






RESEARCH

Open Access



Egocentric video: a new tool for capturing hand use of individuals with spinal cord injury at home

Jirapat Likitlersuang^{1,2} , Elizabeth R. Sumitro^{1,2}, Tianshi Cao² , Ryan J. Visée^{1,2} , Sukhvinder Kalsi-Ryan^{2,3}  and José Zariffa^{1,2*} 

Abstract

Background: Current upper extremity outcome measures for persons with cervical spinal cord injury (cSCI) lack the ability to directly collect quantitative information in home and community environments. A wearable first-person (egocentric) camera system is presented that aims to monitor functional hand use outside of clinical settings.

Methods: The system is based on computer vision algorithms that detect the hand, segment the hand outline, distinguish the user's left or right hand, and detect functional interactions of the hand with objects during activities of daily living. The algorithm was evaluated using egocentric video recordings from 9 participants with cSCI, obtained in a home simulation laboratory. The system produces a binary hand-object interaction decision for each video frame, based on features reflecting motion cues of the hand, hand shape and colour characteristics of the scene.

Results: The output from the algorithm was compared with a manual labelling of the video, yielding F1-scores of 0.74 ± 0.15 for the left hand and 0.73 ± 0.15 for the right hand. From the resulting frame-by-frame binary data, functional hand use measures were extracted: the amount of total interaction as a percentage of testing time, the average duration of interactions in seconds, and the number of interactions per hour. Moderate and significant correlations were found when comparing these output measures to the results of the manual labelling, with $\rho = 0.40, 0.54$ and 0.55 respectively.

Conclusions: These results demonstrate the potential of a wearable egocentric camera for capturing quantitative measures of hand use at home.

Keywords: Tetraplegia, Upper extremity, Outcome measure, Egocentric vision, Community-based rehabilitation, Rehabilitation engineering

Introduction

Upper extremity (UE) impairment can severely limit individuals' ability to perform activities of daily living (ADLs). The recovery of hand function is consequently of great importance to individuals with cervical spinal cord injuries (SCI) [1].

In order to assess new interventions and improve upon currently existing rehabilitation approaches, outcome

measures that can accurately quantify hand function in a natural context are needed. The majority of currently available assessment tools measure impairment or functional limitation and rely on direct observation by a clinician or investigator [2]. On the other hand, limited tools are available to describe how individuals with SCI use their hands in their usual environment or community, restricting our understanding of how changes in UE function impact activity and participation. Often measures that describe function in the community are restricted to potentially biased self-report to try to estimate independence in ADLs [3–6].

* Correspondence: jose.zariffa@utoronto.ca

¹Institute of Biomaterials & Biomedical Engineering, University of Toronto, Toronto, Ontario, Canada

²KITE, Toronto Rehabilitation Institute, University Health Network, Toronto, Ontario, Canada

Full list of author information is available at the end of the article



In an attempt to gauge hand function at home and in the community, multiple studies have explored the use of wearable sensors, particularly accelerometers or inertial measurement units (IMUs). While these devices are small and easily worn, they lack the resolution to capture the complexity of functional hand use. These wearable sensors are limited to capturing arm movements. In studies with hemiparetic stroke survivors, accelerometry has most typically been used to examine the ratio of activity between the impaired and unimpaired arms [7, 8]. Recently in SCI, accelerometers have been used to measure wheeling movements and to assess the laterality of the injuries [9, 10]. Although these studies demonstrated a relationship between accelerometry measures and independence, this approach is not able to reveal direct information about how the hand is used in functional tasks [10]. A recent study in stroke survivors found that improvements in motor function according to clinic-based outcome measures (capacity, as defined in Marino's modification of the International classification of functioning, disability and health (ICF) model [11]) do not necessarily translate into increased limb use in the community (performance), as measured by accelerometry [12]. These findings point to the need for novel outcome measures that can directly measure performance and better describe the impact of new interventions on the daily lives of individuals with SCI.

A system based on a wearable camera has the potential to overcome the limitations of existing outcome measures for UE function. First-person (egocentric) cameras record the user's point of view. Unlike a wrist-worn accelerometer that can only capture arm movement information, an egocentric video provides detailed information on hand posture and movements, as well as on the object or environment that the hand is interacting with. Multiple studies have explored the use of computer vision techniques to extract information about the hand in egocentric videos, though typically not in the context of rehabilitation. Key problems to be solved include hand detection (locating the hand in the image) as well as segmentation (separating the outline of the hand from the background of the image) [13–22]. Beyond hand detection and segmentation, there have also been attempts to use egocentric videos for activity recognition and object detection in ADLs [23–28]. However, generalizability can be a challenge in such systems, given the large variety of activities and objects found in the community.

In our previous work, we proposed to detect interactions of the hand with objects using egocentric videos. This binary classification (interaction or no interaction) is intended to form the basis for novel outcome measures to describe hand function in the community. We have demonstrated, in the able-bodied population, the

possibility of a hand-object interaction detection system, where a system can detect and log whether or not the hand is manipulating an object for a functional purpose. An interaction between an object and the hand is only considered to happen when the hand manipulates the object for a functional purpose; for example, resting a hand on the object would not constitute an interaction [28]. The present work expands the development of the hand-object interaction detection system by describing a novel complete algorithmic pipeline and evaluating for the first time its application to individuals with SCI.

