INTERNATIONAL DESIGN CONFERENCE – DESIGN 2022
https://doi.org/10.1017/pds.2022.50

Pure Vision-Based Motion Tracking for Data-Driven Design - A Simple, Flexible, and Cost-Effective Approach for Capturing Static and Dynamic Interactions

S. H. Johnston ✉, M. F. Berg, S. W. Eikevåg, D. N. Ege, S. Kohtala and M. Steinert

Norwegian University of Science and Technology, Norway
✉ Sondrhjo@stud.ntnu.no

Abstract

This paper presents an exploratory case study where video-based pose estimation is used to analyse human motion to support data-driven design. It provides two example use cases related to design. Results are compared to ground truth measurements showing high correlation for the estimated pose, with an RMSE of 65.5 mm. The paper exemplifies how design projects can benefit from a simple, flexible, and cost-effective approach to capture human-object interactions. This also entails the possibility of implementing interaction and body capturing in the earliest stages of design, at minimal effort.

Keywords: design methods, interaction design, vision-based body-tracking, data-driven design

1. Introduction

Capturing human-object interaction within design science enables designers to understand complex prototype interactions better (Moggridge, 2007). These interactions might include location tracking and dynamics of the human body, often demanding high fidelity to obtain meaningful insights. However, capturing such data is not a trivial challenge, as many current technologies used for tracking interactions are expensive or lack the required fidelity. These technologies also confine testing to lab- environments, making it unsuitable for testing in real-world scenarios. Designers might want to track interactions between prototypes and users to determine how changing ergonomics affects use and output, examine complex design team scenarios, or better understand space utilisation. Manual video coding, in which interactions are filmed and processed post-test, can provide researchers with reliable data on team- and prototype interactions but is generally a tedious process (Wulvik, Erichsen and Steinert, 2016). Filming users also poses privacy challenges and potential behavioural changes that researchers must address.

1.1. Methods for Analysing Human-Object Interaction

Wearable-based methods can track human-object interactions but are often costly, lack flexibility, and are only useable in a lab environment. Wearable trackers include reflector-based motion capture systems that use infrared sensors capturing reflective markers attached to the participant, whose position can be referenced to acquire body composition and movement (Ma, Paterson and Pollick, 2006). Designers also have the choice of using an inertial sensor-based motion capture system that relies on inertial measurement units placed on specific body parts to track acceleration and angular position (Zhu and Zhou, 2004; Zhou et al., 2008; Roetenberg, Luinge and Slycke, 2013). Sjöman et al. (2015) used a system of wearable devices as proxies for capturing team dynamics by measuring the relative distance between objects to analyse interactions over a period. This system, however, also

https://doi.org/10.1017/pds.2022.50 Published online by Cambridge University Press

lacks flexibility and fidelity when examining specific prototype interactions. Time-consuming and expensive testing is restrictive in the initial design stages, where speed and agility are essential to progress development. Therefore, a low-cost, flexible, and simple system for capturing human-object interaction is needed. An alternative to manual video coding and wearable systems is vision-based body capturing, in which cameras are used to produce anonymous skeletal models of the subject (Sudderth, 2006), or more general conveyors without proper anatomical association interpreting movement (Kurakin, Zhang and Liu, 2012).

Previous studies have described vision-based body capturing as lacking fidelity (Mitra and Acharya, 2007; Wulvik, Erichsen and Steinert, 2016), but recent advances in open-source software have increased both fidelity and accessibility for design researchers. If the accessibility and accuracy of these open-source software are acceptably high, they could potentially replace the tedious process of manual video coding in some cases. Furthermore, privacy issues can be neglected since pose information can be collected in real-time without storing and collecting images of people.

Video-based body capturing could have a profound effect on design activities. It can capture prototype interaction in a low-cost and simple way, thus making it available for implementation in the earliest stages of design. Sports equipment development is an example that could benefit from utilising this technology, where technique-based sports, such as running (Glover, Kakar and Chaudhari, 2021), skiing (Yoshioka et al., 2018), rowing (Severin et al., 2021), and cycling (Bini, Daly and Kingsley, 2020) involves complex product interactions that should be captured in the early stages of design.

1.2. Vision-Based Pose Tracking

MediaPipe (Lugaresi et al., 2019) is an open-source, cross-platform framework for applying machine learning (ML) solutions to live media. It contains a state-of-the-art human pose estimation solution (MediaPipe Pose, 2020) based on BlazePose (Bazarevsky et al., 2020), which can run in real-time. Occlusion is a fundamental issue for vision-based pose estimation, in which the subject is not in full view of the camera. Recent studies resolve this by training pose estimation and shape feature prediction algorithms (Hou et al., 2020). BlazePose, the foundation for MediaPipe Pose, is an ML-solution trained with substantial occlusion augmentation (Bazarevsky et al., 2020), making it robust in heavily occluded scenarios. MediaPipe Pose tracks the whole body by predicting the location of up to 33 3D landmarks in an image or a sequence of video frames, with each landmark corresponding to a specific point on the body. Tracking fine movements, such as hands and facial landmarks, can be obtained using other specific algorithms implemented in MediaPipe. The intended use in this paper is however at a gross motor level, and these functions are therefore not considered in this paper. The MediaPipe algorithm consists of two main steps: (1) detecting the region of the frame containing a person, and (2) estimating and tracking the landmarks based on this region. Each landmark consists of the x- and y-coordinates relative to the image frame, thus normalised to a value between zero and one. The z-coordinates represent the depth relative to the origin at the midpoint between the hips. The z-coordinates decrease towards the camera and are scaled similarly to the x-coordinates. Additionally, the real-world coordinates can be estimated in meters using the midpoint between the hips as the origin.

