



# Immersive analysis of user motion in VR applications

Simon Kloiber<sup>1</sup> · Volker Settgast<sup>2</sup> · Christoph Schinko<sup>2</sup> · Martin Weinzerl<sup>3</sup> · Johannes Fritz<sup>3</sup> · Tobias Schreck<sup>1</sup> · Reinhold Preiner<sup>1</sup>

Published online: 12 August 2020  
© The Author(s) 2020

## Abstract

With the rise of virtual reality experiences for applications in entertainment, industry, science and medicine, the evaluation of human motion in immersive environments is becoming more important. By analysing the motion of virtual reality users, design choices and training progress in the virtual environment can be understood and improved. Since the motion is captured in a virtual environment, performing the analysis in the same environment provides a valuable context and guidance for the analysis. We have created a visual analysis system that is designed for immersive visualisation and exploration of human motion data. By combining suitable data mining algorithms with immersive visualisation techniques, we facilitate the reasoning and understanding of the underlying motion. We apply and evaluate this novel approach on a relevant VR application domain to identify and interpret motion patterns in a meaningful way.

**Keywords** Virtual reality · Immersive analytics · Movement analysis · Trajectory visualisation · Virtual training evaluation

## 1 Introduction

As consumer-level virtual reality (VR) devices become more powerful, affordable and widely available, VR-based experiences find ever new applications in entertainment, gaming and industry. Nowadays, VR devices provide an integrated high-end tracking of head and bimanual motion off the shelf, without resorting to any additional motion capture gear. These three basic tracking points alone already provide rich information about a user's motion and actions performed in a VR experience. In fact, the incidental motion data tracked during a VR session is of strong interest for a multitude of VR-based motion analysis tasks, which are relevant for various private, commercial, but also medical and scientific use cases. Examples include the performance analysis and optimisation in task-oriented VR scenarios, but

also the evaluation and assessment in virtual training sessions. The specific key quality of these kinds of use cases is the strong link between rich and complex motion data and the context of the virtual environment it was recorded in. This poses a new challenge for the analysis of human motion recorded in virtual reality.

In this paper, we present an *immersive* approach to analyse such VR motion data. Our key intuition is that such motion data should be analysed in the same perceptive context in which it was recorded, that is, in the same immersive environment. This co-location of data acquisition and analysis allows for several applications, ranging from accelerated improvement cycles in VR training to real-time collaborative remote analysis for VR-based task or game design. Our concept is a first step towards co-located trajectory-based visual analysis of user behaviour in VR environments.

A typical approach to provide a still overview of motion data for analysis is to visualise the space-time trajectories of tracked reference points. However, human motion data tracked in VR sessions comprises multiple semantically linked trajectories (head and hands) per person, which is particularly challenging for comparative side-by-side analysis of multiple user sessions. Moreover, the mapping of a user's real-world motion to its spatial embedding in a virtual world is characteristically not restricted to usual continuity constraints. Typical established VR motion assisting paradigms


**Electronic supplementary material** The online version of this article (<https://doi.org/10.1007/s00371-020-01942-1>) contains supplementary material, which is available to authorized users.

✉ Simon Kloiber  
s.kloiber@cgv.tugraz.at

<sup>1</sup> Institute of Computer Graphics and Knowledge Visualisation, Graz University of Technology, Graz, Austria

<sup>2</sup> Fraunhofer Austria Research GmbH, Graz, Austria

<sup>3</sup> AVL List GmbH, Graz, Austria



**Fig. 1** An immersive visualisation of human motion data that were recorded from multiple VR sessions. Human motion paths are visualised as trajectories, while specific temporal points of interest are automatically detected and depicted as numbered motion key frames.

like teleporting or erratic re-orientation introduce additional complexity and semantics to the recorded motion data that need to be accounted for. At the same time, an immersive setting for the analysis process requires suitable user interaction techniques that support exploration, detail inspection and comparison tasks with a suitable reduction of the data complexity and visual clutter.

In our work, we create an immersive system for the analysis of VR motion data that allows for a seamless integration into existing VR engines and thus an immediate switch between simulation or gaming sessions, where user motion is recorded, and immersive visual analysis sessions, where the motion data is evaluated. Our system specifically aims at comparing the motion of multiple VR sessions, allowing, e.g. assessment of performance improvements in training scenarios or identification of efficiency bottlenecks in virtual design reviews. We address the visual analysis problem both by tools providing a suitable data overview over the entire temporal domain and specific interaction techniques allowing for a detailed assessment and exploration of the data. To this end, we use a specific hybrid of trajectory-based motion path visualisation and an avatar-based visualisation of key events along the time line (*key frames*) to create an immersive 3D *storyboard* of the motion sessions (Fig. 1). We demonstrate and evaluate our immersive motion analysis system at the example of a VR assembly game together with expert users. The example shares many characteristics both with classic objective-oriented VR games as well as serious VR-based industrial training applications. Please note that the system and its interactions are also demonstrated in an accompanying video to this paper (Online Resource 1).

Via a settings tablet, effective interactive exploration and filtering of the complex motion data is supported. Our approach allows for the visual analysis and comparison of VR motion data, facilitating the understanding and evaluation of user behaviour patterns

## 2 Related work

The visual analysis of human motion for performing various tasks is a problem that goes back to the early twentieth century [30]. Frank and Lillian Gilbreth invented Chronocyclographs for the optimisation of manufacturing in factories. They captured the motion of workers by creating long-exposure images showing workers' trajectories, highlighted by worn rings with miniature electric lights. This enabled them to retrace the workers' movements and analyse their timing and efficiency. In accordance with the early approach of Gilbreth et al., a fundamental technique for visualising the space-time path of tracked objects is the depiction as *trajectories*.

Existing visual analytics methods for trajectory data are often based on 2-dimensional data displays. Application areas cover, e.g. movement analysis in team sports [31], traffic analysis [19] or time-dependent multivariate data analysis [33], to name a few. Main analysis tasks in trajectory data analysis include finding aggregations and generalisations of movement patterns, to compare them, and to distinguish groups and outliers. For an overview of methods for visual trajectory analysis please see [4].

### 2.1 3D trajectory visual analysis

Buschmann et al. [10] note that many 3D trajectory visualisation approaches use two-dimensional movement data, while using the third dimension to display additional information [3,24,36].

An important source of real 3D data are flight paths. Multiple methods visualise these paths in three dimensions [2,10]. These approaches face similar problems as us concerning visual clutter and conveying the direction of movement. While 3D approaches need visual cues to improve position and height perception [10], VR environments can make use of stereoscopic vision.

