of clusters (K), mean viewport overlap computed within each cluster composed by at least three users, and viewport overlap within the most populated cluster, that we refer to as the main cluster. The viewport overlap within a cluster is the joint overlap across all users’ viewports in the cluster. The mean overlap is computed by averaging the viewport overlap of all clusters with at least three users identified at a given frame. In Table 1, we also provide the percentage of users covered by clusters. The proposed algorithm always ensures the highest viewport overlap (on average always over 50%) with respect to the other methods. This is due to the implicit constraint that is imposed by the clique-based detection of the clusters. This constraint leads to the identification of clusters that are populated and yet meaningful (i.e., with large viewport overlap among users). For example, in Rollercoaster at frame 40s, our algorithm identifies a main cluster grouping 35% of the population with a viewport overlap of 58.33%. This is much higher than the overlap of 24.20% (0%) in the main cluster identified by the Louvain (K-means) method. Beyond the accuracy, another important parameter is the percentage of the population that is covered by clusters with a significant number of users. These clusters are the most useful ones to allow predictions. For instance in Timelapse at frame 50s, our method identifies a large number of clusters (29), which also includes single users clusters. Nevertheless, half of the population (51.70%) belongs to clusters with more than 3 users with high value of joint overlap (71.40%).

Trajectory-based clustering. Second, we test the proposed algorithm over a time-window with $T = 3s$ and $\tau = 1.8s$. In this case we compare the proposed solution with algorithm SC [24]. The algorithm SC is applied to trajectories spanning the entire video, as in [24], as well as consecutive time windows of 3s. We also consider the case of SC in which the number of clusters K is not evaluated from their affinity matrix but it is imposed as the K obtained from our solution. We label this clustering (“SC - K given”). Figure 5 shows results in terms of overlap among viewports clustered together in both Rollercoaster (a) and Timelapse (b). In more details, all users are clustered over consecutive time-windows of T seconds each. Then, for each frame the viewport overlap among all users within one cluster is evaluated and averaged across clusters. The mean overlap (solid line) and the variance (shaded area) is finally depicted in the figure. Moreover, the mean value of joint overlap in clusters with more than three users across the entire video is showed in the legend. Our solution outperforms SC in terms of mean overlap but also in terms of variance. The latter shows the stability of



(a) Rollercoaster video - $T = 3s$.



(b) Timelapse video - $T = 3s$.

Fig. 5. Mean and variance of the joint overlap across clusters over time. In the legend, the mean value of joint viewport overlap of clusters with more than three users performed across the entire video.

our clustering method ensuring for each cluster a consistent overlap over time. Finally, the performance gain is significant also in terms of overlap in the most populated clusters (value provided in the legend).

5. CONCLUSIONS

In this paper, we proposed a novel graph-based clustering strategy able to detect meaningful clusters, i.e., group of users consuming the same portion of a virtual reality spherical content. First, we derived a geodesic distance threshold value to reflect the similarity among users and then we built a clique-based clustering based on this metric. Results on a set of publicly available VR user navigation patterns show that the proposed method identifies more meaningful clusters with respect to other state-of-the-art clustering methods. The associated code has been made publicly available for future comparisons. Future works will focus on the application of our method in the framework of adaptive streaming of VR videos and for the prediction of user navigation patterns.

