ClimbSense - Automatic Climbing Route Recognition using Wrist-worn Inertia Measurement Units

Felix Kosmalla
German Research Center for Artificial Intelligence (DFKI)
felix.kosmalla@dfki.de

Florian Daiber
German Research Center for Artificial Intelligence (DFKI)
Lancaster University
florian.daiber@dfki.de

Antonio Krüger
German Research Center for Artificial Intelligence (DFKI)
krueger@dfki.de

ABSTRACT
Today, sports and activity trackers are ubiquitous. Especially runners and cyclists have a variety of possibilities to record and analyze their workouts. In contrast, climbing did not find much attention in consumer electronics and human-computer interaction. If quantified data similar to cycling or running data were available for climbing, several applications would be possible, ranging from simple training diaries to virtual coaches or usage analytics for gym operators.

This paper introduces a system that automatically recognizes climbed routes using wrist-worn inertia measurement units (IMUs). This is achieved by extracting features of a recorded ascent and use them as training data for the recognition system. To verify the recognition system, cross-validation methods were applied to a set of ascent recordings that were assessed during a user study with eight climbers in a local climbing gym. The evaluation resulted in a high recognition rate, thus proving that our approach is possible and operational.

Author Keywords
Climbing; sports technologies; inertial sensors; machine learning.

ACM Classification Keywords
H.5.2 User Interfaces: Input devices and strategies

INTRODUCTION
Tracking of sports like running and cycling with smartphones and special sport devices is becoming increasingly popular [10]. These devices allow the recognition of tracks ran or distances cycled. In contrast, it is only possible to assess the routes which were climbed during a climbing session by hand. Common methods are noting the climbed routes and their difficulty levels in a book, spreadsheet, or low fidelity smartphone apps. Existing methods and applications for climbing tracking lack user-friendliness and additional value. Ladha et al. [11] developed a system which assess the user's climbing style and performance by analyzing accelerometer readings assessed by sensors which were worn on both wrists during an ascent.

CHI 2015, April 18 - 23 2015, Seoul, Republic of Korea

ACM 978-1-4503-3145-6/15/04…$15.00
http://dx.doi.org/10.1145/2702123.2702311

Figure 1. ClimbSense wearable sensors to automatically detect climbed routes with wrist-worn Inertia Measurement Units.

In this paper we introduce ClimbSense, a system that is able to record and automatically recognize the route which a user climbed during a climbing session (see Figure 1). A climbing gym usually provides a large amount of different climbing routes that consist of several artificial holds mounted to the climbing walls. One particular route consists of a certain set of holds and only these are allowed to be used as either a hand- or foothold (see Figure 2). Climbing routes usually do not have one single possible sequence of movements to master an ascent, which makes it hard to recognize the climbed routes. In our approach the climber is tracked by wrist-worn Inertia Measurement Units (IMUs). A corpus of climbing data was collected in an initial user study which should investigate the general feasibility of an automatic route recognition system. Eight climbers of different skill levels participated in the study. The features that were extracted from the collected data were used as training data for the recognition system. Our recognition approach is based on feature extraction and string comparison. To verify the recognition system, cross-validation methods were applied to a set of the corpus data. This evaluation resulted in a recognition rate of 90.19% ($SD = 5.39$).

We envision the following scenario that shows the potential of automatic climbing route recognition using wrist-worn IMUs reaching from self tracking for training and coaching to analytics for climbing gym maintenance. Our ClimbSense approach does not fully implement this vision yet. However, an automatic route recognition and tracking system, as it is described in this paper, could be the first step towards such applications.