Other specific examples of depth perception used in body-tracking include Microsofts Kinect (Morrison et al., 2016), Lidar technology (Shimizu et al., 2016), and stereo camera setups. One principle this study attempts to maintain is an easily accessible setup, for which MediaPipe's versatility in hardware makes it available to implement with any camera, given it has an appropriate resolution and framerate. With its cross-platform support, MediaPipe can be implemented on mobile and desktop devices and supports multiple programming interfaces, including Python, C++, and Java. The computer vision library OpenCV (Bradski and Kaehler, 2000) is used for capturing frames from a webcam, while MediaPipe provides the ML solutions for running inference on the captured frames. MediaPipe Pose also provides several configuration options for adjusting the processing time and accuracy, including detection confidence (between zero and one) and model complexity.

https://doi.org/10.1017/pds.2022.50 Published online by Cambridge University Press

1.3. Aim of Study

This paper presents an exploratory case study where video-based pose estimation is used to capture motion tracking data for a data- driven design project. Hence, it exemplifies how design projects can benefit from a simple, flexible, and cost-effective approach to capture interactions reliably. The paper provides two example use cases, one where working ergonomics are analysed, and one where an adjustable rowing seat is tested. Results from the rowing seat test are compared to a ground truth measurement with three different camera angles to determine the reliability of the software as a tool for designers. The paper aims to demonstrate how design researchers can benefit from using open-source software for tracking user interactions in real-world scenarios without the restrictions of invasive or expensive equipment. It also provides limitations and implications of using the software.

2. Method

2.1. Experiment Setup on Processing Procedures

Two experiments were conducted to test the vision-based body-tracking system in different use cases. A Logitech C920 HD-pro webcam was used as the input device. The processing was performed on a PC (MSI Prestige 15 A10SC) with an intel core i7 10710u processor and 16GB 2666 MHz of RAM. We used Python to create a script for capturing frames (OpenCV) and extracting the body pose (MediaPipe), with the procedure shown in Figure 1. The frames were recorded at a resolution of 1920x1080 pixels, providing sufficient details for estimating the pose. Videos were recorded with a sampling rate of ~30Hz, with pose-estimation analysis performed in post-processing. Figure 2 shows an example of an analysed frame, with points 1 to 5 representing detected landmarks and a skeletal model visualising the resulting pose. The angle of the hip-, shoulder-, and elbow-joints were calculated using trigonometry based on the landmark positions (normalised x- and y-coordinates) 1 to 5 in Figure 2 and stored as a CSV file. Ground truth measurements were performed for one test, using a ruler attached to the rowing machine, as shown in Figure 2.

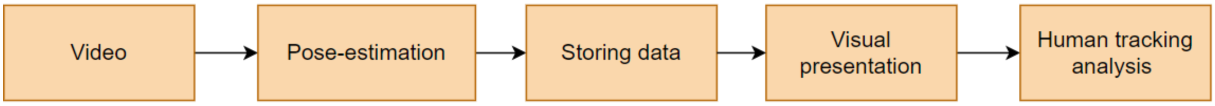

Figure 1. The process model of the visual-based tracking program

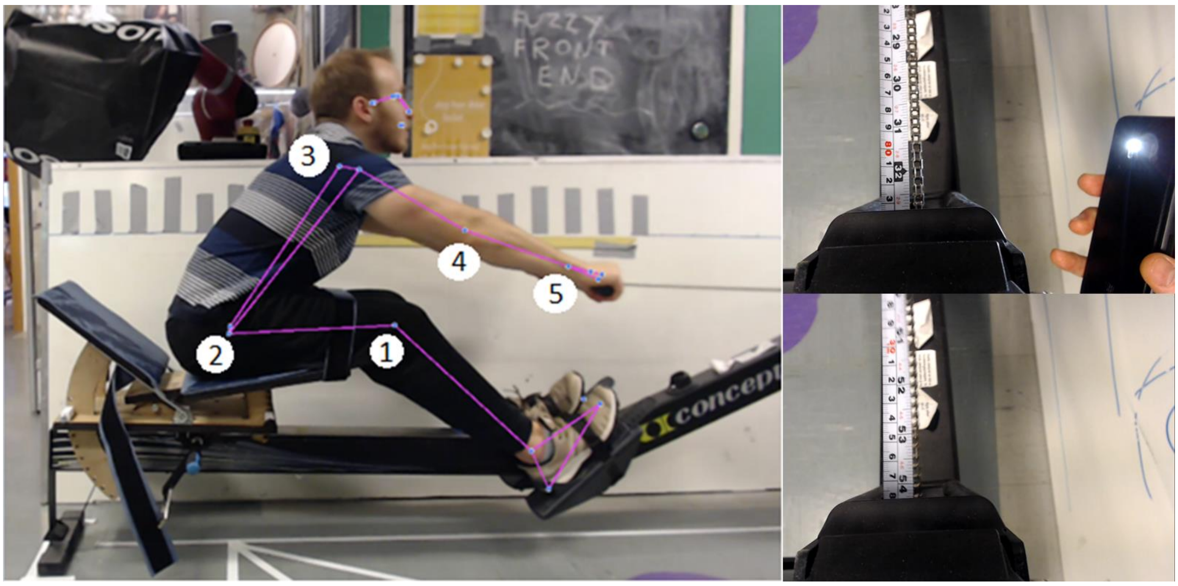

Figure 2. Overlay of pose-estimation with nodes marking the joints used for measurements. A measuring tape was used for ground truth measurements with ~2 mm accuracy, where each camera is synchronised by a flashing light

https://doi.org/10.1017/pds.2022.50 Published online by Cambridge University Press

