

RePiX VR - Learning environment for the Rendering Pipeline in Virtual Reality

B. Heinemann¹  and S. Görzen¹  and U. Schroeder¹ 

¹RWTH Aachen University, Germany

Abstract

Virtual reality can be used to support computer graphics teaching, e.g. by offering the chance to illustrate 3D processes that are difficult to convey. This paper describes the development and first evaluations of RePiX VR a virtual reality tool for computer graphics education, which focuses on the teaching of fundamental concepts of the rendering pipeline and offers researchers the opportunity to study learning in VR by integrating learning analytics. For this, the tool itself is presented and the evaluation, which uses quantitative methods and learning analytics to show the effectiveness of the tool. The first evaluations show that even learners without prior knowledge can use the VR tool and learn the first basics of computer graphics.

CCS Concepts

• **Computing methodologies** → *Virtual reality; Computer graphics*; • **Software and its engineering** → *Software prototyping*; • **Social and professional topics** → *Computer science education*; • **Applied computing** → *Interactive learning environments; E-learning*;

1. Motivation

Early on teachers and researchers investigated the technical support for teaching computer graphics (CG). [SWLR19] have created a review of all computer graphics education tools from 1991 to 2018. They were able to show that we are not only using many different tools but also that many of these tools were not originally developed for education. So far, no fixed standards for computer graphics education have been developed, even though there is a general trend towards web-based educational tools. This trend can also be observed in computer graphics [VPVG21], as well as another trend, which is currently being adopted in computer graphics: Mixed reality applications [SWLR19].

Virtual reality (VR) marks a point on the reality–virtuality continuum [MTUK95]. The potential of VR in teaching is manifold and many researchers showed positive effects on learning, e.g. [MGMADGM17, Pan09, PDS*19, SMBP19, ZKFL19, ANCT21]; nevertheless, an immediate comparison with non-VR variants is difficult to make because many factors come together. The question of *when best to use which technology* is being studied and is difficult to answer conclusively [SIM19]. Specific guidelines for the design of learning environments in VR are still not available [Fow15].

In addition to traditional approaches, such as interviews, learning analytics offer the opportunity to gain insight into the learning process and quantify the effectiveness of learning settings [GSP*20]. All definitions describing this field, "share an emphasis on converting educational data into useful actions to foster learning" [CDST12]. The integration of learning analytics into CG education

tools could be one step to fill a gap that not only [SWLR19] stated in the already mentioned review, but also [RH00] in a more generalised overview. The success of the pedagogy is not systematically evaluated in CG education tools and "more rigorous evaluations should be conducted to determine if the tools are effective for the purpose" [SWLR19].

In this paper, we present RePiX VR, an interactive theory-based VR application, which explains a central topic in introductory CG: the rendering pipeline [BWF17]. Learners can use it to interactively learn the different steps of the graphics pipeline and how images are generated from models. The presented open-source tool offers the possibility to investigate the understanding process of CG concepts and other influencing factors like motivation, spatial ability and the effects of feedback in more detail, while the integration of multi-modal learning analytics also provides opportunities for researchers, learners and teachers alike.

2. Related Work & Foundations

The computing curricula overview report 2020 designates virtual reality and 3D graphics as emerging trends in computing curricula. "Graphics and Visualization" is an element of the software fundamentals [AI20]. [AI13] assigns "affine and coordinate system transformations", "polygonal representations" and "the graphics pipeline" to basic rendering and fundamental graphics techniques. [BWF17] appropriately shows in a review of 20 curricula that these topics are covered by 95% of the CG courses.




Figure 1: Stage 0: A tutorial explaining the selection and locomotion techniques.

The amount of educational CG applications that have been developed for VR, (augmented reality) AR or mixed reality (MR) is still limited. [PVT08] have developed a platform called "Mental Vision", which combines VR and CG. They use modules and tutorials to facilitate the learning of the content. Unlike their approach, we have not used a CAVE and mobile devices, nor do we want to focus on teaching programming skills. Another related approach is the very recent web-based application presented by [VGV*20] in 2020 (and also in [VPVG21]), which aims to teach students a WebGL-based programmable graphics pipeline. Another project researching immersive VR is "GetiT - Gamified Training Environment for Affine Transformations" [OL19], where you can learn about affine transformations (ATs). [OL19] found indicators that VR is beneficial compared to the control group learning with desktop-3D. Regardless of the exact learning content, learning in VR can engage students and shift their role from passive to active learner [Pan09].