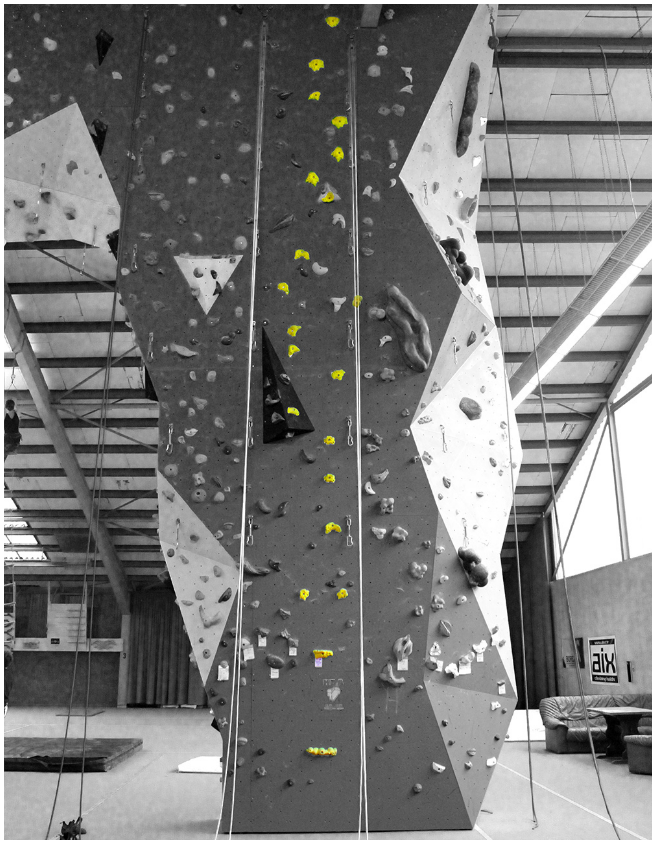

Figure 2. A climbing route consists of a certain set of holds (often denoted by color). Only these holds are allowed to be used during an ascent.

Scenario

When Paul enters the climbing gym, his phone notifies him that, according to his training plan, today is a cardio day. Since Paul uses his sensor armbands that automatically record every ascent he does, the system adapts to him. It knows that Paul usually climbs in the grades around VII+[1] and likes overhanging routes with big holds. Thanks to crowd sourced information every route is tagged with information about the holds, if dynamic moves are required or which inclination the wall has.

The system suggests a cardio training, which requires Paul to climb five different routes with three repetitions, within 30 minutes. To make things easier, five routes are proposed which are one level easier as Pauls average climbing skill. The routes which Paul climbed in the past were automatically assessed by his climbing wrist bands. Pauls climbing partner Sarah, who also uses the system, is also doing cardio training. After an exhausting training session, Paul and Sarah sit in the lounge and compare each others training progress on their smartphones. The system summarizes the last months sessions and visualizes the training progress. Back in the climbing gym office, Michael, the manager of the gym opens the dashboard of the climbing gym management software. The dashboard displays operating numbers like the number of routes, the average route age, and the route distribution. Thanks to the climbing history data he automatically receives from his customers, he can compare the present difficulty distribution with the skill levels of his customers. He sees that many new customers started climbing which require routes in the easier difficulty levels. To counterbalance, he schedules ten new routes for beginners that should be created during the next route setting session.

To sum up, this paper contributes a method to detect features in individual recordings of climbing route ascents based on absolute arm orientations which are assessed by wrist worn IMUs. A modified edit distance algorithm makes it possible to compare two recordings of ascents and returns a measure of similarity. Furthermore we cast a vision of how this technique could be used to enhance the climbing experience for both, climbers and climbing gym operators.

RELATED WORK

Our work builds on methods for (1) sports and activity tracking and (2) climbing technologies and research.

Activity Tracking

Today, a number of devices and smartphone apps are available which enable the user to track various parts of her live. This ranges from sports tracking, through activity tracking, to sleep cycle logging. These examples show that not only sports can be tracked, but also daily activities such as counting steps, recording sleep cycles, and calculating the number of calories burned. Many of these trackers incorporate a corresponding smartphone app that visualize the recorded data, rewards the user with achievements, acts like a mobile coach, or makes it possible to share data within social networks.

The motivations for users to track activities are diverse. According to Rooksby et al. [18] the reasons for activity tracking could be directive, documentary or diagnostic, but also motivated by collectable rewards or simply for the sake of the fetishized gadget. Ojala and Saarela [14] have shown that sharing also plays an important role for the motivation of the users. By interviews they identified seven main categories that motivated them to use the online communities.

