Independent Control of Supernumerary Appendages Exploiting Upper Limb Redundancy

Hideki Shimobayashi
Information Somatics Lab, RCAST, The University of Tokyo
Tokyo, Japan
shimobayashi@star.rcast.u-tokyo.ac.jp

Tomoya Sasaki
Information Somatics Lab, RCAST, The University of Tokyo
Tokyo, Japan
sasaki@star.rcast.u-tokyo.ac.jp

Arata Horie
Information Somatics Lab, RCAST, The University of Tokyo
Tokyo, Japan
arata.horie@star.rcast.u-tokyo.ac.jp

Riku Arakawa
Information Somatics Lab, RCAST, The University of Tokyo
Tokyo, Japan
arakawa-riku428@g.ecc.u-tokyo.ac.jp

Zendai Kashino
Information Somatics Lab, RCAST, The University of Tokyo
Tokyo, Japan
kashino@star.rcast.u-tokyo.ac.jp

Masahiko Inami
Information Somatics Lab, RCAST, The University of Tokyo
Tokyo, Japan
drinami@star.rcast.u-tokyo.ac.jp

ABSTRACT

In the field of physical augmentation, researchers have attempted to extend human capabilities by expanding the number of human appendages. To fully realize the potential of having an additional appendage, supernumerary appendages should be independently controllable without interfering with the functionality of existing appendages. Herein, we propose a novel approach for controlling supernumerary appendages by exploiting upper limb redundancy. We present a headphone-style visual sensing device and a recognition system to estimate shoulder movement. Through a set of user experiments, we evaluate the feasibility of our system and reveal the potential of independent control using upper limb redundancy. Our results indicate that participants are able to intentionally give commands through their shoulder motions. Finally, we demonstrate the wide range of supernumerary appendage control applications that our novel approach enables and discuss future prospects for our work.

CCS CONCEPTS

• **Human-centered computing** → *Human computer interaction (HCI)*; • **Computer systems organization** → *Real-time operating systems*; *Robotic control*.

KEYWORDS

Supernumerary Appendages, Supernumerary Robotic Limbs, Independent Control, Human Body Redundancy, Wearable Sensing

AHs '21, February 22–24, 2021, Rovaniemi, Finland

ACM ISBN 978-1-4503-8428-5/21/02…$15.00
https://doi.org/10.1145/3458709.3458980

ACM Reference Format:
Hideki Shimobayashi, Tomoya Sasaki, Arata Horie, Riku Arakawa, Zendai Kashino, and Masahiko Inami. 2021. Independent Control of Supernumerary Appendages Exploiting Upper Limb Redundancy. In *Augmented Humans International Conference 2021 (AHs '21), February 22–24, 2021, Rovaniemi, Finland*. ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3458709.3458980

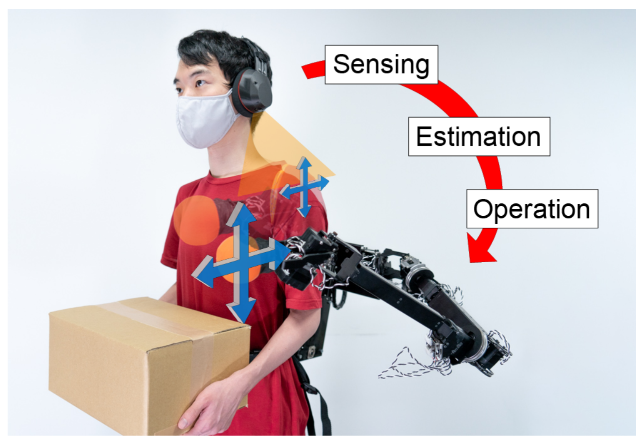

Figure 1: Our proposed system for allowing users to operate supernumerary appendages without disrupting existing appendages by exploiting upper limb redundancy.

1 INTRODUCTION

One aspect of human augmentation is the enhancement of human physical ability through various technologies. One approach to physical augmentation is to expand the number of appendages beyond those which we are naturally born with (i.e., adding supernumerary appendages). To date, research in this area has focused on developing supernumerary robotic appendages (SRAs) to assist in task performance in the physical world [22, 25, 30, 31, 33–35, 38, 38, 40, 42] and adding supernumerary appendages to virtual avatars [21]. For example, existing works have proposed adding a robotic third or fourth arm [33, 38], a 6th finger [35, 42], a tail [25, 30], and a pair of wings to a healthy human body. Similarly, in research pertaining to VR, significant research has been conducted on avatars with supernumerary virtual appendages [21].

While research into the development of such supernumerary appendages is quite active, operation of the appendages remains a significant challenge of the domain. In many cases, appendages are operated using hand motions (e.g., joysticks and VR controllers). These methods of operation, however, are far from ideal. The appendages that we are naturally born with are able to function more or less independently of each other. For example, we are able to move our right arm without having to occupy our left hand with the task of moving the right arm. Ideally, supernumerary appendages would similarly be usable without sacrificing the independence of any other appendages.

Past work has proposed a variety of operation methods that realize hands-free operation of supernumerary appendages. These include operation methods that use feet [38], head direction [23], facial expressions [16], and electromyographic signals from the abdominal muscles [34] to operate supernumerary appendages. However, while these methods do not occupy the hands, they still occupy and limit the functionality of commonly used body parts. For example, we use our feet for walking, our faces and head directions expressing ourselves, and abdominal muscles for breathing and balancing. For supernumerary appendages to be used as natural appendages and realize integrated human augmentation, there is a need to enable operation of the new appendages with minimal interference to our natural body's functionality.

Minimal interference in appendage functionality could be realized by exploiting the redundancy which exists in the human body. Here, redundancy refers to having more controllable degrees of freedom than the total number of degrees of freedom. The human arm, for example, is a highly redundant appendage and allows significant motion even when the end effector (the hand) remains fixed in place. Herein, we propose to utilize these degrees of freedom for supernumerary appendage operation without interfering with appendage functionality.

Namely, we present a system for measuring motions in the shoulder and translating them into control inputs to realize supernumerary appendages operation without disrupting the functionality of any major bodily functions (e.g., manipulation with the hands). Our system consists of a headphone-style wearable visual sensing device and a convolution neural network (CNN) trained to recognize shoulder motions as inputs. Through user experiments, we showed the feasibility and potential of our method. Through a set of demonstrations, we illustrate that our system has a broad range of applications operating a number of supernumerary appendages.

Thus, the contributions of this paper are as follows: (1) A novel input system making use of the redundancy present in the upper limbs; (2) A demonstration of independent hand-based tasks and supernumerary appendages control.

2 RELATED WORK

To enable independent operation of supernumerary appendages (e.g., wearable robotic limbs), we constructed a wearable system for shoulder sensing. Herein, we provide an overview of operation methods for wearable robotics and wearable methods for sensing body movement.

2.1 Operation Methods for Wearable Robotics

One class of wearable robotics, designed with the objective of enhancing human physical abilities in all manner of scenarios [40], is that of supernumerary robotic appendages (SRAs), also known as supernumerary robotic limbs (SRLs). The term SRA refers to wearable robotic systems that add independent appendages to the human body which are not present in the majority of the population. This is in contrast to exoskeletal robotic systems which typically encase and enhance the existing human body. Examples include systems like supernumerary robotic limbs (which may or may not be human-like in structure) [31, 33, 34, 38], supernumerary robotic fingers [22, 35, 42], and robotic tails [25, 30].