VR offers some possibilities and advantages. VR (especially in combination with learning analytics) promotes engagement of the learner with the learning material. VR can implement game-based learning or simulations or exergaming, e.g. [KMD*21]. In our case, VR offers the chance to make abstract domain knowledge interactive (compared to other more action-oriented applications, e.g. teacher training or hard disk changes in [ZKFL19]). Our hypothesis is that VR provides easy access to basic computer graphics topics, e.g., the rendering pipeline. Another point that motivated us is that VR gives a special insight into 3-dimensional learning topics. [AWF10] show there are indicators that the influence of control and active learning on learning outcomes was moderated by spatial ability, while active learning and control itself showed a positive effect on learning outcomes.

In already mentioned publications, there are discussions about the disadvantages and limitations of VR. VR with headset needs the equipment, there is the problem that some learners are affected by problems like simulator sickness [PDS*19]. An even further point is that VR is of course a very visual learning medium and hereby also possible difficulties can arise.

In summary, from the above examples and the discussions, it can be concluded that VR is suitable for computer graphics education and it supports a constructivist approach.

3. Didactic Conceptualization & Development

The systematic procedure we used to develop the open-source tool is described in another paper (actually under review). The didactical concept is based on an analysis of the learners, while the learning analytics concept is based on the investigation of requirements of different stakeholders, namely developers, researchers, students and teachers. Due to a lack of official guidelines, we followed common best-practice recommendations, e.g. [HL18]. These guidelines help to design effective VR applications, which use the potentials of VR and exploit the possibilities of constructivist learning theories as well as situated learning, active learning and creative learning. We used an iterative development process proposed by [Jer15]. Introductory courses at universities and interested high school students are the targeted audience.

The result of the systematic preparation was to design a guided tour for the rendering pipeline, in which learners accompanied by a (virtual) host, get to know the different parts and steps of a simplified version of the rendering pipeline, which is part of the course Learning Technologies at the RWTH Aachen University. This chosen model simplifies especially the *Rasterizer* step with only four simpler stages: *Rasterization*, *(Local) Lighting*, *Texturing* and *Visibility Test*. Fig. 2 shows the short overview, which will be explained in detail in section 4.

4. Technical Realization & Learning Analytics

The planning of the learning data analyses was based on the reference model for learning analytics by [CDST12], which provides an orientation framework to develop planning from different perspectives and helps to set weights for certain aspects. We choose "Experience API" (xAPI) as data format and "Learning Locker" as a compatible Learning Record Store for the data collection [xAPI18]. The xAPI specification allows describing user activities in human-readable statements formatted as a {actor, verb(action) and object(activity)} triplet, e.g. "Diana fixated pointable object", which together with the stored extension, contains the information that our subject Diana looked at the Stanford bunny's mesh for 3,4 seconds - we included eye tracking as one modality for learning analytics. We are not only collecting eye movement but also controller action and head movements, e.g. nodding and shaking. One advantage is that verbs and objects in the xAPI data format are strictly defined and uniquely identifiable. This way, we can make the collected data available to the research community in an interoperable format. Thus, the data can also be useful for comparison between different applications and for sharing it with other researchers.

We provide a web application called *Researcher Configuration Panel* that manages connections to the Learning Record Store (Learning Locker), Unity and various settings for data collection. VR clients (instances of our VR learning application) connect via web sockets to this panel and provide their trackable scene objects. The researcher selects relevant scene objects for tracking, they can start and stop tracking sessions on the panel and start the calibration, for the (optional) eye-tracking. In addition to the interaction data, we can collect eye-tracking data and initial gesture recognition of the HMD. The collection of multimodal learning data helps to provide a comprehensive picture of the learning processes [Wor18].




Figure 2: Short overview of realized Rendering Pipeline stages (Version Jan. 2022): (1) Application, (2) 3D Geometry, (3) 3D Transformation, (4) Clipping, (5) Rasterization, (6) (Local) Lighting, (7) Texturing, (8) Visibility Test and (9) Image

RePiX VR is developed for the HTC Vive Pro Eye. The studies were conducted using a Schenker Compact 15 laptop with an Intel Core i7-8750H CPU (2.20GHz), 64 GB RAM Memory and NVIDIA GeForce RTX 2070 and Windows 10 Education 64-bit. Unity (with SteamVR) was used as the development environment of the VR application, Microsoft Visual Studio 2019 for C# code and JetBrains WebStorm 2020.2 and NodeJS and Nuxt.js for web development handling the LA part.

RePiX VR: The Guided Tour

The result is a guided tour that leads a learner, accompanied by an avatar as host through the various steps of a simplified rendering pipeline in a large neutral environment (short demo: <https://youtu.be/U77hR7udyak>). The host is visible in Fig. 1 and some subfigures of fig.2, especially prominent in Fig. 2.7. The tour begins with a small tutorial, see Fig. 1, where the learner gets introduced to the controller triggers and buttons, its laser pointer to select scene objects (selection technique) and teleport action (locomotion technique). The tutorial is followed by the nine stages of the simplified rendering pipeline shown as previews in Fig.2.