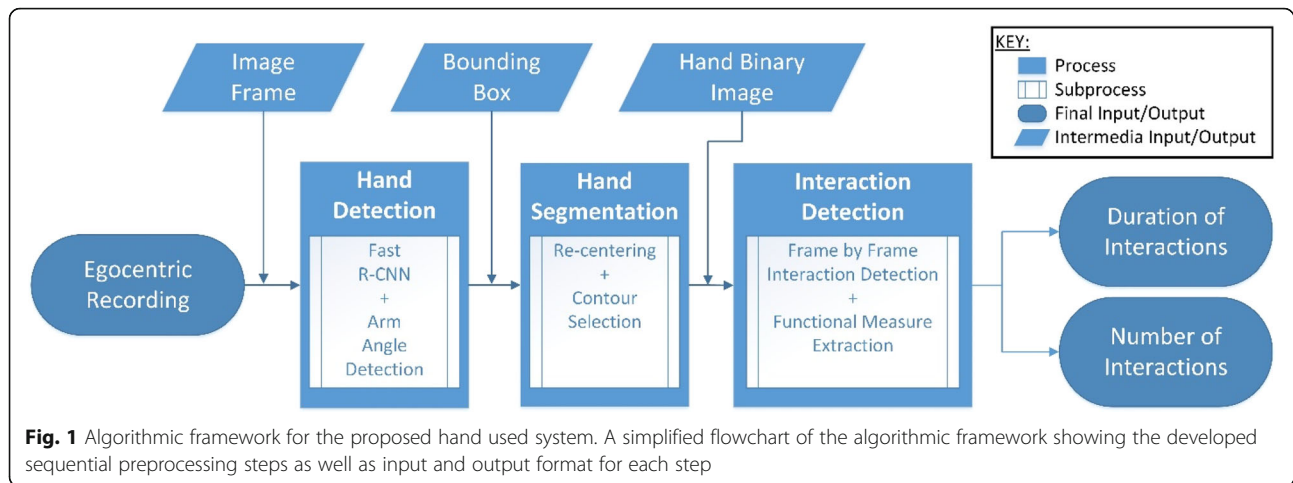
Methods

Dataset and participants

A dataset from participants with cervical SCI was created, the Adaptive Neurorehabilitation Systems Laboratory dataset of participants with SCI ("ANS SCI"). The ANS SCI dataset consists of egocentric video recordings reflective of ADLs obtained using a commercially available egocentric camera (GoPro Hero4, San Mateo, California, USA) worn by the participant overhead via a head strap. The video was recorded at 1080p resolution at 30 frames per second. However, a reduced resolution of 480p was used for analysis in order to reduce computation time, given that higher resolutions are not necessary for our application. The data collection was performed in a home simulation laboratory. Specifically, this study involved recording from 17 participants with SCI, performing different common interactive ADL tasks identified by the American Occupational Therapy Association (AOTA) as important (for example, personal care, eating, and leisure activities) [29]. Each participant performed a total of approximately 38 ADL tasks in several environment (kitchen, living room, bedroom, bathroom). Participants were also asked to perform non-interactive tasks, which involved hand at rest and hand being waved in the air without any interaction with an object.

Algorithmic framework

In order to capture the hand-object interactions, the framework developed consisted of three processing steps. First, the hand location was determined in the form of a bounding box. Next, the bounding box was processed for hand segmentation, where the pixels of the hand were separated from the non-hand pixels (i.e. the background). With the hand being located and segmented, image features including hand motion, hand shape and colour distribution were extracted for the classification of hand-object interaction. The flowchart in Fig. 1 summarizes the algorithmic framework. Note that our previous work in [28] focused on the interaction detection step, whereas in the present study the hand detection and segmentation steps have been



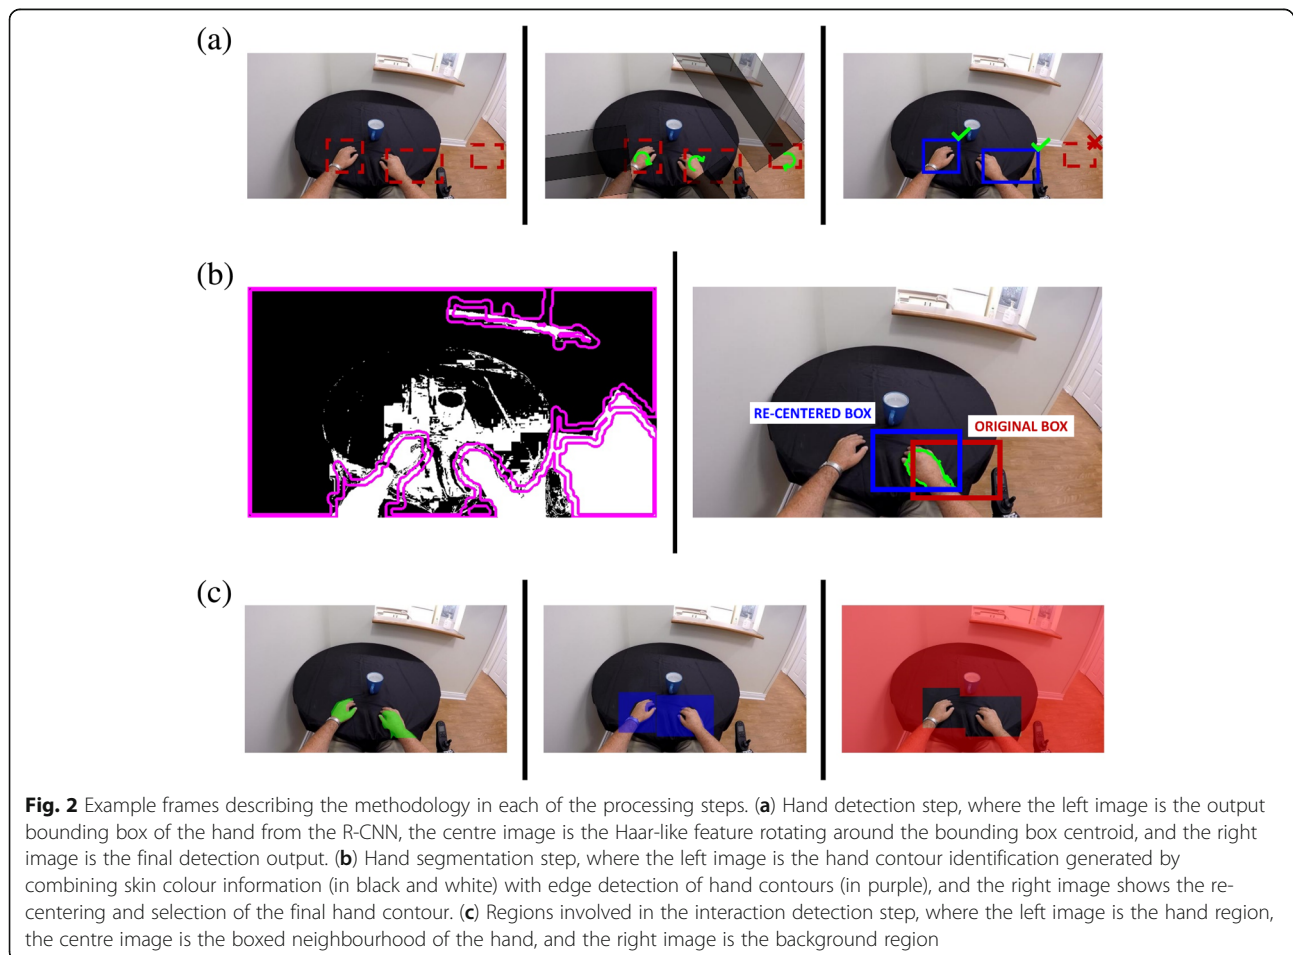
developed and included, allowing us to evaluate the complete pipeline from raw video to interaction metrics.