A great many approaches to enabling effective use of supernumerary robotic appendages, including shared, assistive, and direct control, have been explored [41]. Regardless of the approach, however, designing an effective interface which allows users to operate wearable robotic systems remains a topic of great interest in this field. Past work has proposed a variety of approaches for making use of body motions and gestures as inputs when operating SRAs. For example, Sasaki et. al. [38] have proposed making use of foot motions for operating a wearable robot arm while in a seated position. Other works have suggest making use body parts other than appendages, such as facial expressions [16], facial direction [23], and electromyographic signals from the abdominal muscles [34], to operate SRAs. However, these methods all make use of body parts which are used regularly in our daily lives, limiting the scenarios in which these methods can be used to operate SRAs. To make full use of SRAs for enhancing physical abilities in humans, it is desirable to minimize the impact that operating the SRA has on the operation of existing human body parts.

Herein, we focus on exploiting mechanical redundancies in the human body to enable SRA operation without affecting natural human performance. Our proposed system exploits the mechanical redundancy of the human body's upper limbs to allow effective operation of SRAs without impacting the performance of the human hands. Namely, we present a method for sensing shoulder postures and using shoulder motion for operating SRAs while simultaneously and independently using the hands for manipulation.

2.2 Wearable Sensing of Body Movements

SRAs are wearable devices that are used in a variety of situations, ranging from daily activities to industrial tasks [40]. It is, thus, desirable for sensing modules, SRA interfaces, to also be wearable.

Research into developing wearable sensing systems that act as novel input interfaces is highly active in the field of human-computer interaction (HCI). Methods for sensing hand movements have a particularly long history of investigation [14]. The most common approach here is attaching sensors directly to the hand. However, there are issues with directly attaching a sensor to the body. The sensor, or the way it is attached, can interfere with important functions of the body part. It can, for example, make it difficult to feel fine sensations and impede freedom of movement.

To resolve the problems with direct attachment, researchers have tried indirect approaches. In these approaches, a body part's motion is estimated from the movement of related (e.g., adjacent) body parts. For example, previous works have developed systems for estimating hand posture from a sensing device wrapped around the wrist or arm. These past works make use of a variety of sensors including force sensing resistors (FSRs) [11, 29], electromyography (EMG) [29, 37], piezoelectric sensors [13], electrical impedance tomography [36, 43], and infrared reflectance or transmittance [15, 28]. In other examples, several works achieved sensing face movements through ear-mounted devices [3, 6, 26].

Another approach to resolving issues with direct attachment has been to move even further away from the body part and use computer vision (CV) to sense body movements. Recent advancements in CV have enabled HCI researchers to capture motions and postures of a part or a whole of human body by utilizing wearable cameras. Sensing the hand position and shape is a common application for this approach [4, 8, 19]. For example, Harrison et al. [19] used a body-worn camera for tracking hand touch input on the surface of the body. Chan et al. [8] presented a ring-style fisheye camera that can classify hand gestures. Arakawa et al. [4] used a hand-mounted spherical camera to simultaneously detect hand gestures and hand location around the body.

Similar approaches have been applied to sensing face [10] and whole-body movements [1, 9, 39]. C-Face [10], for example, is a headphone-shaped device consisting of two ear-mounted miniature cameras that can reconstruct face keypoints. Shiratori et al. [39] developed a motion capture system using body-mounted cameras and Chan et al. [9] proposed a single-piece wearable motion-capture device. Recently, Ahuja et al. [1] integrated hemispherical mirrors into a VR headsets for body pose estimation.

Along with the approaches to hand sensing, these two approaches, direct attachment and CV-based approach, could be used for sensing shoulder motion. However, there remain challenges in directly attaching sensors to the shoulder area due to its complex structure and its large range of motion and deformation. CV-based approaches are comparably easier to adopt due to their less restrictive placement requirements. The existence of wearable CV-based approaches for sensing the motion of large body parts has inspired us to adopt a similar CV-based approach for capturing the users' shoulder posture. In particular, we take a data-driven approach to construct a recognition model as previous works suggest that this can result in robust posture classification across users [4, 8–10]. In the next section, we describe our system. This includes our design rationale, our hardware prototype, and the recognition model we used.

3 PROPOSED SYSTEM

Herein, we propose a wearable shoulder sensing system which aims to exploit redundant degrees of freedom in the upper limbs to enable SRA operation without impeding the motion of the hands. The system consists of a headphone-style visual sensing device and a CNN-based recognition model.

3.1 Rationale for Our Design

Many appendages in the human body contain redundant degrees of freedom (DoF). The arm, for example, contains a total of 7DoF (3 in the wrist, 1 in the elbow, and 3 in the shoulder) for controlling the 6DoF (3 for position and 3 for orientation) of the hand [24]. Combined with the shoulder, an upper human limb contains more than 8DoF for controlling the position and orientation (pose) of the hand. These redundant degrees of freedom are typically used to achieve the same hand pose with a range of arm and shoulder postures. This is useful when, for example, reaching into and working in narrow and complex spaces. Herein, we seek to exploit the redundant DoFs of the upper body limbs to enable independent and parallel operation of SRAs and the hands. Specifically, our proposed system aims to detect shoulder postures and use them as inputs to operate SRAs.

Although various techniques have been proposed, there are challenges in applying these existing methods to shoulder sensing, which we need in this study for the purpose of utilizing upper limb redundancy for control of supernumerary appendages. First, attaching sensors directly to shoulder areas is unsuitable. It may interfere with the user's shoulder movement, but also mounting sensors on shoulders stably is a challenge, because of the large displacement of shoulder areas. Next, it is more difficult to apply the existing estimation approaches to sensing shoulder motions. Since human shoulders have a complicated structure and occupy a large area, it is challenging to estimate the movements of shoulders from the movements of the related body parts (e.g., an arm or a neck). Furthermore, just as human bodies come in a variety of shapes and sizes, human shoulders are likely to also vary significantly from person to person in terms of their size, shape, and flexibility. For these reasons, we thought sensing shoulder movements by attaching on-body sensors around the shoulder area could be difficult. Instead, we decided to take a CV-based sensing approach (Section 2.2).

Figure 2: Our proposed wearable visual sensing device.

3.2 Hardware Design

To sense shoulder motion, we developed a headphone-shaped sensing device (Figure 2). The device contains two downwards facing cameras embedded in the headphones to capture top-down video of the shoulders. It also has two speaker units, which were used to give voice instructions to participants during the experiments. Other mounting locations for the shoulder sensing cameras, such

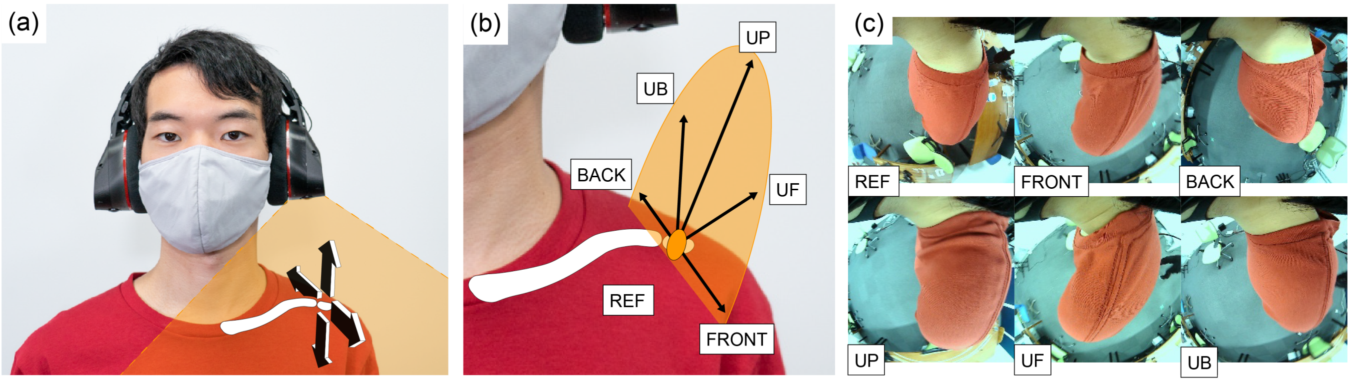

Figure 3: The proposed headphone-shaped visual sensing device designed to acquire the shoulder motion. (a) A user wearing the sensing device with directions for shoulder motion under consideration indicated with arrows. (b) The shoulder postures selected for recognition based on preliminary results. (c) Sample images captured from the left camera for each posture.