2.2. Experiment 1: Human-Product Interaction Tracking to Evaluate Working Ergonomics

The aim of experiment 1 is to analyse a user interacting with a keyboard on an office chair that is configured to different heights resulting in different postures. In the context of chair design, understanding posture can help combat unhealthy habits (Gerr, Marcus and Monteilh, 2004), where Faucett and Rempel (1994) found that having a higher elbow height in respect to the wrist is linked to reduced neck, shoulder, and arm discomfort. For simplification, we have classified three different zones with different criteria, as our experiment is only meant to test the applicability of MediaPipe Pose for tracking purposes rather than implying which positions are best. The three scenarios are shown in Figure 3, where MediaPipe is used to detect if the posture is good or not based on whether the elbow is 5% lower (red) or higher (green), or in between (yellow) relative to the wrist.

The experiment focuses on gross motor movement, on an overall body level. A camera was set up 0.8 m above the ground, approximately 2 m from the sagittal plane of the participant. A video was then captured while the participant used a keyboard and adjusted the seat. The experiment lasted approximately 20 seconds with three different chair positions being captured, as seen in Figure 3. While the person's input images are shown to demonstrate the use cases, with the pose-estimation overlaid, we only used the landmark coordinates of the participants' right elbow and wrist for interpreting each posture.

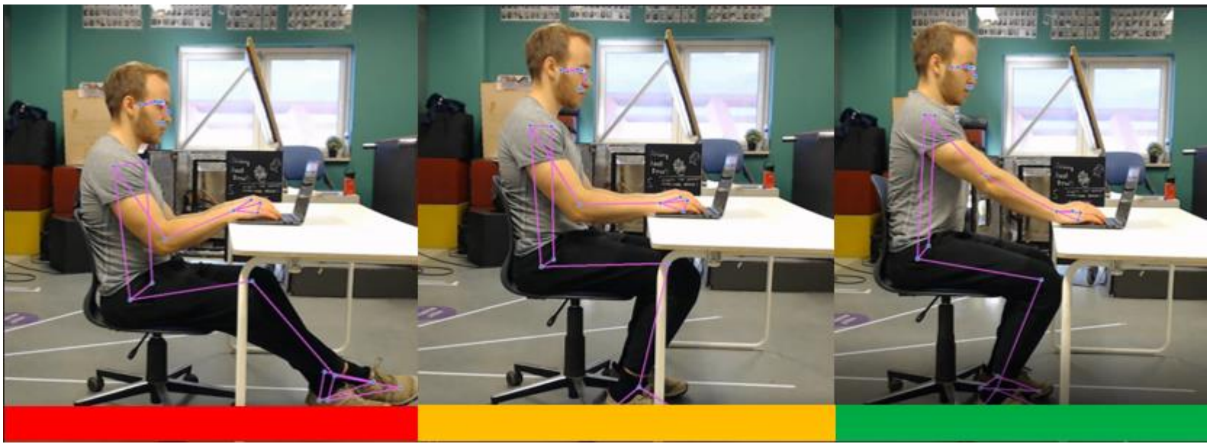

Figure 3. Static images of three different postures, ranging from bad too good, with the pose-estimation overlaid

2.3. Experiment 2: Dynamic Tracking for Movement Analysis

By capturing how humans interact dynamically with equipment in a data-driven design process, we can better understand how design choices affect the overall solution. Machine rowing is an example where the seat configuration can affect the performance of its user. The seating position can therefore be optimised by analysing the movement of a rower (Eikevåg et al., 2020; Severin et al., 2021). Since it involves cycles of repeated motion, it is ideal for assessing the reliability of MediaPipe Pose for movement analysis.

A prototype of an adjustable rowing seat, fixed to a Concept 2 rowing ergometer, was developed to test different pre-determined angles with different ranges of motion and analyse how the body adjusts to various parameters throughout an exercise. The seat has three adjustable parameters, two angular and one linear. Linear movement C in Figure 4 (a) is the horizontal configuration of the seat on the ergometer. The angular parameters A and B have eight and ten individual positions, respectively, as shown in Figure 4. The degrees depicted are the seat and back angles when the pin is locked in the corresponding hole, shown as blue squares. Angle A spans a total of 24 degrees in eight increments to allow for extensive testing of total hip angle movement, while the back-rest B spans a total of 56 degrees in ten increments. The combined angle of A and B defines the full seat angle "h".

https://doi.org/10.1017/pds.2022.50 Published online by Cambridge University Press

