Visualization of User's Behavior in Indoor Virtual Environments through Interactive Heatmaps

Luca Chittaro and Marta Serafini

HCI Lab, Department of Mathematics, Computer Science and Physics,
University of Udine, via delle Scienze 206, 33100 Udine, Italy
serafini.marta@spes.uniud.it

Abstract. Three-dimensional virtual environments (VEs), such as those used in video games and virtual reality experiences, pose new challenges to the study of user's behavior. This paper proposes a system based on the combination of two heatmaps for the analysis of user's movement in the VE and of the areas looked at by him/her. It also describes a pilot study aimed at assessing the efficacy of the system. Results of the study indicate that the system can effectively support analysts in identifying user's look-at behaviors as well as navigation strategies, patterns, and coverage of specific areas during movement.

Keywords: virtual environment, information visualization, heatmap, user behavior, visual analytics.

1 Introduction

Three-dimensional virtual environments (VEs), such as those used in video games and virtual reality experiences, pose new challenges to the study of user's behavior. Recent years saw a growth in interest in tools that support the analysis of user behavior in VEs by visualizing data collected during usage [1, 6, 8-10, 12-15, 18, 19, 21-23, 25, 26]. For example, visualizations can highlight the most frequented areas of the VE and the paths followed by users, helping analysts understand users' navigation behavior and how different design choices can affect it. In the literature, research on tools for the analysis of user's behavior in VEs has typically focused on two-dimensional representations such as maps [1-3, 6, 10-12, 14, 18, 19, 24-26], while solutions that visualize behavioral data directly in the three-dimensional VE are still rare, motivating further research.

This paper focuses on an innovative combination of existing ideas to propose a system that, given a VE that represents an indoor, real or imaginary, environment (indoor VE), provides a visualization of user's movement together with a visualization of how much different parts of the VE have been looked at by the user. We also evaluate the effectiveness of the proposed system by conducting a pilot study with a group of experts in the field of virtual reality and games.

The paper is organized as follows: we first discuss the related work in Section 2, then Section 3 illustrates the proposed system, while Section 4 presents the pilot study and its results. Finally, Section 5 discusses results and outlines future work.

2 Related work

In the literature, tools for visual analysis of user's movement typically employ two alternative types of techniques. The first type visualizes the path followed by each user on a 2D map of the VE [2, 3, 6, 12, 14, 18, 19, 26] to support detailed analysis. The second type is the 2D heatmap, which highlights the most frequented areas of the VE [4, 6, 24]. Most studies of spatiotemporal user's behavior have focused on 2D top-down representations of user's movement, but some of them have started to allow the analyst to examine the visualization from different viewpoints [8, 9, 13, 15, 16, 21, 23]. Bidimensional heatmaps are widely used in the literature also to visualize the frequency of specific events (e.g., user's death, gunshot) [10, 12, 26].