as along the jaw or back, were considered but eliminated due to the lack of space, being too close to the shoulder, or not offering a complete view of the shoulder across all possible shoulder motions. The headphone form factor was also selected as past work has shown that it is a form factor that can be worn with relatively small variations in worn location and orientation [10].

The two cameras we used to capture the shoulders were both ELP USB8MP02G-L180 RGB cameras with fisheye lense (1024*768 pixels, 30 fps). Other sensors we considered during development included depth and infrared cameras, and optical distance sensors. However, due to the measurement target (i.e., the shoulder), its location and proximity, the workspace size and vantage point, and sensor size constraints, we opted to use an RGB camera with a fisheye lense. The wide field of view enabled using the fisheye lense contributed greatly to the continued sensing of the shoulder across their wide range of motion. As mentioned earlier in Section 2.2, similar cameras have been used for on-body sensing and have been proven to give robust sensing and recognition results when used in conjunction with machine learning based recognition systems [4]. Also, we show the installation of the device in Figure 3a.

3.3 Recognition Model

Here, we describe the recognition model used in our proposed system to identify shoulder postures. Inspired by other works which utilize fisheye cameras, we opted to use a convolutional neural network based on the 50-layer residual network (ResNet-50) [20] architecture. This architecture is known for having skip connections and performing well on image recognition tasks. We pre-trained the model on the ImageNet [12] public dataset and then trained the network further on a data set we collected ourselves. To achieve fast recognition results in anticipation of running the system in real-time, we chose to first center crop the input images from a size of 1024×768 pixels to 768×768 pixels and then resize the cropped image to be 224×224 pixels. Manual confirmation of the resulting images showed that the shoulder remained in view across shoulder movements and that the resolution was sufficient for shoulder recognition. From preliminary results, we found that there were little differences between the rest position and movements towards the downward direction. Therefore, we chose six shoulder postures for recognition: "REFERENCE (REF)", "UP", "FRONT", "BACK", "UPANDFRONT (UF)", and "UPANDBACK (UB)". These postures can be seen in Figure 3b. A set of sample images indicating how each posture looks from the left camera are shown in Figure 3c. The output layer of the network was, thus, adjusted to have 6 outputs. For training the network, we used the Adam optimizer with a learning rate and weight decay of 1.0×10^{-3}, 1.0×10^{-2} respectively. In the interest of time, training was performed on a PC with a GeForce RTX 3090 GPU. Real-time recognition, however, is possible with even a Quadro T2000 (i.e., GTX 1060 mobile or equivalent). Finally, the model was developed using the PyTorch framework in Python.

4 FEASIBILITY EXPERIMENTS

In this section, we present 3 user experiments we conducted to demonstrate the feasibility of our proposed shoulder input system. In the first experiment, we evaluated the feasibility of the proposed system under controlled conditions. In the second experiment, we examined whether participants could use the system for real-time operation and learn to blindly use the system without feedback. Finally, in the third experiment, we investigated how the participants' body posture affected the recognition results and revealed that shoulder input has the potential to control supernumerary appendages regardless of what the participant's hands are doing. The results obtained from our experiments showed that our system can recognize shoulder postures in a variety of conditions and that a human operator is capable of operating supernumerary appendages using our system in a variety of body postures. It should be noted that the experiments we conducted are feasibility experiments and not evaluations of robustness.

As the purpose of this experiment was to test the feasibility of our system, all experiments were conducted in a room with controlled fluorescent lighting with participants wearing predetermined clothing. This ensured consistency in the data set during our experiments. We acknowledge that there is a need for the system to be generalizable if it is to be used in a variety of real-world settings. Namely, shoulder posture detection using images obtained from our device will need to be robust against changes in factors like the environment, size and shape of shoulder, body type, and clothing color and texture. However, as other works have shown that generalization across such factors can be achieved with a sufficiently large dataset [4, 27], we choose to assume that the same holds for us. Additionally, for simplicity, only one camera (the left camera, and thus the left shoulder) was used for shoulder posture recognition throughout our experiments.

4.1 Experiment 1: Shoulder Posture Recognition

Our first experiment begins with collecting images of the participants' shoulders using our proposed device. The data collected was then used to train our recognition model and test the feasibility of our system.

4.1.1 Participants. We collected data for shoulder posture recognition from nine participants (nine males, 22 to 25 years old). No participants had any history of shoulder diseases.

4.1.2 Procedure. We began this experiment with the acquisition of shoulder images for six different shoulder postures for each participant. We developed a system to automatically give voice instructions to the participants and collect data. Voice instructions were conveyed to each participant through the speakers embedded in our headphone-shaped device and image data was simultaneously acquired from the cameras in the same device.

Prior to the experiment, participants were briefed on the kinds of shoulder postures they would be asked to take and were allowed to practice the postures. When briefing the participants the experimenter did not give them precise instructions as to the posture they should take. Instead, participants were allowed to take on the posture they thought best fit the description they were given. At this stage, participants were also instructed to not turn their head during the experiment and to move with their head and body orientation locked together. Because our proposed system observes the shoulders from the side of the head, any relative motion between the head and the shoulders can cause a problem in shoulder posture recognition. For example, preliminary experiments by the authors showed motion blur in the shoulder images during head rotation, which rendered the shoulder unrecognizable. After the explanation, participants were asked to wear the sensing device as they would a pair of headphones. The experimenter visually confirmed that the participants were wearing the device appropriately before beginning the experiment to minimize extreme variations in the way the device was worn.

In each recording session, the participant was asked to perform the six different shoulder postures in turn with both shoulders while moving freely inside the room with hands relaxed at their sides. They were asked to walk around in order to obtain data with many different backgrounds. It should be noted that participant body postures did not change significantly during each experiment (i.e., no participants sat or crouched). In each recording session, the shoulder postures were presented to the participant in a predefined order (REF, UP, FRONT, BACK, UPANDFRONT, and then UPANDBACK). The participant kept each shoulder posture for approximately 17 seconds (500 frames at 30 frames per second) for a total recording time of approximately 100 seconds. Ten recordings were collected for each participant. This resulted in a total of 30,000 image frames being collected per participant. It should be noted that the device was not taken off at any time during the experiment.

Video footage was collected using a laptop connected via USB to the cameras. To avoid having participants trip over wires, the experimenter followed the participant while holding the laptop. The USB cable connecting the cameras to the laptop was made long enough to not impede the motion of the participants. After collecting video footage, we extracted images from the video streams. After completing the experiment, we had obtained a total of 270,000 images from each camera constituting 45,000 images for each shoulder posture. The neural network was then trained on this data and tested to demonstrate the feasibility of our proposed system.