Besides consumer-ready devices and services, a variety of research exists on activity recognition and tracking. Kranz et al. [10] introduced GymSkill, an app-based personal trainer which assesses the exercises performed on a balance board. They proposed an assessment algorithm that relies on the accelerometer of a smartphone and evaluated it with the assessment of a professional coach as ground truth. Swimming as a sport got also some attention in research. Daukantas et al. [5] made a first step towards measuring swimming performance of butterfly strokes with inertial sensors attached to the swimmers' spine. Their work reveals several problems when trying to assess the moving speed from acceleration data. Stamm et al. [20, 19] also tried to measure the swimming velocity with

[1]For climbing, many different grading systems exist varying according to country. The UIAA grading system is used within the scope of this work because the participants were familiar with this system. The grade reach from I (easy) to XII- (hard). Average skilled and sportive people might quickly reach the V range while climbing a VI requires some training. Climbing in grades above VII make a regular and rigor training regime necessary.

the help of an accelerometer that was attached the to swimmers sacrum. An evaluation indicated a good match between the proposed accelerometer-based velocity profile and a tethered speed probe. Kooyman et al. [9] presented a gyroscope-driven system to improve motor skills while performing a golf putt. A user study showed that experienced and inexperienced participants improved their putting performance after using the feedback GUI. Morris et al. [13] introduced RecoFit, a system for automatic tracking of repetitive exercises. They used an arm worn IMU to recognize fitness exercises. In an evaluation study 20 participants had to perform two rounds of a four-exercise course which resulted in an average recognition accuracy of 99.3%.

The proposed papers suggest that the usage of accelerometers and gyroscopes attached to the sport device or the user herself is a common and proven method for activity tracking. In addition to the research in tracking technology, the related work showed that using a tracking system could improve the performance of the user. We envision this for climbing and assume that a virtual climbing coach could be built that is based on the climbing data, in terms of the determination of climbed routes, assessed by the system proposed in this paper.

Climbing and HCI

Sport climbing can be performed indoors as well as outdoors and requires protection such as harnesses and ropes. Indoor climbing gyms provide the climber with a variety of routes in different difficulty levels. Artificial holds are mounted to the climbing wall which allows a large variety in different route styles and difficulties. A climber can decide whether she wants to climb *top rope* or *lead*. In the first case, the climber is tied on a rope that leads from the bottom to the top, passing through carabiners and is held fast by a second person, the belayer, with the help of a belay device. In case of a fall the belayer prevents the climber from hitting the ground. *Lead climbing* describes the type of sport climbing where the rope is carried up by the climber who has to clip the rope into carabiners. Bouldering is another discipline of climbing which requires neither a rope nor a harness. Bouldering is climbing in low heights that is performed indoors as well as outdoors. Artificial bouldering walls are surrounded by thick mats that prevent a climber from severe injuries in case of a fall. As of recently, climbing got some attention in research maybe because of its growing popularity. Climbing-related research includes work on technology aided route creation, instrumented climbing walls, and automated skill assessment.

Pfeil et al. [16] propose a climbing route designer that aims to enable even non-expert route setter to create quality climbing routes. An informal user study showed some limitations that were explained by the missing specification of certain climbing moves and feedback on why a move fails. They concluded that this system could be used by experienced and novice climbers to design routes for children. Daiber et al. [4] investigated handheld augmented reality for collaborative boulder training. They present a mobile augmented reality application to define, document and share boulder problems. Kajastila and Hämäläinen [8] also explored augmented reality for climbing walls but they directly augmented the wall with a projector. A preliminary Wizard-of-Oz study with six interaction prototypes and structured interviews showed that users liked the system.

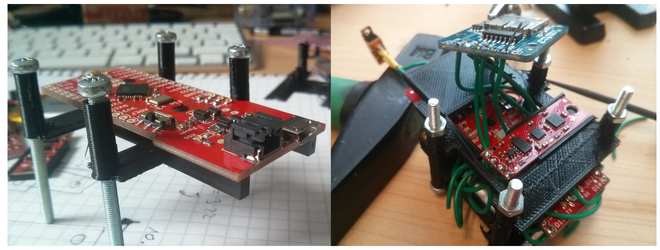

Figure 3. Mounting for the sensor modules. Left: Arduino Fio. Right: SDCard writer and sensor stick.

