

PerSiVal: On-Body AR Visualization of Biomechanical Arm Simulations

Xingyao Yu^{1,4}, David Rosin^{2,4}, Johannes Kässinger^{3,4}, Benjamin Lee^{1,5}, Frank Dürr^{3,4}, Christian Becker^{3,4}, Oliver Röhrle^{2,4}, Michael Sedlmair^{1,4},

¹ Visualization Research Center (VISUS), University of Stuttgart, 70569, Germany,

² Institute for Modelling and Simulation of Biomechanical Systems (IMSB), University of Stuttgart, 70569, Germany,

³ Institute for Parallel and Distributed Systems (IPVS), University of Stuttgart, 70569, Germany

⁴ Stuttgart Center for Simulation Science - SimTech, University of Stuttgart, 70569, Germany

⁵ Global Technology Applied Research, JPMorganChase, NY, 10017, United States of America

Abstract—In this work, we explore different combinations of techniques for an interactive, on-body visualization in augmented reality (AR) of an upper arm muscle simulation model. In terms of data, we focus on a continuum-mechanical simulation model involving five different muscles of the human upper arm, with physiologically realistic geometry. In terms of use cases, we focus on the immersive illustration, education, and dissemination of such simulation models. We describe the process of developing six on-body visualization prototypes over the period of five years. For each prototype, we employed different types of motion capture, AR display technologies, and visual encoding approaches, and gathered feedback throughout outreach activities. We reflect on the development of the individual prototypes and summarize lessons learned of our exploration process into the design space of situated on-body visualization.

Biomechanical simulation plays a crucial role in the modern scientific exploration of biological systems and is widely used in fields such as sports science, biomechanical engineering, and rehabilitation medicine [1]. Computer simulation allows for investigations into the human body's inner workings that, due to ethical concerns, are preferable to invasive procedures. These simulations have greatly enhanced our understanding of human motor control, injury, microscopic and macroscopic tissue biomechanics, and the design of, for example, prostheses.

Simulating parts of the human body in detail is a complex task and often demands substantial computational resources, as is the case with musculature. When simulating the movement of muscle tissue and its geometry, one method is to use forward simulation with muscle activation level (MAL) as input. MAL represents the degree to which muscles are activated by the central nervous system, thereby reflecting the extent of muscle contraction and force output. Continuum-

mechanical musculoskeletal system models provide a framework to describe these processes in 3D space, using prior knowledge about the material properties, structure, and morphology of muscles and joints. This enables a physics-based simulation of the muscles as a 3D organ, in combination with the finite element (FE) method. However, with the granularity needed for physiological geometries, this approach is computationally expensive even by high-performance computing standards [2]. Moreover, when multiple muscles act on the same joint, various MAL combinations may lead to the same simulated joint angle [3]. Optimization is therefore needed to determine the optimal MALs for a given joint angle based on predefined criteria, requiring frequent evaluation of an already expensive model. In 2018, however, Valentin et al. [4] achieved real-time optimization of arm muscle activation states by using gradient-based B-spline optimization techniques on sparse grids. This served as the foundation for fast 3D biomechanical simulations, further making the real-time visualization of continuum-mechanical models such as 3D human musculature feasible.

To capitalize on this newfound possibility, we at the Cluster of Excellence for Data-Integrated Sim-

ulation Science (EXC SimTech)¹ formed an interdisciplinary team of researchers from three different academic backgrounds: biomechanical simulation, distributed systems, and visualization. The objective was to make the simulation and visualization of arm biomechanics both accessible and flexible such that it can be used by anyone, anytime, and anywhere, and even on mobile devices such as smartphones, tablets, and mobile head-mounted displays (HMDs). In other words, we sought to achieve **Pervasive Simulation and Visualization** (*PerSiVal*). Through this, we aim to make biomechanical simulation data — traditionally confined to high-performance computing devices and professional settings — accessible to a broader user base. At the same time, we seek to enhance users' understanding of biomechanical processes, particularly that of muscle activity, in scenarios such as physiotherapy and specialized educational environments.

So far, we have demonstrated advancements in both **biomechanical simulation** and **distributed systems** over the past five years. In particular, we made it possible to use a combination of surrogate modeling and distributed processing techniques to predict arm muscle deformation based on movement data in real time [5] and to have these run on resource-constrained mobile devices including the Apple iPad, iPhone, and Microsoft HoloLens via distributed processing [6], [7], [8]. These advancements have opened the door for a broader audience to access biomechanical simulations. However, many of these users may lack the domain-specific expertise necessary to effectively interpret and understand traditional biomechanical visualizations, which are typically derived from post-processed data. As a result, there is a need for a novel, real-time visualization framework that can run on resource-constrained devices, enabling a more intuitive and accessible interpretation of biomechanical data.

To address this, we visualize the corresponding biomechanical simulation data directly on the human arm using *Augmented Reality* (AR), creating an **on-body visualization**. By *on-body visualization*, we refer to using the moving human body as the physical referent for the visualization, thus forming a unique type of *situated visualization* (SV) [9]. In this context, the spatial alignment between the biomechanical data and the human body enhances the comprehensibility of the information and fosters a more immersive and engaging learning experience. However, despite its potential, the technical feasibility of delivering real-time biomechanical simulations through AR on resource-

constrained devices has yet to be fully explored. Key challenges include determining how to efficiently acquire motion capture data on mobile devices to serve as input for biomechanical simulations, balancing the trade-off between simulation efficiency and visualization quality, and selecting appropriate devices to establish a functional setup. Addressing these issues is critical for realizing the practical application of on-body AR biomechanical visualizations in everyday environments.

In this paper, we describe our parallel efforts toward developing an AR application for on-body visualization that takes a user's motion as input and outputs a visualization of biomechanical arm simulations in real-time, representing the third pillar of *PerSiVal* project. To establish a foundational understanding of the design of on-body AR visualization of upper arm biomechanical simulations, we designed and implemented six prototypes — each a different instantiation of the vast design space. For each prototype, we elicited feedback from simulation scientists as well as the public through outreach activities, demos, and discussions. Through this, we sought to understand the benefits and tradeoffs of five key design dimensions: motion capture solution, display type, arm representation, viewing perspective, and data encodings. We describe our process and learnings from each of the six prototypes and summarize our findings and reflections, which may help guide the design and development of future on-body AR visualizations of biomechanical simulations. To the best of our knowledge, our work is the first to integrate human motion capture, realistic biomechanical simulation, and on-body visualization into AR devices in real-time.

Background: Simulation Data and Processing

In the following, we summarize the generation and processing of the simulation data that feeds into our on-body AR visualizations. Details on biomechanical modeling and distributed algorithms can be found in previous publications [5], [6], [7], [8].