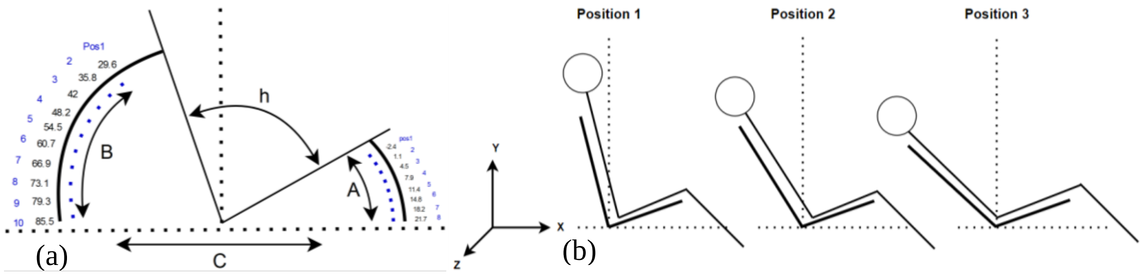

Figure 4. (a) The adjustable parameters of the equipment, (b) The different positions used in dynamic testing.

The experiment was performed in the three positions shown Figure 4 (b) with a fixed seat angle A and increasing back angles B. The participant adjusted the horizontal positions C subjectively as the feet have a fixed position on the machine. The participant rowed for approximately 40 seconds at a constant intensity, of which equal intervals were extracted for processing and to compare results. We have simplified the angular measurements as the depth is disregarded from this view, thus calculating the angles described in Table 1 based on the x- and y-coordinates. Stroke distances were analysed directly from the estimated world coordinates from MediaPipe. Three cameras were set up 210 cm from the rower at 90, 60, and 30 degrees, respectively, as shown in Figure 5 (a). Camera position 1 and 3 were located 80 cm above the ground, while position 2 was elevated to 120 cm providing an isometric view. Pearson correlation coefficients were then calculated between the different camera measurements for each seating position, to evaluate the robustness and reliability of MediaPipe.

To evaluate the accuracy of the MediaPipe, an additional test was captured with the camera in position 1 with a measuring tape fixed to the rowing machine and rowing handle. A second camera was pointed directly down on the measuring tape and synchronized to the first camera, where we manually noted the maximum and minimum stroke distances frame by frame. Root-mean-square error (RMSE) was calculated between the ground truth and MediaPipe stroke distance measurements to evaluate accuracy.

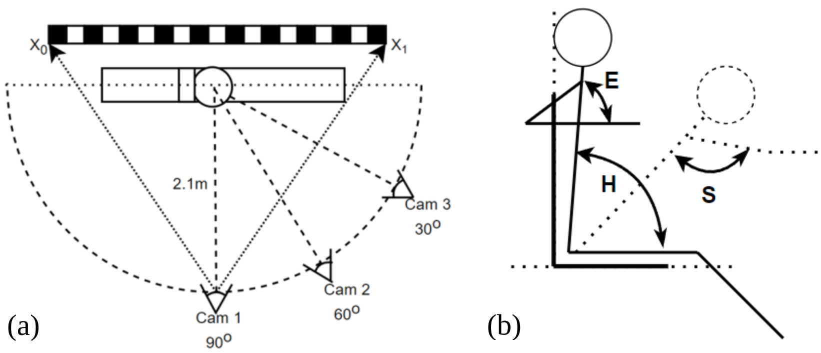

Figure 5. (a) Setup with three camera positions and (b) the angles measured in the dynamic experiment

3. Results

3.1. Experiment 1: Analysis of Different Working Ergonomics

Figure 6 shows the subject's elbow and wrist positions based on the y-coordinates from the estimated world coordinates, where the origin is set between the hips and the y-axis pointing towards the ceiling. We can interpret these landmark-coordinates without using the actual images of the person considering the points have IDs corresponding to their body location, thus maintaining the participants' privacy. Furthermore, the proposed method can be used to analyse posture continuously, allowing the natural sitting habit of a person to be captured without invading their privacy, which can provide valuable insights that are difficult to measure in a controlled lab setting. The graph in Figure 6 shows the duration in which the user has bad, moderate, and good posture, in addition to showing

https://doi.org/10.1017/pds.2022.50 Published online by Cambridge University Press

when the seat configuration is changed. The relative height difference between the wrist and elbow were verified visually, which was captured correctly by the predicted pose.

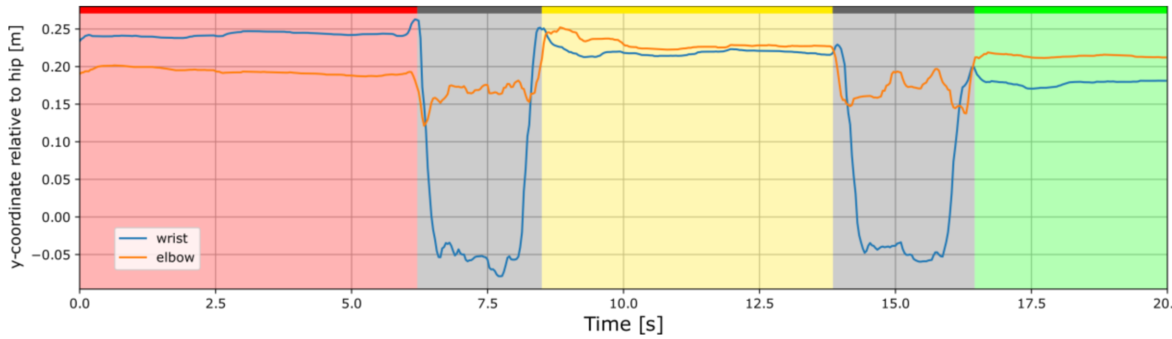

Figure 6. Comparison between right elbow and wrist height with bad (red), mediocre (yellow), and good (green) posture highlighted, corresponding to the postures illustrated in Figure 3.