6. REFERENCES

- [1] A. Serrano, V. Sitzmann, J. Ruiz-Borau, G. Wetzstein, D. Gutierrez, and B. Masia, "Movie editing and cognitive event segmentation in virtual reality video," vol. 36, 2017.
- [2] S. Rossi and L. Toni, "Navigation-Aware Adaptive Streaming Strategies for Omnidirectional Video," in *IEEE 19th International Workshop on Multimedia Signal Processing (MMSP)*, 2017.
- [3] X. Corbillon, G. Simon, A. Devlic, and J. Chakareski, "Viewport-adaptive navigable 360-degree video delivery," in *IEEE International Conference on Communications (ICC)*, 2017.
- [4] S. Petrangeli, V. Swaminathan, M. Hosseini, and F. De Turck, "An HTTP/2-Based Adaptive Streaming Framework for 360 Virtual Reality Videos," in *Proceedings of the 2017 ACM on Multimedia Conference*, 2017.
- [5] C. Fan, J. Lee, C. Lo, W. and Huang, K. Chen, and Cheng-H. H., "Fixation Prediction for 360 Video Streaming in Head-Mounted Virtual Reality," in *Proceedings of the 27th Workshop on Network and Operating Systems Support for Digital Audio and Video*. ACM, 2017.
- [6] M. Yu, H. Lakshman, and B. Girod, "A Framework to Evaluate Omnidirectional Video Coding Schemes," in *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2015.
- [7] M. Broeck, F. Kawsar, and J. F. Schöning, "It's all around you: Exploring 360-degree video viewing experiences on mobile devices," in *Proceedings of the 2017 ACM on Multimedia Conference (MM)*, 2017.
- [8] V. Sitzmann, A. Serrano, A. Pavel, M. Agrawala, D. Gutierrez, B. Masia, and G. Wetzstein, "Saliency in VR: How Do People Explore Virtual Environments?," vol. 24, no. 4, April 2018.
- [9] K. Srivastava, RC Das, and S Chaudhury, "Virtual reality applications in mental health: Challenges and perspectives," *Industrial psychiatry journal*, vol. 23, 2014.
- [10] E. Upenik and T. Ebrahimi, "A simple method to obtain visual attention data in head mounted virtual reality," in *IEEE International Conference on Multimedia Expo Workshops (ICMEW)*, July 2017.
- [11] B. Hu, I. Johnson-Bey, M. Sharma, and E. Niebur, "Head movements during visual exploration of natural images in virtual reality," in *IEEE 51st Annual Conference on Information Sciences and Systems (CISS)*, 2017.
- [12] C. Wu, Z. Tan, Z. Wang, and S. Yang, "A Dataset for Exploring User Behaviors in VR Spherical Video Streaming," in *Proceedings of the 8th ACM on Multimedia Systems Conference (MMSys)*, 2017.
- [13] X. Corbillon, F. De Simone, and G. Simon, "360-degree video head movement dataset," in *Proceedings of the 8th ACM on Multimedia Systems Conference (MMSys)*, 2017.
- [14] C. Lo, W. and Fan, J. Lee, C. Huang, K. Chen, and C. Hsu, "360 Video Viewing Dataset in Head-Mounted Virtual Reality," in *Proceedings of the 8th ACM on Multimedia Systems Conference (MMSys)*, 2017.
- [15] S. Fremerey, A. Singla, K. Meseberg, and A. Raake, "AV-track360: an open dataset and software recording people's head rotations watching 360 videos on an HMD," in *Proceedings of the 9th ACM Multimedia Systems Conference (MMSys)*, 2018.
- [16] A. Singla, S. Fremerey, W. Robitza, P. Lebreton, and A. Raake, "Comparison of Subjective Quality Evaluation for HEVC Encoded Omnidirectional Videos at Different Bit-rates for UHD and FHD Resolution," in *Proceedings of the ACM Conference on Multimedia (MM) Thematic Workshops*, 2017.
- [17] A. Duchowski and G. Marmitt, *Modeling Visual Attention in VR: Measuring the Accuracy of Predicted Scanpaths*, Ph.D. thesis, 2002.
- [18] E. J. David, J. Gutiérrez, A. Coutrot, M. Da Silva, and P. Le Callet, "A Dataset of Head and Eye Movements for 360° Videos," in *Proceedings of the 9th ACM Multimedia Systems Conference (MMSys)*, 2018.
- [19] Y. Rai, P. Le Callet, and P. Guillotel, "Which saliency weighting for omni directional image quality assessment?," in *International Conference on Quality of Multimedia Experience (QoMEX)*, May 2017.
- [20] I. Bogdanova, A. Bur, and H. Hugli, "Visual Attention on the Sphere," *IEEE Transactions on Image Processing*, vol. 17, no. 11, Nov 2008.
- [21] I. Bogdanova, A. Bur, H. Hügli, and P. Farine, "Dynamic visual attention on the sphere," *Computer Vision and Image Understanding*, vol. 114, no. 1, 2010.
- [22] M. Xu, J. Song, Y. and Wang, M. Qiao, L. Huo, and Z. Wang, "Modeling Attention in Panoramic Video: A Deep Reinforcement Learning Approach," *arXiv preprint arXiv:1710.10755*, 2017.
- [23] L. Xie, X. Zhang, and Z. Guo, "CLS: A Cross-user Learning based System for Improving QoE in 360-degree Video Adaptive Streaming," in *ACM Multimedia Conference on Multimedia Conference (MM)*, 2018.
- [24] S. Petrangeli, G. Simon, and V/ Swaminathan, "Trajectory-Based Viewport Prediction for 360-Degree Virtual Reality Videos," in *IEEE conference on Artificial Intelligence and Virtual Reality (AIVR)*, 2018.
- [25] C. Bron and J. Kerbosch, "Algorithm 457: finding all cliques of an undirected graph," *Communications of the ACM*, vol. 16, no. 9, 1973.
- [26] S. Atev, G. Miller, and N. P. Papanikolopoulos, "Clustering of vehicle trajectories," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, 2010.
- [27] V. D. Blondel, J. Guillaume, and E. Lambiotte, R. and Lefebvre, "Fast unfolding of communities in large networks," *Journal of statistical mechanics: theory and experiment*, , no. 10, 2008.
- [28] J. Hartigan and M. Wong, "Algorithm as 136: A k-means clustering algorithm," *Journal of the Royal Statistical Society*, vol. 28, 1979.