ARrow: A Real-Time AR Rowing Coach

E. Iannucci [1,2], Z. Chen[1], I. Armeni[2], M. Pollefeys[2], H. Pfister[1], J. Beyer[1]

[1]Harvard University
[2]ETH Zurich

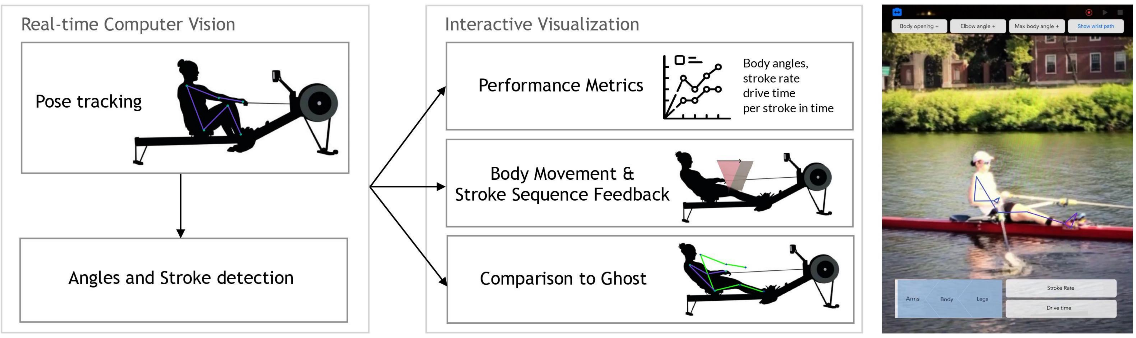

Figure 1: ARrow *pipeline. Left: Using state-of-the-art computer vision, we track the rower's pose and detect inidividual stroke cycles. Middle: We interactively visualize extracted performance metrics, give situated visual feedback on the rower's body movement and stroke sequence, and allow the comparison of the rower's stroke to a previously recorded 'ghost' rower. Right: User interface of* ARrow.

Abstract

Rowing requires physical strength and endurance in athletes as well as a precise rowing technique. The ideal rowing stroke is based on biomechanical principles and typically takes years to master. Except for time-consuming video analysis after practice, coaches currently have no means to quantitatively analyze a rower's stroke sequence and body movement. We propose ARrow, *an AR application for coaches and athletes that provides real-time and situated feedback on a rower's body position and stroke. We use computer vision techniques to extract the rower's 3D skeleton and to detect the rower's stroke cycle.* ARrow *provides visual feedback on three levels: Tracking of basic performance metrics over time, visual feedback and guidance on a rower's stroke sequence, and a rowing ghost view that helps synchronize the body movement of two rowers. We developed* ARrow *in close colaboration with international rowing coaches and demonstrate its usefulness in a user study with athletes and coaches.*

CCS Concepts

• ***Computing methodologies*** → ***Mixed / augmented reality;*** • ***Human-centered computing*** → ***Visualization;***

1. Introduction

Rowing is widely considered one of the hardest sports. It is not only physically demanding, using 85% of the body's musculature during each rowing stroke, but also technically challenging, requiring expert body control, proprioception, timing, and rhythm. While indoor rowing machines (i.e., *ergs*) automatically report performance metrics such as stroke rate and speed, providing similar feedback for on-the-water rowing requires expensive sensors on the boat [NKb, NKa]. Furthermore, these performance metrics do not give rowers actionable feedback on how to improve their stroke sequence, body position, or stroke effectiveness, which typically requires one-on-one feedback from experienced coaches.

In recent years, augmented reality (AR) has become a viable path for sports analysis and coaching. Advances in computer vision (CV) have made it possible to achieve real-time tracking of objects, augmenting them in AR, and overlaying visualizations into live videos of the real world. AR can provide athletes with real-time information [LSY*21, Put], but it is equally important for providing coaches with quantitative data that they might miss otherwise.

To the best of our knowledge, there is currently no AR coaching system for rowing that tracks body movement and provides real-time feedback. Current AR approaches for rowing are limited to showing simple performance metrics (captured with external sensors) in an AR heads-up display [Rowb]. Post-hoc video analysis is becoming popular for more quantitative and visual feedback for rowers. However, none of the available tools work in real-time, and most require time-consuming post-processing [Rowa].