3.2. Experiment 2: Dynamic Motion Tracking

The hip, elbow, and shoulder angles are shown in Figure 7 for seating positions 1, 2 and 3, measured directly from MediaPipe in camera position 1. The S angle bounces back up when the shoulder joint moves behind the back and reaches the zero-degree limit. The resulting relationships between angles are shown in Table 1, with the mean and standard deviation for both the maximum and minimum H angle throughout the exercise. By using body-tracking, human motion dynamics can be analysed to explain if the user consciously or subconsciously exceeds the equipment's boundary conditions. For example, in position 1, the participant's max H angle is 114.8 degrees, which is 19.3 degrees more than the seat angle (h). The participant pushes his lumbar and pelvic away from the seat, rendering the full seating surface unused, potentially causing discomfort and suboptimal performance.

Lastly the furthest point forward in the movement is different in the three positions even though the seat angle remains the same and the participant has little to no restrictions from the seatback. This indicates a correlation between the participant's ability to stretch further forwards and one of two factors. The first is the participant's leg position, which is adjusted subjectively for comfort, and the second is how much momentum the participant has gathered in the movement leading up to the final forward stretch pushing him further. Both factors could also work alongside each other.

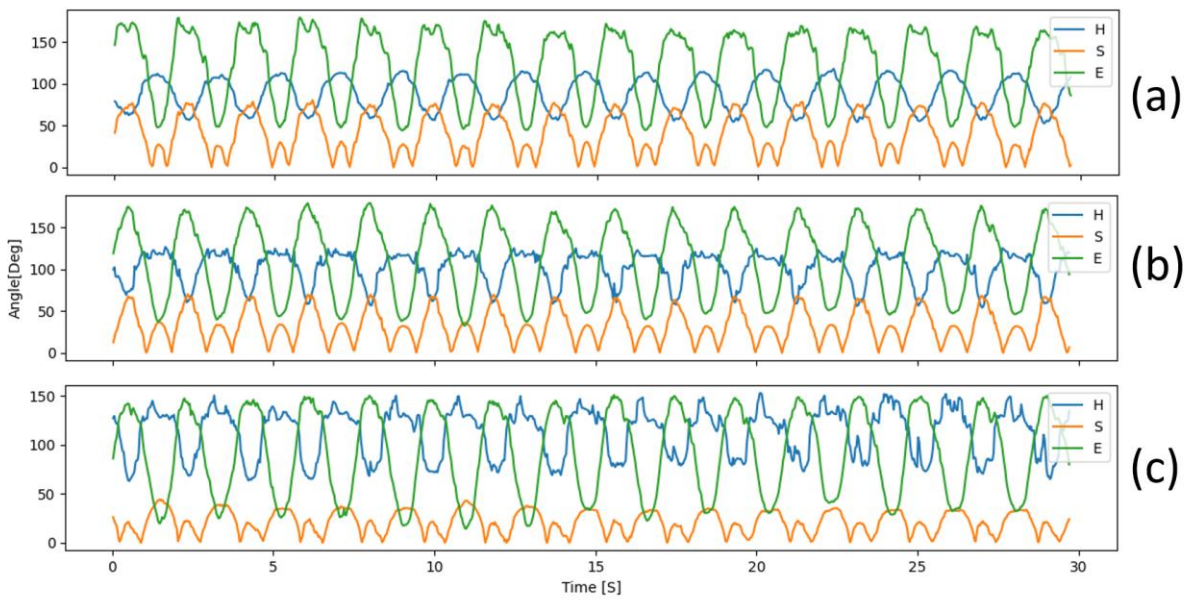

Figure 7. Movement pattern (a) position 1, (b) position 2, (c) position 3.

https://doi.org/10.1017/pds.2022.50 Published online by Cambridge University Press

Table 1. Adjusted seat angles and the measured pose angles in degrees

Variables	Position 1	Position 2	Position 3
Seat bottom angle A	11.4	11.4	11.4
Seatback angle B	73.1	54.5	35.8
Full seat angle (h)	95.5	114.1	132.8
Max H	114.8 ± 3.4	122.4 ± 5.9	132.8 ± 3.4
Min H	56.5 ± 2.3	69.0 ± 3.5	61.9 ± 3.3
Total hip angle (H)	58.3	53.4	70.9

Figure 8 shows the absolute stroke distance, measured from the right wrist landmark in world coordinates directly from MediaPipe, with the origin set at the initial forward-most wrist position. The measurements are shown for each seating position (a-c) and each camera position (1-3), synchronized by a flashing light. The Person correlation coefficients between cameras range from 0.81-0.98, as shown in Table 2, indicating that the camera angle does not substantially affect the robustness of the measurements, although calibration is needed to improve the correctness of the measurements for various camera angles. The lowest correlations at ~0.8 can be explained by camera 3 in seat position 2, where the origin is not captured at the outer position for the wrist landmark for this camera angle.