1. *Application*: An introduction into the matrix components of objects in a 3D world. As example, a red cube transforms matrix by matrix (translation, rotation, scale) and results in its' *TRS matrix* (Fig. 2.1). After this demonstration, the learner can practise this knowledge by combining the matrices of a small solar system.

The matrices are coloured (e.g. the rotation matrix is green) to identify its' responsibility.

2. *3D Geometry*: The learner gets introduced to the basics of a 3D mesh of an object. For example, a Stanford Bunny mesh (without surfaces) rotates in front of the learner. Subsequently, the learner has to create the first triangle by placing some vertices (represented by small spheres) and finally connecting them into a triangle and more complex objects ((Fig. 2.2). The colour and position of the vertices and, of course, of the triangle are chosen by using the VR controller (a gradient colour picker is attached to it). The vertex indices are represented by a text label above each vertex.
3. *3D Transformation*: This stage explains the frustum-based projection and its' influencing factors (far plane, near plane, the field of view and aspect ratio). Thereafter, a semi-transparent, blue pyramid-shaped mesh appears, representing the frustum view. A few 3D objects inside this pyramid represent a 3D scene. The learner is able to move the planes (representing the far and near plane) along the z-axis by grabbing/selecting them with the (laser-)pointing ((Fig. 2.3). The values of the field of view and aspect ratio can be changed with sliders. Both factors define the width and height of the pyramid-shaped mesh.
4. *Clipping*: The clipping method of the rendering pipeline gets explained here. To illustrate the normalized device coordinates concept, a unit sized semi-transparent cube appears with transformed copies of the 3D objects of the previous stage repre-

senting the transformed scene. Primitives outside this unit-sized cube are demonstratively tinted cyan (Fig. 2.4). Those can be toggled to transparent by pressing a button on the VR controller. The transformation of the 3D objects depends, like in a real-world rendering pipeline, on factors of the frustum. Thus, in this stage, the learner is again able to see the results of the frustum-based transformation by interacting with the pyramid-shaped mesh. Transformed primitives outside the unit-sized cube get clipped (or tinted) in real-time and the learner can observe clipped primitives live (see Fig. 3).

5. *Rasterization*: A textual explanation provides the principle of the Rasterization stage. The explanation is supported by an image showing the detection of fragments (potential pixels) of 3D objects (Fig. 2.5).
6. *(Local) Lighting*: In this stage, the tour guide explains roughly the idea of Local-Lighting. For demonstration, the directional light source gets turned on and off. Afterwards, the learner is invited to seek for the light source in the sky. Further, the learner is invited to place flashlights and light bulbs with different colours to discover the influence of different light sources (Fig. 2.6).
7. *Texturing*: For the texturing stage the tour guide refers to a "naked" (without texture) copy of itself. Its texture is placed on a board next to it spanned on a UV coordinate system (Fig. 2.7). The learner can go nearer to the UV map and see the 2D mesh grid (UV map) lying on the texture.
8. *Visibility Test*: The textual explanations contain the concept of the z-buffer and its' responsibility to decide between visible fragments (pixels on screen) and not visible (discarded) fragments). To illustrate the final chosen fragments, the resulting image appears. It is placed where the unit-sized cube stood in the stage before (Fig. 2.8).
9. *Image*: The final stage offers to interact with the frustum again. The learners can not only see the result of the clipping stage, but they can also see the final rendered image of the transformed scene in front of the unit-sized cube (Fig. 2.9). By (laser-)pointing on this image plane, the image gets semi-transparent. By standing directly in front of the image, the learner can directly compare the image with its' transformed 3D geometry laid behind. Clipped primitives can be tinted to directly observe how it fits with the final image.

Overall every stage has explanations, animations and some provide VR interactions. For orientation, a panel is placed in front of the tour guide to provide an overview of all stages. Each stage has a button that can start the learning lesson of a stage. The learner can only select the already visited stages or the next one on the tour agenda. This way we ensure linear guidance without disallowing the exploration of several stages.

The interactive stages are designed to invite trial and exploration, following the approach of constructivist learning theories. The learner gets explanations about the possible interactions, receives small interaction tasks to get familiar with them. Finally, the tour guide invites to explore the stage interactions freely. This way the learner can practise the learned theory provided by the tour guide.

Depending on the task, further control panels are used, e.g. in the 3D Geometry stage. It is used to switch between functions such

as "place vertices", "build triangles" (connect vertices), "move vertices" and "remove vertices".