Biomechanical simulation, at a high level, provides two direct data sources that we intend to visualize: muscle activation levels (MALs) and muscle geometry. Physical movement data is used to derive the MALs (Step 1), which can then be used to calculate the muscle geometry in each simulation cycle (Step 2). This latter stage, depending on its complexity, may be computationally expensive, and so it may optionally be offloaded for distributed processing (Step 3).

¹ <https://www.simtech.uni-stuttgart.de/>

Step 1: From Movement Data to MALs

Movement data of human bodies generally includes joint angles that can be used to calculate MALs [10], and is a commonly used method in bio-robotics. Like Valentin et al. [4], we used an optimization based on a sparse grid surrogate using B-splines, to determine a unique connection between joint angle and MALs. However, this optimization was initially too computationally expensive for a mobile device, and therefore was not suitable for a pervasive real-time visualization. To address this, we ran this optimization offline and exported all possible pairs between joint angles and MALs into a training dataset for a deep-learning model [5]. As the end result, we receive activation values for each muscle in the simulation that range from 0% to 100% contractile force.

Step 2: From MALs to Muscle Geometry

To then generate the muscle geometry based on the obtained MALs, we used a continuum-mechanical simulation with a finite element (FE) model to calculate muscle deformation. Our model focuses on the human upper arm, and consists of three bones (humerus, radius, and ulna), five muscles (biceps, triceps, brachialis, brachioradialis, and anconeus), and one mechanical degree of freedom at the elbow. However, this only yields results corresponding to a limited number of sampled combinations of MALs. We, therefore, interpolated the results using a sparse grid surrogate, which not only increases the number of available samples but also finds MALs optimal for reaching a given joint angle. To ensure real-time performance across this series of simulations, we employed a deep learning surrogate [5] which is trained on the MALs and muscle geometry. The result of this step is a 3D mesh of vertices that form the shape of the arm and its muscle deformation, which is dependent on the input MALs.

Step 3: Distributed Processing

Distributed processing can be employed as an optional step to ensure pervasive simulation on mobile devices. The muscle geometry can be reduced to a representative mesh node subset using a genetic optimizer, and the deep learning architecture for muscle geometry can then be split using a distributed systems approach [6]. In this manner, the majority of mesh computation will happen on a server. The server then transmits the optimal subset of the mesh nodes to the local mobile device, which uses this data to reconstruct and render the muscle geometry mesh. Additionally, we utilize forecasting methods that help reduce the

impact of latency and potential sample loss (e.g., from motion capture) [7].

PerSiVal: Pervasive Simulation and On-body AR Visualization

Our on-body visualization allows people to make use and sense of biomechanical simulation — the third pillar of *PerSiVal*, and the main focus of this paper. Our overarching goal was to use pervasive mobile computing devices, most notably AR, to visualize this simulation data of muscle activation level and muscle geometry directly on the body of a human user.

In this section, we first list our design goals, which we identified early on in the project and which helped to narrow our scope in our design explorations. We then describe the technical pipeline for how our biomechanical simulation functions and is used to produce the on-body AR visualizations. Throughout this pipeline, we also highlight the components wherein five key design choices need to be made, which would likely affect the usability and effectiveness of the visualization. These design choices provided the structure that guided our exploration into the design space of on-body AR visualization via our six prototypes, as described in the next section.

Design Goals

As described earlier, the *PerSiVal* team was initially formed as an interdisciplinary group of researchers at EXC SimTech from three domains: biomechanical simulation, distributed systems, and visualization (i.e., the authors). Because of this, it was necessary for us to clarify the requirements, design goals, and technical feasibility of on-body AR visualization among ourselves — both early and often. During this process, we also chose to interview a physiotherapist to understand the possible applications and use cases of *PerSiVal*. These resulted in six design goals which served to motivate our work and refine our scope, which are as follows.

G1 On-body augmented reality. AR is capable of rendering muscle geometry in 3D and can overlay it directly above the corresponding positions of a person. This helps establish a direct relationship between the person's body and the biomechanical simulation [9].

G2 Virtual arm representation. Biomechanical simulation provides both muscle geometry and MAL data. A virtual arm representation is needed to facilitate an easier understanding of this data, with each muscle showing its own geometry and activation level. Virtual muscles can also facilitate interactive visualizations,

such as by color-coding MALs [10] that update in real-time.

G3 High level of realism. An extension of G2, a highly realistic virtual arm should be used to allow users to accurately discern the shape and deformation of individual muscles.

G4 Accurate and stable motion capture. Motion capture allows for human bodies to be tracked in real-time, which can serve as realistic limb movements for the biomechanical simulation to acquire finer muscle activity data [10]. It also should be accurate and stable to facilitate the positioning of the visualization on-body in AR, as per G1.

G5 Pervasiveness. The visualization should be usable in a wide range of different scenarios and applications. By extension, the visualization and the required simulation should be capable of running on a single mobile device.

G6 Understandability. The application and its visualization should be easily usable and understood by anyone, especially non-experts who may be viewing biomechanical simulations and visualizations for the first time.

Pipeline & Design Choices

Figure 1 illustrates the pipeline of how our prototypes, at their core, compute and render biomechanical simulations in real-time. Throughout this pipeline lie several design choices that influence the understandability and effectiveness of the resulting visualization, which we later explored as variations in each individual prototype.

First, some movement data is needed so that it can be fed into the simulation. *Motion capture* is therefore needed to track the arm movements [10] of a subject of observation (SOB) — that is, a human person. These arm movements are then inputted into the biomechanical simulation, which, as mentioned earlier, is a deep learning surrogate [5] that is lightweight enough to run locally on the mobile device. Note that a part of the simulation can be configured to run on a distributed system in case it is too computationally expensive to do so locally. This simulation returns the MALs and muscle geometry, which, along with the original movement data, can be used as the *data mapping* to whichever visual channels are deemed necessary. These visual channels are then applied to and rendered on some virtual *arm representation*, thus forming the visualization. The arm may either be static (i.e., a pre-defined 3D model) or dynamic (i.e., the mesh is simulated and processed in real-time). This rendering is finally projected onto an *AR display*, which may vary in its *viewing perspective*.

These five design choices — motion capture, data mapping, arm representation, display, and viewing perspective — were what we explored throughout the six prototypes in the *PerSiVal* project. Note that the other components relating to the underlying biomechanical simulation remained conceptually the same throughout the development process.

Technical Prototypes

Our exploration into on-body AR visualization of biomechanical arm simulations was conducted through the creation of six individual prototypes (P1~P6). Each prototype explores a different configuration of the five aforementioned design choices, seeking to understand their benefits and tradeoffs and thus provide a more holistic understanding of this design space.