Several task-specific analysis frameworks for motion data have been presented in the past, requiring specifically tailored tracking systems. Anagnostakis et al. [1] analyse motion by visually tracking a stylus with reflective markers. By testing on a constrained task, user actions can be classified and visualised as a Chronocyclegraph. These graphs are used as simple visualisation tools and offer limited to no interaction or filtering approaches to address visual clutter and exploration. Tashiro et al. [35] use magnetic tracking to record the movement of surgical instruments, to compare the performance of experienced surgeons and trainees through trajectory visualisation. Both approaches work well for short tasks, while longer tasks would require addressing the problem of increasing visual clutter.

Visual analysis of motion data tracked in immersive environments has been addressed in various previous works. Sas et al. [32] visualise and cluster three-dimensional trajectories in two dimensions by creating self-organising maps. The maps create clusters of movement patterns, allowing to classify user movement on-line. Covaci et al. [15] use the captured data to compare basketball throws in the real and the virtual world, as well as the effectiveness of guidance in the virtual world. Büschel et al. [9] use 3D trajectories to analyse the movement of analysts using a 3D in-place data visualisation on a mobile device.

Virtual training environments for manufacturing can also benefit from motion analysis during training. As shown by Gomes De Sá [16], these environments have been of interest for commercial applications for several years. Multiple assembly training systems have used Chronocyclegraphs to visualise the motion of users in their immersive training set-up [17,18,26]. These works focus on creating and evaluating haptic VR assembly training systems built for specific tasks, but provide only limited insight in their visualisation and do not discuss more elaborated interaction with the data.

## 2.2 Immersive trajectory visualisation

An emerging research area for virtual reality is immersive analytics (IA) applications [11]. As outlined in Marriott et al. [27], virtual reality provides an immersive experience with better spatial clues than 3D and with potential for more intuitive interaction. Previous work has found that performance in VR IA applications is similar to desktop 3D, despite the unfamiliarity of test subjects. IA environments have also shown greater immersion and a lower mental workload due

to more natural interactions [6,22,37,38]. Classic VR-based IA research concerned with human motion analysis and trajectory visualisation is receiving its data from outside the virtual environment and mostly covers a single motion path per actor in a large-scale context.

Zhang et al. [40] perform immersive trajectory stacking, using the third dimension to encode additional information into two-dimensional trajectories. The approach reduces over-plotting issues present in traditional 3D renderings. Similarly, Wagner Filho et al. [38] use the third dimension to encode time in GPS tracking data to create a space-time cube of trajectories. While our focus lies on virtual reality, augmented reality research is also looking into space-time cubes and is evaluating fitting interaction metaphors [34].

Hurter et al. [22] have created an interactive VR environment for air-traffic visualisation, which uses trajectories linked with traditional information visualisation graphs. Hurter et al. have identified several significant challenges for the visualisation of trajectory datasets. Some of these challenges are also relevant for human motion analysis: trajectories are inherently multidimensional and therefore two-dimensional projections might lose important features; trajectory visualisations can become dense and tangled, necessitating mechanisms for reducing this visual clutter. Nguyen et al. [29] use movement data from bee tracking to visualise their flight paths as trajectories. The trajectories are embedded in the natural 3D geo-spatial context of the bees.

In contrast to the discussed methods, human motion tracked in VR produces multiple trajectories per actor (head and hands), which poses the additional challenge of maintaining the visual association between corresponding trajectory points. Since commercial head-mounted displays (HMDs) have head and hand tracking built-in, we tailor our application to these VR devices. Additionally, Cordeil et al. [14] have shown that HMDs are a viable tool for (collaborative) immersive analysis.

## 2.3 Motion tracking for motion analysis

With full-body tracking, more elaborate movement analytics is possible (e.g. [5,7,41]). These methods track many points on the body, and while their analysis can be more comprehensive, tracking set-ups for VR are more complex and need additional hardware and wearable accessories, mostly to mount multiple markers. In contrast, we present an accessible consumer-level solution for which users need not change their virtual reality set-up, which allows to make use of the data that is already inherently tracked during a typical VR session, i.e. the headset (head) and both controllers (hands).

### 3 Immersive motion analysis concept

In the following, we describe our concept for immersive visual analysis of VR motion data, supporting analysis and comparison of motion patterns of users in task-oriented VR applications. Our analysis design is applicable to a wide range of training and/or learning applications, in which movements of VR users need to be understood. In this paper, we will consider a VR-based automotive assembly/repair task.

#### 3.1 Analysis objectives and challenges


The analysis of motion from gaming and training tasks in VR environments share similar goals. We distinguish three main high-level analysis tasks, which in our experience are important for understanding user behaviour in VR-based training applications:

- T1 Analysis of change in behaviour of repeated motion
- T2 Detecting anomalies in task solving
- T3 Finding variations in user motion within single sessions or between multiple session

Inspecting the motion data should reveal patterns and variances in repeated motion or highlight anomalies in how users tackled the task. Furthermore, a fundamental task is the comparison of different sessions. Therefore, we need a way to see variations in the movement of different users, like different timings and sequences for common operations.

The realisation of the above-stated goals in an immersive analytics system comes with a set of challenges. As motion data are acquired in 3D VR space, we propose to do the visual analysis also in 3D VR space, for (1) not having to transform the data to, e.g. a 2D display and (2) to retain the application context including the 3D application environment. Direct visualisation of the raw trajectory data can quickly become cluttered (cf. Fig. 2). To address this issue, we allow for exploring the data both by giving a suitable *overview* (Sect. 3.2) and by enabling a *detailed* view (Sect. 3.3) and implement a set of visualisation options to further reduce visual clutter. The provided tools and options need to support the exploration and provide an intuitive interaction with the data. Their utilisation must be clear, and they have to be easily accessible.

Motion data are created by recording and tracking the movement of individual users during a VR session. Since flexibility and ease-of-use is important, we rely on common sources of motion for commercial VR set-ups: the headset for head motion and two controllers for hand motion. The movement in VR environments comprises fluent motion and teleport actions.



**Fig. 2** Trajectory visualisation can be a powerful tool to gain insight into human motion in VR environments. However, appropriate processing and abstraction of trajectories is needed as visual clutter arises when showing multiple trajectories


#### 3.2 Overview visualisation

To provide an overview of the entire spatio-temporal motion data domain, we employ a hybrid approach of trajectory visualisation and avatar-based key frame depiction (Fig. 3). Our approach allows analysts to navigate through the data in the same immersive environment it was recorded in. Alternatively, they can observe the entire spatial embedding of the data. For this, it is important to create elevated vantage points in the environment.