Figure 3: An early result from clipping stage.

5. Evaluation Methodology

In this paper, we describe two evaluations that took place at different points in the development. It is important to note that interactivity was increased in the different versions. Two larger-scale studies in 2020 and 2021 were conducted to improve the usability, the learning process and the learning analytics components. In addition, further smaller tests were conducted while developing. Due to the special circumstances in the academic years of 2020 and 2021, the user studies had to be done under strict safety precautions. Special admittance regulations to the university premises have been in place and our hygiene standards involved avoiding the repeated usage of a VR HMD as well as additional precautions to avoid the presence of multiple people in the institute. Eight persons participated in this first evaluation all in the VR condition. Ten VR users and fourteen desktop users took part in the second evaluation.

Each test session took between 35 and 75 minutes. All participants were between 25 and 40 years old. The first task was to explore the guided tour of the rendering pipeline (almost) independently. We provided help with questions about the operation of the controllers. The participants were instructed to communicate their thoughts while exploring. The instruction took place in the sense of the think-aloud method, followed by a free interview [ES84, GFA09]. A photograph taken in the experiment is shown in fig. 5. The stage is the first version of the *Application* stage.

5.1. First Evaluation

The first evaluation aimed to find the possible technical problems, to test the usability and to check the difficulty of the content. Four test subjects, who had limited prior knowledge, were tested in particular for usability and interactions. These four test subjects also had no experience with VR. To reduce the pressure, we assured the subjects that they did not need to understand every detail. In order to meet the required conditions our test persons were not only learners with no prior knowledge, we also asked four domain experts to

evaluate the guided tour. The four expert participants were experienced in user studies and VR, they provided feedback from different perspectives. They gave feedback concerning learning material, usability, didactics and the learning- and user experience.

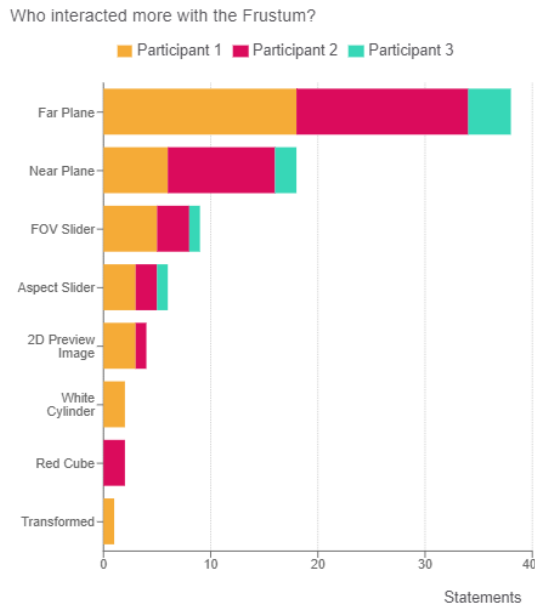


Figure 4: Visualisation of learning analytics data.

Four novices and four experts with different stakeholder perspectives gave feedback and confirmed the progress of the development and gave indications for the next iteration.

The Researcher Configuration Panel was evaluated by two of the above-mentioned persons since this will only be used by people with some prior experience anyway.

5.2. Second Evaluation

For the second evaluation, we tested learners, studying computer science, or a similar program. All participants took part in the master's course Learning Technologies at the RWTH Aachen University. Because of pandemic limitations, we created a second variant of the pipeline on short notice, which can be used from a desktop computer. In future iterations, this can be used for control groups or educators without time or resources to use the VR version. We collected interaction data in two platforms for the possibility of comparative evaluation. We will refer here in particular to the VR students, four of the ten had prior knowledge of computer graphics or virtual reality.

6. Results and discussion

Due to the number of participants, we analysed the logged interaction, gaze and motion data in an exploratory and descriptive manner without statistical reliability. Especially the evaluation of the think-aloud protocol and the interviews provided useful feedback for the already ongoing developments.

The user impressions were consistently positive. The virtual tour host was experienced as helpful. The application can be helpful to understand fundamental concepts of the rendering pipeline and the visualisations are very descriptive. The small video explanations, which demonstrate the next interaction in some stages, were helpful to prevent usability problems with approaching and orienting the objects, a problem for example described in [SK00]. The walkthrough questions explained in [SK00] could be answered positively for both evaluations.

Results of the pretests

The dialogue system had to be adjusted early on, during the pre-trial runs. In the very first version, it was equipped with autoplay. In particular, different prior knowledge does not allow to set a common pace for the dialogues, which is why they can now be controlled manually by each user. These results, among other findings, coincide with the various discoveries of researchers like J. Nielsen, which has led us to check everything again with heuristics for usability, see [Nie94].