Hand detection

The detection of the hand was accomplished using a convolutional neural network (CNN). Specifically, the faster regional CNN (“faster R-CNN”) architecture [30]

was selected. The CNN outputs the coordinates of the box surrounding the detected hand (Fig. 2a). The CNN was trained on the even frames of eight participants in our dataset (participants # 10–17) resulting in 33,256 frames manually labelled with hand bounding boxes.

The output bounding box from the CNN was further processed to help eliminate false positives and to extract



arm angle information, using a rotating Haar-like feature. This feature consists of three adjacent parallel rectangles extending from the centre and sides of the bounding box to the border of the image; this was inspired by the Haar-like feature used in real-time face detection [31]. This selected centre rectangular region was compared to the two parallel regions on either side, by taking the average of the difference between the coefficient of variation in the centre rectangle and in each of the two side rectangles; this difference is expected to be greatest when the center rectangle is aligned with the arm. This Haar-like feature was rotated 360 degrees around the centroid of the hand bounding box. The output from the rotating feature, computed in 1-degree steps was then summed over bins of 5 degrees to obtain a feature vector with 72 values. This vector was used as input to a random forest classifier with binary output, to confirm whether the bounding box truly included a hand (Fig. 2 a). The training of this classifier was performed on the odd frames of eight participants (i.e. participants # 10–17, the same ones used to train the CNN), which again consisted of 33,256 frames labelled with bounding boxes.

Additionally, given the egocentric viewpoint, the arm angle was used to determine which hand was being used (i.e. user's left or right hand, as well as a hand belonging to another individual). This was established by summing the values in the Haar-like feature vector in the top two quadrants (Quadrants I + II, 0–180 degrees), bottom right quadrant (Quadrant III, 180–270 degrees), or bottom left quadrant (Quadrant IV, 270–360 degrees). The hand was determined to be the left hand, right hand or other person's hand if the quadrant with the highest sum was IV, III or I + II, respectively.

Hand segmentation

The output of the hand detection stage was processed to segment the hand (i.e. identify pixels belonging to the hand). The segmentation process consisted of the following steps:

1. Identify candidate hand contours. This step combined colour and edge information. For colour, the RGB image was back-projected using a histogram obtained from a generic mixture-of-Gaussians skin colour model [32]. The back-projected image was thresholded at 0.75 of the maximum value (this value was selected empirically and may change depending on the camera system). To obtain edge information, a Structured Forests edge detection [33] method was used, which was specifically trained on hand images to preferentially identify hand edges. For the purposes of training this model, a publicly available dataset [34] with pixel-level hand annotations [35] was used. The output of the edge detection was thresholded at 0.05 of the maximum value (again

selected empirically), and morphological operations (dilation followed by erosion) were used to remove small gaps in the contours detected. Lastly, the edge information was used to improve the delineation of the hands in the colour-based segmentation (Fig. 2 b).

2. Re-centre the bounding box and select the final hand contour. The bounding box from the hand detection step was applied to the image from step 1. Within this box, we sought to identify the contour most likely to be the hand and re-centre the box around it to minimize the occurrence of truncated hands. The determination of this final hand contour was based on shape, again using the information from the edge detection. From the list of contours obtained in the previous step from the combined colour and edge information, we selected the one that had the highest overlap with the edge image, filled in using dilation. Prior to this determination, any contour whose area was less than 2% or more than 75% of the bounding box area was eliminated; similarly, contours with arc lengths between 90 and 110% of the bounding box perimeter were removed. Once the hand contour was determined, a new box and associated centroid was selected from the mean of the hand contour's centroid and top pixel. This step promotes maximum coverage of the hand and not the arm, which is often located below the hand in an egocentric view.

Interaction detection

The interaction detection is built on our previous work [28], with additional colour features included. Furthermore, the interaction detection for this study also supported multiple hands, with the user's left or right hand as well as the hands of other individuals being identified using the arm angle obtained via the hand detection step. Three categories of features were extracted in the interaction detection:

1. Object Motion. Motion features assume that an object being held in the hand will be moving with a similar direction and speed as the hand. Conversely, an item in the frame that is not being interacted with is more likely to have a motion similar to that of the background. To capture this distinction, a dense optical flow map [36] was separated into three regions: the segmented hand, the bounding box around the hand, and the background (Fig. 2 c). Note that the bounding box size and location are equivalent to the segmentation step after re-centering. The dense optical flow from each region was summarized into respective histograms of magnitude and direction, each with 15 normalized bins. The final feature consisted of two vectors: the subtraction of the histograms of the bounding box near the hand from those of the hand, and the

- subtraction of the histograms of the bounding box near the hand from those of the background.
2. Hand shape. Certain characteristics of hand shape are indicative of hand activity (i.e. grip type). The hand shape was represented using histograms of oriented gradients (HOG), implemented as in our previous work [28]. The HOG features were extracted from the same bounding box used in the motion feature analysis above. The selected image regions were then resized to 10% of the frame height and 15% of the frame width to guarantee identical dimensions before principal component analysis (PCA) was applied. The HOG feature vector dimension was reduced from 960 to 60 to keep the dimensions identical to that of the motion cues feature.
 3. Object Colour. The objects near the hand may have a different colour than the background and the hand. The closer the objects to the hand, the greater the likelihood of interaction. HSV colour histograms were extracted from the same three regions as in the motion feature analysis described earlier. The two comparison scores are extracted based on the Bhattacharyya distance between the histograms of the bounding box and the hand, as well as those of the bounding box and the background.

Finally, the combined features were input to a random forest binary classifier (see “System Evaluation” below). The random forest used 150 trees, following the work in [28].

The classifier was trained using data manually labelled by a human observer, where each frame was either classified as interaction or no interaction. An interaction between the object and the hand was only considered to happen when the hand manipulates the object for a functional purpose. For example, resting a hand on the object or moving a hand through space would not constitute an interaction. Labelling was performed on a frame-by-frame basis, with no bounding box or segmentation shown to the annotator. An annotator was instructed to label the user’s left hand, the user’s right hand, and other people’s hands separately.