However, when the analyst needs to study where users looked at in the VE, this kind of heatmaps is not sufficient because the information on the third dimension is lost. For this reason, heatmaps superimposed on the entire VE, called 3D heatmaps, have been proposed as an alternative to the traditional 2D heatmap. For example, Stellmach et al. [22] proposed two different techniques for visualizing looked-at objects of the VE with an eye tracker. The first technique (object-based attentional map) considers the entire surface of each object in the VE and assigns to it a color that indicates how much the user looked at the object. Since this technique considers the object as a whole, it provides an overview of the looked-at objects but does not give detailed information on which areas of the object have been looked at. The second technique (surface-based attentional map, also called triangle-based attentional map) visualizes observational data directly on the surface of each object in the VE using vertex-based mapping (i.e., it first colors the vertices based on looked-at data and then the triangles of the 3D model are colored by interpolation). This heatmap should allow the analyst to see which areas of the object have been looked at but is affected by two issues. First, the accuracy of the visualization depends on the resolution of the triangulation (i.e., objects with a small number of vertices will not produce accurate attentional maps). Second, when triangles are colored, if the object is relatively small, the coloring may extend to parts of the object that were not visible from the user's viewpoint. It must also be noted that an object is considered to be looked-at when the eye-tracker detects a fixation on that object, and both visualizations ignore the other objects in close proximity, possibly leading to a lack of coloring of nearby objects in the user's field of view. Maurus et al. [17] proposed a 3D heatmap that considers occlusions of the user's view by objects, and, as Stellmach et al., uses an eye tracker to collect user's precise gaze data. Unlike [22], they perform the computation of the looked-at areas at the pixel level instead of the vertex level. Pfeiffer et al. [20] presented a 3D heatmap based on special textures, called attention textures, applied to the surfaces of objects and created by mixing the colors of the heatmap with those of the original object so that its original texture remains distinguishable. Kraus et al. [16] proposed a visualization technique of looked-at areas that wraps all the surfaces to be monitored in the VE with a uniform grid. Each vertex of the grid increases a counter every time a user looks at it. In this way, the visualization allows to clearly distinguish among looked-at, not looked-at and occluded areas by coloring the grid in green, red, and blue respectively according to a threshold value.

To summarize, current techniques show that heatmaps are able to properly represent both user's movement and looked-at areas, but their combinations are still rare in 3D VEs. An exception is the recent work by Kepplinger et al. [15] who proposed a combination of visualizations that shows in the 3D VE the user's path and the objects (s)he looked at. The former is represented by connected segments while the latter is represented by a gradient heatmap on the contours of the looked-at objects of interest.

In this work, we aim at integrating a 3D heatmap of looked-at areas with a heatmap of movement, making them viewable from any angle inside the VE. The combination of these two visualizations can help the analyst to better understand user behavior, e.g. by highlighting what is the object the user looked at from each of his/her positions.

3 The proposed system

The system we propose aims at helping in the analysis of single user's behavior in indoor VEs by visualizing usage data directly within the VE and supporting exploration of the visualization from arbitrary view angles. As shown in Fig. 1, the system integrates two heatmaps into an interactive interface, using two different gradients. To visualize user's movement, we adopted the 2D heatmap proposed by Kraus et al. [16], with a black-to-yellow gradient. To visualize looked-at areas, we combined instead the surface-based attentional map [22] and the 3D heatmap by Maurus et al. [17], with a black-to-red gradient. The resulting heatmap operates on pixels instead of vertices of 3D objects, thus becoming independent from the number of vertices used to model the object. Moreover, it is able to exclude occluded objects [17], with the possible exception of some parts of the objects that are not visible from the user's viewpoint, as in surface-based attentional maps [22].

To determine looked-at objects and areas in the VE, the system uses a raycast that points to the center of the scene viewed by the user and colors a circular area around that point on the surface of the object. Furthermore, unlike Stellmach et al., it does not consider as looked-at only the object in the center of the field of view, but spreads the attention to the objects in its close proximity under a distance threshold, obtaining a more accurate result. Furthermore, the system works without the need for an eye tracker, supporting collection of data from large numbers of online users who do not have special hardware.

The interactive interface of the system adds the possibility for the analyst to select objects of interest: they will be circled in the VE to highlight their location and facilitate visual retrieval. The analyst can turn on or off each visualization as well as object-circling. In addition, (s)he can apply a temporal filter to visualize only data in a time interval. For example, in Fig. 2 time is restricted to the first 30 seconds of data, allowing the analyst to discover the following behaviors. From the movement visualization, the analyst can conclude that the user walked only along the upper corridor and moved continuously instead of using teleportation. From the visualization of looked-at objects, the analyst can conclude that the user: (i) passed close to the "Luggage" object but did not look at it because (s)he was looking towards the middle row of seats, and (ii) saw the “Ball” object even though it was on the corridor (s)he did not travel.