Results of the first evaluation

We collected feedback in different categories. The learning material and content was perceived differently by the test persons. We were also able to get a lot of valuable remarks on the contents in order to further improve the content and visualisations in the next cycle.

During the final on-site conduct, an average of 4.500 xAPI statements were collected for each of the eight test persons in the Learning Record Store. Even though this data is not enough to give well-grounded predictions about the learning content from a didactical perspective, we were able to see trends and test first statistical methods. All available objects were selected at least once and the triangle building seems to be the most catchy interaction. The more complex interactions and visualisations were visited the most, which also fits with the users' verbal statements. The more interactive content was prepared, the more likely they were to spend more time there and to engage beyond the target. Fig. 4 shows a subset of the collected data (for better readability). The third person (turquoise) shows significantly lower interactions with Frustum than the other two persons. Some functions were only found or used by some persons, e.g. the turquoise person did not interact a single time with the resulting 2D image. Notable is that different VR experience levels during the studies had a clear impact on the behaviour of the subjects. While people without prior experience were more static and did not move and interact much, experienced users moved around the objects and tried different buttons. In total, five people were very eager to experiment and continued to stay in the stages even after completing the tasks, for example, to shape the frustum to exclude the bunny from the visible area.

Results of the second evaluation

In one case, we learned through the Think-Aloud protocol that it can be difficult for users with impaired vision to read the text boxes and that the robot's small movements can make reading difficult in such cases. We are currently working on a redesign of the scene that addresses this need. Reading was found to be exhausting by one



Figure 5: A test person watching the animated cube and the transformation matrix composition in the Application. From left to the right: Translation, Rotation and Scale multiplied into TRS matrix. Individual matrices were later coloured to allow better recognition.

more person and, in line with the principles of multimedia learning, will be supplemented by voice output in the next version [May01].

Students also want opportunities for interaction for the stages that are not yet interactive and provided suggestions directly once they explored the content and stage animations. Examples are to make the light source movable to change the shadow cast or a connection between the UV map and the model without texture (draw on the model and see the result on the map), which are already in development. Another point is that students very quickly wish they could explore the scene from all angles. More teleportation points could invite a closer examination of the 3D scene. In order to investigate these points more systematically, the positions in the desktop variant will be logged in the future, to compare them with the VR data and we allow free teleportation for the study of the stages.

We made an unintended observation at the end of the 2021 summer semester. The data collection took place at the beginning of the semester. In the exams for the lecture, many students were able to recollect details from the application to provide them as examples for exam questions. This indicates a positive learning outcome, also after a longer period, which has not been explicitly researched yet (this was not a computer graphics introduction).

We collected around 147k statements from 24 participants in the Learning Record Store. The evaluations show that VR users spent a longer overall time within the application (29 minutes on average). Desktop users spent an average of 9 minutes less and showed less variance in usage time than VR users. An initial evaluation of head movements shows that users nod significantly more often than shake their heads, a gesture that we could interpret as an indicator of understanding under certain conditions, e.g. when they have previously read a text and solved a task. Shaking the head is currently not an indicator that can be evaluated without error, since some stages invite you to move your head from right to left several

times, e.g. stage 5 - *Rasterization*, when you follow the lines from the camera to the objects.

The desktop version had a mixed reception. Some interactions are easier to use, while others clearly show the advantages of the VR application. For example, desktop users created fewer triangles on average in the *3D Geometry* stage because they are more difficult to place (four vs ten triangles). Another factor besides the usability could be the environment, which is more stimulating for the users of the VR application. On the one hand, the environment itself might be more motivating to experiment with the triangles, and on the other hand, the overall learning context is different, i.e. only these students were asked to Think Aloud in the VR condition.

The data also show that the stages are used for different lengths of time depending on the degree of interactivity, but also that the particularly three-dimensional tasks are used significantly more intensively by the VR learners. VR learners spend on average six times as long in the geometry stage and even 7 times as long in the Image Stage. It is particularly interesting that they are much longer in the Image Stage, since comparatively little new content is explained there. Already learned concepts, e.g. the interactions with the frustum are repeated and the changes in the resulting image are observed. The stage where the users of the desktop version took longer was the tutorial. They also needed a little longer in the rasterization stage. When looking at the movement data, it is noticeable that the 3D transformation animated both groups to move in 3D space.

One unanswered question that emerged from the log data analysis is why users interacted with the "far plane" significantly more in both evaluations, an exemplary representation of this from the first evaluation can be found in Fig. 4.