As these prototypes were developed sequentially over the span of five years, each prototype also became increasingly targeted toward meeting our design goals. In particular, P1 and P2 act as foundational explorations into on-body AR biomechanical visualization, P3 and P4 apply these learnings to physiotherapy as a real-world application, and P5 and P6 move closer toward a true pervasive simulation and visualization. This process was guided at each stage with outreach activities, where we sought feedback from our expert physiotherapist, simulation scientists, and members of the general public. It should be noted, however, that these prototypes were developed in parallel with the other two pillars of *PerSiVal* — biomechanical simulation and distributed systems — and thus, our prototypes vary based on the presence or absence of this real-time *simulation data*. All prototypes were created using the Unity3D Game Engine, which is a popular software for creating AR applications.

P1: Foundational Exploration of Concepts

At the project's outset in 2019, we first sought to validate *PerSiVal* as a proof-of-concept prototype. The goal was, therefore, to use a straightforward visualization with a basic motion capture setup.

Materials

The real-time pervasive simulation had yet to be developed this early into the project, which meant we could not run any biomechanical simulations directly on the mobile device. We instead pre-processed and exported simulated MALs as a hash table which was indexed by the elbow angle. Since we also did not have the muscle geometries, we used a static arm model which featured only the biceps, triceps, and a few upper-arm muscles

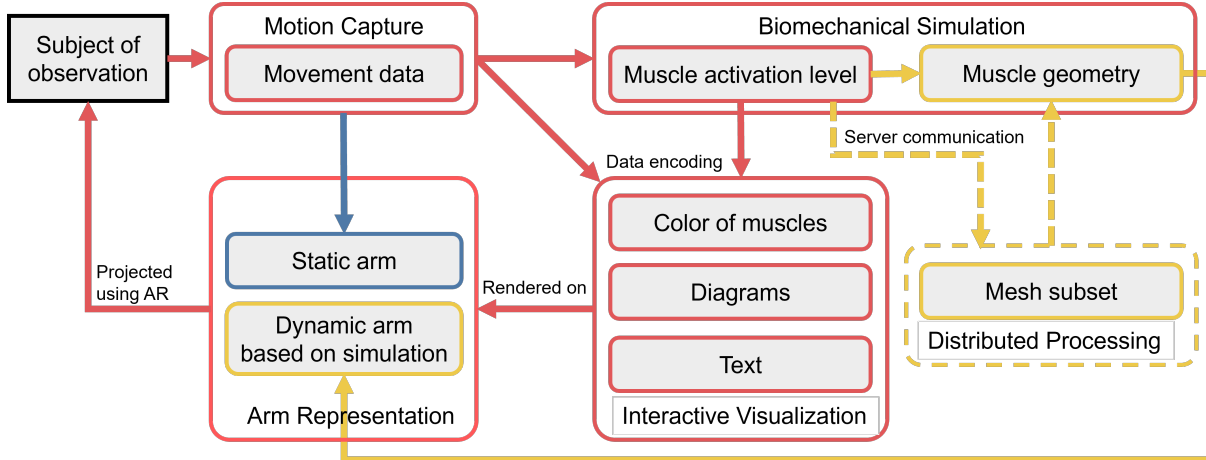


FIGURE 1. The framework of PerSiVal system. The red parts represent the basic pipeline for the system: (1) the movement data is captured by a real-time MoCap system and then fed into a biomechanical simulation to obtain MALs. (2) the MAL will be live encoded into the color of muscles. Text and diagrams can be used, e.g., to show the history of MALs or to additionally encode joint angles. (3) the arm representation could be either static (in blue), which is the static arm model directly manipulated by the movement data, or dynamic (in yellow), which is derived from simulations and optionally lightweight through the distributed process.

(G2). While this model is capable of elbow flexion, it could not depict muscle deformation.

Implementation & Features

We chose the Microsoft HoloLens 1 as it was the state-of-the-art AR HMD back in 2019. We used Vuforia to facilitate motion capture of the SOB, wherein three image targets had to be affixed to their arm joints: wrist, elbow, and shoulder. Based on the positions of these three joints, the elbow joint angle is calculated, and its corresponding MAL is retrieved from the hash table, which is then encoded onto the corresponding muscles using color [10]. This is shown in Figure 2. The arm model is scaled, and the elbow angle is adjusted for the whole to better align with the actual arm of the SOB. Due to the low FOV of the HoloLens 1, looking at the visualizations on your *own* body in first-person was impractical, meaning that the SOB needed to be a different person from the person wearing the HMD (i.e., third-person).

Outreach Activity

We showcased P1 as a demonstration at the SimTech Status Seminar in 2019 — an annual event that brings together about 150 members of EXC SimTech. In each seminar, we were visited by researchers who have expertise in simulation technology and received feedback, comments, and suggestions from approximately 10 of them. During our session that year, it was quickly apparent that the poor tracking performance of Vuforia




FIGURE 2. P1 and P2 share a similar design: the movement of the SOB is captured by the image target, the MAL is color-coded on the muscles, and the static arm representation provides no deformation of muscles.

made it challenging for people to even look at the visualization. In particular, there was noticeable latency causing the virtual and real arms to be misaligned when moving, and tracking was frequently lost in the low indoor lighting conditions.

P2: Improving the Simulation and Setup

While the COVID pandemic impacted internal team communication and development efforts, the increasingly available HoloLens 2 provided a clear opportunity for a second prototype. P2 was, therefore, a more fully realized version of P1, improving on the latter's technical limitations.

Materials

By this point in time, we had made it possible to run the simulation of the MAL using a lightweight neural network, which was capable of running on mobile devices in real-time. Thus, the color-coded MALs were now based on simulation data and not on the pre-computed hash table.

Implementation & Features

We upgraded to the Microsoft HoloLens 2, which is overall a much more powerful AR HMD than its predecessor. This hardware upgrade also made it feasible to run the aforementioned neural network to compute MALs. We also replaced Vuforia with ARToolkit to track the arm joints of the SOB, which, from our testing, performed much better when tracking image targets. Otherwise, the visualization remained the same as in P1 and Figure 2.

Outreach Activity

P2 was presented in the SimTech Status Seminar in 2021. As expected, the use of ARToolkit made the prototype much more reliable and reduced the risk of tracking loss. Most people we spoke to found it novel to be able to run and see the results of biomechanical simulations on a resource-constrained mobile device like the HoloLens 2, expressing anticipation and enthusiasm for the future development of *PerSiVal*. This excitement was made especially clear as we were awarded Best Poster² by all participants of the seminar.

P3 & P4: Physiotherapy Scenarios

After having confirmed the feasibility of *PerSiVal*, our attention shifted to the use case of physiotherapy. In doing so, we also considered the other possible configurations of design choices in terms of both the technical setup and the visualizations shown in AR. This was achieved by developing two different prototypes that are tailored toward two contrasting users: the patient with P3, and the physiotherapist with P4.