Both visualizations use gradients that scale on brightness because color maps such as the classic rainbow map: (i) are not perceptually ordered, (ii) are not suitable for presenting small details because the human visual system cannot perceive small hue changes which can thus obscure information, and (iii) can introduce artifacts into the visualization. The last problem arises from the fact that sharp transition between the different hues of the rainbow color map can be perceived by the analyst as a sharp difference in the visualized data, leading to misinterpretations [5].

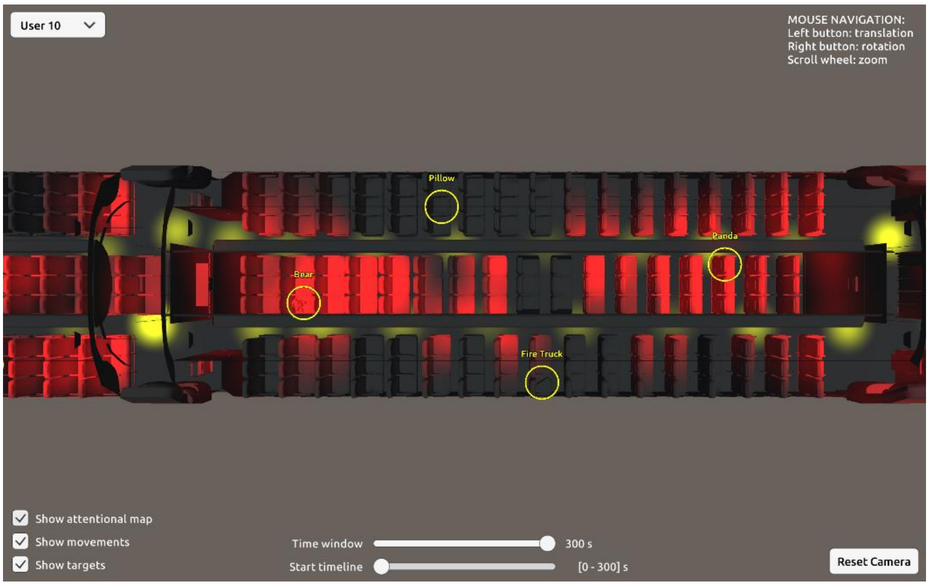

Fig. 1. Top-down view with both visualizations activated. The black-to-yellow gradient on the floor allows to quickly notice that this user’s navigation strategy is based on teleportation. Visualization of looked-at areas uses instead a black-to-red gradient: black coloring of the “Pillow” and “Fire truck” objects indicates they have not been looked at by the user. For the interpretation of color in this figure, the reader is referred to the web version of this article.

Our proposed system shows user's behavior directly in the VE, representing user's movement and the observed areas with a 2D and 3D heatmap, respectively. In this way, unlike Keppliger at al. [15], it is possible to recognize the navigation strategy used (e.g., to distinguish continuous motion from teleportation), as well as identify which parts of the objects' surfaces were looked at, by highlighting them directly on the 3D objects, and examining the objects from different angles. For example, in Fig. 2., the arbitrary viewpoint allows one to see that even though the user did not look at the top row of seats, (s)he still looked at the sides of seats facing the corridor, including their armrest.