4.1.3 Usability Survey. After each participant completed all of their recordings, we informally asked them if any of the shoulder postures were particularly difficult to assume. All participants answered that none of the shoulder postures were particularly difficult. However, some participants noted that they felt little difference between the UPANDFRONT and FRONT posture and the BACK and UPANDBACK postures. Also, some participants reported that our headphone-shaped device shifted slightly during data collection, indicating that the wearability of the prototype needs to be improved in future work.

4.1.4 Result. To evaluate the feasibility of our proposed system and its ability to recognize shoulder motions in unknown users, we trained and evaluated our recognition model using the leave-one-out cross-validation approach. Namely, we selected one participants' data to be the test set and trained/validated the model with data obtained from the other participants. After selecting the testing participant, the other participants' data was splitted according to a training:validation data ratio of 9:1. This was repeated nine times, once for each participant's data set being the test set. In each case, we first trained the neural network model for up to 20 epochs. At the end of the training, we chose to make use of the historically best model, evaluated based on its accuracy on the validation data. Finally, we tested the selected model on the test data to obtain a measure of accuracy for the model at hand. Since the training data did not include any test data, the result of the final accuracy test can be seen as a measure of the robustness of the model to unknown subject data.

The results, summarized in Figure 4 as confusion matrices, indicate that there is significant variation in the recognition results. The average recognition accuracy across participants was 69% (53.8% - 83.8%, sd: 9.0%). The values shown in the confusion matrix on the left are averaged values across all participants. The middle and right matrices are for the best and worst performing participants, respectively. The confusion matrices show that the system had the greatest success detecting the REF posture but had the lowest accuracy with the UP posture. We also see that UPANDFRONT had a tendency to get mislabeled as UP or FRONT and that UPANDBACK had a tendency to get mislabeled as UP or BACK.

We suspect that this issue arises from variations in shoulder motion and size, the device location, and neck length. The fact that the worst overall accuracy was observed when the test participant was the one whose body shape was most different from

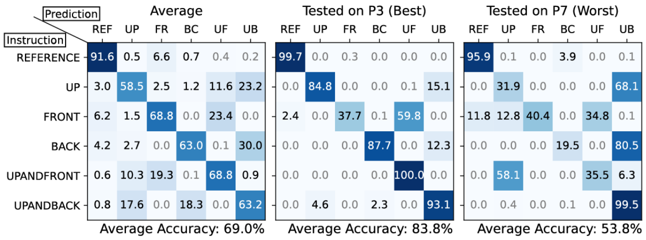

Figure 4: Confusion matrices showing model recognition accuracy using leave-one-out cross-validation. Shown are the average performance matrix (left), the best performance matrix (center), and worst performance matrix (right).

the other participants supports this theory. The shifts in the device during experiments may also have contributed classification errors. Finally, a manual check of the images showed that there were some instances where the participant turned their head independent of their body despite the instructions to the contrary. Images taken during these instances may have contributed to lowering the recognition accuracy.

4.2 Experiment 2 : Real-Time Shoulder Input

In the second experiment, we made use of an estimation model, trained using all of the data acquired in Experiment 1, and investigated whether real-time shoulder inputs were possible. We also sought to identify whether participants could learn how to perform shoulder-based operations without visual feedback.

4.2.1 Participants. There were twelve participants in this study (11 males and 1 female, aged between 22 to 34 years old). Nine of the participants also participated in Experiment 1. This experiment was conducted on a different day than Experiment 1.

4.2.2 Procedure. As with Experiment 1, this experiment began with a briefing on the shoulder postures the participants would be asked to assume. Again, participants were not given precise instructions but were allowed to take on the posture they thought best fit the description they were given. A detailed explanation of the experimental system was then given to the participants. The experimental system consisted of the headphone-shaped sensing device, a computer, and a display (See Figure 5). Participants followed instructions on display to take the displayed shoulder posture. Participants were also told that they would need to focus their gaze on the screen they were seated in front of, for the duration of the experiment. This instruction was given to prevent head motion and motion blur.

For every trial during the experiment, the system visually presented words representing the desired shoulder posture to the participant on the display and presented corresponding voice commands through the speakers embedded in the headphone-shaped device as well. The participants were allowed three seconds to assume the desired posture. A command to relax was then given to the participants (both visually and audibly). The next posture command was given another three seconds later and marks the start of a new trial. Finally, the system could optionally also display the real-time shoulder posture recognition result, giving the participant

Figure 5: A participant performing Experiment 2. Participants followed the displayed instructions and took the appropriate shoulder posture. The picture shows the participant in the second session, where the system indicated the real-time estimated posture.

visual feedback regarding the posture they were taking. We note that the frame rate of our recognition system is above 9 fps. This is enough to allow participants to take the indicated postures and control supernumerary appendages in real-time, as is shown in our demonstrations (Section 5).

A total of three sessions, each containing 60 trials were conducted per participant. The trials in each session were such that each of the six postures would be done ten times each, but in a random order. The first session had the participants conducting trials without real-time visual feedback. In the second session, the participants were provided with visual feedback and instructed to use the feedback to adjust their shoulder posture. Finally, in the last session, the visual feedback was turned off again. After completing all three sessions, participants were given a simple questionnaire asking how difficult they felt it was to take on shoulder postures and whether the shoulder posture recognition results seemed to match their intent.

4.2.3 Result. By comparing the participant shoulder posture data acquired to the desired postures displayed, we were able to identify whether participants were successfully able to assume the desired shoulder posture. We chose to deem a trial a success if the shoulder posture immediately before the relaxation command matched the previously given shoulder posture command. Figure 6 shows the distribution of success rates across participants for each of the three sessions.

The average success rate in each session was 73.5%, 84.9%, and 78.9% respectively. A one-way ANOVA test showed that there was a significant difference between the results obtained across the three sessions ($F = 4.16, p = 0.024$). A Tukey-Kramer test showed that a significant difference only existed between the first and second sessions. Results from the questionnaire showed that the participants found it difficult to assume postures that raised the shoulder, such as UPANDFRONT and UPANDBACK, while seated. Additionally, there were several comments indicating that the system would

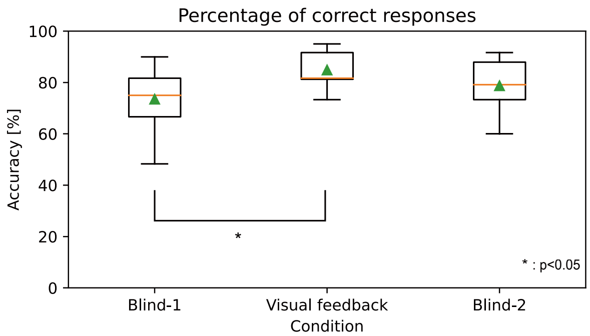

Figure 6: Results from Experiment 2, real-time shoulder posture classification, in three conditions: without real-time visual feedback [Blind-1]; with feedback [Visual feedback]; without feedback after having visual feedback [Blind-2].

almost never recognize certain shoulder postures they took. The UPANDBACK posture was often mentioned as being difficult to assume.

In this experiment, we showed that our system can perform accurate, real-time shoulder posture classification. Namely, participants were able to achieve accurate control inputs without significant training. As noted earlier, this was the first experiment some of the participants participated in. Even then, the lowest overall accuracy recorded in the with-feedback part of this experiment was 70%. While lower accuracies were observed when no visual feedback was presented, the results still indicate that participants were able to achieve the intended results. This was all despite the shoulder recognition model being a general one not tailored to the body shape of any user. These results, as well as the fact that participants of Experiment 1 and 2 likely did not wear the device identically in both experiments, suggest that our system has robust recognition capabilities and could be used in a practical setting. However, the system is not without its flaws. Some participants commented that they were unable to perform shoulder postures such as UPANDBACK. As in Experiment 1, this may be due to physical differences among the participants.