System evaluation

The interaction detection was evaluated using Leave-One-Subject-Out cross-validation. We applied our system to 9 participants from our dataset (participants # 1–9), for whom interactions had been manually labelled and whose data had not been used to train the hand detection algorithms. In the cross-validation process, each participant in turn was left out for testing while the other 8 participants were used for training. Depending

on the participant being left out, the training set of 8 participants on average consisted of $28,057 \pm 2,334$ frames (935 ± 77 s) of interaction (48%) and $30,850 \pm 3,369$ frames ($1,028 \pm 112$ s) of no interaction (52%). The participant in the testing set consisted on average of $3,507 \pm 3,369$ frames (116 ± 112 s) of interaction (48%) and $3,856 \pm 2,334$ frames (128 ± 77 s) of no interaction (52%). The classification was compared with manually labelled data.

Extraction of functional measures

In order to translate the frame-by-frame results into more easily interpretable measures, the binary output of hand-object interaction detection was processed to extract: the amount of total interaction as a percentage of testing time, average duration of individual interactions in seconds, and number of interactions per hour (i.e., number of interactions normalized to video segment duration).

Hand-object interaction outputs from the algorithm were assigned to one of two timelines, depending on arm angle: user’s left hand and user’s right hand. Failures in hand detection or segmentation could result in missing frames at the interaction detection stage. To address this issue, interactions were prolonged for 90 frames (3 s) if the hand was suddenly lost. Outside of this range, the frame was classified as a non-interaction. Furthermore, a moving average filter was applied to the binary frame-by-frame timelines of interaction. The moving average promotes temporal smoothness in the output, reduces the impact of labelling errors in the start and stop of interaction, and corrects minor faulty hand detection and segmentation. The filter window was chosen to have a length of 120 frames (corresponding to 4 s), with equally weighted samples. This duration was optimized empirically on the basis of its ability to meaningfully summarize the number and duration of underlying activities. The moving average was similarly applied to the manually labelled timelines.

The output of the moving average was then normalized by subtracting the minimum value over the entire video and dividing by the difference of maximum and minimum values. The result was thresholded at 0.5 (i.e. values greater than 0.5 were considered to be interactions). Number and duration of interactions were extracted from this filtered output.

For each of the three metrics, the correlation between the algorithm output and the manually labeled data over the 9 tested participants was computed. The resulting correlations were tested against a hypothesis of no correlation using a one-tailed (right) test. Significance was set at $p = 0.05$. Where data were normally distributed, a Pearson Correlation was chosen, otherwise a Spearman correlation was used.

Results

Training and testing sets

The demographic and injury characteristics of the participants are provided in Table 1. The interaction detection was evaluated on participants # 1–9. This testing set consisted of 8 males and 1 female with an average age of 52 ± 13 years and an average upper extremity motor score (UEMS) of 17 ± 4 . As mentioned previously, participants # 10–17 were used for the training of the preprocessing steps, namely the CNN-based hand detection and the classifier based on the Haar-like arm angle features. This hand detection training set consisted of 7 males and 1 female with an average age of 47 ± 11 years and an average UEMS of 20 ± 3 .

This separation of the dataset was designed such that the final interaction detection evaluation was conducted on participants (# 1–9) whose data had not been used in training any parts of the preprocessing algorithms. The two subsets have similar characteristics in terms of gender, age and UEMS.

Interaction detection

The filtered frame-by-frame interaction detection results are quantified using classification accuracy and F1-score (harmonic mean of precision and recall) for each participant; these are shown in Table 2. The average accuracy and F1-score are calculated for all the participants in the interaction detection evaluation. Over the 9 participants, the F1-scores were 0.74 ± 0.15 for the left hand and $0.73 \pm$

Table 2 F1-Score and accuracy for left (L) and right (R) hand for each participant as well as the average for each of the features

	F1-score		Accuracy	
	L	R	L	R
Participant				
1	0.54	0.53	0.42	0.42
2	0.6	0.75	0.73	0.79
3	0.86	0.73	0.79	0.63
4	0.85	0.58	0.79	0.59
5	0.72	0.67	0.75	0.7
6	0.55	1	0.48	0.99
7	0.84	0.8	0.82	0.69
8	0.78	0.63	0.67	0.51
9	0.93	0.91	0.88	0.88
Mean \pm S.D.	0.74 ± 0.15	0.73 ± 0.15	0.70 ± 0.16	0.68 ± 0.18
Feature				
Optical Flow	0.73 ± 0.14	0.70 ± 0.13	0.68 ± 0.16	0.66 ± 0.15
HOG	0.72 ± 0.12	0.72 ± 0.14	0.69 ± 0.12	0.68 ± 0.15
Colour Histogram	0.70 ± 0.12	0.66 ± 0.17	0.68 ± 0.10	0.66 ± 0.16

0.15 for the right hand. The accuracies were 0.70 ± 0.16 for the left hand and 0.68 ± 0.18 for the right hand.

Functional hand use

The three proposed interaction metrics were compared between the automated output and the manual observations,

Table 1 Participant Demographics and Injury Characteristics

Participant	Age (Years)	Sex	Level of Injury	AIS grade	Traumatic (T)/ Non-traumatic (NT)	Time since injury (Years)	Upper Extremity Motor Score (UEMS)
1	63	Male	C5-C6	A ^a	T	8	15
2	58	Male	C3-C5	D	T	1	24
3	59	Male	C2-C6	D	T	1	20
4	55	Male	C7-T1	C/D ^a	T	4	18
5	56	Male	C2-C7	D	T	2	19
6	56	Male	C5-C6	D	T	2	16
7	20	Male	C5	B	T	4	9
8	58	Male	C5	C/D ^a	T	32	13
9	44	Female	C6-C7	A	T	20	20
10	51	Male	C4-C6	D	T	1	22
11	34	Male	C5-C6	C	T	5	21
12	40	Female	C2-T1	D	NT	2	20
13	70	Male	C4-C6	C	T	1	24
14	42	Male	C4-C6	B	T	0.4	16
15	56	Male	C1-C6	D	NT	0.3	23
16	44	Male	C4-C5	B ^a	T	21	21
17	41	Male	C6-C7	A ^a	T	20	14

^aThese AIS grades are based on self-report

and the results are shown in Fig. 3. For all three metrics (proportion of interaction time, average duration of interactions, and number of interactions per hour), a moderate, significant correlation was found between the actual and predicted metrics ($\rho = 0.40$, $p = 0.04$; $\rho = 0.54$, $p = 0.01$; $\rho = 0.55$, $p = 0.01$, respectively).