The main problem for providing a suitable data overview and insight are different sources of visual clutter appearing both in the data and in the embedded environment: (i) the motion data of a single session already consists of three trajectories per session (head and two hands), (ii) during comparison, multiple sessions may be shown at once, (iii) multiple visits to the same locations will lead to entangled trajectories, iv. the environment itself can obfuscate the visualisation especially when textured with high-frequency patterns. Figure 2 shows a direct visualisation of raw trajectory data that exhibits these types of clutter. The following will describe our set of tools and measures that address these problems for an improved data overview.

**Visual Representation of Trajectories** The tracked data comes in varying intervals and originates from natural motion in a VR application context. To approximate the original motion, we visualise the trajectories as tubes based on centripetal Catmull–Rom splines [39]. They interpolate the original data without creating loops and follow it more tightly when compared to other Catmull–Rom schemes. These properties and no pre-processing ensure that we retain the essential trajectory information. To visually group the different trajectories, we uniquely colour each session. Colours are chosen via a repeated clockwise hue shift. The first four colours are distributed over 90° angles, the next eight fill the angular sections in between, and so on. Perceiving the trajectory directions is important when directly looking at the trajectories to gain a quick insight into the movement of the user. We texture each trajectory tube with glyphs flowing in the






**Fig. 3** Pose triangle in a cluttered environment (left) and its particle trail during the animation with hidden trajectories (right). The head is represented as a headset, and the hands have corresponding hand meshes. The particles of the trail spawn in fixed intervals. Their density indicates the speed of the animated motion and their colour shows their age

direction of movement, which we found the most fitting due to its simplicity and low visual complexity. The glyphs' speed varies based on the speed of the original motion, and it encodes its motion source using a small symbol for head or left/right hand. We refrain from using more complex colouring or patterns (e.g. [24]), since we found that they are too hard to see on thin trajectories in a room-scale set-up.

A particular type of clutter stems from teleportation events present in the data, which leads to many long linear trajectory segments (see left half of Fig. 2). To visually discriminate these path segments from the actual human motion data, we reduce their tubular diameter and exclude them from the spline interpolation. Analysts can also hide the tubes of teleportation events completely, if desired.

**Average Trajectory** Visualising all three sources of motion (head, hands) is important to fully understand the motion of users in interaction. However, when only the coarse movement in the environment is important, it is beneficial to use a more compact representation with only one representative trajectory per motion session. To this end, we allow analysts to switch to only one trajectory resulting from the average motion of all three motion sources. This allows for a less cluttered overview of the locations and order of user movement. We use the average motion to ensure that we can still see movement when one motion source moves vastly different from the others. Reaching far forward with one hand or bending down to change perspective, while remaining mostly stationary with the rest of the body, are common examples for this.

**Key Framing and Storyboard** To find a middle-ground between animation and static trajectories, we use a key framed storyboard metaphor, similar to [5,12,21]. The story-



**Fig. 4** A comparison between uniformly timed key frames (top) and key frames based on our clustering (bottom). Pose triangles at the clustering-based key frames have less overlap and have more geometrical spread than the uniform key frames. The time axis on the bottom left of each image shows the differing distributions over time

board comprises a set of key frames defined by a timestamp. These key frames are displayed for all loaded sessions as static pose triangles at the corresponding locations (cf. Fig. 3). Pose triangles match the session's colour and consist of three meshes to represent the position and rotation of head and hands, with a thin semi-transparent triangle connecting them.

Key frames are chosen at timestamps around which motion tracking points exhibit a high spatially concentrated activity. This ensures a better spread between them (cf. Fig. 4). We compute the clusters of concentrated activity through centroid linkage agglomerative clustering [20] on the 4d space-time locations of the tracking points, using the Euclidean distance metric. Key frames are then defined by the median timestamps of each cluster in the selected hierarchy level. Since the clustering is fully automatic, we allow analysts to choose the displayed level of the hierarchy (i.e. the number of key frames).

For facilitating the comparison of timing and order between sessions, we compute key frames for one user-defined reference session and use them for all sessions. Pose triangles for each loaded session are displayed with an overhead number showing the index of the corresponding key frame. These indices allow analysts to easily find corresponding pose triangles from different sessions that represent the same key frame.

Furthermore, analysts can apply a time-offset to all key frames for additional insight into the movement around each key frame. Figure 1 shows a storyboard for three sessions with red as the reference session. It also shows how combining averaged trajectories and hidden teleportation segments



**Fig. 5** Difference between an unchanged environment (left) and an environment with simplified texturing (right)

can highlight the storyboard visualisation. To the lower right, we can see the settings tablet with an option to switch between our clustered storyboard and the uniform one shown in Fig. 4. *Environment* Displaying the environmental context of the movement is important for understanding. However, high-frequent textures or colourful backgrounds in environments can contribute to visual clutter. To simplify the environment, we allow analysts to set all texturing to a uniform white colouring. This does not destroy the context for the movement since shadows and shapes are still perceivable. Figure 5 shows the difference between a normal and a simplified environment, demonstrating the above-stated problems: the high-frequent floor texture introduces visual unrest while the red car body partially hides some trajectories.


### 3.3 Detail visualisation

While an exploratory visual analytics process needs an overview of the data, measures for a more detailed, in-depth view are equally important.

*Filtering* Motion can exhibit regions of locally overlapping and entangled patterns, especially in scenarios with repeated tasks in the same location. To address this, our analytics system allows the filtering of trajectory ranges by time.

*Detail Information* If analysts want to see more detailed information at a certain point on the trajectory, they can use a detail popup. It displays information for a tracked object (head, hand) at the specified location. It shows the session ID, the interpolated position and rotation, the current time and the source of motion (head, right or left hand).

*Motion Animation* Visualising motion as trajectories is a static representation of spatio-temporal data. To enable better insight into the movement of users, our system can show a simplified replay of the movement. Similar to the storyboard, we use pose triangles to animate the motion. Each of the meshes in the triangle has a short-lived trail of particles when moving, to indicate the path it has taken. Figure 3 shows the pose triangle and the trail of particles. The triangle should give analysts enough insight to understand the motion of users and the relative positioning of head and hand, but be simple enough to not contribute to even more visual clutter. Analysts can modify the animation speed but can also set the time of the animation manually.