Materials

As we still did not have the real-time muscle geometry simulation, we opted to use a more detailed anatomical masculine arm model³ to improve the realism of the visualization. While still static in nature, it carries more muscles than the one previously used in P1 and P2. To further enhance the model, we used Autodesk 3ds Max to add skeletal animations, showing simple muscle deformations in response to changes in elbow and shoulder joint angles.

Implementation & Features (P3, Magic Mirror AR)

For patients receiving physical therapy, we explored the use of a magic mirror AR (MMAR) setup. The use of the mirror metaphor allows the patient to naturally see their own bodies overlaid with visualizations on a large screen, which is advantageous as their own body is at the center of attention. The patient's movements are captured using a Microsoft Azure Kinect DK. Using the workflow shown in Figure 1, all of the visualizations in this setup are rendered on the reflection of the user on the screen.

Once again, we color-coded the MALs of five arm muscles for the on-body visualization. We also explored *off-body* visualizations as a means to supplement the on-body visualizations. This included dynamically updating bar charts and line charts which show the exact values and change over time of MALs in a more familiar manner. As suggested by the physiotherapist, we also included numeric text values for four main upper-arm angles: the shoulder angle projected in the sagittal, frontal and transverse planes, and the elbow angle. All of these can be seen in Figure 3(a).

From our earlier prototypes, we also learned that on-body visualizations suffer when the virtual and real arms do not perfectly align, causing visual clutter and confusion. This is further exacerbated by render latency and tracking drift. We therefore experimented with *diminished reality* [11] as a means to dynamically "remove" the actual arm from the MMAR video feed, thus leaving only the virtual arm present. This was done using inpainting based on a prerecorded model of the background.

Implementation & Features (P4, Head-Mounted AR)

The physiotherapist would instead inspect the motion of a patient with the aid of an on-body visualization on the patient's body. They, therefore, used a head-mounted AR (HMAR) setup, wearing a HoloLens 2 to see the visualizations in a third-person perspective (i.e., visualizations on another person). We also

²<https://www.simtech.uni-stuttgart.de/press/Best-Poster-Award-goes-to-Project-PerSiVal/>

³<https://assetstore.unity.com/packages/3d/characters/arm-muscles-motion-104538>




FIGURE 3. P3 and P4. MALs are shown through the color of the muscle models, and the dynamic bar charts and line charts. The muscle name labels and user interface are located on the 2D panel. (a) P3-MMAP: Mirror-based setup with diminished reality removing the patient's arm from the mirror video stream; the degrees of the joint angles are listed directly on the right panel. (b) P4-HMAR: View from a physiotherapist onto a patient through the HoloLens; this setup additionally shows the joint angles around the patient's arm (red for sagittal, blue for frontal, green for transverse and yellow for elbow angle).

wanted to ensure that the tracking was as accurate and stable as possible. We, therefore, used an OptiTrack motion capture system, equipping the patient's arm with the required trackers on their arm joints.

We added a floating 2D user interface that can be seen in the right of Figure 3(b). This provides the basic interactions, such as muscle selection and tracking calibration, and shows the same bar charts, line charts, and upper-body angles as was the case in P3. In addition to this, we leveraged the HoloLens 2's spatial capabilities by embedding visualizations of the angles at the corresponding joints of the patient, as can be seen from the colorful arcs. To then allow the patient to understand and be aware of what the physiotherapist is seeing and doing, the physiotherapist's viewpoint is streamed onto the TV screen. Note that this stream is different from the MMAR view of P3 — it is simply the physiotherapist's perspective.

Exploratory Study

We conducted a small-scale exploratory study to gauge the usability of the two prototypes. We first collected feedback from the same aforementioned physiotherapist who tried both prototypes, with the experimenter playing the role of the patient. He formulated a typical physiotherapy task sequence during this process. We then collected feedback from six students from our university (three females, 22–28 years old) who also tried both prototypes. They, however, always played the role of the patient, with the experimenter the role of the physiotherapist who employed the formulated task sequence. When student participants tried P3 (MMAR), they could see the on-body visualizations on themselves on the TV screen.

When they tried P4 (HMAR), they saw the viewpoint of the experimenter (playing the role of a physiotherapist) on the same TV screen whilst wearing the OptiTrack trackers. We did not get them to use the HMD due to its complexity. Further details are included in the supplemental material.

From the physiotherapist's perspective, he saw the potential for both MMAR and HMAR to be used for educational purposes, but only if the muscles and their deformation are as "lifelike" as possible. The physiotherapist particularly appreciated the arcs in HMAR, as they helped him to convey movement instructions and reduced the need to physically hold and guide the patient's arms. However, the physiotherapist had clarified that for real-world scenarios, the actual MALs of someone who is injured, for example, will inevitably be different from someone who is healthy. Thus, personalization is required to ensure an accurate biomechanical simulation.

For the MMAR from the patient's perspective, four participants found it comfortable as they did not need to wear any devices to use the magic mirror, though two participants complained that the visualization would "jump" onto the physiotherapist's body whenever the Kinect accidentally tracked the wrong person. All participants appreciated the use of diminished reality and agreed that it reduced distractions and helped them focus more on the visualization.

For the HMAR from the patient's perspective, we generally observed participants could easily follow the viewpoint of the physiotherapist. They all liked the visualized arcs for joint angles as it helped them understand the physiotherapist's orders to move their arms.

However, three participants complained about the jitter of the virtual arm, and two felt the field of view was too narrow.

Regarding the visualization itself for both prototypes, there was a tendency for participants to focus on the colors of the on-body visualization, with the off-body charts only being consulted sporadically. Three participants were also dissatisfied with the lack of wrist twists in the virtual arm model, indicating that the arm needed to be more realistic. Personalization was, again, raised as an important topic, as participants wanted to see their own personal data and muscle activity.

P5: Mesh-Based Biceps

For this prototype, we were able to implement the aforementioned dynamic arm representation based on *pervasive simulation* of muscle geometry, though this was limited to just the bicep muscle. Based on the feedback from the study with P3 and P4, we opted to move away from the HoloLens 2 due to its narrow field of view, relatively low processing power, and the need to use an external tracking solution (i.e., image targets and OptiTrack system). To explore a different variety of AR, we decided to use the Apple iPad Pro, which featured the M1 chip and LiDar sensor, giving access to both good performance and depth-based motion capture, which no longer requires markers or trackers.

Materials

The main addition to P5 is the real-time muscle geometry simulation, with the same MAL simulation being used as before. This prototype, therefore, becomes close to delivering on *PerSiVal* being a pervasive simulation that runs on mobile devices, with both MALs and muscle deformations being visualized in real-time. Note that whilst the muscle geometry simulation could be offloaded for distributed processing in the case of low-power devices, we found that the iPad Pro was actually powerful enough whereby any performance differences were hardly perceptible to us, at least in the case of the single bicep muscle. We, therefore, chose not to include this distributed processing in our prototype.