Features analysis

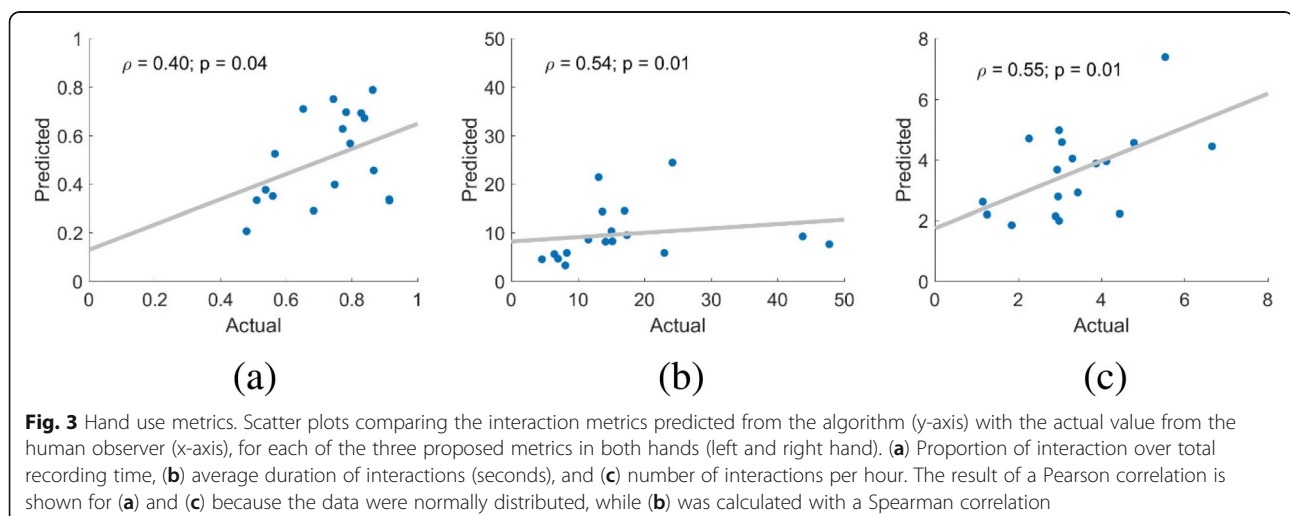
Moreover, we sought to understand the importance of each feature for hand-object interaction. This involved selecting only one type of feature at a time (optical flow, HOG and colour histogram) in the classifier. The overall average accuracy and F1-score for all subjects are articulated in Table 2. There was a general trend towards relatively slight decreases in performance when features were removed, with the resulting F1-scores ranging from 0.66 ± 0.17 to 0.73 ± 0.14 , and accuracies ranging from 0.66 ± 0.15 to 0.69 ± 0.12 .

Discussion

Outcome measures capable of capturing the hand function of individuals with SCI in their home environments are needed to better assess the impact of new interventions. Current methods are limited to self-report or direct observation. In this study, we demonstrated the feasibility of measuring the functional hand use of individuals with UE dysfunction resulting from cervical SCI using a novel wearable system. The system is based on a wearable camera and a custom algorithmic framework capable of automatically analyzing hand use in ADLs at home. The frequency and duration of hand-object interaction could potentially serve as the basis for new outcome measures that provide clinicians and researchers with an objective measure of an individual's independence in UE tasks.

We demonstrated the criterion validity of an interaction-detection system by comparing the algorithm outputs to manually labelled data, taken here to be the gold standard. The binary hand-object interaction classification was shown to be robust across multiple ADLs, environments, and individuals, with an average F1-score of 0.74 ± 0.15 for the left hand and 0.73 ± 0.15 for the right hand after the moving average (Table 2). For the system evaluation, we sampled activities at random among participants; thus, each participant had a unique set of activities during the leave-one-subject-out cross-validation, meaning that the system was tested on ADLs that it had not necessarily seen before. The evaluation was performed in a non-scripted manner; however, the tasks were specified (i.e. participants were only given the name of the tasks with no instruction on how these should be performed). The dataset generated was balanced between the proportion of interaction and non-interaction frames (48 and 52% respectively).

As visible in Table 2, performance varied between participants and sometimes between hands. It is important to note that since subsets of the data were labelled and used for the interaction detection evaluation, participants were not all evaluated on the same activities. This difference in activities may account for variations in F1-score and accuracy across individuals. For example, in participant # 1, the data used consisted largely of activities in the living room and involved finer manipulation activities such as eating potato chips and picking up a coin. However, perhaps the main factor that resulted in faulty hand detections was scenes consisting of wooden objects and flooring. Based on qualitative observation the R-CNN tended to produce false positives on wood. While the Haar-like methods are able to eliminate the majority of the false positives, their impact is greater in certain environments and scenes than others. The



inconsistency observed is also related to differences in classifier performance between the left and right hand in some cases, for example in participant # 4. Direct observation of the frames again point to the poor detection of the right hand, which is closer to the wooden floor, leading to some missed identifications. It is therefore a point of limitation given that the hand-object interaction directly relies on this pre-processing step due to the serial architecture of the pipeline. We anticipate that improvements in hand detection will improve the performance of the overall pipeline. Although minor, variations in task and injury characteristics may also contribute to lower accuracy. For example, participant # 6 used a joystick to control their wheel chair. Such activities involved a relatively small movement, which can be difficult to reliably identify in the egocentric video. Aspects of participant postures post-injury, such as pronounced curving of the back, were also found to reduce the Field of View (FOV) of the camera, resulting in the hand not appearing in the frames.

The functional measures of interest are the number and duration of hand-object interactions. Given the variability in activities and labelled video durations between participants, the measures were summarized using the amount of total interaction normalized to testing time, the average duration of individual interactions, and the number of interactions per hour. A moving average was applied to reduce noise from short-term fluctuations caused by a potentially faulty preprocessing step (e.g. hand detection or segmentation), misclassification, or fast interaction sub-tasks within an activity. Perhaps most important is that the moving average highlights important long-term trends that are obscured by a frame-wise analysis of long recordings. Consider the timeline in Fig. 4, where the moving average of the binary interaction graph is plotted against time. Here, the transitions between tasks are clearly observed as changes between the interaction and non-interaction classes. This example illustrates the value of capturing these metrics for the measurement of hand use at home.