**Fig. 6** Local interaction ball and its popup menu (left) and the settings tablet together with the laser pointer (right)

### 3.4 Interaction

The described visualisation provides a variety of tools analysts can use for the exploration and analysis of the motion data. To provide a flexible interaction concept, we offer three types of interaction: pointing-based interaction, local interaction and menu-based interaction. They cover different needs and provide a mixture of shared and exclusive functionality. *Pointing-based Interaction* This interaction uses the metaphor of a laser pointer. Analysts can interact with user interface elements or show detail information of specific points on trajectories. The interaction with user interface elements works well for all ranges, but accurately pointing the laser at trajectories from range can be difficult. It can be seen in Fig. 6 right and is triggered with its own button on the controllers.

*Local Interaction* Local interaction gives analysts a more tangible control when standing in front of a trajectory [13]. When gripping a trajectory, an interaction ball appears on the trajectory. After being grabbed, the ball highlights the trajectory it lies on and it can be moved along the trajectory by hand movement or via thumbstick tilt. This gives a fine control with the hands, but also allows for moving the ball further than the personal reach of the analyst. By touching the ball, analysts can open a menu to show detail information or to display a pose triangle on the position of the ball. When grabbed by both hands, the ball splits into two and creates a filter to show only the trajectories in between the two balls. This filter applies to all trajectories. The local interaction is especially useful for many trajectories, where the laser pointer cannot reach inside a cluster. Figure 6 shows the local interaction ball and its menu.

*Menu-based interaction / Settings Tablet* Some introduced tools are independent of positioning. For these tools, a tablet is following the analysts until they grab and place it in the world. When ‘thrown away’ with enough force, it will resume following. This tablet is set up similar to the settings menu of a video game for controlling the tools. Analysts can set trajectory visualisation options, control the global time filter and the global animation (disable/enable, speed and time) and

set the hierarchy level and a time-offset for the storyboard key frames. They can change settings either with the laser pointer or by touching the controls with their virtual index fingers. Figures 6 and 1 show it in use.

## 4 Implementation

We used the Unreal Engine<sup>1</sup> to create our system. Ideally, the VR environment producing the data is the same as in the analysis process. However, using a game engine also allows exporting and importing an environment from other sources. We used the Oculus Rift CV1<sup>2</sup> and the HTC Vive Pro<sup>3</sup> in our experiments.

**Session Data** We have created our own logging data format and data logger. Since logging movement can lead to large amounts of data, adaptive logging intervals reduce the number of needed log entries by pausing the logging of objects that are not moving. In VR, analysts can start and stop session tracking, load and unload sessions, and can switch between analysis and game. This functionality is provided by stationary user interface elements in the environment that can be controlled either with the laser pointer or by touch.

**Locomotion** An important ergonomic aspect in any immersive application is to minimise the amount of discomfort via locomotion. In our system, analysts can either physically move around or use teleportation. This ensures a minimal amount of nausea when moving further distances and follows the recommendation of Jerald [23]. Analysts who prefer sitting and are less prone to nausea can fly around using the keyboard. We do not support locomotion methods that warp the space or otherwise alter the correspondence between real and virtual movement, like re-orientation or redirected walking, because we want to preserve the relation between real and virtual movement for our industrial application context.

**Data Processing** The trajectories are computed when the data is loaded. We use a custom mesh with one continuous tube for each trajectory, which allows us to use lighting and shadowing. While it is not the most efficient method for rendering trajectories, we found that this approach represents a fitting trade-off between fidelity and performance for our visualisation.

Trajectories have a mesh structure and collision information. Generally, their generation is executed within less than half a second. When analysts load multiple sessions at once, they might experience a short stutter. Since hierarchical clustering is a relatively slow algorithm ( $O(n^2)$  [28]), we compute and save it the first time a session is loaded. For a session with 1400 tracking entries spanning over 13 min,

creating the tubes and the collision takes 100 ms in total, while the agglomerative hierarchical clustering takes about 5 seconds (on an Intel I7-4930K CPU).

## 5 Exemplary analysis workflow

To demonstrate the use of our system, we outline an exemplary workflow in an exploratory analysis session based on a task-oriented car-shop game. The goal of the game is mounting four wheels as fast and as precise as possible onto a race car. This scenario expectantly produces interesting motion patterns with different focus regions (wheels, wheel pickup) and varying dwell times at different locations throughout the session. The common objective allows for a comparable set of motion data from different sessions, while giving users enough freedom (e.g. mounting order) to observe significant differences between individual motion sessions. A visual outline of the workflow is given in Fig. 7.

First, to gain an overview, the analyst teleports to an overhead location and looks at the storyboard. From there, differences in order and movement-based anomalies can easily be seen. In Fig. 7a, we can see that the blue user has a different timing for key frame 1 and seems to have a different order according to key frame 4, while green is venturing off the usual path ( $T_2$ ,  $T_3$ ). Then, to gain more insight into the observed motion, the analyst looks at the animation for the three sessions (b). The animation further supports the observations concerning speed and order. Next, the analyst can observe certain locations of interest in more detail. With the local interaction, a more thorough investigation of the assembly behaviour at wheel assembly spots can be performed (c). In the shown case, the analyst has filtered the trajectories to the time frame of the assembly of a specific wheel. With a pose triangle, the movement can be studied, while the filtering shows the locations of the other sessions in the same timeframe ( $T_3$ ). Finally, to observe the pickup behaviour of the different users, the analyst shifts their attention to the pickup spot (d). From the different heights of the pickup, it appears that the users have different body heights and change their stance when picking up the wheel ( $T_1$ ). Using the laser pointer or local interaction, the analyst can examine the timeframes of each pickup and the animation can give more insight into the precise movement of the users.