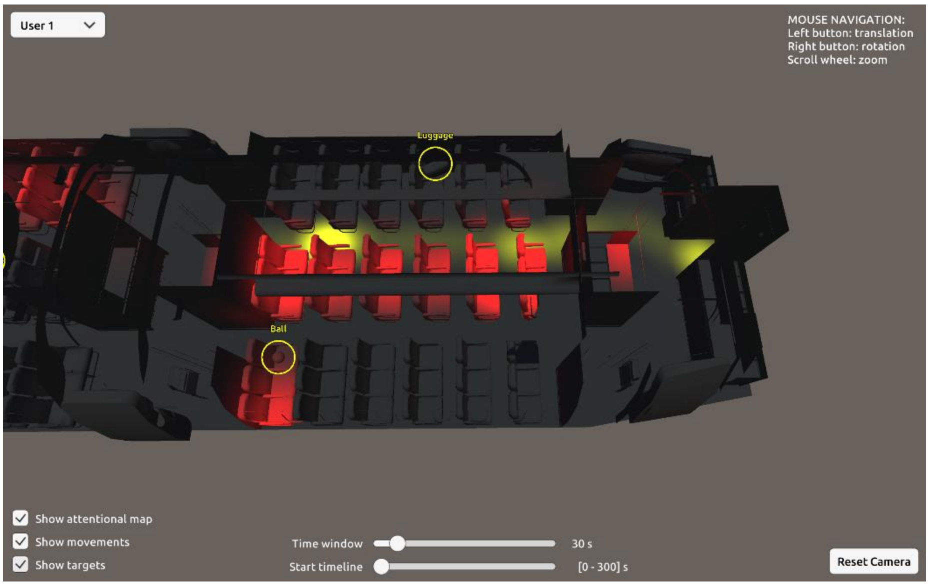

Fig. 2. The system with arbitrary view angle active, both visualizations active, and a time interval set to 30 seconds since the beginning of the session. The black-to-yellow gradient on the floor highlights that the user moved continuously instead of using teleportation. The black-to-red gradient highlights that the user looked at the "Ball" object but not at the "Luggage" object. For the interpretation of color in this figure, the reader is referred to the web version of this article.

4 Pilot study

4.1 Objectives, dataset and participants

A pilot study was conducted with a group of experts in VEs to evaluate the effectiveness of the interactive visualization system. We used a dataset of real users' behavioral data that was collected in an experiment whose purpose was to test users' ability to remember the position of target objects placed on some seats of an indoor VE that could be navigated with a continuous or teleportation strategy. The indoor VE was a reproduction of the cabin of a Boeing 777 twin-aisle aircraft, and users carried out a visual inspection that flight attendants have to perform with no passengers on board. We chose this dataset because it is representative of the typical datasets concerning users moving in a VE and at the same time puts central aspects of our system under test. Indeed, it allows to evaluate the effectiveness of: (i) the visualization of movement, which should allow analysts to recognize the general and detailed patterns of movement followed by each user, and (ii) the visualization of looked-at areas, which should allow analysts to recognize which target objects were looked at or not by each user.

The pilot study involved four experts in the field of VEs. They were members of the HCI Laboratory at the University of Udine, and none of them was involved in the design of the system or had previously used it. They were asked to carry out four different analysis tasks on the data of each of 10 users taken from the dataset. We chose the 10 specific cases from a total of 127 in a way that covered different possibilities such as: users who relied on teleportation vs. users who did not, users who visited the entire VE vs. users who left some areas unexplored, users who looked at all target objects vs. users who looked only at some target objects.

4.2 Procedure

Since the study was carried out during the Covid-19 lockdown, it was conducted remotely. Each expert was sent a link to download the system and a document that described the visualizations and the interface. They used the system on their own computer and completed a questionnaire on Google Forms. This actually contributed to make the context more naturalistic, since analysts are going to use this kind of tools on their computers, in the places where they work.

The effectiveness of the system in supporting the experts was measured by asking them to perform four different analysis tasks on each of the 10 users, resulting in 40 analyses performed by each expert. The questionnaire was organized in 10 identical sections, one for each of the 10 cases. Each section was organized into 4 subsections, one for each task. The aim of task T1 was to recognize the general user's navigation pattern, and experts were asked to identify whether the user moved in the VE using a continuous movement strategy or teleportation. In T2, the goal was to detect whether the user followed a specific path: experts were asked to report whether or not the user had followed a path that was suggested to him/her, i.e. traveling from the starting point (identical for all users) to the opposite side of the VE through only one aisle, and then returning to the starting point through only the other aisle. T3 focused on the ability to detect whether the user had explored a specific wide area of the VE, and experts were asked to report whether or not the user had completed at least one full tour of the VE. Experts were informed that a tour of the VE was to be considered complete if the user had visited the entire VE without leaving areas unexplored, thus fully using both corridors in the VE. T4 focused on the ability to identify which objects were looked at and which were not by the user: for each of the 8 target objects in the VE, experts were asked to indicate whether or not it had been looked at by the user.