The comparison of the predicted and actual values for the interactions per hour, average duration of interactions, and proportion of interaction time for every participant are summarized in Fig. 3. These results indicate to what extent the conclusions about hand use provided by the automated system correlate to the results from a human observer. We found moderate and significant correlations, which provides evidence for the viability of this approach, while highlighting the need for continued improvements in the algorithms used for interaction detection. Deviations between the predicted and actual metrics are likely a combination of errors in hand detection, interaction detection, and the filtering process, providing multiple targets to increase performance. It is also

possible that outliers in the limited number of participants in this evaluation may have partially skewed the results, and that a larger dataset would yield a more robust outcome (e.g., consider the influence of the two data points with high actual values but low predicted values in Fig. 3b).

A challenge in extracting hand use metrics based on frame-by-frame interaction detections is that complex timelines of interactions (Fig. 4) need to be simplified into simple outcome units or scores. Further exploration is warranted in how to best summarize the interactions taking place over time among a variety of activities, and in optimizing the normalization used for comparing different recording times. The correlations obtained do suggest the potential of such metrics for summarizing the interaction detection, but further evaluation on larger and more diverse dataset is required.

To the best of our knowledge, the hand-object interaction problem has not been studied before in individuals with neurological injuries. Thus, we sought to understand what types of information were most beneficial to accurate interaction detection. This understanding is crucial in that hand movements or postures may vary substantially between individuals with SCI having different patterns of injury or impairment. The analysis in Table 2 revealed that all the features (optical flow, HOG and colour histogram) were able to contribute to the hand-object interaction classification. The combined features produced the overall highest performance, but the gain over using a single feature type was minimal. The average F1-scores when using only optical flow or when using only HOG were very close, and then followed by colour histogram. In contrast, in our previous work with able-bodied participants [28], using only HOG was found to be more useful than using only optical flow. The fact that this finding no longer held true with participants with SCI suggests that relying on shape (i.e., hand posture) may be less beneficial in the presence of varying levels of impairment and compensatory postures. Classifiers tailored to different types and severities of injury could potentially increase the performance achievable with shape features.

Our validity evaluation focused on a comparison with manually labelled interaction data, and we did not attempt to correlate the results with existing outcome measures. The collected dataset consisted of standardized and pre-determined activities. As such, it is not representative of the frequency of hand use in a natural home environment and includes tasks that some participants may not normally perform on their own. While we expect that individuals with better hand function may independently perform UE tasks more frequently in their daily life, evaluating this relationship will require collecting data in the home in an unscripted manner. In the

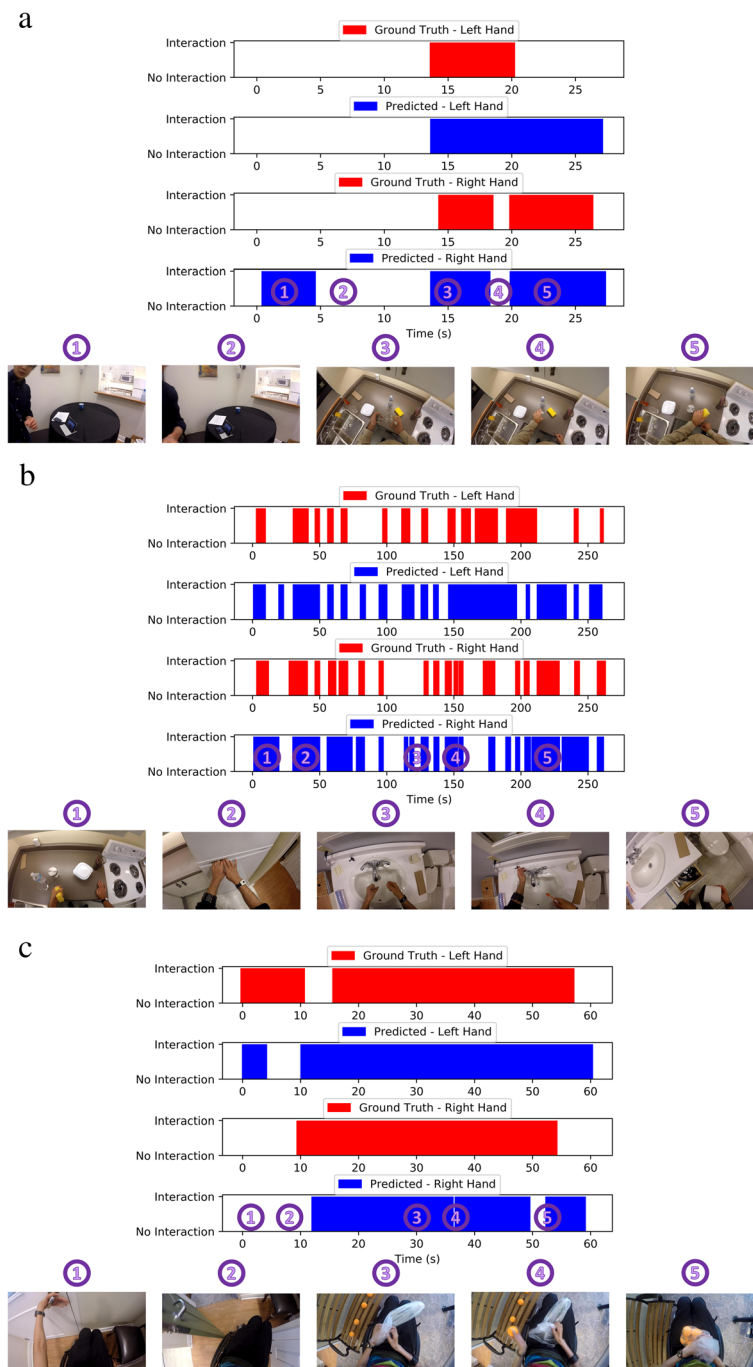


Fig. 4 Example binary hand-object interaction graphs of 3 participants. The graphs compare the predicted interactions from the algorithm output to the actual interactions from the manually labeled data, after applying the moving average filter. Example frames of the activities in different segments of videos are shown underneath. **(a)** Participant # 2. **(b)** Participant # 5. **(c)** Participant # 9. Note that in some cases the videos were briefly paused in between the activities shown

present study, UEMS data was collected for demographic and study inclusion purposes, but its relationship to the interaction metrics was not investigated, for the reasons above. With appropriate data collection in the home, an important extension of our analysis will

include the comparison of the interaction-detection metrics with other UE outcomes measures. For example, greater independent use of the hand would be expected to correlate with the Graded Redefined Assessment of Strength, Sensibility and Prehension (GRASSP), which is

sensitive to fine gradations in hand function after SCI, or the Spinal Cord Independence Measure III (SCIM), which aims to quantify independence in the community. Of note, the actual amount of hand use at home (corresponding to Performance in the revised ICF model [11]) represents a different construct than hand impairment or the capacity to perform activities, which are more closely captured by the GRASSP and SCIM. Nonetheless, some degree of relationship is still expected. While these investigations will be the subject of future work in parallel with continued algorithm improvements, the present study was required to validate the feasibility of the interaction detection step before metrics derived from it (i.e., number and duration of interactions) can be meaningfully compared to other measures.