Implementation & Features

We used the aforementioned iPad Pro as the mobile AR device, and used the Apple ARFoundation/ARKit in Unity 3D to drive the LiDar sensor for markerless motion capture. We purposefully omitted the off-body visualizations from P3 and P4, allowing us to focus solely on this new muscle geometry and ensure its




FIGURE 4. P5: tracker-free motion capture, biomechanical simulation, rendering calculation, and display integrated in a mobile device. The virtual biceps will deform based on elbow angle, angular velocity, and angular acceleration.

rendering quality and performance. This muscle deforms based on the elbow angle, angular velocity, and angular acceleration of the real arm of the SOB. As with the other prototypes, the MAL of the bicep is also encoded using color, and can be seen in Figure 4.

Outreach Activity

P5 was presented at the SimTech Status Seminar in 2022. We observed that the motion capture of the iPad was generally more accurate and much easier for the SOB to use, as they did not need to wear any markers or trackers. However, when multiple bodies were within the camera's field-of-view, it failed to identify a unique tracking subject, causing the virtual arm representation to switch erratically between bodies, which was not an issue in the one-on-one indoor physiotherapy scenario used in P3. Regarding the visualization, people told us that the muscle deformation was easy to observe and very realistic despite there being only a single bicep muscle. Even still, some appeared to become scared, describing the muscle as being "overly realistic and bordering on gory". In general, people again showed great interest in *PerSiVal* and suggested some future topics including that of personalization.

P6: More Muscles and Particle Rendering

P6 was, again, an evolution of its predecessor, especially as the muscle geometry simulation was now extended to five muscles: biceps, triceps, brachialis, brachioradialis, and anconeus. This brought with it several implementation challenges regarding the performance of the simulation and rendering.

Materials

As the number of muscles in the simulation increased, so too did the number of vertices of the muscle's mesh that needed to be simulated and rendered locally, going from 2809 from the biceps to 18,641 for all




FIGURE 5. P6: rendering the simulation of all five muscles simultaneously based on the mesh vertices, in the form of complete surfaces (left) and the particle (right).

five muscles. This meant that it was non-trivial to go from P5 to P6, as the iPad Pro was now not able to keep up with this increased load. Whilst distributed processing of the simulation was the intended solution, the author responsible for this had since left the team after graduating, making this option unavailable to us. To compensate for this, we decided to optimize the rendering component of the prototype, switching from the surface mesh that was previously used in the visualization to a particle-based rendering technique, as shown in Figure 5. In our testing, this approach reduced the number of computed triangles by 28% and reduced the time per frame by 0.1 ms. We also hoped that this new approach would (purposefully) make the muscles less realistic, thus making them less scary to people.

Implementation & Features

Being an evolution of P5, P6 once again used the iPad Pro with its LiDar camera for motion capture. Each of the five muscles were again color-coded based on their MAL, with the muscles also deforming as per the simulation.

Outreach Activities

We presented P6 at the SimTech Status Seminar in 2023, and also at the Simulated Reality Science Exhibition⁴. The latter was a public event where we gathered feedback from approximately 20 visitors, including non-experts. Many people who we showed the prototype to had once again expressed interest and excitement about the work. As expected, people felt a noticeable lag as five muscles were simulated and rendered at the same time using the surface-based rendering, but this lag was absent in the particle-based rendering. More so, no one commented on the particle-based rendering being overly realistic or gory.

⁴<https://www.project.uni-stuttgart.de/simulierte-realtaet/en/>

Lessons Learned & Reflections

Whilst by no means a complete nor perfect project, our final prototype, P6, we believe does meet our design goals of *PerSiVal*. In particular, it was capable of using AR to present visualizations on-body (G1) of the simulated arm (G2) that was highly realistic with muscle deformations (G3). It did so using real-time motion capture (G4) that was capable of running on a mobile device along with the simulation itself (G5). Over the course of the project, more than 70 people were able to get a basic grasp of the biomechanical process and related muscle morphology using our prototypes (G6).

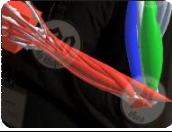
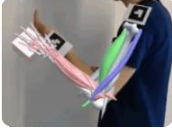
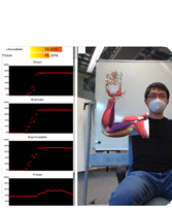
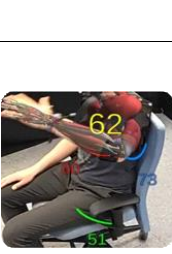
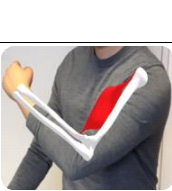
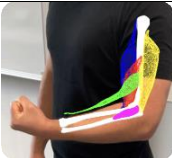
We had, along the way, developed six prototypes that accumulatively represent an exploration into the different opportunities and design choices for on-body AR biomechanical simulation and visualization. In this section, we discuss and reflect on the lessons learned throughout this design space exploration, providing recommendations for future work in this area. Table 1 provides an overview of the design and evaluation of all six prototypes with respect to the five design choices.

Motion Capture

Precise motion capture is the cornerstone of *PerSiVal*. This ensures the accuracy of both biomechanical simulation [10] and 3D registration of AR. During our development, we have experimented with various motion capture strategies, including marker-based approaches (Vuforia and ARToolkit), an outside-in large-scale system (OptiTrack), and depth sensors (Azure Kinect and Apple LiDar). We acknowledge that there are more advanced motion capture methods available. However, we did not consider those methods due to the use of resource-constrained mobile devices and the necessity of also allocating computational resources to simulation and rendering.

While marker-based approaches have been widely used in AR scenarios, they are rarely applied to track dynamic movements of deformable objects, such as a human arm. However, marker-based tracking is simple to implement and easy to deploy in a proof-of-concept setting like ours. As such, we chose it as a low-cost starting point for our first two prototypes by attaching markers directly to the arm joints of the SOB. However, marker-based tracking comes with considerable drawbacks in our biomechanical use case. First, the need to simultaneously capture three image targets located at the shoulder, elbow, and wrist necessitates the camera to be held at a far enough distance away from the SOB (>1.1 meters in our case). At this distance, however, tracking is frequently lost as the captured targets are

TABLE 1. The overview of all the six prototypes, including the design choices and feedback from outreach. More details can be found in our supplemental video.