4.3 Results

The results obtained with the proposed visualization system were compared with those that would be obtained if the contribution of the visualization was not better than chance. For each of the four tasks, the proportion of correctly identified user behaviors was thus subjected to the appropriate non-parametric test (binomial test), after checking its assumptions were met, following Cohen [7]. Since the outcome of each answer obtained from experts was dichotomous, the expected proportion of correct answers if the effects of the visualization were not different from chance is 50%. For each task, the obtained proportions were instead ideal (100%) or close to ideal (95%), and the binomial test indicated that they were significantly higher than those expected. The obtained results for each task are illustrated respectively by Tables 1, 2, 3 and 4, for each expert as well as for all experts combined. For all 10 users, all experts correctly recognized the type of navigation (T1) and identified whether users had followed the suggested path (T2). Experts were able to recognize if users had covered a specific area (T3) for a number of users that was 9 or 10 (M=9.5, SD=0.58). Finally, in T4 they were able to correctly identify which target objects were looked at or not for a number of objects that ranged from 72 to 80 (M=76.0, SD=4.08).

Table 1. Results for task T1

Expert	Right answer	Wrong answer	Significance test
E1	10	0	p<0.01
E2	10	0	p<0.01
E3	10	0	p<0.01
E4	10	0	p<0.01
All	40	0	p<0.001

Table 2. Results for task T2

Expert	Right answer	Wrong answer	Significance test
E1	10	0	p<0.01
E2	10	0	p<0.01
E3	10	0	p<0.01
E4	10	0	p<0.01
All	40	0	p<0.001

Table 3. Results for task T3

Expert	Right answer	Wrong answer	Significance test
E1	10	0	p<0.01
E2	10	0	p<0.01
E3	9	1	p<0.05
E4	9	1	p<0.05
All	38	2	p<0.001

Table 4. Results for task T4

Expert	Right answer	Wrong answer	Significance test
E1	72	8	p<0.001
E2	80	0	p<0.001
E3	73	7	p<0.001
E4	79	1	p<0.001
All	304	16	p<0.001

5 Discussion and Conclusions

This paper proposed a visualization system to support the analysis of user's behavior in indoor VEs. The results obtained in the pilot study provide an indication that the system can effectively support analysts in detecting navigation strategies, patterns, and coverage of specific areas in terms of movement as well as look-at behaviors. By discussing with experts about the errors made with the movement visualization (a total of 2 errors, both in task T3), it came out that one expert had considered a tour complete also when the user reached the opposite side of the VE and then returned to the starting point, even if the user did not walk through all the in-between areas. The discussion with experts also allowed us to identify the four reasons behind the relatively few errors made in identifying which target objects were looked at or not by users. First, in 4 of the 16 errors, the expert concluded that the target object was looked at by the user even if the visualization had correctly painted the object in black, e.g., the target "Fire Truck" in Fig.1. This happened because the expert reasoned that, since nearby objects had been looked at (e.g., some seats near "Fire Truck" in Fig. 1), then the user might have looked at the target too. Second, 2 errors were due to a slip by the expert in selecting the answer in the questionnaire. Third, for 2 errors, the reason was particularly interesting because it highlighted a difficulty in interpretation that may arise in a specific case, i.e. when the visualization paints a small object in dark red because it has been briefly looked at by the user, but the small object is placed on a large object which has mostly not been looked at and is thus painted in black. This highlights the need for further refinement of 3D heatmaps of the areas looked at by the user. For example, it should be evaluated whether changing the size of the colored area for each looked-at point on objects could make it easier to clearly recognize that the target object has been seen in the above described case. The remaining 8 errors revealed that some experts concluded that the user had not looked at the target object even if it was painted in dark red. This coloring reflected the fact the user had looked at the target briefly compared to the total duration of the session. However, the expert reasoned that, since the object was painted in very dark red, it was unlikely that the user had looked at it.