This experiment also demonstrated that participants can significantly improve their accuracy given visual feedback. The participants were able to use the visual feedback to adjust their shoulder postures during the trial to match the desired posture. However, while some improvement was observed after the visual feedback was removed, it was not a statistically significant improvement over pre-feedback accuracy. From this, we infer that there is a possibility that a training session could improve participant accuracy, but further experimentation and analysis is necessary to reach a definitive conclusion.

4.3 Experiment 3 : Robustness to Body Posture Changes

Through Experiments 1 and 2, we showed that it is possible for our system to estimate the user's shoulder posture when they are seated and their arms are at their side. In this third experiment, we investigated how the participants' body posture can affect the estimation results. In particular, we focused on several body postures where the hands were fixed in space (as they might be during a task). By examining the system's performance in these postures, we saw how our system could perform in parallel task scenarios.

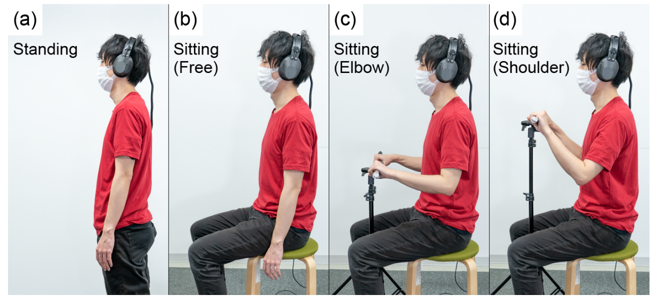

Figure 7: The four different body postures in Experiment 3: (a) Standing; (b) Sitting; (c) Sitting with hands fixed at elbow height; (d) Sitting with hands fixed at shoulder height.

4.3.1 Procedure. The participants in this experiment were identical to those in Experiment 2. The experimental system used for this experiment was also identical to the one developed for Experiment 2. The variable in this experiment, however, was the participant's body posture. Namely, we collected data for the four postures shown in Figure 7 in a random order while participants followed on-screen instructions to assume the indicated shoulder posture. The body postures were 1) standing, 2) sitting, 3) sitting with hands fixed at elbow height, and 4) sitting with hands fixed at shoulder height. The latter two postures were selected to resemble postures a user might take when doing a manual task (e.g., while typing on a keyboard or holding a box). For the fixed hand postures, participants were asked to hold onto handles placed in front of them. The distance to the handles was set based on each participant's arm length (i.e., 70% of the acromion to ulnar styloid process distance). It should also be noted that the data obtained in Experiment 2 was used for the sitting posture.

As in Experiment 2, participants were given three seconds to assume the posture and 3 seconds to relax before being instructed to take the next posture. Participants performed 60 trials (six postures times ten repetitions in a randomized order) in each body posture. It should be noted that participants were provided with visual feedback and were instructed to use the feedback to adjust their shoulder posture. After each experiment, participants were asked to fill out a questionnaire asking how difficult it was to take each shoulder posture and whether the system appeared to recognize their intended posture.

4.3.2 Result. As in Experiment 2, we checked whether the participants' shoulder postures estimated by our estimation model matched the instructions displayed on the screen. The percentage of correct responses for each of the four body postures examined in this experiment are shown in Figure 8. The percentage of correct responses obtained in the second phase of Experiment 2 was used as the accuracy of the sitting and idle condition. The mean accuracy of the standing and seated postures, when hands were not in use, were 90.1% and 84.9%, respectively. The mean percentages of correct responses were 75.6% and 61.1% for the seated participants with hands fixed at elbow and shoulder height, respectively.

A one-way ANOVA was used to determine whether there was a significant difference in the percentage of correct responses given between the four body postures. A significant difference was found with statistics $F = 33.9, p = 1.6 * 10^{-11}$. A post-hoc analysis using the Tukey-Kramer test revealed that there was no significant difference in accuracy between the seated and standing postures where the hands were free. All other differences were significant. Responses to the questionnaires revealed that participants had difficulty assuming shoulder postures where the shoulders needed to move away from where their hands were fixed in space (i.e., the BACK and UPANDBACK postures).

In previous experiments, we demonstrated the possibility of giving control inputs using shoulder postures. In this experiment, we investigated whether participants can assume various shoulder postures, and thus give control inputs, even when in restrictive postures (e.g., the hand position is fixed). The results, Figure 9, showed that participant accuracy was high for certain shoulder postures (e.g., REFERENCE, UP, FRONT postures), and that redundancy of upper limbs allows the control to be performed even when the hand was fixed. On the other hand, participant's accuracy was significantly lower in the two conditions where the hands were fixed than in the other conditions.

We believe the main reasons for the deterioration in accuracy in the fixed-hand postures lie in limitations of body structure. Specifically, we believe restricting the hand position resulted in a reduced range of motion and a shift in the resting position of the shoulder (i.e., there was a change in the mean and reduction of the variance of the shoulder position). A reduced range of motion resulting from the restricted hand position decreased the participant's response accuracy overall as they were unable to assume the same shoulder postures they exhibited when training the recognition model (see the lower two confusion matrices in Figure 9). Restricting the hands had a particularly large effect on the UPANDBACK and BACK postures, which required shoulder movements away from the hands, due to the relatively long distance we fixed the hands at (70% of the acromion - ulnar styloid process distance). Participant comments and the confusion matrices indicating many BACK-REF errors and UPANDBACK-UP errors support this hypothesis. A shift in the resting shoulder position, in contrast, was only clearly observed in when the hands were fixed at shoulder height. In this body posture, the shoulder moves forward and is typically recognized as FRONT when the participant is at rest (i.e., when they attempt to present the REF state). The confusion matrix we obtained for the "Sitting - Shoulder" posture, bottom right in Figure 9, confirms that the REF-FRONT and UP-UPANDFRONT errors are prevalent in this position. Further investigation will be needed to investigate when control input using redundancy is possible.

Further investigation is also necessary to determine whether a unique estimation model will be required for each individual user. Comparing the confusion matrices from Experiment 1, Figure 4, and from the standing posture of Experiment 3, Figure 9, we can see that the response accuracy of UPANDBACK posture did not improve despite the availability of visual feedback. From this, and participant comments indicating that some shoulder postures

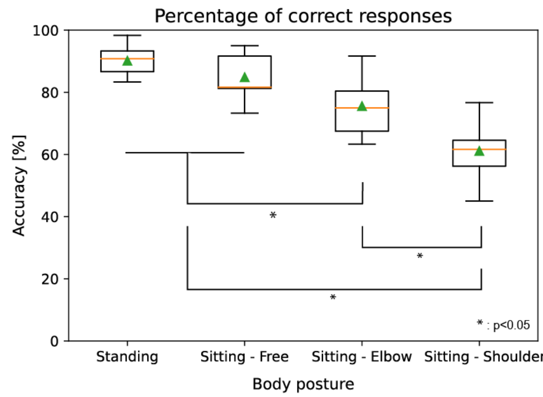

Figure 8: Results from Experiment 3: real-time shoulder posture classification accuracy in four body postures. Green triangles represent the mean values.