7. Discussion

The first part of this section deal with the results of the evaluations presented here that have already been incorporated into the iterative development. Subsequently, larger changes and other future work are explained as well as developments that still have to be discussed with more stakeholders.

7.1. Integrated results

Particularly interesting was the experts' wish to expand the first stage further and to let learners know more about and experiment with the creation of the model/TRS matrix in the *3D Geometry* stage. These findings are also in line with [HL18], who request comprehensive interactivity in VR, and supports [OL19]. The result of this request can already be seen in the description of the learning application (section 4). The latest version includes a solar system where the learners can manipulate the matrices themselves.

The next goals, currently in the *Make* phase of the iterative design cycle of [Jer15], are the extension of the *Lighting* and the *Texturing* stage. These stages are also to be designed interactively matching the results of our evaluations.

7.2. Future directions

The idea to extend the application to participatory VR in the future (guests experience the tour in a shared VR environment via interaction), see [VGCL19].

[HL18] argues, that VR also offers great potential for cooperative and collaborative learning. We will address this point in one of the next versions. Participatory VR - guests experience the tour in a shared VR environment via interaction [VGCL19] - needs further didactical decisions. How can we adapt the tasks so that they are collaboratively actionable?

The usability has been iteratively improved, but there are still points we need to improve. For example, the colour of the vertices can be changed via the trackpad of the controller. This functionality was not immediately discovered by everyone. Another issue is the pointing pattern, which we used as an interaction technique. It is experienced very intuitively, but for some interactions, like adjusting the far and near plane the interaction was not accurate enough. To improve these rather small points (as well as the usability of the desktop version) short development cycles with quick tests will be used.

There is no shortage of remaining research questions: Does the VR application take longer because the learners engage more intensively with the material? Do they learn more as a result? To what extent might the longer processing time correlate with enthusiasm, which (at best) has a moderating effect on the overall course learning outcomes? Another one, do we help specific groups of learners more than others, e.g. learners with a low visual-spatial ability like [GCSNAM19] indicate?

8. Conclusion

This paper presented the promising results of a guided tour, which introduces the concepts of the rendering pipeline interactively. Both

evaluations show that an iterative approach as suggested in [Jer15] leads to user-friendly VR applications.

But as motivating as our preliminary results are, we need to be tested with a broader audience. Getting more participants is complicated by points such as the comparatively high acquisition costs and, for example, Cybersickness [LaV00, SIM19]. Another issue we need to resolve is the difficulty of generating truly comparable control groups for some questions in order to accurately determine the effects of the VR learning environment.

Nevertheless and regardless of our direction, we hope that our decision to provide our environment as an open-source project will encourage educators and contributors to either use the application or to cooperate with us. From a teaching perspective, the modular development of the individual stages leaves us with many options that we can discuss. E.g. currently a more advanced version of the light and shadow calculation is developed.

As [FP13] showed, diverse audiences are an issue in introductory graphics courses. One approach is to not only begin with the fundamentals of the field, as the rendering pipeline but also by paying special attention to the concepts in the exams. This way they were able to distribute grades without correlation to the heterogeneous pre-knowledge of their students. The presented tool might help in that regard by assisting future students, i.e. by providing more information on foundations "on-demand" as well as going more in-depth if requested (like planned with the light and shadow stage). Thus, an encouraging self-regulated learning approach could create adaptive learning experiences levelling out differences in prior knowledge.

References

- [AI13] ACM ASSOCIATION FOR COMPUTING MACHINERY, IEEE COMPUTER SOCIETY (Eds.): *Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science*. ACM, Inc, New York, 2013. doi:10.1145/2534860. 1
- [AI20] ACM ASSOCIATION FOR COMPUTING MACHINERY, IEEE COMPUTER SOCIETY (Eds.): *Computing Curricula 2020: Paradigms for Global Computing Education*. ACM, 2020. doi:10.1145/3467967. 1
- [ANCT21] ASAD M. M., NAZ A., CHURI P., TAHANZADEH M. M.: Virtual Reality as Pedagogical Tool to Enhance Experiential Learning: A Systematic Literature Review. *Education Research International* 2021 (2021), 1–17. doi:10.1155/2021/7061623. 1
- [AWF10] AI-LIM LEE E., WONG K. W., FUNG C. C.: How does desktop virtual reality enhance learning outcomes? A structural equation modeling approach. *Computers & Education* 55, 4 (2010), 1424–1442. doi:10.1016/j.compedu.2010.06.006. 2
- [BWF17] BALREIRA D. G., WALTER M., FELLNER D. W.: What we are teaching in Introduction to Computer Graphics. In *EG 2017 - Education Papers* (2017), Bourdin J.-J., Shesh A., (Eds.), The Eurographics Association. ISSN: 1017-4656. doi:10.2312/eged.20171019. 1
- [CDST12] CHATTI M. A., DYCKHOFF A. L., SCHROEDER U., THÜS H.: A reference model for learning analytics. *International Journal of Technology Enhanced Learning* 4, 5/6 (2012), 318. doi:10.1504/IJTEL.2012.051815. 1, 2
- [ES84] ERICSSON K. A., SIMON H. A.: *Protocol analysis: Verbal reports as data*. Protocol analysis: Verbal reports as data. The MIT Press, Cambridge, MA, US, 1984. Pages: 426. 4