In this paper, we introduce *ARrow*, an AR system that provides real-time visual feedback for rowing training in a situated AR environment. *ARrow* was developed in close collaboration with international rowers and coaches with the goal of providing useful metrics and quantitative feedback that is typically hard to catch by sight and without the need for expensive external sensors. Our system is built on CV methods for real-time 3D pose estimation and rowing stroke detection. *ARrow* provides three levels of coaching feedback: the display of basic performance metrics; visual feedback on a rower's stroke sequence and body movement; and visual feedback on the synchronization between rowers using a *visual ghost*. We demonstrate the usefulness of *ARrow* in a user study with rowers and coaches, and report on their positive qualitative feedback.

2. Related Work

Embedded visualizations can blend data visualizations into physical scenes and have been widely used in sports domains where data is inherently associated with physical environments. Prior research has mainly focused on the use of embedded visualizations for post-game analysis [SJL*18, SBH*18, YCC*21, CXY*22] or presentation [CYC*22, CYX*23]. However, AR can also be used to give real-time feedback to athletes during practice [LSY*21]. This work focuses on using embedded visualizations through AR to provide quantitative real-time feedback for rowing training.

AR is often used to provide real-time feedback on tracked objects (e.g., a ball) in sports training. For instance, DribbleUp [Dri] is an application for children to learn new sports, count repetitions, and correct their ball movements by automatically tracking the ball while they dribble. Similarly, HomeCourt [Hom] records and tracks basketball shots, makes, misses, and locations by tracking the ball and the basket. To track ball trajectories in AR, Kahrs et al. [KRM*], for example, developed an early AR prototype for visual guidance in basketball shooting that shows a 3D trajectory to the center of the rim. Lin et al. [LSY*21] investigated situated AR visualization using the Hololens for basketball free-throw training, and found the situated 3D visualization more effective than a co-located 2D visualization. PuttView [Put] uses AR to help golfers identify the ideal ball path in putting practice by visualizing the ball's path and pace on the green. These examples demonstrate the potential of AR to enhance sports training by providing visual feedback and previously unavailable insights.

While these systems have proven useful, very few examples explore body-centered sports, such as rowing and running. The challenges in these sports lie in estimating the player's pose and visualizing data on top of the human body, which can be significantly different from visualizing data for a ball. VCoach [LSC*22] uses pose estimation for running pose analysis, but it focuses on post-analysis based on videos rather than real-time feedback with AR. Holofit [Boj] is another relevant example that allows rowers to track and improve their progress in virtual worlds, but it uses Virtual Reality (VR) instead of AR. Our work aims to fill this gap by exploring the use of AR for real-time feedback in rowing, a body-centered sport that has not been extensively explored in literature.

3. Goals and Tasks Analysis

We interviewed three rowing coaches from the US national team and one Austrian national team rower (also a co-author) to identify domain goals and requirements for *ARrow*. The coaches were interested in obtaining more quantitative feedback on a rower's stroke and body movement, as subtle changes in body positioning and movement sequencing can significantly affect stroke efficiency and boat speed. They emphasized the importance of tracking body angles at different points during the stroke cycle and performance metrics such as stroke rate, drive time, and body angle range. Based on the interviews, we identified four main goals: **G1)** Real-time tracking and quantitative evaluation of a rower's performance and technique; **G2)** Intuitive visual feedback on the erg and on the water for both coaches and athletes; **G3)** Tracking performance and body movement over several stroke cycles; and **G4)** Visual feedback on the synchronization of rowers in a team boat.

Based on the above goals, we identified three main visualization tasks: **T1)** Analyzing a rower's performance metrics at the current moment and over several stroke cycles. **T2)** Analyzing a rower's motion during the rowing stroke cycle, such as the time of the leg-drive, body opening, elbow bending, and handle height. **T3)** Comparing a rower's timing and motion to other rowers.

4. AR Guided Rowing Feedback

We have designed *ARrow* as a mobile AR app. We use a tablet's camera to capture a live video of a rower and overlay embedded visualizations on the video in real-time. A second display device can also be connected to give live feedback to athletes while rowing. Figure 1 shows the system design of *ARrow*. We perform 3D human pose estimation on the video feed and use the detected key points for subsequent stroke detection. *ARrow* supports three main types of coaching feedback: tracking of basic performance metrics (over several stroke cycles), tracking the rower's stroke sequence and providing visual feedback on body positioning and movement, and a ghost view that gives feedback on the synchronization between rowers. All visualizations are designed to give immediate feedback to rowers and coaches without distracting them.