Our research will now continue along two main lines, also considering the above mentioned issues. First, we will introduce new features into the system to increase the level of support offered by the interface to the analyst. For example, the system could semi-automatically propose an identification of the objects that were looked at based on thresholds defined by the analyst. This would also help preventing the last source of error described above. Second, we will assess the effectiveness of the system in supporting the detection of VE design problems, e.g. finding the VE areas where the user got stuck, due to the used navigation techniques or the architectural modeling of the VE. In addition, we will compare our system with visualization tools that use 2D representations, such as those presented in Section 2, to highlight the strengths and weaknesses of both solutions. A larger sample of participants will also be considered.

References

1. Ahmad, S., Bryant, A., Kleinman, E., Teng, Z., Nguyen, T. H. D., Seif El-Nasr, M.: Modeling Individual and Team Behavior through Spatio-temporal Analysis. In Proceedings of the Annual Symposium on Computer-Human Interaction in Play, 601–612 (2019), https://doi.org/10.1145/3311350.3347188
2. Almeida, S., Mealha, O., Veloso, A.: Video game scenery analysis with eye tracking. Entertainment Computing 14, 1–13 (2016), https://doi.org/10.1016/j.entcom.2015.12.001
3. Aung, M., Demediuk, S., Sun, Y., Tu, Y., Ang, Y., Nekkanti, S., Raghav, S., Klabjan, D., Sifa, R., Drachen, A.: The trails of Just Cause 2: spatio-temporal player profiling in open-world games. In Proceedings of the 14th International Conference on the Foundations of Digital Games, 1–11 (2019), https://doi.org/10.1145/3337722.3337765
4. Barros, V. P., Notargiacomo, P.: Big data analytics in cloud gaming: Players' patterns recognition using artificial neural networks. In 2016 IEEE International Conference on Big Data (Big Data), 1680–1689 (2016). doi: 10.1109/BigData.2016.7840782
5. Borland, D., Taylor II, R.M.: Rainbow color map (still) considered harmful. IEEE Computer Architecture Letters 27(02), 14–17 (2007)
6. Chittaro, L., Ranon, R., Ieronutti, L.: Vu-flow: A visualization tool for analyzing navigation in virtual environments. IEEE Transactions on Visualization and Computer Graphics 12(6), pp. 1475–1485 (2006). doi: 10.1109/TVCG.2006.109
7. Cohen, B. H. (2013). Explaining Psychological Statistics, 4th Edition, John Wiley & Sons.
8. Dixit, P.N., Youngblood, G.M.: Understanding information observation in interactive 3d environments. In: Proceedings of the 2008 ACM SIGGRAPH symposium on Video games. pp. 163–170 (2008), https://doi.org/10.1145/1401843.1401874
9. Dixit, P. N., & Youngblood, G. M.: Understanding playtest data through visual data mining in interactive 3d environments. In 12th international conference on computer games: AI, animation, mobile, interactive multimedia and serious games (CGAMES), 34–42 (2008)
10. Drachen, A., Canossa, A.: Analyzing spatial user behavior in computer games using geographic information systems. In: Proceedings of the 13th international MindTrek conference: Everyday life in the ubiquitous era. pp. 182–189 (2009), https://doi.org/10.1145/1621841.1621875
11. Drachen, A., Canossa, A.: Towards gameplay analysis via gameplay metrics. In: Proceedings of the 13th international MindTrek conference: Everyday life in the ubiquitous era. pp. 202–209 (2009), https://doi.org/10.1145/1621841.1621878
12. Drachen, A., Canossa, A.: Evaluating motion: Spatial user behaviour in virtual environments. International Journal of Arts and Technology 4(3), 294–314 (2011), https://doi.org/10.1504/IJART.2011.041483
13. Drenikow, B., Mirza-Babaei, P.: Vixen: interactive visualization of gameplay experiences. In: Proceedings of the 12th International Conference on the Foundations of Digital Games. pp. 1–10 (2017), https://doi.org/10.1145/3102071.3102089
14. Hoobler, N., Humphreys, G., Agrawala, M.: Visualizing competitive behaviors in multi-user virtual environments. In: IEEE Visualization 2004. pp. 163–170. IEEE (2004). doi: 10.1109/VISUAL.2004.120
15. Kepplinger, D., Wallner, G., Kriglstein, S., Lankes, M.: See, Feel, Move: Player Behaviour Analysis through Combined Visualization of Gaze, Emotions, and Movement. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, 1–14 (2020), https://doi.org/10.1145/3313831.3376401