Prototype	Simulation Data		Technical Setup		Visualization Design			Outreach & Feedback
	Muscle Activation Level (MAL)	Muscle Geometry	Motion Capture	AR Display	Arm	Perspective	Data Encoding	
 <p>P1</p>	Precomputed MAL	---	Vuforia (markers)	Microsoft HoloLens 1	Static arm	Third-person	MAL : • Color	<i>SimTech Status Seminar 2019 (n ≈ 10):</i> • Vuforia had poor latency and tracking
 <p>P2</p>	Simulated real-time MAL	---	ARToolkit (markers)	Microsoft HoloLens 2	Static arm	Third-person	MAL : • Color	<i>SimTech Status Seminar 2021 (n ≈ 10):</i> • ARToolkit was more performant with better tracking • People expressed enthusiasm for PerSiVal
 <p>P3 (MMAR)</p>	Simulated real-time MAL	---	Azure Kinect (depth)	TV Screen	Static arm with more muscles and animated deformation	Mirror	MAL : • Color • Bar charts • Line charts Joint angles : • Text	<i>Exploratory Study (n = 7):</i> <i>HMAR & MMAR (Expert Physiotherapist):</i> • Joint angles in HMAR made giving oral instructions about specific motions easier • Both prototypes have potential in physiotherapy education if muscles and their deformations were truly accurate • Personalization is required for actual use <i>MMAR (Patient Perspective):</i> • Comfortable due to no wearable trackers • Visualization would "jump" away from body • Diminished reality reduced distractions and helped improve focus
 <p>P4 (HMAR)</p>	Simulated real-time MAL	---	OptiTrack	Microsoft HoloLens 2	Static arm with more muscles and animated deformation	Third-person	MAL : • Color • Bar charts • Line charts Joint angles : • Text • Arcs	<i>HMAR (Patient Perspective):</i> • Unstable tracking caused frequent jitter of virtual arm • This tracking provided higher degree of freedom at the cost of comfort • Field of view of virtual content was limited • On-body joint angles helped patient understand orders from the physiotherapist <i>General Feedback for Visualizations:</i> • On-body visualizations were used more than off-body visualizations • Line charts were harder to understand than bar charts
 <p>P5</p>	Simulated real-time MAL	Simulated real-time geometry for one muscle	LiDar (depth)	iPad Pro	Dynamic arm based on simulated muscle geometry	Third-person	MAL : • Color	<i>SimTech Status Seminar 2022 (n ≈ 10):</i> • LiDar was efficient, accurate, and comfortable • LiDar was susceptible to interference from bystanders • Arm muscle was very realistic to the point of being scary to some people
 <p>P6</p>	Simulated real-time MAL	Simulated real-time geometry for five muscles	LiDar (depth)	iPad Pro	Dynamic arm based on simulated muscle geometry	Third-person	MAL : • Color	<i>SimTech Status Seminar 2023 (n ≈ 10) & Simulated Reality Exhibition (n ≈ 20):</i> • Particle rendering was more performant than surface rendering • No reports of the muscle visualization being scary for the particle rendering

not as sufficiently clear, exacerbated by low camera resolution and/or low ambient brightness. Second, as all image targets are required to be facing the camera, tracking is again frequently lost whenever the SOB extends their arms, unintentionally rotating their elbow fossa and thus the image target.

The large-scale motion capture system allows for a much larger motion area and theoretically provides the most accurate tracking of all available options. While full-body motion capture is possible, given the scope of *PerSiVal*, we opted only to track the arm of the SOB

in P4. Despite being a smaller subset of points to track on the SOB, because the HMAR required there also be trackers on the HMD, jitter would occur whenever the physiotherapist moved too close to the tracked arm. Furthermore, the software required to operate these motion capture systems can only run on a PC, requiring movement data to be transmitted via a local area network, which increases the noticeable latency of the visualization.

In contrast to the above two methods, depth sensors have the advantage that the SOB does not need

to wear any trackers. According to Puthenveetil et al. [12], this allows for a higher degree of freedom for limb movements, thus making it easier and more comfortable to use. However, depth sensors cannot accurately identify a certain single object when multiple people are in the camera's field-of-view, making their use in public spaces challenging. Therefore, depth sensors are likely only practical when used in intimate or individual settings, such as the physiotherapy scenarios in P3 and P4.

To summarize, whilst real-time on-body AR biomechanical visualization is highly dependent on *motion capture accuracy*, trying to maximize this, in turn, sees trade-offs in the *comfort* of the SOB and the *pervasiveness* of the devices. Moreover, when using motion capture solutions and setups that require the use of more than one device (e.g., PC and HMD as in P4), this further incurs an increased *latency* cost which would impair the user experience of the simulation and visualization.

Display & Perspective

The choice of AR display and its accompanying perspective has a direct influence on how users perceive the on-body visualization.

Many of our prototypes used a third-person perspective. That is, the person with the AR display observes a separate person who acts as the SOB. Within this, we note that while the iPad Pro was the most powerful and capable of running the simulations of all our devices, the use of the HoloLens HMDs allowed for hands-free viewing of the AR visualizations. This was especially useful when other interactions were needed, such as interacting with a user interface or physically holding and guiding patients in rehabilitation exercises (P4). We note that we had initially explored the use of *first-person* perspective with an AR HMD (i.e., HoloLens 1 & 2). However, similar to Hammady et al. [13], we immediately found that the limited field of view would easily lead to the loss of virtual content. Specifically, our focus on the upper arm muscles meant that the entire visualization could not be fully observed unless the arm was raised high enough. Thus, we turned to the third-person perspective when using AR HMD, supporting an observer to view the visualization on someone else's arm.

In addition to HMDs, we also experimented with a magic mirror perspective on a stationary TV screen. Most of our participants found this setup familiar and easy to use. The use of diminished reality was particularly praised for reducing the interference and visual clutter caused by misalignment between the virtual and

real arms. As highlighted by Siltanen et al. [11], the audience reported that this reduction in visual distractions enhances the immersive experience of biomechanical visualization. Whilst diminished reality can also be employed on mobile AR devices like HMDs or on the iPad Pro, the technical requirements grow dramatically due to the constantly changing background.

Realism

Muscle deformation is a crucial part of biomechanical simulation and visualization, which requires a high level of realism to convey properly.

For our earlier prototypes (P1–P4), we needed to employ simple static arm representations that did not allow for realistic muscle deformations. This, according to the feedback we received from our outreach activities, detracted from the understanding of muscle behaviors. While we did add animated deformation to the static arm in P3 and P4, our physiotherapist quickly recognized the lack of accuracy and realism in the visualization, thus leading him to emphasize this as an important requirement of future biomechanical visualizations. Furthermore, as we focused primarily on the upper arm, we did not add support for wrist twists or animations, which some participants who used P3 and P4 found disappointing. This made it so that certain movements, such as full arm stretches, could not be accurately visualized and thus further compromising the realism and utility of the visualization.