4.1. Human Pose Estimation & Stroke Detection

3D Human Pose Estimation. We perform 3D human pose estimation using BlazePose [BGR*20], a lightweight convolutional neural network architecture for human pose estimation. It is tailored for real-time inference on mobile devices, calculating the 3D coordinates of 33 keypoints in a human skeleton at over 30 fps.

Body Angles Computation. Next, we compute body angles throughout the rower's stroke, including the angle of the upper body (i.e., leaning forward or backward), hip, knee, and elbow. We first compute angles using the 3D coordinates we get from BlazePose. However, since the depth coordinates are not always accurate, we can optionally compute angles in just the image plane (as long as the camera direction is perpendicular to the rowing direction) and take an average of those two computed angles.

Stroke and Stroke Phase Detection. Upon obtaining the angles, we use a heuristic to detect the different phases in the stroke cycle,

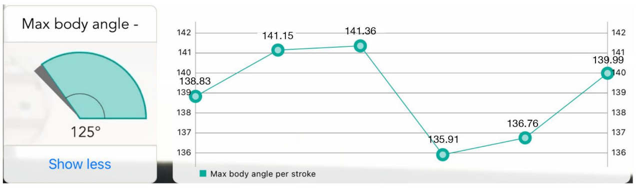

Figure 2: Finish body angle. *Left: The arc diagram shows the current finish angle, highlighting the difference to the previous stroke's finish angle (grey area). Right: Finish body angle over time.*

including catch, leg drive, body opening, arm bending, and finish. We first detect the start and finish of each stroke based on handle direction changes and then extract the different stroke phases based on body angles and the biomechanics of rowing [Kle20].

4.2. Visualizing Performance Metrics

Stroke Rate and Drive Time. The most important performance metrics for rowing are stroke rate (i.e., how many strokes a rower performs in a minute) and speed. We obtain the stroke rate by averaging the last three stroke intervals. Speed, on the other hand, cannot be extracted from video alone but needs external sensors (e.g., GPS, or the force at which the erg handle is pulled). Therefore, we focus on drive time (i.e., the duration of the drive phase of the stroke), which is a good indicator of stroke rhythm and efficiency.

ARrow displays the current stroke rate and drive time as permanent visual overlays (see Fig. 3, c), as rowers need constant feedback on these metrics. Additionally, *ARrow* can display a temporal line chart of these metrics. The chart can be scrolled to view earlier numbers to evaluate technique changes and to help coaches visualize data they might have missed while observing the rower.

Finish Body Angle. *ARrow* can also display the angle of the upper body at the catch and the finish. These angles are usually challenging to quantify, especially for less experienced coaches, but are critical components of the rowing stroke.

We offer three different levels of detail to display the finish angle: a compact label, a custom body angle view (Fig. 2, left), or a temporal line chart (Fig. 2, right). The custom body angle view displays an arc that represents the maximum body angle achieved in the last stroke. We also display the arc of the previous stroke to emphasize whether the angle has increased or decreased.

4.3. Visualizing Stroke Sequence Feedback

The rowing stroke involves multiple body movements in a specific sequence. The drive phase of a stroke starts with leg pushing, followed by opening the upper body, and finally bending the elbows to complete the stroke. The timing and sequencing of these movements are crucial for an efficient stroke, but difficult to master. For example, based on biomechanic principles, the upper body should not begin to open until the legs have reached a 90° angle, and the arms should not start pulling until the legs are straight and the upper body is perpendicular to the floor [Kle20]. Currently, rowers only receive verbal feedback on their stroke sequence. *ARrow* provides four different visualizations to help identify and correct any incorrect sequence of movements.

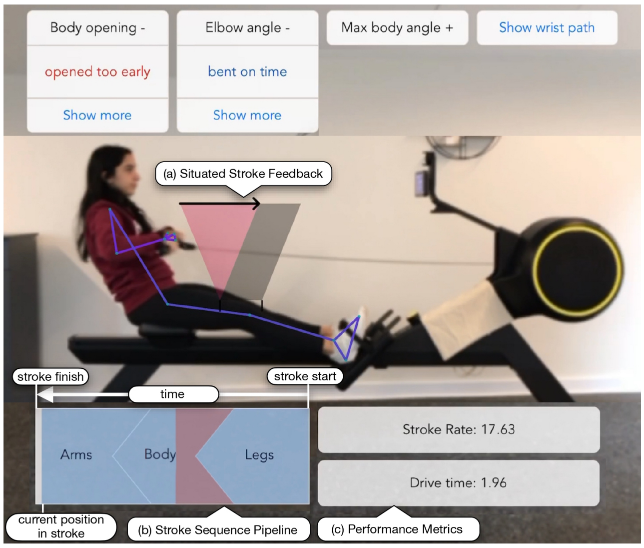

Figure 3: *Embedded Stroke Sequence Feedback (a). The rower opens her upper back too early during the leg drive. The red wedge indicates the correct back position at the beginning of the* body-opening *phase (right side of red wedge) vs. the actual rower's position (left side of red wedge). The stroke sequence view (b) tracks the current position in the stroke (blue bar, filled from right to left) and highlights incorrect stroke phase transitions in red.*