Another limitation in this study is that the hands of individuals other than the user were not analysed. At home, it would be of great interest to quantify caregiver assistance by tracking their manipulations of objects in front of the individual with SCI. The system described supports the detection of caregiver hands and their associated interactions (a valuable metric for reliance on care). However, limited data of this type was collected in this study, because participants most often chose to complete the tasks on their own or skip them completely, instead of asking for assistance. This behaviour can likely be explained by the lack of familiarity with the researchers and the pressure to perform in a research setting. In the home setting, however, we expect that more caregiver actions will be captured by the system, which will provide important information about independence.

There are remaining technical challenges to be explored beyond this study, namely the improvement of the preprocessing steps and computation time. The detection and segmentation of the hand are based on skin and hand shape characteristics, which can be influenced by glare and objects with a similar shape or colour to the hand (e.g. a wooden floor, table, or door). The performance of the overall hand-object interaction system is dependent on the accuracy of the hand detection and segmentation steps and can be expected to improve with them. For example, in a preliminary analysis, only manually selected frames with good hand segmentation were used, and the interaction detection performance was found to be 0.81 for F1-score in 15,471 frames from 3 participants [37]. Secondly, to avoid privacy and usability concerns for the user [38], the ideal system should be mobile and process the video in real-time, storing only the extracted metrics and not the video. Unfortunately, the proposed algorithm remains computationally expensive for a mobile system. It takes on average 2.70 s per frame to process hand-object interaction from the input image frames to the output metrics (Intel- i7-8700

k@4.8GHzOC, DDR4-16GB@3200MHzOC, GTX1080Ti OC-GDDR5X-11GB, Ubuntu16.04 LTS 64-bit). Lastly, the performance of the algorithms will need to be evaluated in a wider range of environments, with challenges that may include imperfect lighting, differences in camera orientation, and more diverse tasks. Thus, a more varied dataset is needed both for evaluation purposes and to ensure that model parameters that were selected empirically here can be refined to yield robust performance. Improvements using recent computer vision techniques for hand detection and segmentations [39], as well as better feature selection, have the potential to improve performance and speed.

Conclusions

In this study, we demonstrated the potential of an ego-centric wearable camera system for capturing an individual's functional hand use in the home environment. Novel outcome measures based on this system have the potential to fill the research gaps in home-based assessment of the UE in neurorehabilitation, which currently relies heavily on self-report. We have demonstrated the feasibility of the interaction-detection process and illustrated how this concept can be used to derive meaningful outcome measures, such as the number and duration of independent, functional object manipulations. More broadly, our study provides a framework for future research in UE assessment within the broader community.

Abbreviations

ADLs: Activities of daily living; ANS SCI: Adaptive Neurorehabilitation Systems Laboratory dataset of participants with spinal cord injury; AOTA: American Occupational Therapy Association; CNN: Convolutional neural network; cSCI: Cervical spinal cord injury; faster R-CNN: Faster regional convolutional neural network; GRASSP: Graded Redefined Assessment of Strength, Sensibility and Prehension; HOG: Histograms of oriented gradients; ICF: International classification of functioning, disability and health; PCA: Principal component analysis; SCI: Spinal cord injury; SCIM: Spinal Cord Independence Measure III; UE: Upper extremity

Acknowledgments

The authors would like to thank Pirashanth Theventhiran for his valuable assistance in the data labelling process. The authors would also like to thank all the participants of the study.

Authors' contributions

JL, SKR and JZ designed the study. JL and ERS carried out the data collection. JL, TC, and RJV participated in algorithm development and implementation. JL and JZ drafted the manuscript. All authors read and approved the final manuscript.

Funding

This study was supported in part by the Natural Sciences and Engineering Research Council of Canada (RGPIN-2014-05498), the Rick Hansen Institute (G2015-30), and the Ontario Early Researcher Award (ER16-12-013).

Availability of data and materials

The dataset will be shared for academic purposes upon reasonable request. Please contact author for data requests.

Ethics approval and consent to participate

The study participants provided written consent prior to participation in the study, which was approved by the Research Ethics Board of the institution (Research Ethics Board, University Health Network: 15–8830).

Consent for publication

All study participants consented to have non-identifying videos and images used for publication.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Institute of Biomaterials & Biomedical Engineering, University of Toronto, Toronto, Ontario, Canada. ²KITE, Toronto Rehabilitation Institute, University Health Network, Toronto, Ontario, Canada. ³Department of Physical Therapy, University of Toronto, Toronto, Ontario, Canada.