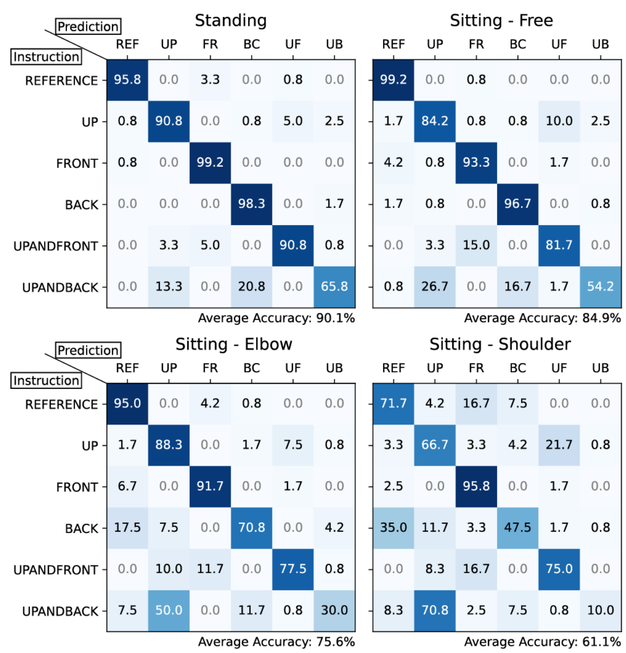

Figure 9: Confusion matrices from Experiment 3 showing recognition accuracy using our estimation model in four different body postures.

were not recognized at all, we conjecture that some postures are difficult to reach even without any restrictions on body posture. Variations in the range of motion should, thus, be taken into account in the design of future shoulder-based input systems. One approach to resolving this issue is to achieve input with minimal shoulder motion. Another solution is to acquire more data across a range of body types, ages, and genders, and use them to train a, possibly more sophisticated, recognition model. However, it should be noted it is not the intent of this study to propose a robust input method. We only show the feasibility and possibilities of shoulder-based input. Future work will endeavor to further develop our approach into a robust control method by, for example, extending the data collection to a wider range of upper limb postures and body types in an effort to achieve robust shoulder posture estimation.

5 APPLICATION & DEMONSTRATION

In this work, we have shown that it is possible to make use of upper limb redundancy to enable control input without interfering in hand operations. In this section, we suggest four example applications of supernumerary appendages which become feasible given our system, shown in Figure 10. These applications were implemented and tested to verify that they are realizable.

Figure 10: Photos of applications and demonstrations: (a) Robotic Arms; (b) Robotic Fingers; (c) Robotic Tentacles; (d) Virtual Appendages.

5.1 Robotic Arms (SRLs)

First, we show a demonstration wherein a user controls a set of SRLs. Figure 10a shows our system is being used to operate an SRL to press an elevator button while the user's hands are busy holding a box. This is one scenario in which SRLs are commonly expected to be useful [17]. Our implementation of this application enabled a user to move a robot arm up-down and forward and backward using the shoulder.

5.2 Robotic Fingers

Figure 10b shows a demonstration of a robotic finger attached to the human hand. This finger extends the functions and capabilities of the human hand as a sixth finger [35, 42]. This finger can provide additional dexterity as it can be manipulated independently of the user's other finger movements. In our working prototype, the two axes of the robotic finger were manipulated using the front, back, and diagonal shoulder postures as input.

5.3 Robotic Tentacles

Figure 10c shows an application wherein our system is being used to manipulate robotic tentacles. Such robotic tentacles composed of several serially connected actuators suggest the possibility of body extension unconstrained by existing body structures [2, 32]. In our case, we placed tentacles on the back of the hand to allow for additional object grasping. As in the other demonstrations, this tentacle can be used while walking. To enable operation of the tentacle, we mapped forward and backward movement of the shoulders to the opening and closing of the robotic tentacle. Using the tentacle, we allowed the user to grasp objects with the tentacle without occupying the hands.

5.4 Virtual Appendages

Figure 10d shows our system being used to operate a set of virtual wings, a kind of virtual appendages. The wings flap in response to shoulders movements picked up by our system. In this scenario, the wings are being operated while the user is performing complex maneuvers with their hands (i.e., juggling a set of three balls). This demonstration suggests that our system could be applied for performance purposes in addition to practical uses. Furthermore, this demonstration shows that our system is compatible with use in real, virtual, and mixed reality scenarios.

6 DISCUSSION

In this section, we will first discuss the conceptual contributions of our proposed system. Then, we will discuss the limitations of our system as well as future work that might develop from the system and the concepts it embodies.

6.1 Conceptual Contributions

In this paper, we proposed a novel wearable sensing system which enables the independent operation of supernumerary appendages. The system takes a powerful approach which takes advantage of the mechanical redundancy present in the human body structure. While our system focuses on exploiting upper limb redundancy and the degrees of freedom it gives to the shoulder, we believe that this approach could be applied to the elbows, head, or lower limbs. Our first conceptual contribution is, thus, highlighting the potential of exploiting mechanical redundancy in the human body. Namely, that it can be exploited to enable supernumerary appendage control that minimizes the performance impact on the rest of the body.

Our second conceptual contribution is showing that our approach to supernumerary appendage operation is not task or appendage dependent. Our applications and demonstrations showed that our approach could be used to operate a variety of physical and virtual supernumerary appendages. This is in contrast to previous works which typically focus on task-dependent inputs [17].

Finally, our work has shown the potential of shoulder-based input. The shoulder has many aspects of expression which humans can intentionally control due to the presence of the scapula. For example, besides moving the whole shoulder, it is possible to slide and reorient the scapula as well as tense the surrounding muscles. Being able to sense these expressions could enable delicate, precise, and complex inputs. While we have only focused on sensing the externally visible shape of the shoulder in this work, future work may consider developing methods for extracting more information from the shoulder by, for example, applying on-body sensing method to enable at will control of supernumerary appendages.

6.2 Limitations

Despite the potential of our current system, it still has multiple issues and limitations which must be resolved. First, there is the issue of mapping shoulder motion to supernumerary appendage control inputs. In more traditional methods of appendage control, it is common to have six input dimensions for controlling the six degrees of freedom (three for position and three for orientation) of the end effector. Having six reliable input dimensions is challenging given our current system. Since we were able to realize six inputs, we were only able to demonstrate operation of a robotic arm along two positional dimensions independent of the rest of the body. Additional issues may arise due to all of our control inputs being in the saggital plane. Namely, control of appendages along the coronal plane may require significant training before they become usable by users. To enable control of supernumerary appendages with higher degrees of freedom, we may consider applying control dimensionality reduction techniques such as those applied to prosthetic and robot arm control [5, 7]. Also, in this study, it is necessary to hold the input posture when giving a command. This can be tiring when using our system for an extended period of time. There are several ways to address this problem. One solution for this issue is to implement trigger inputs that hold a command until an opposite input is given. Another solution is to minimize the amount of shoulder motions during input.

The second major issue is regarding the limitations on the types of SRAs which can be used with our system. Since our method uses visual data of the shoulder to infer control inputs, any SRA which may cover the shoulder cannot be used with the system. SRAs, for example, which are worn like a backpack can significantly impact the input accuracy as the shoulder straps interfere with shoulder sensing. For such SRAs, attaching a marker to the strap and utilizing an algorithm that detects the marker pose may result in improved performance.