Stroke Sequence Overview - Legs/Body/Arms. The drive phase consists of three main activation phases: leg straightening, body opening, and arm contraction. These movements are performed sequentially, with a slight overlap between two consecutive phases. *ARrow* uses a compact horizontal pipeline view to visualize this sequencing (Fig. 3, b), with overlapping sections to indicate the sequential overlap of consecutive phases. The pipeline gradually fills from the right (catch position) to the left (finish position) to display the rower's position throughout the stroke. In addition, *ARrow* highlights parts of the stroke sequence in red when the athlete starts muscle activation too early, which is a common mistake.

Leg Drive and Upper Body Opening. We show additional visual feedback on a rower's stroke sequence by overlaying it onto the live video stream. We highlight the timing of the leg drive and body opening to identify if a rower is opening up too early (Fig. 3, a). We start by indicating the position of the hip at the beginning of the leg drive and the position of the hip where the body opening should begin (when the knees are 90°bent [Kle20]). Then, we display the body angle at both of these positions. The body angle should ideally remain constant, forming a parallelogram, which we display as a gray overlay. Moreover, we use a red wedge to display the hip position when the rower *actually* begins to open their back and the amount they opened it too early. Athletes can also choose to track the timing of their back opening over time in a line chart if desired.

Elbow Bending. *ARrow* also provides visual feedback on the timing of elbow bending in a rowing stroke. If a rower bends their elbows too early, a red arc is displayed over their elbow (see Fig. 4, left). Moreover, *ARrow* offers a temporal line chart that displays when the elbows began to bend for each previous stroke.

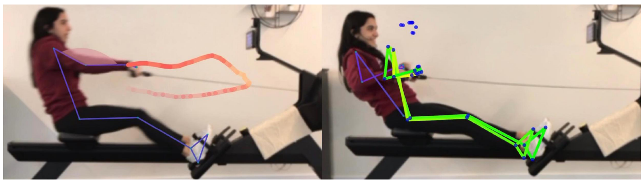

Figure 4: *Left: Wrist path and elbow feedback. The wrist path shows a rise at the beginning of the drive but consistent handle speed on the drive. Right: The ghost rower overlay (green skeleton) highlights their difference at the finish position.*

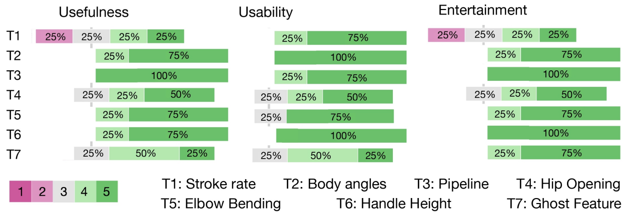

Figure 5: *Survey results on usefulness, usability and entertainment of each feature, from 1 (red, negative) to 5 (green, positive).*

Handle Height. *ARrow* provides feedback on handle height and speed throughout the stroke cycle (see Fig. 4). Hands should be kept at a relatively constant height rather than dipping too low before the catch. *ARrow* visualizes handle height by tracing the path that the rower's wrist follows throughout the stroke, gradually fading over time. Handle speed is encoded from slow (yellow) to fast (red).

4.4. Visualizing Synchronization with Ghost Rower

ARrow provides a visual *ghost rower overlay* for athletes on the erg. This functionality is mainly aimed at rowers in team boats, to aid the synchronization of rowers in a boat (e.g., by making sure their body opening starts at the same time). First, *ARrow* allows one to record the movement of a rower and to extract the rower's pose in terms of keypoints. Next, we can overlay the recording onto the current rower by aligning both skeletons at their feet, which is the only part of the body that is not moving during a stroke. Furthermore, we automatically adjust the playback speed of the recording to match the current stroke rate. This allows us to highlight differences in stroke sequencing while keeping the same stroke rate.