Received: 25 January 2019 Accepted: 25 June 2019

Published online: 05 July 2019

References

- Anderson KD. Targeting recovery: priorities of the spinal cord injured population. *J Neurotrauma*. 2004;21:1371–83.
- Kalsi-Ryan S, Curt A, Verrier MC, Fehlings MG. Development of the graded redefined assessment of strength, sensibility and Prehension (GRASSP): reviewing measurement specific to the upper limb in tetraplegia. *J Neurosurg Spine*. 2012;17:65–76.
- Linacre JM, Heinemann AW, Wright BD, Granger CV, Hamilton BB. The structure and stability of the Functional Independence Measure. *Arch Phys Med Rehabil*. 1994;75:127–32.
- Heinemann AW, Linacre JM, Wright BD, Hamilton BB, Granger C. Relationships between impairment and physical disability as measured by the functional independence measure. *Arch Phys Med Rehabil*. 1993;74:566–73.
- Catz A, Itzkovich M, Agranov E, Ring H, Tamir A. SCIM–spinal cord independence measure: a new disability scale for patients with spinal cord lesions. *Spinal Cord*. 1997;35:850.
- Itzkovich M, Gelernter I, Biering-Sorensen F, Weeks C, Laramie MT, et al. The spinal cord Independence measure (SCIM) version III: reliability and validity in a multi-center international study. *Disabil Rehabil*. 2007;29:1926–33.
- Noorkoiv M, Rodgers H, Price CI. Accelerometer measurement of upper extremity movement after stroke: a systematic review of clinical studies. *J Neuroeng Rehabil*. 2014;11:144.
- Lemmens RJM, Timmermans AAA, Janssen-Potten YJM, Janssen-Potten SANTDP, et al. Accelerometry measuring the outcome of robot-supported upper limb training in chronic stroke: a randomized controlled trial. *PLoS One*. 2014;9:e96414.
- Broggioli M, Schneider S, Popp WL, Albisser U, Brust AK, et al. Monitoring upper limb recovery after cervical spinal cord injury: insights beyond assessment scores. *Front Neurol*. 2016;7:142. <https://www.frontiersin.org/article/10.3389/fneur.2016.00142>.
- Broggioli M, Popp WL, Albisser U, Brust AK, Frotzler A, et al. Novel sensor technology to assess independence and limb-use laterality in cervical spinal cord injury. *J Neurotrauma*. 2016;33:1950–7.
- Marino RJ. Domains of outcomes in spinal cord injury for clinical trials to improve neuro- logical function. *J Rehabil Res Dev*. 2007;44:113.
- Waddell KJ, Strube MJ, Bailey RR, Klaesner JW, Birkenmeier RL, et al. Does task-specific training improve upper limb performance in daily life post-stroke? *Neurorehabil Neural Repair*. 2017;31:290–300.
- Li C, Kitani KM. Pixel-level hand detection in ego-centric videos. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2013. p. 3570–7.
- Li C, Kitani KM. Model recommendation with virtual probes for egocentric hand detection. In: Proceedings of the IEEE International Conference on Computer Vision; 2013. p. 2624–31.
- Serra G, Camurri M, Baraldi L, Benedetti M, Cucchiara R. Hand segmentation for gesture recognition in ego-vision. In: Proceedings of the 3rd ACM international workshop on interactive multimedia on mobile & portable devices: 31–36ACM; 2013.
- Betancourt A, Morerio P, Marcenaro L, Barakova E, Rauterberg M, Regazzoni C. Towards a unified framework for hand-based methods in first person vision, in Multimedia & Expo Workshops (ICMEW), 2015 IEEE International Conference on:1–6IEEE2015.
- Fathi A, Rehg JM. Modeling actions through state changes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2013. p. 2579–86.
- Likitlersuang J, Zariffa J. Arm angle detection in egocentric video of upper extremity tasks, in World congress on medical physics and biomedical engineering, June 7–12, 2015. Toronto, Canada: Springer; 2015. p. 1124–7.
- Bambach S, Lee S, Crandall DJ, Yu C. Lending a hand: detecting hands and recognizing activities in complex egocentric interactions. In: Proceedings of the IEEE International Conference on Computer Vision; 2015. p. 1949–57.
- Zariffa J, Popovic MR. Hand contour detection in wearable camera video using an adaptive histogram region of interest. *J Neuroeng Rehabil*. 2013;10:114.
- Khan AU, Borji A. Analysis of Hand Segmentation in the Wild, arXiv preprint arXiv:1803.03317; 2018.
- Fathi A, Farhadi A, Rehg JM. Understanding egocentric activities, in Computer vision (ICCV), 2011 IEEE international conference on:407–414IEEE; 2011.
- Fathi A, Ren X, Rehg JM. Learning to recognize objects in egocentric activities. In: Computer vision and pattern recognition (CVPR), 2011 IEEE conference on:3281– 3288IEEE; 2011.
- Ishihara T, Kitani KM, Ma WC, Takagi H, Asakawa C. Recognizing hand-object interactions in wearable camera videos. In: Image processing (ICIP), 2015 IEEE international conference on:1349– 1353IEEE; 2015.
- Matsuo K, Yamada K, Ueno S, Naito S. An attention-based activity recognition for egocentric video. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops; 2014. p. 551–6.
- Pirsiavash H, Ramanan D. Detecting activities of daily living in first-person camera views. In: Computer vision and pattern recognition (CVPR), 2012 IEEE conference on:2847– 2854IEEE; 2012.
- Ren X, Philipose M. Egocentric recognition of handled objects: Benchmark and analysis. In: Computer vision and pattern recognition workshops, 2009. CVPR work- shops 2009. IEEE computer society conference on:1–8IEEE; 2009.
- Likitlersuang J, Zariffa J. Interaction detection in egocentric video: towards a novel outcome measure for upper extremity function. *IEEE J Biomed Health Inform*. 2016;22(2):561–9.
- Association American Occupational Therapy. Occupational therapy practice framework: domain and process. *Am J Occup Ther*. 2008;62:625–83.
- Ren S, He K, Girshick R, Sun J. Faster r-cnn: towards real-time object detection with region proposal networks. In: Advances in neural information processing systems; 2015. p. 91–9.
- Viola P, Jones MJ. Robust real-time face detection. *Int J Comput Vis*. 2004; 57:137–54.
- Jones MJ, Rehg JM. Statistical color models with application to skin detection. *Int J Comput Vis*. 2002;46:81–96.
- Dollár P, Zitnick CL. Structured forests for fast edge detection. In: Proceedings of the IEEE International Conference on Computer Vision; 2013. p. 1841–8.
- Betancourt A, Lopez MM, Regazzoni CS, Rauterberg M. A sequential classifier for hand detection in the framework of egocentric vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops; 2014. p. 586–91.
- Cartas A, Dimiccoli M, Radeva P. Detecting hands in egocentric videos: towards action recognition. *Computer Aided Systems Theory -- EUROCAST 2017*. Cham: Springer International Publishing. 2018. p. 330–8.
- Farneback G. Two-frame motion estimation based on polynomial expansion. *Image Analysis*. Berlin: Springer. 2003. p. 363–70.
- Likitlersuang J, Sumitro ER, Pi T, Kalsi-Ryan S, Zariffa J. Abstracts and Workshops 7th National Spinal Cord Injury Conference November 9–11, 2017 Fallsview Casino Resort Niagara Falls, Ontario, Canada. *J Spinal Cord Med*. 2017;40:813–69.
- Likitlersuang J, Sumitro ER, Theventhiran P, Kalsi-Ryan S, Zariffa J. Views of individuals with spinal cord injury on the use of wearable cameras to monitor upper limb function in the home and community. *J Spinal Cord Med*. 2017;40:706–14.
- Redmon J, Farhadi A. YOLO9000: better, faster, stronger. In Proceedings of the IEEE conference on computer vision and pattern recognition. 2017. p. 7263–71.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.