## 6 Evaluation and results

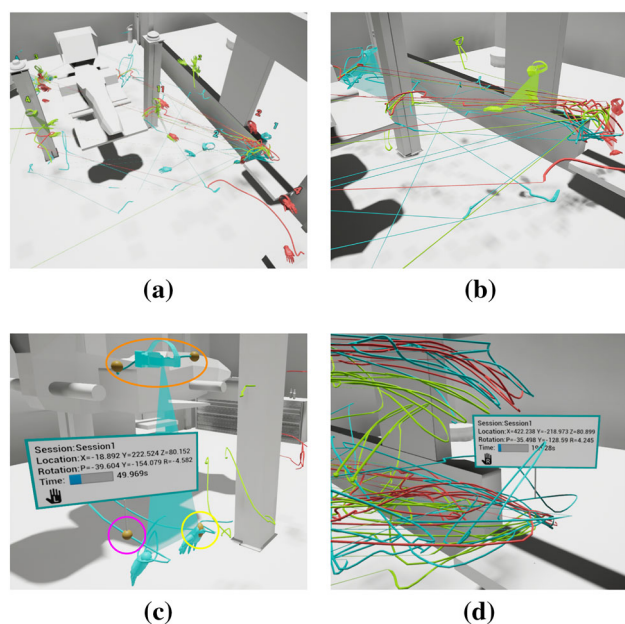
We evaluate the validity of our approach based on a user study to analyse our design choices and whether the proposed immersive analytics system can support the visual analysis process to meet the analysis objectives stated in Sect. 3.1.

<sup>1</sup> Unreal Engine [www.unrealengine.com](http://www.unrealengine.com).

<sup>2</sup> Oculus Rift [www.oculus.com/rift/](http://www.oculus.com/rift/).

<sup>3</sup> HTC Vive Pro [www.vive.com/uk/product/vive-pro/](http://www.vive.com/uk/product/vive-pro/).





**Fig. 7** Examples for different stages in using our tool: **a** overview with storyboard; **b** animation view; **c** local interaction animation (yellow, right) with filter (orange, top), and detail (violet, left); **d** detail inspection of wheel pickup


## 6.1 Evaluation

The data for the study comes from the simplified industrial car-shop use-case mentioned in Sect. 5. Based on previous demonstrations of the car-shop application at various events to introduce students to VR, we have identified a set of interesting motion patterns and features for our user study from logged movement data. These include different assembly order, mistakes when assembling, different behaviour and stances while assembling and picking up a wheel, simulated issues with the headset, and different behaviour concerning movement / teleporting. To create a controlled dataset for the user study, we have simulated and captured a set of assembly sessions containing these patterns and features. One of these sessions is shown in Fig. 8. In the middle, we can immediately make an observation pertaining to high-level task  $T2$  by noticing the outlier in the visualisation, which depicts the movement of a user re-adjusting their headset.

**Evaluation tasks** The primary evaluation of the system comprised four analysis tasks: two detail tasks and two comparison tasks. Detail tasks include the analysis of the motion for one session and comparison tasks compare three sessions with each other:

### Detail

1. How do repeated actions change over time?
2. Are there any outliers in the movement patterns of the user? What happened/why?



**Fig. 8** One of the session datasets used in the evaluation. In the middle (green) we can quickly see an anomalous trajectory where all three trajectories meet at head-level. In this case, it is from a user that is adjusting their headset

### Comparison

3. How does the order of actions for the three sessions differ?
4. Can you find differences in how the users performed the assembly task? Which was the best strategy?

These tasks cover our range of high-level tasks ( $T1$ ,  $T2$ ,  $T3$ ) in Sect. 3.1 and are a mixture of different levels of specificity. This way, we can see how well experts find patterns with varying guidance and how well patterns can ‘pop out’ within the scope of less specific analysis tasks. To analyse the effect of the simplified environment, it was disabled in Task 3 without giving the experts an explanation why this was done.

## 6.2 Evaluation set-up

The evaluation is a qualitative, thinking-aloud user study with five experts: three experienced VR users from different fields of VR research, one augmented reality researcher and one expert in visual data analysis. All experts had previous experiences with VR and quickly familiarised themselves with this new VR environment. For an insight of where the data came from, the experts first had a short session in the car shop, in which they carried out the assembly for themselves. Next, they were given a short presentation of the system and the implemented tools. This presentation was followed by an interactive session in the system in which they were guided to try out all system features. We kept this introduction relatively short, to gauge the intuitiveness of the system and to evaluate the clarity of use for our implemented tools. The main part of the evaluation comprised the four analysis tasks to be solved. After completing these steps, the experts filled in a System Usability Score [8] and we discussed their experience. On average, one evaluation session lasted for 1 h. Solving each task ranged from 5 to 10 min and was cut short



when experts could not find all objectives of the task after 15 min.

### 6.3 Evaluation results

This section will discuss the results and observations of the evaluation. First, we discuss how well experts were able to solve our tasks, followed by the results of the System Usability Score. Next, we present our observations of how they solved the given tasks and discuss the received feedback concerning the implemented tools.

**Task Completion** Experts were able to find differences in wheel assembly order (*T3*) and positional outliers (*T2*) consistently and quickly. Positional outliers were identified by the trajectory visualisation itself, while assessing the wheel assembly order could not rely on only the trajectory visualisation. Experts needed to use our implemented tools for further insight to see the assembly order, like the motion animation. Timing and related changes in wheel assembly behaviour (e.g. wheel falling to the ground, or removing and re-assembling a wheel) could also reliably be detected by the experts (*T1*, *T2*). More subtle differences in stances and specific movement patterns for single and multiple sessions (*T1*, *T2*, *T3*) were less easy to detect, considering the more general nature of these observations. An example is the way wheels are assembled where some users do not bend down to assemble the wheel, some are standing to the side and others bend down with either their knees or upper body for more precise fitting of the wheel. These differences in behaviour can influence the speed and precision of assembly between users. When asked more specific questions, experts were able to also find such more subtle differences. Finding a reason for outlier behaviour was prone to misinterpretations. While all experts found the outlier where one user re-adjusted their headset, only one expert stated the correct interpretation for this behaviour. Finally, the experts also found other, non-planned, interesting movement patterns like a change in handedness for picking up the wheel (*T1*).

**System Usability** Overall, the experts stated that the system is intuitive after an initial familiarisation and expressed an interest in continued use of the system for motion analysis tasks. They judge that the immersive virtual environment is advantageous for understanding the spatial data and that it succeeds in immersing them into the motion. The System Usability Score (SUS) supports these statements with an average of 76.25 points ( $\sigma = 4.87$ ), which is slightly above *average* compared to SUS scores in general [25]. A prominent observation during the evaluation is that the implemented system is created for analysts and not for novice users. The experts expressed a need for a longer introductory period than the given short introduction to fully grasp the system. There were also minor comments on improvements for some interactions, but they were not listed as major influ-

ences on the usability score. None of the experts had issues with discomfort in the VR environment.