Finally, there are some issues that must be resolved with the sensing device. While our camera-based approach resulted in robustness against differences between participants, it carries with it many problems which are common to this kind of approach. Using a camera can lead to, for example, privacy violation if a person is caught camera and recognition being sensitive to changes in lighting. Since SRAs are intended to be used in a variety of situations [40], our sensing device should be applicable both in a variety of scenarios and for a variety of people. Another issue with the sensing device is that its field of view changes with head rotations. This shifts the shoulder position in the field of view and introduces motion blur which can be problematic for detecting shoulder postures. This latter issue could be resolved by increasing the camera framerate or by adding an IMU to detect head motions. The other issues may be resolved by introducing methods from previous works which encountered similar issues [4, 10].

6.3 Future Work

There are several possible directions for future research. One direction is to further our measurement system and its capabilities. An important issue to be addressed in future work is the generalization of the sensing system. To increase the generalized performance of the system, it is necessary to acquire training data from a variety of users and usage situations. This might include data from users in different clothing, reattachment positions, head postures, and various arm postures. Using this data, we believe that we can construct a robust recognition system and expand the situations in which we could allow for SRA operation. Another direction for furthering our system is enabling independent operations with each shoulder as, in this work, it was assumed that both shoulders moved in tandem. This would increase the degrees of freedom available to the user and enable more complex control. Another extension of the current work would be looking at continuous shoulder position inputs instead of discrete ones. Finally, the addition of other sensors (e.g., IMUs and EMGs) could be used to extend the sensing capabilities of our camera-based approach to sensing shoulder posture.

We would also like to conduct a usability study of SRA operations using shoulder input. In this study, we showed the possibility of multitasking with SRA operations by using shoulder motions. However, it has not been sufficiently investigated whether user tasks and SRA tasks can actually be performed independently. In future work, it will be necessary to investigate the psychological and physical loads on the user and ease of use during multitasking.

Another topic for future work is to investigate the human side of redundancy-based input. We could, for example, investigate how participants learn to use the system. We observed some learning through the presentation of visual feedback in Experiment 2. However, learning could be accelerated by providing multimodal feedback, as is suggested by other works investigating SRLs and extra fingers [18, 22].

Finally, we could explore other applications for our proposed sensing system. While we have focused on controlling systems related to human augmentation (i.e., SRAs and VR avatars), our system could be applied more broadly in the field of HCI. It could, for example, be used in the context of rehabilitation, as a general purpose input system, or for the control of prosthetics.

7 CONCLUSION

In this study, we proposed a novel approach for realizing independent control of supernumerary appendages. The most novel part of our approach is its exploitation of upper limb redundancy to achieve control inputs without interfering with the existing appendages. Our approach consists of capturing top-down visual data of the shoulder and recognizing shoulder postures using a neural network classifier. To capture the visual data, we created a headphone-shaped sensing device making use of two fisheye lens cameras to capture a large field of view. The neural network we used for classification was based on ResNet-50. Through user experiments, we demonstrated the feasibility and potential of our approach. Furthermore, through a set of demonstrations, we have shown a wide range of human augmentation applications that can be realized with our approach. Future work will further explore independent control methods that exploit mechanical redundancies present in the human body.

ACKNOWLEDGMENTS

This work was supported by JST ERATO Grant Number JPMJER1701, Japan (INAMI JIZAI Body Project).

REFERENCES

[1] Karan Ahuja, Chris Harrison, Mayank Goel, and Robert Xiao. 2019. MeCap: Whole-Body Digitization for Low-Cost VR/AR Headsets. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. ACM. https://doi.org/10.1145/3332165.3347889

[2] Mohammed Al-Sada, Thomas Höglund, Mohamed Khamis, Jaryd Urbani, and Tatsuo Nakajima. 2019. Orochi. In *Proceedings of the 10th Augmented Human International Conference 2019 on - AH2019*. ACM Press. https://doi.org/10.1145/3311823.3311850

[3] Toshiyuki Ando, Yuki Kubo, Buntarou Shizuki, and Shin Takahashi. 2017. CanalSense: Face-Related Movement Recognition System based on Sensing Air Pressure in Ear Canals. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology - UIST '17*. ACM. https://doi.org/10.1145/3126594.3126649

[4] Riku Arakawa, Azumi Maekawa, Zendai Kashino, and Masahiko Inami. 2020. Hand with Sensing Sphere: Body-Centered Spatial Interactions with a Hand-Worn Spherical Camera. In *Symposium on Spatial User Interaction*. ACM. https://doi.org/10.1145/3385959.3418450

[5] P.K. Artemiadis and K.J. Kyriakopoulos. 2010. EMG-Based Control of a Robot Arm Using Low-Dimensional Embeddings. *IEEE Transactions on Robotics* 26, 2 (April 2010), 393–398. https://doi.org/10.1109/tro.2009.2039378

[6] Abdelkareem Bedri, David Byrd, Peter Presti, Himanshu Sahni, Zehua Gue, and Thad Starner. 2015. Stick it in your ear: building an in-ear jaw movement sensor. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers - UbiComp '15*. ACM Press. https://doi.org/10.1145/2800835.2807933

[7] S. Bitzer and P. van der Smagt. 2006. Learning EMG control of a robotic hand: towards active prostheses. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*. IEEE. https://doi.org/10.1109/robot.2006.1642128

[8] Liwei Chan, Yi-Ling Chen, Chi-Hao Hsieh, Rong-Hao Liang, and Bing-Yu Chen. 2015. CyclopsRing: Enabling Whole-Hand and Context-Aware Interactions Through a Fisheye Ring. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST '15*. ACM Press. https://doi.org/10.1145/2807442.2807450

[9] Liwei Chan, Chi-Hao Hsieh, Yi-Ling Chen, Shuo Yang, Da-Yuan Huang, Rong-Hao Liang, and Bing-Yu Chen. 2015. Cyclops: Wearable and Single-PieceFull-Body Gesture Input Devices. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*. ACM Press. https://doi.org/10.1145/2702123.2702464

[10] Tuochao Chen, Benjamin Steeper, Kinan Alsheikh, Songyun Tao, François Guimbretière, and Cheng Zhang. 2020. C-Face: Continuously Reconstructing Facial Expressions by Deep Learning Contours of the Face with Ear-mounted Miniature Cameras. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. ACM. https://doi.org/10.1145/3379337.3415879

[11] Artem Dementyev and Joseph A. Paradiso. 2014. WristFlex: low-power gesture input with wrist-worn pressure sensors. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology - UIST '14*. ACM Press. https://doi.org/10.1145/2642918.2647396

[12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE. https://doi.org/10.1109/cvpr.2009.5206848

[13] Travis Deyle, Szabolcs Palinko, Erika Shehan Poole, and Thad Starner. 2007. Hambone: A Bio-Acoustic Gesture Interface. In *2007 11th Annual International Symposium on Wearable Computers - ISWC '07*. IEEE. https://doi.org/10.1109/iswc.2007.4373768

[14] L. Dipietro, A.M. Sabatini, and P. Dario. 2008. A Survey of Glove-Based Systems and Their Applications. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38, 4 (July 2008), 461–482. https://doi.org/10.1109/tsmcc.2008.923862

[15] Rui Fukui, Masahiko Watanabe, Tomoaki Gyota, Masamichi Shimosaka, and Tomomasa Sato. 2011. Hand shape classification with a wrist contour sensor: development of a prototype device. In *Proceedings of the 13th International Conference on Ubiquitous Computing - UbiComp '11*. ACM Press. https://doi.org/10.1145/2030112.2030154