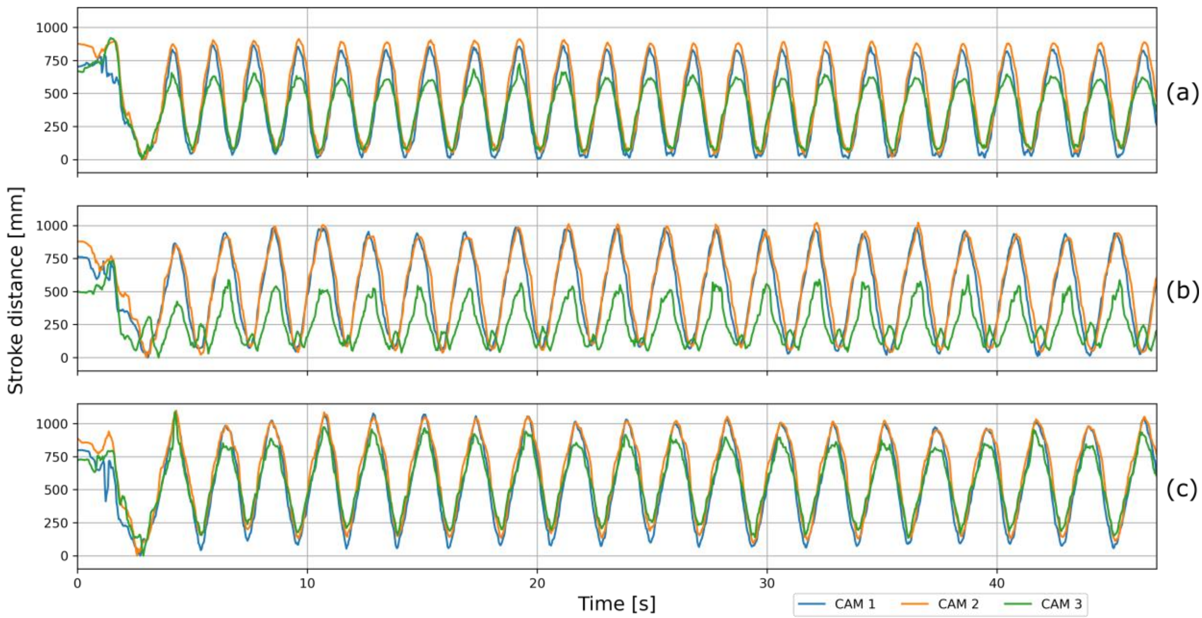

Figure 8. Absolute stroke distance measured from the right wrist landmark in seating position (a) 1, (b) 2, and (c) 3, for each camera angle.

Table 2. Average stroke distance with SD, and Pearson correlation coefficients between camera measurements.

Position	Avg. stroke distance ± SD [mm]			Pearson correlation coefficients		
	CAM 1	CAM 2	CAM 3	CAM 1-2	CAM 1-3	CAM 2-3
1	821.4 ± 16.6	834.4 ± 22.3	564.5 ± 31.8	0.98	0.96	0.97
2	903.7 ± 32.6	901.3 ± 35.6	496.5 ± 25.4	0.97	0.82	0.81
3	948.7 ± 45.8	880.5 ± 43.7	722.4 ± 70.5	0.98	0.98	0.98

Figure 9 shows the stroke distance measured through MediaPipe from camera position 1 for one test scenario, with manual ground truth measurements using a ruler (see Figure 2). Although MediaPipe

https://doi.org/10.1017/pds.2022.50 Published online by Cambridge University Press

and the ground truth measurements have a high correlation (>0.99), the RMSE is 65.5 mm. After fitting a simple linear regression model (resulting in an intercept of -26.2 mm and a coefficient of 1.09) using the scikit-learn python library, with only two samples as input (2nd and 3rd ground truth point with the corresponding MediaPipe measurements), the RMSE is reduced to 28.5 mm, showing that the accuracy of MediaPipe for real world measurements can be improved with relatively simple steps.

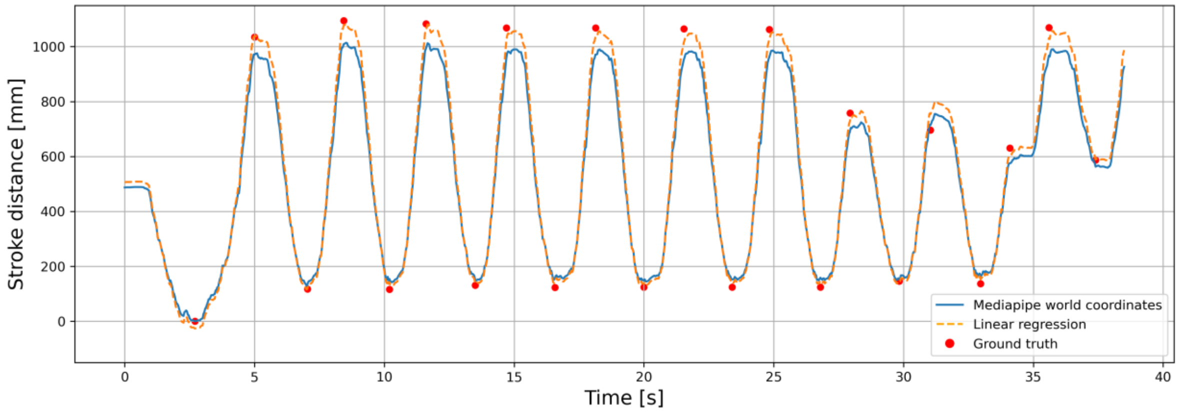

Figure 9. Ground truth stroke distance measurements compared with MediaPipe and linear regression.