5. Implementation

ARrow is implemented for the iPad Pro using Xcode, Swift 5, and the Charts library [Gin]. For pose estimation, we use MLKit [Goo] and manually changed the pose tracking network to a heavier but more accurate model from Mediapipe [Med]. Finally, we use the iOS AirPlay feature to stream to other iOS devices positioned so that rowers can get feedback while rowing.

6. Evaluation

Here, we report on the accuracy of our stroke detection and angle computation, and on our qualitative user study.

Quantitative Analysis. We evaluated the accuracy of our stroke detection and angle computation by analyzing four different videos: two videos of athletes on ergs facing the same direction, and two videos of rowers on the water facing different directions. *ARrow* accurately detected all 60 strokes in the analyzed videos. We then randomly selected 60 frames and manually measured angles using a protractor and compared them to the computed angles. In 93% cases, the difference between the actual and computed angles was less than 10° and in 85% cases, it was less than 5°.

Qualitative User Feedback. We conducted a small user study with four participants (F: 2, M: 2; age 20-35), including three rowers (one of whom had coaching experience) and one coach. After introducing *ARrow* and its features, we asked the participants to freely explore *ARrow* by recording a rower on the erg. Following their exploration, we asked them to fill out a post-study survey.

The feedback we received indicated that our participants were most interested in getting quantitative feedback on their body angles and stroke sequencing because there is currently no other way to get accurate and real-time feedback on those measures. Participants further reported that visualizing numerical performance metrics over time is a valuable feature (with an average rating of 4.75/5). The usefulness, usability, and fun factor of each feature were rated on a Likert scale, and Fig. 5 shows the results. All participants reported that they would like to use the app in the future.

The *ghost rower overlay* was evaluated by two male rowers in their twenties who row in doubles. They rated it 5/5 in terms of usefulness and reported that there is currently no other way to get feedback on the synchronization of their movements.

Discussion. *ARrow* is designed to give live rowing feedback. However, the application could also be used on prior video recordings of athletes in an offline setting. One limitation of our user study is that it was conducted entirely indoors due to weather constraints during winter. Although *ARrow* can also be used for on-water rowing, when capturing video from a motor launch following a rower, it can be challenging to achieve a camera angle approximately perpendicular to the rowing direction. Additionally, participants suggested integrating feedback on power application per stroke, which would require incorporating external sensor information into *ARrow*. Finally, all participants found quantitative and visual feedback on body position to be helpful in guiding athletes toward a more efficient stroke, as *"Seeing yourself do something wrong can sometimes be even more instructive than a coach telling you."*

7. Future Work and Conclusion

ARrow is a first real-time AR rowing coach that provides embedded visualizations to give direct and quantitative feedback on a rower's stroke sequence and body movement. This allows rowers to immediately adjust their strokes and see the result of their changes. Our work demonstrates the potential of AR-based sports coaching based on human pose tracking and biomechanical principles. In the future, we plan to incorporate automatic oar detection for on-the-water rowing, which would enable us to provide additional details on catch and finish oar angles, stroke length, and the efficiency or *slip* of the oar. We also aim to extend *ARrow*'s capabilities to provide simultaneous feedback on multiple rowers in a team boat. We hope our work inspires future research in this area.

References

[BGR*20] BAZAREVSKY V., GRISHCHENKO I., RAVEENDRAN K., ZHU T., ZHANG F., GRUNDMANN M.: BlazePose: On-device Real-time Body Pose tracking, June 2020. arXiv:2006.10204 [cs]. URL: http://arxiv.org/abs/2006.10204, doi:10.48550/arXiv.2006.10204. 2

[Boj] BOJANA: HOLOFIT - VR Fitness App for Your Rowing Machine, Bicycle or Elliptical. last accessed 2/27/2023. URL: https://www.holodia.com/. 2

[CXY*22] CHU X., XIE X., YE S., LU H., XIAO H., YUAN Z., CHEN Z., ZHANG H., WU Y.: TIVEE: Visual Exploration and Explanation of Badminton Tactics in Immersive Visualizations. *IEEE Transactions on Visualization and Computer Graphics 28*, 1 (Jan. 2022), 118–128. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi:10.1109/TVCG.2021.3114861. 2

[CYC*22] CHEN Z., YE S., CHU X., XIA H., ZHANG H., QU H., WU Y.: Augmenting Sports Videos with VisCommentator. *IEEE Transactions on Visualization and Computer Graphics 28*, 1 (Jan. 2022), 824–834. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi:10.1109/TVCG.2021.3114806. 2