16. Kraus, M., Kilian, T., Fuchs, J.: Real-time gaze mapping in virtual environments. In: EUROVIS 2019: 21st EG/VGTC Conference on Visualization (2019), https://dx.doi.org/10.2312/eurp.20191135
17. Maurus, M., Hammer, J.H., Beyerer, J.: Realistic heatmap visualization for inter- active analysis of 3d gaze data. In: Proceedings of the Symposium on Eye Tracking Research and Applications. pp. 295–298 (2014), https://doi.org/10.1145/2578153.2578204
18. Moura, D., el Nasr, M.S., Shaw, C.D.: Visualizing and understanding players' behavior in video games: discovering patterns and supporting aggregation and comparison. In: ACM SIGGRAPH 2011 game papers, pp. 1–6 (2011), https://doi.org/10.1145/2037692.2037695
19. Mueller, S., Solenthaler, B., Kapadia, M., Frey, S., Klingler, S., Mann, R. P., Sumner, R. W., Gross, M.: HeapCraft: Interactive data exploration and visualization tools for understanding and influencing player behavior in Minecraft. In Proceedings of the 8th ACM SIGGRAPH Conference on Motion in Games, 237–241 (2015), https://doi.org/10.1145/2822013.2822033
20. Pfeiffer, T., Memili, C.: Model-based real-time visualization of realistic three-dimensional heat maps for mobile eye tracking and eye tracking in virtual reality. In: Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications. pp. 95–102 (2016), https://doi.org/10.1145/2857491.2857541
21. Schertler, R., Kriglstein, S., Wallner, G.: User guided movement analysis in games using semantic trajectories. In: Proceedings of the Annual Symposium on Computer-Human Interaction in Play. pp. 613–623 (2019), https://doi.org/10.1145/3311350.3347156
22. Stellmach, S., Nacke, L., Dachselt, R.: 3d attentional maps: aggregated gaze visualizations in three-dimensional virtual environments. In: Proceedings of the international conference on advanced visual interfaces. pp. 345–348 (2010), https://doi.org/10.1145/1842993.1843058
23. Stellmach, S., Nacke, L., Dachselt, R.: Advanced gaze visualizations for three-dimensional virtual environments. In Proceedings of the 2010 symposium on eye-tracking research & Applications, 109–112 (2010), https://doi.org/10.1145/1743666.1743693
24. Tremblay, J., Torres, P., Rikovitch, N., Verbrugge, C.: An exploration tool for predicting stealthy behaviour. In: Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment. vol. 9 (2013)
25. Wallner, G., Kriglstein, S.: Multivariate Visualization of Game Metrics: An Evaluation of Hexbin Maps. In Proceedings of the Annual Symposium on Computer-Human Interaction in Play, 572–584 (2020), https://doi.org/10.1145/3410404.3414233
26. Wallner, G., Kriglstein, S.: Plato: A visual analytics system for gameplay data. Computers & Graphics 38, 341–356 (2014), https://doi.org/10.1016/j.cag.2013.11.010