There are various possibilities to track a climber such as body-worn sensors, image processing, or instrumented climbing walls. Liljedahl et al. [12] proposed Digiwall, which consists of holds that can sense the climber's position with built-in capacitive sensors and provide subtle feedback with LEDs. The focus of this work is gaming, competitions, and challenges that can be rather used for playful activities than rigorous training. A very similar instrumentation was done by Ouchi et al. [15] whereas the goal of their work was to model play behavior of children. The used climbing holds that incorporated a LED and a strain gauge. Their work aims to improve the design of age-appropriate and safer playground equipment. Aladdin and Kry [1] proposed an instrumented climbing wall for static pose reconstruction. They use holds equipped with 6-axis force torque sensors that were used to reconstruct the climber's pose during an ascent. An evaluation showed that dynamic motions and higher errors coincide. Fuss and Niegl [6] also used torque sensors in instrumented climbing holds to measure the performance of a climber. Data collected on three climbing events was be segmented into the three phases of contact: set-up phase, crank phase, and lock off.

While the previous methods required an instrumented climbing wall, Ladha et al. [11] used wrist worn accelerometer sensors to assess the climbing performance of the user. An evaluation of the system during a climbing competition resulted in a positive correlation between the predicted and the actual score of the participants.

Although our system is closely related to the work of Ladha et al. [11] we follow a slightly different approach. The purpose of our system is not to assess the skills of the climber, but rather to automatically detect which routes she is climbing during a session. We further used sensor-fused orientation information and accelerometer data of a 9 degrees of freedom (DOF) IMU instead of accelerometer-only data. Our approach is motivated by self-tracked sports activities and thus requires no additional instrumentation of the environment. Other approaches that instrument the climbing wall (e.g. with RFID/NFC technology) instead of the climber are indeed possible, but we think that instrumenting solely the climber is the better option.

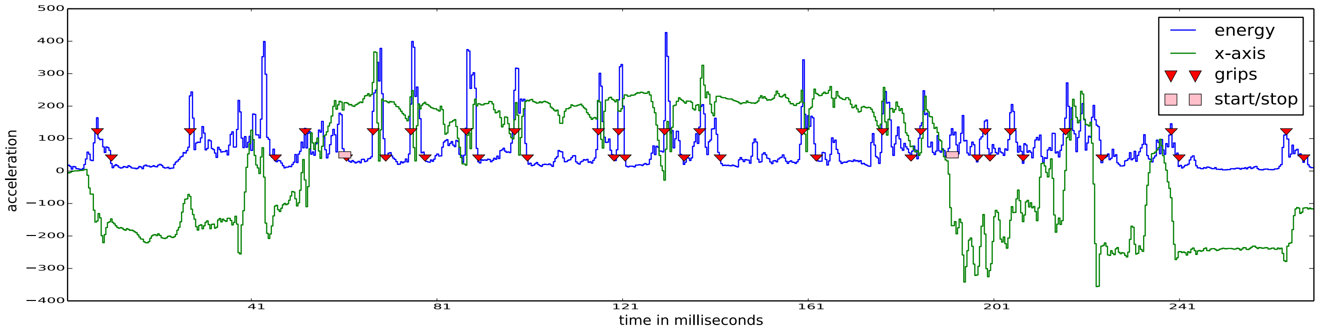

Figure 4. Recorded data (left hand) of an ascent. Blue: sum of standard deviations, green: mean of x-axis acceleration, red triangles: thresholds for move segmentation, pink squares: start and stop of climbing sequence.

CLIMBING DATA COLLECTION

Wrist-worn inertia measurement units (IMUs) were used to automatically recognize climbed routes. After introducing the hardware prototype a data collection study with eight climbers of different skill levels is presented. The aim of this study was to collect a corpus of climbing data to extract features as training data for the recognition system.

Climbing Sensor Hardware

Commercial accelerometer driven products like the fitbit or the Nike+ FuelBand do not provide access to the raw data collected by the device [3]. The missing API of commercial devices and the need for additional data like the orientation of the sensor required to build a custom prototype.