- [Fow15] FOWLER C.: Virtual reality and learning: Where is the pedagogy?: Learning activities in 3-D virtual worlds. *British Journal of Educational Technology* 46, 2 (2015), 412–422. doi:10.1111/bjjet.12135. 1
- [FP13] FAIRÉN M., PELECHANO N.: Introductory Graphics for Very Diverse Audiences. In *Eurographics 2013 - Education Papers* (2013), Bourdin J.-J., Cerezo E., Cunningham S., (Eds.), The Eurographics Association. ISSN: 1017-4656. doi:10.2312/conf/EG2013/education/009-010. 7
- [GCSNAM19] GONZÁLEZ CAMPOS J. S., SÁNCHEZ-NAVARRO J., ARNEDO-MORENO J.: An empirical study of the effect that a computer graphics course has on visual-spatial abilities. *International Journal of Educational Technology in Higher Education* 16, 1 (2019), 41. doi:10.1186/s41239-019-0169-7. 7
- [GFA09] GROSSMAN T., FITZMAURICE G., ATTAR R.: A Survey of Software Learnability: Metrics, Methodologies and Guidelines. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2009), CHI '09, ACM, pp. 649–658. doi:10.1145/1518701.1518803. 4
- [GSP*20] GEDRIMIENE E., SILVOLA A., PURSIAINEN J., RUSANEN J., MUUKKONEN H.: Learning analytics in education: Literature review and case examples from vocational education. *Scandinavian Journal of Educational Research* 64, 7 (2020), 1105–1119. doi:10.1080/00313831.2019.1649718. 1
- [HL18] HUANG H.-M., LIAW S.-S.: An Analysis of Learners' Intentions toward Virtual Reality Learning Based on Constructivist and Technology Acceptance Approaches. *International Review of Research in Open and Distributed Learning* 19, 1 (2018), 91–115. 2, 7
- [Jer15] JERALD J.: *The VR Book: Human-Centered Design for Virtual Reality*. Association for Computing Machinery and Morgan & Claypool, 2015. 2, 7
- [KMD*21] KELLER N., MCHENRY N., DUNCAN C., JOHNSTON A., WHITTLE R. S., KOOCK E., BHATTACHARYA S. S., DE LA TORRE G., PLOUTZ-SNYDER L., SHEFFIELD-MOORE M.: Augmenting Exercise Protocols with Interactive Virtual Reality Environments. In *2021 IEEE Aerospace Conference (50100)* (2021), IEEE, pp. 1–13. 2
- [LaV00] LAVIOLA J. J.: A discussion of cybersickness in virtual environments. *ACM SIGCHI Bulletin* 32, 1 (2000), 47–56. doi:10.1145/333329.333344. 7
- [May01] MAYER R. E.: *Multimedia Learning*. Cambridge University Press, 2001. URL: <https://www.cambridge.org/core/books/multimedia-learning/E9595926786F5DEA326A3774D2F23DB2>, doi:10.1017/CBO9781139164603. 6
- [MGADM17] MARTÍN-GUTIÉRREZ J., MORA C. E., AÑORBE-DÍAZ B., GONZÁLEZ-MARRERO A.: Virtual technologies trends in education. *EURASIA Journal of Mathematics, Science & Technology Education* 13, 2 (2017), 469–486. Publisher: EURASIA. doi:10.12973/eurasia.2017.00626a. 1
- [MTUK95] MILGRAM P., TAKEMURA H., UTSUMI A., KISHINO F.: Augmented reality: A class of displays on the reality-virtuality continuum. In *Telem manipulator and Telepresence Technologies* (1995), Das H., (Ed.), SPIE, pp. 282–292. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=981543>, doi:10.1117/12.197321. 1
- [Nie94] NIELSEN J.: Enhancing the explanatory power of usability heuristics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 1994), CHI '94, Association for Computing Machinery, pp. 152–158. doi:10.1145/191666.191729. 5
- [OL19] OBERDÖRFER S., LATOSCHIK M. E.: Knowledge encoding in game mechanics: Transfer-oriented knowledge learning in desktop-3d and VR. *International Journal of Computer Games Technology* 2019 (2019), 7626349. Publisher: Hindawi. doi:10.1155/2019/7626349. 2, 7