However, we note that — should the visualization be fully realistic — users may exhibit the opposite response. That is, many people whom we showed P5 found the realistic anatomical deformation to be unsettling and eerie. We note that all people who commented on that did not have backgrounds in biomechanics or anatomy. In contrast, due to their professional training and occupational exposure, professionals such as physiotherapists and biomechanical researchers are more familiar with the appearance of red-colored organs and tissues, which reduces their likelihood of experiencing such discomfort [14]. In fact, they preferred detailed visualizations that accurately present the subtleties of biomechanical processes and bring them deeper insights into the dynamics of human movement and anatomical functions.

To summarize, there appears to be a varying threshold for what should be considered the “appropriate” amount of realism for biomechanical visualization, which is seemingly related to the user's expertise. Although *PerSiVal* aims to provide highly realistic biomechanical visualizations, it's important to allow users to adjust the level of realism out of their expertise and

needs. This flexibility ensures that the visualization is both accessible and relevant to a diverse range of users.

Data Encoding

In addition to the aforementioned muscle deformation, MALs and joint angles could also be visualized to the user through some form of data encoding.

MALs were, most commonly, encoded directly onto the muscles of the on-body visualization using color. We also experimented with presenting MALs on off-body bar charts and line charts. According to participants, color was the easiest to observe as it coincides with muscle deformations. The off-body visualizations, on the other hand, were found to be not as useful, especially with the line charts being both harder to understand and seen as redundant by participants.

For the joint angles, we visualized the angle of the upper arm projected into the three body planes (frontal, sagittal, and transverse) and elbow angle as text on the 2D panel in P3 and P4. According to our physiotherapist, this information would help him devise treatment plans and physiotherapy exercises for patients. In addition, we also included arcs in AR that show the same four angles directly on the body of the SOB. Consistent with the findings of Xie et al. [15], having explicit visual cues made it easier for participants to follow the instructor's auditory directions. For instance, as illustrated in Figure 3, the verbal instruction "raise your arm and make the blue angle 73 degrees" would prompt participants to follow the blue arc and laterally abduct their arm 73 degrees on the frontal plane.

Limitations & Outlook

We highlight the three most prevalent steps towards future work.

First, while we have demonstrated the technical feasibility of improving the understandability of biomechanical visualizations on resource-constrained devices, formal validation of this conceptual application is still lacking, particularly in terms of usability testing across different user groups, such as students, physiotherapists, and non-expert individuals. Future work should focus on user studies to assess the understandability and usability of the application. Toward this goal, we, for instance, plan to run a comparative study based on P6, with participants carrying different levels of domain expertise. Other interesting factors that would benefit from further evaluation are the number of rendered muscles and the different visualization styles (as shown in Figure 5). Quantitative measures, such

as subjective NASA Task Load Index (TLX) [16] and System Usability Scale (SUS) [17] scores, could help to more systematically characterize, test, and compare different approaches.

Second, whilst our biomechanical simulations were able to run locally on a mobile device, thus enabling their pervasiveness, we note the limited number of muscles included in our work. As the number of simulated muscles increases, so too will the processing requirements for the simulation to still perform in real-time. Therefore, further optimizations will continuously need to be made, in terms of the distributed processing of the simulation, and the visualization and rendering of the virtual arm representation. Besides, our explorations on pervasive simulation and visualization provide a foundation for integrating real-time live muscle activity data into the system. This will increase the complexity of our input data and hence require explorations of more machine learning models and architectures to still ensure a lightweight surrogate.

Third, whilst the on-body AR visualization is intended to promote a closer association between the visualization itself and the SOB's body, we acknowledge the incongruity between the simulation and the SOB themselves. That is, the continuum-mechanical musculoskeletal system model used in the simulation is based on a healthy young male human and is not representative of the actual SOB, which may be misleading, particularly if they significantly deviate from this baseline (e.g., due to fatigue or medical conditions). Future work will seek to implement personalized simulations for users, using techniques such as ultrasound scanning [18] and incorporating factors including muscle injuries [19], which may require personalized visualizations.

According to our physiotherapist, the latter two steps — incorporating more muscles and a wider variety of muscle states — are prerequisites for the application of *PerSiVal* in physiotherapy education. With the addition of more accurate motion capture, it is also then possible for *PerSiVal* to be used for real-time muscle state presentation in physiotherapy diagnoses or during workouts. Furthermore, during actual treatment, this enables patients to gain a more intuitive understanding of their muscle condition and will make it easier for doctors to explain movements and therapies, making patient-doctor communication easier. It may also provide users a complete end-to-end process of self-accessing their muscle state, performing an exercise [20], and then re-assessing their new muscle state — all in the absence of a physically present coach or physiotherapist.

Conclusion

In this work, we presented an exploration of the design of on-body AR biomechanical visualizations. Taking advantage of our interdisciplinary team, our visualizations integrated biomechanical simulations that could be run on lightweight mobile devices in real-time, thus achieving pervasive simulation and visualization — *PerSiVal*. Throughout this five-year project, we developed six prototypes investigating the benefits and trade-offs of different design choices for on-body AR visualization, collecting feedback through outreach activities along the way. Based on feedback we received, we believe that applications like *PerSiVal* demonstrate the utility of AR and can directly help with, for example, providing guidance for rehabilitation exercises whilst informing the patient about their muscle status [20]. Thus, we argue that establishing a better understanding of how best to visualize biomechanical simulations is of high importance. Further studies should seek to validate the effectiveness of these on-body visualizations for practical usage, guided by our own conceptual explorations and outlook that were presented in this work.

ACKNOWLEDGMENTS

This work was funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 - 390740016. We acknowledge the support by the Stuttgart Center for Simulation Science (SimTech). We also thank Urban Daub from Fraunhofer Institute for Manufacturing Engineering and Automation (IPA) for his expert feedback on physiotherapy. This paper was prepared for informational purposes by the Global Technology Applied Research center of JPMorgan Chase & Co. This paper is not a product of the Research Department of JPMorgan Chase & Co. or its affiliates. Neither JPMorgan Chase & Co. nor any of its affiliates makes any explicit or implied representation or warranty and none of them accept any liability in connection with this paper, including, without limitation, with respect to the completeness, accuracy, or reliability of the information contained herein and the potential legal, compliance, tax, or accounting effects thereof. This document is not intended as investment research or investment advice, or as a recommendation, offer, or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction.