4. Discussion

We have shown how MediaPipe can be used for pose-estimation in the context of design, including both static and dynamic scenarios. Our results demonstrate a simple, flexible, and cost-effective approach for providing new opportunities and insights in design research and product development.

The static analysis demonstrated a simple use case where MediaPipe can automatically monitor posture, which can be used for long periods of time without having to record personal data directly, thus reducing the ethical requirements regarding privacy. The information we extracted could reliably determine if the user had good or bad posture, using measurements based on findings by Faucett and Rempel1994).

In the dynamic experiments, we were able to extract useful information that can be related back to the rowing equipment, to better understand the design requirements and user-interaction. Three cameras with different orientations were used to test how reliably MediaPipe can estimate the pose, showing high correlation between measurements. However, an initial calibration step is needed to improve the correctness of the measurements, as the different camera angles resulted in different world coordinates being estimated. Ground truth measurements were compared to MediaPipe for one test case where the stroke distance was analysed, resulting in a Pearson correlation coefficient of >0.99 and an RMSE of 65.5 mm, which was reduced to 28.5 mm using only two datapoints for fitting a linear regression model.

Capturing dynamic movement in real-time has an upper limit in velocity and acceleration depending on the camera's framerates and the computer's processing performance. Initially, attempts were made to connect all three cameras to one computer for the dynamic experiment, which severely limited the framerate to around 8-12 fps. Capturing videos independently before post-processing with MediaPipe is therefore required in these scenarios. Camera distance and resolution also play a prominent role as the program interprets body parts from only 2D images, where distinct features need to be recognised in order to improve landmark detection. In our experiment, the camera placed with a 30 degree angle to the subject struggled with contrast as the participant wore black trousers. For optimal accuracy of motion and dynamic human-equipment interaction, the contrast between clothing, user, equipment, and background, in addition to lighting, must be considered.

Occlusion is a fundamental challenge for pose-estimation. Using the previous frames in a video sequence can predict occluded body-parts well, which is more challenging when using single images.

https://doi.org/10.1017/pds.2022.50 Published online by Cambridge University Press

The capturing method should therefore be carefully considered based on the context of the study being conducted.

Figure 7 and Table 1 provide insights on areas where vision-based body tracking helps interpret important human-product interactions. Comparing the results given by MediaPipe to the physical changes in position on the rowing ergometer can give clear implications on whether the parameters are restricting or superfluous. It can also show how relations work in conjunction, improving the designer's ability to converge on a position where the output parameters are optimised with minimal limitations on others. This could be directly implemented when trying to understand the best position for a specific athlete to use regarding the back angle of the fixed seat during rowing. Our testing implies that the seat restricts the participant's movement for anything above 54.5 degrees. It also suggests that the optimal back angle where it is neither restricting nor superfluous lies around 35.8 degrees, where further testing around this specific area could give valuable insight on how minor adjustments affect the participant's performance. Table 1 also shows a concrete example of a change in position where one would expect an improvement in range of motion, but the opposite happens. This can give designers indications on when human attributes desire changes in equipment that are hard to distinguish subjectively by the user.

Movement can also alter significantly when switching from lab-setting to field testing and executions. Vision-based body tracking is a simple way of tracking movements during field tests, where only a single camera is required, compared to other solutions where multiple camera systems are often needed. However, further comprehensive studies are required to properly evaluate the millimetre accuracy of MediaPipe compared to state-of-the-art tracking solutions.

5. Conclusion

This paper has presented an exploratory case study where camera-based pose estimation is used to track and analyse human motion to support data-driven design research. The paper provides two example use cases, one where the working ergonomics and static postures were analysed, and one dynamic where an adjustable rowing seat was evaluated. A test case in the dynamic scenario compared the estimated pose information to ground truth measurements, showing a correlation of >0.99 and an RMSE of 65.5 mm. Our results demonstrate how design projects can benefit from a simple, flexible, and cost-effective system using open-source software, while showing that interactions can be captured reliably with some limitations regarding the estimated pose accuracy. The paper also demonstrates how design researchers can benefit from open-source pose-estimation algorithms without the restrictions imposed by invasive or expensive equipment, supporting early adaptation of this technology in design and product development at minimal effort.

Acknowledgement

This research received funding from the Research Council of Norway, project number 321386.

References

Bazarevsky, V. et al. (2020) ‘BlazePose: On-device Real-time Body Pose tracking’, arXiv:2006.10204 [cs] [Preprint]. Available at: http://arxiv.org/abs/2006.10204 (Accessed: 9 November 2021).

Bini, R., Daly, L. and Kingsley, M. (2020) ‘Changes in body position on the bike during seated sprint cycling: Applications to bike fitting’, European Journal of Sport Science, 20(1), pp. 35–42. https://dx.doi.org/10.1080/17461391.2019.1610075.

Bradski, G. and Kaehler, A. (2000) ‘OpenCV’, Dr. Dobb’s journal of software tools, 3.