- [Pan09] PANTELIDIS V. S.: Reasons to Use Virtual Reality in Education and Training Courses and a Model to Determine When to Use Virtual Reality. *Themes in Science and Technology Education* (2009), 12. 1, 2
- [PDS*19] PAPANASTASIOU G., DRIGAS A., SKIANIS C., LYTRAS M., PAPANASTASIOU E.: Virtual and augmented reality effects on K-12, higher and tertiary education students' twenty-first century skills. *Virtual Reality* 23, 4 (2019), 425–436. doi:10.1007/s10055-018-0363-2. 1, 2
- [PVT08] PETERNIER A., VEXO F., THALMANN D.: The mental vision framework - a platform for teaching, practicing and researching with computer graphics and virtual reality. In *Transactions on Entertainment I*, Pan Z., Cheok A. D., Müller W., El Rhalibi A., (Eds.), Lecture Notes in Computer Science. Springer, 2008, pp. 242–260. doi:10.1007/978-3-540-69744-2_19. 2
- [RH00] REEVES T. C., HALL A.: Enhancing the Worth of Instructional Technology Research through "Design Experiments" and Other Development Research Strategies. *Proc. International Perspectives on Instructional Technology Research for the 21st Century* (2000), 15. 1
- [SIM19] STOJŠIĆ I., IVKOV-DŽIGURSKI A., MARIČIĆ O.: Virtual Reality as a Learning Tool: How and Where to Start with Immersive Teaching. In *Didactics of Smart Pedagogy*, Daniela L., (Ed.). Springer International Publishing, Cham, 2019, pp. 353–369. doi:10.1007/978-3-030-01551-0_18. 1, 7
- [SK00] SUTCLIFFE A. G., KAUR K. D.: Evaluating the Usability of Virtual Reality User Interfaces. *Behaviour & Information Technology* 19, 6 (2000), 415–426. doi:10.1080/014492900750052679. 5
- [SMBP19] SCRIVNER O., MADEWELL J., BUCKLEY C., PEREZ N.: Best Practices in the Use of Augmented and Virtual Reality Technologies for SLA: Design, Implementation, and Feedback. In *Teaching Language and Teaching Literature in Virtual Environments*, Carrió-Pastor M. L., (Ed.). Springer, Singapore, 2019, pp. 55–72. doi:10.1007/978-981-13-1358-5_4. 1
- [SWLR19] SUSELO T., WÜNSCHE B. C., LUXTON-REILLY A.: Technologies and Tools to Support Teaching and Learning Computer Graphics: A Literature Review. In *Proceedings of the Twenty-First Australasian Computing Education Conference* (New York, NY, USA, 2019), ACE '19, Association for Computing Machinery, pp. 96–105. doi:10.1145/3286960.3286972. 1
- [VGCL19] VELHO L., GIANNELLA J., CARVALHO L., LUCIO D.: VR TOUR: GUIDED PARTICIPATORY META-NARRATIVE FOR VIRTUAL REALITY EXPLORATION. *Revista GEMInS* 10, 2 (2019), 122–140. 7
- [VGV*20] VITSAS N., GKARAVELIS A., VASILAKIS A.-A., VARDIS K., PAPAIOANNOU G.: Rayground: An Online Educational Tool for Ray Tracing. In *Eurographics 2020 - Education Papers* (2020), Romero M., Sousa Santos B., (Eds.), The Eurographics Association. ISSN: 1017-4656. doi:10.2312/eged.20201027. 2
- [VPVG21] VASILAKIS A. A., PAPAIOANNOU G., VITSAS N., GKARAVELIS A.: Remote Teaching Advanced Rendering Topics Using the Rayground Platform. *IEEE Computer Graphics and Applications* 41, 5 (2021), 99–103. doi:10.1109/MCG.2021.3093734. 1, 2
- [Wor18] WORSLEY M.: Multimodal learning analytics' past, present, and potential futures. In *CrossMMLA@LAK* (2018). 2
- [xAPI18] Documentation of experience api (or xapi): a specification for learning technology, 2018. (visited on 18/01/2022). URL: <https://xapi.com/>. 2
- [ZKFL19] ZENDER R., KNOTH A. H., FISCHER M. H., LUCKE U.: Potentials of Virtual Reality as an Instrument for Research and Education. *i-com* 18, 1 (2019), 3–15. Publisher: De Gruyter Section: i-com. doi:10.1515/icom-2018-0042. 1, 2