**Task observations** Experts showed different strategies in the ways they solved the given tasks. In general, the trajectory visualisation worked well for seeing positional outliers and all experts used the animation for overview and detail observations. One expert focused mostly on the filter tool, while another expert mainly used the detail information via laser pointer for insight into the movement. Due to the short introduction, some hints to existing tools were also necessary. These hints were given as the question ‘would tool *x* help you here?’. Once reminded, they saw its use and used it more often. The settings panel was the primary interaction method for all experts, while the local interaction was mostly unused. This is due to experts forgetting about it and because the settings panel offers similar functionality, albeit through a potentially more convenient interaction metaphor. The storyboard was mostly used to see the order of operations, since it provided a still overview of the clusters. Finally, experts displayed vastly different movement patterns during the analysis themselves. Some remained mostly stationary, others moved at ground level and looked at the data more closely. Two experts used an overhead view, which is a roof-platform surrounding the walls.

**Tool feedback** Experts appreciated the overall interaction and visualisation of the implemented system. Their favourite tool was the animation visualisation. According to them, the combination of head/hand meshes for orientation and the triangle in between is ideal and they do not need a more complex avatar for the animation, since it might contribute further to visual clutter. The experts liked the ability to hide teleports to remove unnecessary visual clutter, since the teleport trajectories represent instantaneous movement only. One expert in particular disliked trajectories passing through the area of their physical body, for whom hidden teleports increased the available space. Our average trajectory visualisation did not have the same acceptance, since only the expert with the previously mentioned dislike used it extensively. While the experts agreed on the potential usefulness of the storyboard, two stated that they did not fully understand how it works and that they had difficulties in interpreting the visualisation. The others like the still-frame visualisation of key moments and find the overhead numbers useful for in-session chronological order and for comparing between different sessions. Experts constrained themselves to the first three levels of the storyboard, as higher levels started to contribute to visual clutter. We switched to the normal environment instead of the simplified one in Task 3 without telling the experts why. When asked, they did not notice a difference between both environment styles. Finally, experts liked the interaction with the settings panel, but small improvements concerning its behaviour were deemed necessary.

## 6.4 Discussion

The presented results show that the implemented system is a step in the right direction for providing an environment for the intuitive and immersive visualisation and exploratory analysis of motion data. Differing analysis behaviours suggest that it is important for such systems to remain flexible and provide users with a wealth of tools to cater to individual preferences. While the storyboard was well received overall, a clustering method that is easier to understand and without the hierarchical nature might be an improvement. The indifferent perception of normal and simplified environment might be due to the relatively simple and reduced texturing of the given car-shop environment. Other VR environments with a more camouflaging and obfuscating appearance might make this texture-reduction measure more essential. Another factor is the way the environment change was set up in the evaluation. A more explicit change might be necessary. Additionally, while our visualisation techniques can reduce visual clutter to some extent, experts felt that visual clutter might become too much for comparing more than the given three sessions. Some experts mentioned they presume that their analysis behaviour might change once they have more experience and a better routine for selecting certain tools for certain tasks.

## 6.5 Changes after evaluation

From the feedback we have received, we added two improvements after the evaluation. First, since the average trajectory might not always be the ideal way of simplifying the trajectories, we allow analysts to hide the visualisation of any motion source (left/right hands and head) for all displayed sessions. This way, they can decide to, e.g. view a single motion source instead of the average trajectory. Second, while the numbers over the storyboard's key-frames indicate the temporal order, they do not convey the time difference, which is of special interest for our clustered, non-uniform distribution of key-frames. Hence, we have added a circular time indicator next to the numbers, to show the elapsed time of each key-frame relative to its session.

## 7 Limitations and future work

We focus on analysis of up to a handful of sessions and have opted in favour of visual fidelity for our trajectories. Since a stable frame-rate is important, displaying more than a dozen sessions would necessitate a different approach in their rendering. Additionally, while we focused on three motion sources in this work, we plan to evaluate whether additional trackers or full body tracking could be supported or if we need a different visualisation method due to too much visual clutter.

To give analysts the needed flexibility in how they interact and visualise the motion, we have implemented several tools and interactions. However, this flexibility necessitates that analysts know the available options to make use of them as there is little guidance concerning the choice of the right tool for the desired exploration. While we focused on the exploratory analysis of motion in this work, we want to add more automated features for analysis and guidance in the future. This includes automatic highlighting of interesting locations and viewpoints, integration of events and tasks for better semantic analysis, and more performance statistics. We focused on clustering the movement within one session, but there are many avenues for clustering different aspects of the motion data, based on different metrics. Clustering complex moves entails a segmentation problem that would need to be solved in an application-dependent way. A potential avenue for the future is the addition of a distributed work environment. With this, users could play or train anywhere in the world, while analysts see ongoing sessions and can analyse them shortly after completion. Such a system can, for example, enable manufacturing training on remote locations, with short response times and feedback from experts and trainers. Finally, we have evaluated our work as a first step towards trajectory-based visual analysis of user behaviour in VR environments and are planning more task-specific designs and more comprehensive evaluations in the future.

## 8 Conclusion

We have created and evaluated an immersive analytics system for the explorative analysis of human VR motion data, focused on task-oriented applications. The system provides analysts with a wealth of tools for gaining an insight into the data by analysing behaviour and finding patterns or outliers either in single sessions or when comparing different sessions. The most prominent challenge in creating such a visualisation is dealing with visual clutter and giving an intuitive way of interacting with the data. To tackle the visual clutter, we have implemented a range of improvements to trajectory visualisations and offer several tools, giving analysts the freedom of interacting in a way that fits their preferences. Our evaluation demonstrates that it is possible to uncover many interesting patterns and behaviours from tracked motion.

**Acknowledgements** Open access funding provided by Graz University of Technology. Thanks to our project partners, division for Visual Computing of Fraunhofer Austria Research GmbH and AVL List GmbH.