Eikevåg, S.W. et al. (2020) ‘DESIGNING AN EXPERIMENT FOR EVALUATING SEATING POSITIONS IN PARALYMPIC ROWING’, Proceedings of the Design Society: DESIGN Conference, 1, pp. 2485–2494. https://dx.doi.org/10.1017/dsd.2020.101.

Faucett, J. and Rempel, D. (1994) ‘VDT-related musculoskeletal symptoms: Interactions between work posture and psychosocial work factors’, American Journal of Industrial Medicine, 26(5), pp. 597–612. https://dx.doi.org/10.1002/ajim.4700260503.

Gerr, F., Marcus, M. and Monteilh, C. (2004) ‘Epidemiology of musculoskeletal disorders among computer users: lesson learned from the role of posture and keyboard use’, Journal of Electromyography and Kinesiology, 14(1), pp. 25–31. https://dx.doi.org/10.1016/j.jelekin.2003.09.014.

Glover, N.A., Kakar, R.S. and Chaudhari, A.M.W. (2021) 'Effects of spinal coupling and marker set on tracking of spine models during running', Journal of Biomechanics, 116, p. 110217. https://dx.doi.org/10.1016/j.jbiomech.2020.110217.

Hou, T. et al. (2020) 'MobilePose: Real-Time Pose Estimation for Unseen Objects with Weak Shape Supervision', arXiv:2003.03522 [cs] [Preprint]. Available at: http://arxiv.org/abs/2003.03522 (Accessed: 9 November 2021).

Kurakin, A., Zhang, Z. and Liu, Z. (2012) 'A real time system for dynamic hand gesture recognition with a depth sensor', in 2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO). 2012 Proceedings of the 20th European Signal Processing Conference (EUSIPCO), pp. 1975–1979.

Lugaresi, C. et al. (2019) 'MediaPipe: A Framework for Building Perception Pipelines', arXiv:1906.08172 [cs] [Preprint]. Available at: http://arxiv.org/abs/1906.08172 (Accessed: 12 November 2021).

Ma, Y., Paterson, H.M. and Pollick, F.E. (2006) 'A motion capture library for the study of identity, gender, and emotion perception from biological motion', Behavior Research Methods, 38(1), pp. 134–141. https://dx.doi.org/10.3758/BF03192758.

MediaPipe Pose (2020) Mediapipe. Available at: https://google.github.io/mediapipe/solutions/pose.html (Accessed: 11 November 2021).

Mitra, S. and Acharya, T. (2007) 'Gesture Recognition: A Survey', IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 37(3), pp. 311–324. https://dx.doi.org/10.1109/TSMCC.2007.893280.

Moggridge, B. (2007) Designing interactions. Cambridge, Mass: MIT Press.

Morrison, C. et al. (2016) 'Vision-based body tracking: turning Kinect into a clinical tool', Disability and Rehabilitation: Assistive Technology, 11(6), pp. 516–520. https://dx.doi.org/10.3109/17483107.2014.989419.

Roetenberg, D., Luinge, H. and Slycke, P. (2013) 'Xsens MVN: Full 6DOF Human Motion Tracking Using Miniature Inertial Sensors', p. 9.

Severin, A.C. et al. (2021) 'Case Report: Adjusting Seat and Backrest Angle Improves Performance in an Elite Paralympic Rower', Frontiers in Sports and Active Living, 3, p. 625656. https://dx.doi.org/10.3389/fspor.2021.625656.

Shimizu, M. et al. (2016) 'LIDAR-based body orientation estimation by integrating shape and motion information', in 2016 IEEE International Conference on Robotics and Biomimetics (ROBIO). 2016 IEEE International Conference on Robotics and Biomimetics (ROBIO), pp. 1948–1953. https://dx.doi.org/10.1109/ROBIO.2016.7866614.

Sjöman, H. et al. (2015) 'Dynamically capturing engineering team interactions with wearable technology', in DS 80-11 Proceedings of the 20th International Conference on Engineering Design (ICED 15) Vol 11: Human Behaviour in Design, Design Education; Milan, Italy, 27-30.07. 15, pp. 153–162.

Sudderth, E.B. (Erik B. (2006) Graphical models for visual object recognition and tracking. Thesis. Massachusetts Institute of Technology. Available at: https://dspace.mit.edu/handle/1721.1/34023 (Accessed: 12 November 2021).

Wulvik, A., Erichsen, J. and Steinert, M. (2016) 'Capturing Body Language in Engineering Design–Tools and Technologies', DS 85-[1] Proceedings of NordDesign 2016, Volume 1, Trondheim, Norway, 10th-12th August 2016, pp. 165–174.

Yoshioka, S. et al. (2018) 'Pose tracking with rate gyroscopes in alpine skiing', Sports Engineering, 21(3), pp. 177–188. https://dx.doi.org/10.1007/s12283-017-0261-y.

Zhou, H. et al. (2008) 'Use of multiple wearable inertial sensors in upper limb motion tracking', Medical Engineering & Physics, 30(1), pp. 123–133. https://dx.doi.org/10.1016/j.medengphy.2006.11.010.

Zhu, R. and Zhou, Z. (2004) 'A real-time articulated human motion tracking using tri-axis inertial/magnetic sensors package', IEEE Transactions on Neural Systems and Rehabilitation Engineering, 12(2), pp. 295–302. https://dx.doi.org/10.1109/TNSRE.2004.827825.