[CYX*23] CHEN Z., YANG Q., XIE X., BEYER J., XIA H., WU Y., PFISTER H.: Sporthesia: Augmenting Sports Videos Using Natural Language. *IEEE Transactions on Visualization and Computer Graphics 29*, 1 (Jan. 2023), 918–928. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi:10.1109/TVCG.2022.3209497. 2

[Dri] DRIBBLEUP: Dribbleup. last accessed 2/27/2023. URL: https://dribbleup.com/hp/default. 2

[Gin] GINDI D. C.: Charts: The Apple side of the crossplatform MPAndroidChart. last accessed 2/27/2023. URL: https://github.com/danielgindi/Charts. 4

[Goo] GOOGLE: Pose detection, ML Kit. last accessed 2/27/2023. URL: https://developers.google.com/ml-kit/vision/pose-detection. 4

[Hom] HOMECOURT: HomeCourt. last accessed 2/27/2023. URL: https://www.homecourt.ai. 2

[Kle20] KLESHNEV V.: *Biomechanics of Rowing: A unique insight into the technical and tactical aspects of elite rowing*. The Crowood Press, Sept. 2020. Google-Books-ID: lzD4DwAAQBAJ. 3

[KRM*] KAHRS L. A., RACZKOWSKY J., MANNER J., FISCHER A., WÖRN H.: Supporting Free Throw Situations of Basketball Payers with Augmented Reality. 2

[LSC*22] LIU J., SAQUIB N., CHEN Z., KAZI R. H., WEI L.-Y., FU H., TAI C.-L.: VCoach: A Customizable Visualization and Analysis System for Video-based Running Coaching, Apr. 2022. arXiv:2204.08805 [cs]. URL: http://arxiv.org/abs/2204.08805. 2

[LSY*21] LIN T., SINGH R., YANG Y., NOBRE C., BEYER J., SMITH M. A., PFISTER H.: Towards an Understanding of Situated AR Visualization for Basketball Free-Throw Training. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, no. 461. Association for Computing Machinery, New York, NY, USA, May 2021, pp. 1–13. URL: https://doi.org/10.1145/3411764.3445649. 1, 2

[Med] MEDIAPIPE: Mediapipe Pose. last accessed 2/27/2023. URL: https://google.github.io/mediapipe/solutions/pose.html. 4

[NKa] NK: EmPower Oarlock. last accessed 2/27/2023. URL: https://nksports.com/empower-oarlock. 1

[NKb] NK: StrokeCoach and Speedcoach. last accessed 2/27/2023. URL: https://nksports.com/category-strokecoach-and-speedcoach. 1

[Put] PUTTVIEW: PuttView - Upgrade your Game. last accessed 2/27/2023. URL: https://www.puttview.com/. 1, 2

[Rowa] ROWERUP: RowerUp - The first AI rowing technique video analysis platform for rowing on a boat and an erg. last accessed 2/27/2023. URL: https://rowerup.com/. 2

[Rowb] ROWINGPERFORMANCE: Quiske Rowing Technique. last accessed 2/27/2023. URL: https://www.rowingperformance.com/. 1

[SBH*18] STEIN M., BREITKREUTZ T., HAUSSLER J., SEEBACHER D., NIEDERBERGER C., SCHRECK T., GROSSNIKLAUS M., KEIM D., JANETZKO H.: Revealing the Invisible: Visual Analytics and Explanatory Storytelling for Advanced Team Sport Analysis. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)* (Oct. 2018), pp. 1–9. doi:10.1109/BDVA.2018.8534022. 2

[SJL*18] STEIN M., JANETZKO H., LAMPRECHT A., BREITKREUTZ T., ZIMMERMANN P., GOLDLÜCKE B., SCHRECK T., ANDRIENKO G., GROSSNIKLAUS M., KEIM D. A.: Bring It to the Pitch: Combining Video and Movement Data to Enhance Team Sport Analysis. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (Jan. 2018), 13–22. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi:10.1109/TVCG.2017.2745181. 2

[YCC*21] YE S., CHEN Z., CHU X., WANG Y., FU S., SHEN L., ZHOU K., WU Y.: ShuttleSpace: Exploring and Analyzing Movement Trajectory in Immersive Visualization. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (Feb. 2021), 860–869. Conference Name: IEEE Transactions on Visualization and Computer Graphics. doi:10.1109/TVCG.2020.3030392. 2