**Funding** This work is supported by the Austria Research Promotion Agency (FFG) under Grant No. 864814.

## Compliance with ethical standards

**Conflict of interest** The authors declare that they have no conflict of interest

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Anagnostakis, D., Ritchie, J., Lim, T., Sung, R., Dewar, R.: Automated coordinate measuring machine inspection planning knowledge capture and formalization. *J. Comput. Inf. Sci. Eng.* **18**(3), 031005 (2018)
2. Andrienko, G., Andrienko, N., Fuchs, G., Garcia, J.M.C.: Clustering trajectories by relevant parts for air traffic analysis. *IEEE Trans. Vis. Comput. Graph.* **24**(1), 34–44 (2018)
3. Andrienko, G., Andrienko, N., Schumann, H., Tominski, C.: Visualization of Trajectory Attributes in Space-Time Cube and Visualization of Trajectory Attributes in Space-Time Cube and Trajectory Wall, pp. 157–163. Springer, Berlin (2014)
4. Andrienko, N., Andrienko, G.: Visual analytics of movement: an overview of methods, tools and procedures. *Inf. Vis.* **12**(1), 3–24 (2013)
5. Assa, J., Caspi, Y., Cohen-Or, D.: Action synopsis. In: ACM SIGGRAPH 2005 Papers, p. 667. ACM Press, New York, USA (2005)
6. Batch, A., Cunningham, A., Cordeil, M., Elmqvist, N., Dwyer, T., Thomas, B.H., Marriott, K.: There is no spoon: evaluating performance, space use, and presence with expert domain users in immersive analytics. *IEEE Trans. Vis. Comput. Graph.* pp. 1–1 (2019)
7. Bernard, J., Wilhelm, N., Krüger, B., May, T., Schreck, T., Kohlhammer, J., Kohlhammer, J.: MotionExplorer: exploratory search in human motion capture data based on hierarchical aggregation. *Trans. Vis. Comput. Graph. (Proc. VAST 2013)* **19**(12), 2257–2266 (2013)
8. Brooke, J.: SUS—a quick and dirty usability scale. *Usability Eval. Ind.* **189**(194), 4–7 (1996)
9. Büschel, W., Reipschläger, P., Langner, R., Dachsel, R.: Investigating the use of spatial interaction for 3D data visualization on mobile devices. In: ISS'17 Proceedings of the 2017 ACM International Conference Interaction Surfaces Spaces, pp. 62–71. ACM Press, New York, USA (2017)
10. Buschmann, S., Trapp, M., Döllner, J.: Animated visualization of spatial-temporal trajectory data for air-traffic analysis. *Vis. Comput.* **32**(3), 371–381 (2016)
11. Chandler, T., Cordeil, M., Czauderna, T., Dwyer, T., Glowacki, J., Goncu, C., Klapperstueck, M., Klein, K., Marriott, K., Schreiber, F., Wilson, E.: Immersive analytics. In: 2015 Big Data Vis. Anal. (BDVA), pp. 529–533. ACM Press, New York, USA (2015)
12. Coffey, D., Korsakov, F., Ewert, M., Hagh-Shenas, H., Thorson, L., Ellingson, A., Nuckley, D., Keefe, D.: Visualizing motion data in virtual reality: understanding the roles of animation, interaction, and static Presentation. *Comput. Graph. Forum* **31**(3pt3), 1215–1224 (2012)
13. Cordeil, M., Bach, B., Li, Y., Wilson, E., Dwyer, T.: Design space for spatio-data coordination: tangible interaction devices for immersive information visualisation. *IEEE Pac. Vis. Symp.* **2017**, 46–50 (2017)
14. Cordeil, M., Dwyer, T., Klein, K., Laha, B., Marriott, K., Thomas, B.H.: Immersive collaborative analysis of network connectivity: CAVE-style or head-mounted display? *IEEE Trans. Vis. Comput. Graph.* **23**(1), 441–450 (2017)
15. Covaci, A., Olivier, A.H., Multon, F.: Visual perspective and feedback guidance for VR free-throw training. *IEEE Comput. Graph. Appl.* **35**(5), 55–65 (2015)
16. Gomes De Sá, A., Zachmann, G.: Virtual reality as a tool for verification of assembly and maintenance processes. *Comput. Graph.* **23**(3), 389–403 (1999)
17. Gonzalez-Badillo, G., Medellin-Castillo, H., Lim, T., Ritchie, J., Garbaya, S.: The development of a physics and constraint-based haptic virtual assembly system. *Assem. Autom.* **34**(1), 41–55 (2014)
18. Gonzalez-Badillo, G., Medellin-Castillo, H.I., Lim, T.: Development of a haptic virtual reality system for assembly planning and evaluation. *Proc. Technol.* **7**, 265–272 (2013)
19. Guo, H., Wang, Z., Yu, B., Zhao, H., Yuan, X.: TripVista: triple Perspective visual trajectory analytics and its application on microscopic traffic data at a road intersection. In: 2011 IEEE Pacific Visual Symposium, pp. 163–170. IEEE (2011)
20. Han, J., Kamber, M., Pei, J.: Data Mining: Concepts and Techniques, 3rd edn. Elsevier Inc., Burlington (2012)
21. Hentschel, B., Tedjo, I., Probst, M., Wolter, M., Behr, M., Bischof, C., Kuhlen, T.: Interactive blood damage analysis for ventricular assist devices. *IEEE Trans. Vis. Comput. Graph.* **14**(6), 1515–1522 (2008)
22. Hurter, C., Riche, N.H., Drucker, S.M., Cordeil, M., Alligier, R., Vuilleumot, R.: FiberClay: sculpting three dimensional trajectories to reveal structural insights. *IEEE Trans. Vis. Comput. Graph.* **25**(1), 704–714 (2019)
23. Jerald, J.: The VR Book. Association for computing machinery and Morgan and Claypool, New York (2015)
24. Lange, D., Samsel, F., Karamouzas, I., Guy, S.J., Dockter, R., Kowalewski, T., Keefe, D.F.: Trajectory Mapper: interactive widgets and artist-designed encodings for visualizing multivariate trajectory data. *EuroVis 2017 - Short Pap* (2017)
25. Lewis, J.R., Sauro, J.: Item benchmarks for the system usability scale. *J. Usability Stud.* **13**(3), 158–167 (2018). <https://doi.org/10.5555/3294033.3294037>
26. Lim, T., Ritchie, J., Sung, R., Kosmadoudi, Z., Liu, Y., Thi, A.: Haptic virtual reality assembly—Moving towards Real Engineering Applications. In: Adv. Haptics, pp. 694–721. IntechOpen (2010)
27. Marriott, K., Schreiber, F., Dwyer, T., Klein, K., Riche, N.H., Itoh, T., Stuerzlinger, W., Thomas, B.H.: Immersive Analytics State-of-the-Art, Lecture Notes in Computer Science, vol. 11190. Springer (2018)
28. Murtagh, F., Contreras, P.: Algorithms for hierarchical clustering: an overview, II. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.* **7**(6), e1219 (2017)
29. Nguyen, H., Wang, F., Williams, R., Engelke, U., Kruger, A., De Souza, P.: Immersive visual analysis of insect flight behaviour. In: Immersive Analytic Work. IEEE VIS, pp. 1–5. IEEE, Phoenix, USA (2017)
30. Price, B.: Frank and Lillian Gilbreth and the manufacturing and marketing of motion study, 1908–1924. *Bus. Econ. Hist.* **18**(c), 88–89 (1989)