[16] Masaaki Fukuoka, Adrien Verhulst, Fumihiko Nakamura, Ryo Takizawa, Katsutoshi Masai, and Maki Sugimoto. 2019. FaceDrive: Facial Expression Driven Operation to Control Virtual Supernumerary Robotic Arms. *ICAT-EGVE 2019 - International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments* (2019). https://doi.org/10.2312/EGVE.20191275

[17] Jacob Guggenheim, Rachel Hoffman, Hanjun Song, and H. Harry Asada. 2020. Leveraging the Human Operator in the Design and Control of Supernumerary Robotic Limbs. *IEEE Robotics and Automation Letters* 5, 2 (April 2020), 2177–2184. https://doi.org/10.1109/lra.2020.2970948

[18] Jacob William Guggenheim and H. Harry Asada. 2020. Inherent Haptic Feedback from Supernumerary Robotic Limbs. *IEEE Transactions on Haptics* (2020), 1–1. https://doi.org/10.1109/toh.2020.3017548

[19] Chris Harrison, Hrvoje Benko, and Andrew D. Wilson. 2011. OmniTouch: wearable multitouch interaction everywhere. In *Proceedings of the 24th annual ACM symposium on User interface software and technology - UIST '11*. ACM Press. https://doi.org/10.1145/2047196.2047255

[20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE. https://doi.org/10.1109/cvpr.2016.90

[21] Ludovic Hoyet, Ferran Argelaguet, Corentin Nicole, and Anatole Lécuyer. 2016. "Wow! I Have Six Fingers!": Would You Accept Structural Changes of Your Hand in VR? *Frontiers in Robotics and AI* 3 (May 2016). https://doi.org/10.3389/frobt.2016.00027

[22] Irfan Hussain, Leonardo Meli, Claudio Pacchierotti, Gionata Salvietti, and Domenico Prattichizzo. 2015. Vibrotactile haptic feedback for intuitive control of robotic extra fingers. In *2015 IEEE World Haptics Conference (WHC)*. IEEE. https://doi.org/10.1109/whc.2015.7177744

[23] Yukiko Iwasaki and Hiroyasu Iwata. 2018. A face vector - the point instruction-type interface for manipulation of an extended body in dual-task situations. In *2018 IEEE International Conference on Cyborg and Bionic Systems (CBS)*. IEEE. https://doi.org/10.1109/cbs.2018.8612275

[24] J. Lenarcic and A. Umek. 1994. Simple model of human arm reachable workspace. *IEEE Transactions on Systems, Man, and Cybernetics* 24, 8 (1994), 1239–1246. https://doi.org/10.1109/21.299704

[25] Azumi Maekawa, Kei Kawamura, and Masahiko Inami. 2020. Dynamic Assistance for Human Balancing with Inertia of a Wearable Robotic Appendage. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2020)*.

[26] Hiroyuki Manabe, Masaaki Fukumoto, and Tohru Yagi. 2013. Conductive rubber electrodes for earphone-based eye gesture input interface. In *Proceedings of the 17th Annual International Symposium on Wearable Computers - ISWC '13*. ACM Press. https://doi.org/10.1145/2493988.2494329

[27] Fabrice Matulic, Riku Arakawa, Brian Vogel, and Daniel Vogel. 2020. PenSight: Enhanced Interaction with a Pen-Top Camera. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems - CHI '20*. ACM. https://doi.org/10.1145/3313831.3376147

[28] Jess McIntosh, Asier Marzo, and Mike Fraser. 2017. SensIR: Detecting Hand Gestures with a Wearable Bracelet using Infrared Transmission and Reflection. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology - UIST '17*. ACM Press. https://doi.org/10.1145/3126594.3126604

[29] Jess McIntosh, Charlie McNeill, Mike Fraser, Frederic Kerber, Markus Löchtefeld, and Antonio Krüger. 2016. EMPress: Practical Hand Gesture Classification with Wrist-Mounted EMG and Pressure Sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM. https://doi.org/10.1145/2858036.2858093

[30] Junichi Nabeshima, MHD Yamen Saraiji, and Kouta Minamizawa. 2019. Arque. In *ACM SIGGRAPH 2019 Posters*. ACM. https://doi.org/10.1145/3306214.3338573

[31] Koki Nakabayashi, Yukiko Iwasaki, and Hiroyasu Iwata. 2017. Development of Evaluation Indexes for Human-Centered Design of a Wearable Robot Arm. In *Proceedings of the 5th International Conference on Human Agent Interaction*. ACM. https://doi.org/10.1145/3125739.3125763

[32] Ken Nakagaki, Sean Follmer, and Hiroshi Ishii. 2015. LineFORM: Actuated Curve Interfaces for Display, Interaction, and Constraint. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST '15*. ACM Press. https://doi.org/10.1145/2807442.2807452

[33] Federico Parietti and Harry H. Asada. 2013. Dynamic analysis and state estimation for wearable robotic limbs subject to human-induced disturbances. In *2013 IEEE International Conference on Robotics and Automation*. IEEE. https://doi.org/10.1109/icra.2013.6631123

[34] Federico Parietti and H. Harry Asada. 2017. Independent, voluntary control of extra robotic limbs. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. https://doi.org/10.1109/icra.2017.7989702

[35] Domenico Prattichizzo, Monica Malvezzi, Irfan Hussain, and Gionata Salvietti. 2014. The Sixth-Finger: A modular extra-finger to enhance human hand capabilities. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*. IEEE. https://doi.org/10.1109/roman.2014.6926382

[36] J. Rekimoto. 2001. GestureWrist and GesturePad: unobtrusive wearable interaction devices. In *Proceedings Fifth International Symposium on Wearable Computers*. IEEE Comput. Soc. https://doi.org/10.1109/iswc.2001.962092

[37] T. Scott Saponas, Desney S. Tan, Dan Morris, Ravin Balakrishnan, Jim Turner, and James A. Landay. 2009. Enabling always-available input with muscle-computer interfaces. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology - UIST '09*. ACM Press. https://doi.org/10.1145/1622176.1622208

[38] Tomoya Sasaki, MHD Yamen Saraiji, Charith Lasantha Fernando, Kouta Minamizawa, and Masahiko Inami. 2017. MetaLimbs. In *ACM SIGGRAPH 2017 Emerging Technologies*. ACM. https://doi.org/10.1145/3084822.3084837

[39] Takaaki Shiratori, Hyun Soo Park, Leonid Sigal, Yaser Sheikh, and Jessica K. Hodgins. 2011. Motion capture from body-mounted cameras. In *ACM SIGGRAPH 2011 papers on - SIGGRAPH '11*. ACM Press. https://doi.org/10.1145/1964921.1964926

[40] Vighnesh Vatsal and Guy Hoffman. 2017. Wearing your arm on your sleeve: Studying usage contexts for a wearable robotic forearm. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE. https://doi.org/10.1109/roman.2017.8172421

[41] Sang won Leigh, Harshit Agrawal, and Pattie Maes. 2018. Robotic Symbionts: Interweaving Human and Machine Actions. *IEEE Pervasive Computing* 17, 2 (2018), 34–43. https://doi.org/10.1109/mprv.2018.022511241

[42] Faye Wu and Harry Asada. 2014. Bio-Artificial Synergies for Grasp Posture Control of Supernumerary Robotic Fingers. In *Robotics: Science and Systems X*. Robotics: Science and Systems Foundation. https://doi.org/10.15607/rss.2014.x.027

[43] Yang Zhang and Chris Harrison. 2015. Tomo: Wearable, Low-Cost Electrical Impedance Tomography for Hand Gesture Recognition. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology - UIST '15*. ACM Press. https://doi.org/10.1145/2807442.2807480