REFERENCES

1. M. Ezati, B. Ghannadi, and J. McPhee, "A review of simulation methods for human movement dynamics with emphasis on gait," *Multibody System Dynamics*, pp. 265–292, 2019.
2. B. Maier, N. Emamy, A. Krämer, and M. Mehl, "Highly parallel multi-physics simulation of muscular activation and emg," in *COUPLED VIII: Proc. International Conf. Computational Methods for Coupled Problems in Science and Engineering*, pp. 610–621, CIMNE, 2019. ISBN:978-84-949194-5-9.
3. O. Röhrle, M. Sprenger, and S. Schmitt, "A two-muscle, continuum-mechanical forward simulation of the upper limb," *Biomechanics and modeling in mechanobiology*, pp. 743–762, 2017.
4. J. Valentin, M. Sprenger, D. Pflüger, and O. Röhrle, "Gradient-based optimization with b-splines on sparse grids for solving forward-dynamics simulations of three-dimensional, continuum-mechanical musculoskeletal system models," *International journal for numerical methods in biomedical engineering*, vol. 34, no. 5, p. e2965, 2018.
5. D. Rosin, J. Kässinger, X. Yu, O. Avci, C. Bleiler, and O. Röhrle, "Persival: Neural-network-based visualisation for pervasive continuum-mechanical simulations in musculoskeletal biomechanics," *arXiv preprint arXiv:2312.03957*, 2023.
6. J. Kässinger, D. Rosin, F. Dürr, N. Hornischer, O. Röhrle, and K. Rothermel, "Persival: Simulating complex 3d meshes on resource-constrained mobile ar devices using interpolation," in *IEEE Int. Conf. Distributed Computing Systems (ICDCS)*, pp. 961–971, IEEE, 2022.
7. J. Kässinger, D. Rosin, F. Dürr, B. Mehler, T. Hubatscheck, and K. Rothermel, "Persival: Using delayed remote updates in a distributed mobile simulation," in *Int. Conf. Computer Communications and Networks (ICCCN)*, pp. 1–10, IEEE, 2023.
8. J. Kässinger, H. Trötsch, F. Dürr, and J. Edinger, "Simedged: Towards accelerated real-time augmented reality simulations using adaptive smart edge computing," in *Proc. Int. ACM Conf. Modeling Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, pp. 181–190, 2023.
9. B. Lee, M. Sedlmair, and D. Schmalstieg, "Design patterns for situated visualization in augmented reality," *IEEE Trans. Visualization and Computer Graphics (TVCG)*, 2023.
10. G. Palmas, M. Bachynskyi, A. Oulasvirta, H.-P. Seidel, and T. Weinkauff, "Movexp: A versatile visualization tool for human-computer interaction studies with 3d performance and biomechanical data," *IEEE*

Trans. Visualization and Computer Graphics (TVCG), vol. 20, no. 12, pp. 2359–2368, 2014.

11. S. Siltanen, "Diminished reality for augmented reality interior design," *The Visual Computer*, vol. 33, pp. 193–208, 2017.
12. S. C. Puthenveetil, C. P. Daphalapurkar, W. Zhu, M. C. Leu, X. F. Liu, A. M. Chang, J. K. Gilpin-Mcminn, P. H. Wu, and S. D. Snodgrass, "Comparison of marker-based and marker-less systems for low-cost human motion capture," in *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 55867, p. V02BT02A036, American Society of Mechanical Engineers, 2013.
13. R. Hammady, M. Ma, and C. Strathearn, "User experience design for mixed reality: a case study of hololens in museum," *International Journal of Technology Marketing*, vol. 13, no. 3-4, pp. 354–375, 2019.
14. V. Bernhardt, H. J. Rothkötter, and E. Kasten, "Psychological stress in first year medical students in response to the dissection of a human corpse," *GMS Zeitschrift für Medizinische Ausbildung*, vol. 29, no. 1, 2012.
15. H. Xie, R. E. Mayer, F. Wang, and Z. Zhou, "Coordinating visual and auditory cueing in multimedia learning.," *Journal of Educational Psychology*, vol. 111, no. 2, p. 235, 2019.
16. S. G. Hart, "Nasa-task load index (nasa-tlx); 20 years later," in *Proceedings of the human factors and ergonomics society annual meeting*, vol. 50, pp. 904–908, Sage publications Sage CA: Los Angeles, CA, 2006.
17. J. R. Lewis, "The system usability scale: past, present, and future," *International Journal of Human-Computer Interaction*, vol. 34, no. 7, pp. 577–590, 2018.
18. A. S. Sahrman, G. G. Handsfield, L. Gizzi, J. Gerlach, A. Verl, T. F. Besier, and O. Röhrle, "A system for reproducible 3d ultrasound measurements of skeletal muscles," *IEEE Transactions on Biomedical Engineering*, 2024.
19. E. Ramasamy, O. Avci, B. Dorow, S.-Y. Chong, L. Gizzi, G. Steidle, F. Schick, and O. Röhrle, "An efficient modelling-simulation-analysis workflow to investigate stump-socket interaction using patient-specific, three-dimensional, continuum-mechanical, finite element residual limb models," *Frontiers in bioengineering and biotechnology*, vol. 6, p. 126, 2018.
20. X. Yu, K. Angerbauer, P. Mohr, D. Kalkofen, and M. Sedlmair, "Perspective matters: Design implications for motion guidance in mixed reality," in

IEEE Symp. Mixed and Augmented Reality (ISMAR), pp. 577–587, IEEE, 2020.

Xingyao Yu is a research associate and doctoral candidate at the University of Stuttgart. His research interests are VR/AR-based human-computer interaction. Contact him at Xingyao.Yu@visus.uni-stuttgart.de

David Rosin is a research associate and doctoral candidate at the University of Stuttgart. His research interest is surrogate modelling in biomechanical context. Contact him at rosin@imsb.uni-stuttgart.de

Johannes Kässinger is a research associate and doctoral candidate at the Institute of Parallel and Distributed Systems (IPVS) of University of Stuttgart. His research topics include pervasive computing, mobile simulations, offloading, and AR applications. Contact him at johannes.kaessinger@ipvs.uni-stuttgart.de.

Benjamin Lee conducted this research whilst he was a postdoctoral researcher at the University of Stuttgart. He is now a Senior Applied Research Associate at JPMorganChase. His research interests include immersive analytics and situated analytics with VR and AR. Contact him at benjamin.lee@jpmchase.com.

Frank Dürr is a lecturer and senior researcher at the Institute of Parallel and Distributed Systems of University of Stuttgart. His research is focused on mobile and pervasive computing and time-sensitive/software-defined networking. Contact him at Frank.Duerr@ipvs.uni-stuttgart.de.

Christian Becker is a professor at the University of Stuttgart, where he works on distributed systems and context-aware computing. Contact him at Christian.Becker@ipvs.uni-stuttgart.de.

Oliver Röhrle is a professor at the University of Stuttgart, where he works at computational biomechanics, finite element and skeletal muscle mechanics. Contact him at roehrle@simtech.uni-stuttgart.de.

Michael Sedlmair is a professor at the University of Stuttgart, where he works at the intersection of human-computer interaction, visualization, and virtual/augmented reality. Contact him at michael.sedlmair@visus.uni-stuttgart.de.