The prototype uses a Razor 9DOF Sensor Stick containing a triple-axis gyroscope, a triple-axis accelerometer, and a triple-axis magnetometer. This sensor allowed to assess the orientation. A microSD card writer was used to store the sensor readings during the climbing sessions. Both components were driven by an Arduino Fio v3 with an ATmega32U4 chip and a rechargeable battery. The sensor box could record approximately three hours of climbing with a fully charged battery.

Since the sensor should be worn on the wrist, a 3D model of a housing was designed using OpenScad[2] and printed with an Ultimaker[3] 3D printer. As it can be seen in Figure 3, stackable frames were used to hold the individual modules in place. This setup also allows later extensions of the prototype by muscle or heart rate sensors. The frames were held in place by off-the-shelf nuts and bolts inside a custom built housing. This housing had the dimensions of 7cm $\times$ 3.5cm $\times$ 5cm. A leather strip with Velcro was used to fix the sensor boxes on the wrist of the climber (see Figure 5): The sensor bands have to be tightly attached to the wrist to avoid slipping of the box during a climb.

Participants and Task

Eight climbers with different skill levels participated in the study, were one participant was female. In average the participants climbed for 7.25 years ($SD = 7.12$), ranging from four to 24 years. The participants' skill levels ranged from IV to

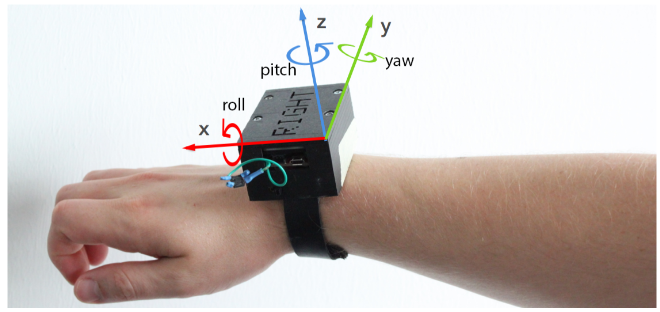

Figure 5. Sensor box attached to wrist. The attached button is used to manually mark timestamps in the recorded data. The x-axis runs parallel to the arm which is later used to determine if the arm is pointing up- or downward. Yaw, pitch, and roll corresponds to the y-, z-, and x-axis.

$VIII+$ (UIAA). Seven of the participants were right-handed. Furthermore the height ($M = 177.38cm$, $SD = 9.40$) and arm span ($M = 183.25$, $SD = 10.61$) of the participants was assessed. To build a data corpus, the participants were asked to climb a set of 5 predefined routes. The difficulty of the routes ranged from $IV-$ to $VI+$ and they contained dynamic and static moves, and different hold types. All routes had the same height and roughly the same number of holds. The routes were straight with no to little overhang. Each route varied in terms of holds in different forms and difficulty to grip which resulted in a broad variety of climbing styles. Most of the routes had more holds than necessary, allowing varying ascents and thus, resulting in differing recordings for every ascend. Furthermore, all routes were top-rope routes. Beside the general security aspect, top-roping allowed the participants to fully concentrate on climbing. The participants were asked to climb each route two times. No participant climbed any of the selected routes before. Both ascents were recorded. Since not every participant was able to climb every route during the recording sessions, the average number of recording per route was 10 ($SD = 3.53$). As a result, in total 50 recordings could be achieved.

Procedure

To assess the climbing data, the following procedure was performed for each climb. First, the participant had to put on the wristbands. Since the recognition of the route relies on the orientations of the arms, it was necessary to put on the sensors

[2] `http://www.openscad.org/`
[3] `https://www.ultimaker.com`

in the same way for each recording (details on the recognition see below). The sensors had labels to distinguish between the left and right hand. Figures 1 and 5 show the sensors, worn by participants. After the participant tied herself in, the recording was started by simultaneously pressing the external buttons of the sensor boxes. This was necessary since the current prototype does not contain clock modules for synchronization. Additionally the climbing tasks were video captured for later analysis. The video footage was mainly used to compare the collected sensor data with the actual movements of the climber.