31. Sacha, D., Al-Masoudi, F., Stein, M., Schreck, T., Keim, D.A., Andrienko, G., Janetzko, H.: Dynamic visual abstraction of Soccer movement. *Comput. Graph. Forum* **36**(3), 305–315 (2017)
32. Sas, C., O'Hare, G., Reilly, R.: Virtual environment trajectory analysis: a basis for navigational assistance and scene adaptivity. *Future Gener. Comput. Syst.* **21**(7), 1157–1166 (2005)
33. Schreck, T., Bernard, J., Tekusova, T., Kohlhammer, J.: Visual cluster analysis of trajectory data with interactive Kohonen Maps. In: 2008 IEEE Symposium on Visual Analytics Science and Technology, pp. 3–10. IEEE (2008)
34. Ssin, S.Y., Walsh, J.A., Smith, R.T., Cunningham, A., Thomas, B.H.: GeoGate: correlating geo-temporal datasets using an augmented reality space-time cube and tangible interactions. In: 2019 IEEE Conference Virtual Reality 3D User Interfaces, pp. 210–219. IEEE (2019)
35. Tashiro, Y., Miura, H., Nakanishi, Y., Okazaki, K., Iwamoto, Y.: Evaluation of skills in arthroscopic training based on trajectory and force data. *Clin. Orthop. Relat. Res.* **467**(2), 546–552 (2009)
36. Tominski, C., Schumann, H., Andrienko, G., Andrienko, N.: Stacking-based visualization of trajectory attribute data. *IEEE Trans. Vis. Comput. Graph.* **18**(12), 2565–2574 (2012)
37. Wagner Filho, J., Rey, M., Freitas, C., Nedel, L.: Immersive visualization of abstract information: an evaluation on dimensionally-reduced data scatterplots. In: IEEE Virtual Reality. 2018, March (2018)
38. Wagner Filho, J.A., Stuerzlinger, W., Nedel, L.: Evaluating an immersive space-time cube geovisualization for intuitive trajectory data exploration. *IEEE Trans. Vis. Comput. Graph.*, pp. 1–1 (2019)
39. Yuksel, C., Schaefer, S., Keyser, J.: On the parameterization of Catmull–Rom curves. In: 2009 SIAM/ACM Joint Conference on Geometric and Physical Modeling, pp. 47–53. ACM Press, New York, NY, USA (2009)
40. Zhang, M.J., Li, J., Zhang, K.: Using virtual reality technique to enhance experience of exploring 3D trajectory visualizations. In: Proceedings of the 8th International Symposium on Visual Information Communication and Interaction—VINCI'15, pp. 168–169. ACM Press, New York, New York, USA (2015)
41. Zia, A., Sharma, Y., Bettadapura, V., Sarin, E.L., Essa, I.: Video and accelerometer-based motion analysis for automated surgical skills assessment. *Int. J. Comput. Assist. Radiol. Surg.* **13**(3), 443–455 (2018)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Simon Kloiber** is currently employed as project assistant and Ph.D. student at Graz University of Technology. He has graduated from the same university with a specialization on computer graphics and algorithms. The topic of his Ph.D. is Immersive Analytics in Virtual Reality, mostly for industrial applications. Within this field, Simon is looking to improve the ways we can train in VR, but also to create data analytics environments by leveraging the improved depth perception of VR.



**Volker Settgast** studied computer science at the Technical University of Braunschweig. After graduating in 2005, he worked as a scientific assistant at the Institute for Computer Graphics and Knowledge Visualization at Graz University of Technology. Since 2006 he has contributed significantly to the development of the visualization center at Graz University of Technology. Starting in 2009 he has been working for Fraunhofer Austria Research GmbH. In 2013 he graduated from doctoral school in computer science and has since been working in the area of Virtual, Mixed and Augmented Reality.



**Christoph Schinko** studied Telematics at Graz University of Technology. He received his doctorate in March 2018 with his work “Shape Processing for Content Generation” on the subject of generative modeling and geometrical optimization. From 2009 to 2013 he worked as a researcher at the Institute of Computer Graphics and Knowledge Visualization (CGV) at Graz University of Technology. Currently, he works for Fraunhofer Austria as a senior researcher in the Visual Computing business area in Graz and for the innovation center KI4LIFE in Klagenfurt. His research focuses on computer graphics, mixed reality, 3D web technologies and machine learning.



**Martin Weinzerl** has studied Mechanical Engineering, specializing in mechatronics at the University of Technology of Graz. During his first job as production planner in the electronic manufacturing he got his basic Knowledge regarding production. As Head of technology and project management at SVI Austria he was responsible for all production, new product introduction and project management processes. At AVL he currently holds the position as Project Manager in the department “Product Quality Assurance and Production Engineering” for Engineering and Technology Powertrain Systems and is responsible for coordination and project management of several national and EU funded research projects.





**Johannes Fritz** has studied mechanical engineering at Graz University of Technology and business informatics at University for Applied Science Campus02 in Graz. He holds the position “Senior Systems Engineering Methods Manager” at AVL List in Graz, Austria, where he is responsible for the Systems Engineering deployment in the field Global Project Operations for Powertrain Test Systems. This includes the development and rollout of state-of-the-art engineering methods. Besi-

des his work at AVL List, he is external lecturer for “Model-Based Engineering” at University of Applied Sciences Campus02 master’s degree program for automation technology.



**Reinhold Preiner** is a Senior Researcher in computer science, specialized in 3D computer graphics and visualization. He received his Ph.D. from TU Wien in 2017 and is now with the Institute of Computer Graphics and Knowledge Visualization at Graz University of Technology, working on projects of applied geometry processing and interactive visualization. His main research fields include geometry processing and representation, rendering and visualization.



**Tobias Schreck** is a professor and head of the Institute of Computer Graphics and Knowledge Visualization at Graz University of Technology, Austria. Previously, he held positions as Assistant Professor with University of Konstanz, Germany, and Postdoc Researcher with TU Darmstadt, Germany. He received a Ph.D. degree in Computer Science from the University of Konstanz in 2006. His research interests include visual data analysis and 3D object retrieval.