Each recording of the ascents lasted about 30 minutes per participant. The participants were allowed to take breaks as often as they wanted. When the participant had finished the route she was lowered. As a final step the recording was stopped by switching of the sensor boxes. The sensor data for each ascent was automatically written in a separate file on the internal SDCard. These files were finally uploaded via an API into the analytics toolkit. We developed the analysis toolkit with the help of the Django web framework [4] which made it easy to implement both, GUIs and CLIs, to access the functions of our system. The analytics toolkit is responsible for the storage and processing of the collected climbing data.

Results

To analyze the recordings, the data was preprocessed by the analytics toolkit. First the orientation of the sensors was calculated with a sensor fusion algorithm[17] using the direction cosine[5]. The algorithm calculates a robust orientation of the device based on the data from the magnetometer, the gyroscope, and the accelerometer. In a second step the data stream was segmented with a sliding window procedure: each window covered $280ms$ of data with $20ms$ overlap on each side. Next, the average and the standard deviation of the three accelerometer values was calculated for each window. The sum of the standard deviations served as a description for the movement energy within the borders of a window.

When visually analyzing the plots from the data described above, the climbing sequence and the individual grips of the hands could be easily identified (see Figure 4). Most of the time during an ascent both hands are pointing upwards. This results in a mostly positive value of the acceleration in the x-axis (green). Transitions from one hold to another can be identified by the peaks of the standard deviation (blue). The valleys in-between the peaks indicate the gripping of a climbing hold.

CLIMBSENSE SYSTEM

The automatic route recognition introduced in this section is based on the assumption that a route can be characterized by a sequence of arm orientations, asessed in the moment of gripping a hold. Our system relies on the assumption that some of the recordings collected during training are similar to the recording to be recognized. Since the setup of a route usually limits the number of possibilities on how a route can be climbed, this assumption should be fulfilled. We chose the arm orientations in the moment of gripping a hold as characterization, since the determination of absolute movement based on the available accelerometer data is hard. This is due to errors based on the double integration of accelerometer readings [22]. Especially during movements like fast hand switching or shaking the hands for relief, using the absolute arm orientations in holding position as opposed to interpreting movements, is a viable approach.

Every aspect of processing and data management is covered by the analytics toolkit. Each recording of an ascent is stored in a so-called recording dataset. This includes the raw data obtained from the sensors as also the processed data, i.e. the extracted features of the ascent.

The process of the route recognition consists of the following steps: (0) preprocessing, (1) climb segmentation and grip detection, (2) feature extraction, and (3) matching. After segmenting the climb within a recording, grips are detected and the arm orientations of each grip are extracted. The orientations are discretized and considered as a sequence of symbols, i.e. a string. Since each route can now be characterized by a string, the identification of new unrecognized recordings is achieved by applying a weighted Levenshtein edit distance [2].

(0) Preprocessing

As mentioned above, preprocessing includes the calculation of the orientation based on a sensor fusion algorithm which uses the sensor readings from the accelerometer, gyroscope, and magnetometer [17]. A sliding window procedure is applied and parameters such as the average acceleration and the standard deviation are calculated. The following steps are applied on the resulting windows.

(1) Climb Segmentation and Grip Detection

After preprocessing the recorded data, the climbing sequence was extracted from each recording dataset. Each recording dataset only contains one climbing sequence. This also includes standing on the ground, the ascent itself, and lowering.

From the accelerometers point of view, a climbing sequence is defined as a sequence of alternations between low and high movement. These alternations arise when gripping a hold and the transition from one hold to another. Analog to this definition, the segmentation of individual movements is achieved by applying two thresholds t_{low} and t_{high} on the energy (i.e. the sum of the standard deviations) for each window. An alternation between those two thresholds indicates a movement that is set as property of the window (*is_gripping_window*). A second pair of thresholds is applied to the x-axis (parallel to the climber's arm) acceleration indicating whether the arm is pointing down (t_{down}) or up (t_{up}). After identifying the movements and arm position in the dataset for the left and right hand, the climbing sequence is narrowed down by observing both arms simultaneously. For this, both window sequences (left hand and right hand) are iterated and flagged with *both_hands_down* or *both_hands_up* whenever in two corresponding windows (two windows correspond if they share the same timestamp) the arms are pointing up or down.

[4] `https://www.djangoproject.com/`

[5] `https://github.com/ptrbrtz/razor-9dof-ahrs`

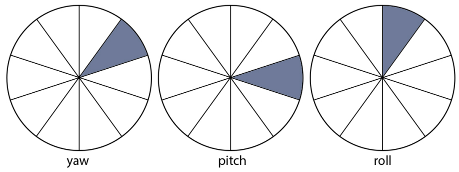

Figure 6. Segmentation of the orientation $yaw = 50°$, $pitch = 125°$, $roll = 10°$. One segment spans an interval of $[i \times 36, (i+1) \times 36 - 1]$. For the sake of simplicity, the number of segments was reduced to 10.

The thresholds for climb segmentation and grip detection were determined in an iterative procedure on data recorded during a pilot study. For this, the actual number of grips was manually transcribed from all videos. Based on this, an average number of grips was calculated for each route. An interval ranging from a relatively low to a relatively high energy was assigned to each threshold. After iterating over all values within these intervals, the script returned the "best" combination of thresholds. The "best" combination is defined as the set of thresholds which result in the smallest delta between the actual number and the detected number of grips. During the later measurement of the recognition performance, data different from the data assessed in the pilot study was used.

In a final step, the climbing sequence within the recording is marked by applying the definition above. Figure 4 shows a plot of a route's recording with the hold transitions and the boundaries of the climbing sequence indicated as red triangles and pink squares. This results in two sequences of groups of windows which were marked as gripping windows. These sequences are then used to extract the features of this ascent.

(2) Feature Extraction

The method of gesture recognition as proposed by [21] uses a codebook to convert the direction vectors which were calculated from accelerometer data into symbols. A codebook is basically a mapping from a set of values to one single value or symbol. A sequence of those symbols forms a string. This string is then considered a word, which can be used to calculate the Levenshtein distance [2] between two "stringified" gestures.

We adapted this method for the route recognition process. Instead of direction vectors, arm orientations, which arise during the ascent of a route, are used. These orientations are fetched whenever a hand grips a hold. An orientation is converted, with the help of a codebook into a symbol representing the yaw, pitch and roll angles. The codebook is generated as follows. First, a circle is divided into 30 segments. The resulting $12°$ circular segments represent a satisfying tradeoff between generalization and sufficient discriminability. This step has to be executed three times, one time for each of the three angles (yaw, pitch, roll).

The following example illustrates this procedure: To convert a specific orientation of the sensor box, the matching segment for each angle is chosen. Supposing the orientation $yaw = 50$, $pitch = 125$, $roll = 10$ the resulting segments for the three axis would have the indices 4, 10, and 0 (see Figure 6). By using this method an orientation of a sensor box can be coded into one of $30 \times 30 \times 30 = 27.000$ resulting symbols. During recognition each orientation is assigned to a combination of segment indices.

(3) Matching

In a final step of the route recognition process, the route is determined based on the calculated strings. The general approach is to compare the incoming dataset with each of the training datasets that are stored in the database for each individual route. As comparison, a weighted Levenshtein distance is used. The weighted Levenshtein distance used by the system is defined by the following equation.

$$m = |u|, \;\; n = |v|$$
$$D_{i,0} = i, 1 \leq i \leq m, \;\; D_{0,j} = j, 1 \leq j \leq n$$
$$D_{i,j} = \min \begin{cases} D_{i-1,j-1} & +0 \text{ if } u_i = v_j \\ D_{i-1,j-1} & +0.1 \times SymDist(u_i, v_j) \\ D_{i,j-1} & +1 \textit{ (Insertion)} \\ D_{i-1,j} & +0.1 \textit{ (Deletion)} \end{cases}$$
$$1 \leq i \leq m, 1 \leq j \leq n$$

$SymDist$ depicts the distance between two symbols, i.e. two orientations and is used when applying a substitution. The symbol distance has to be included in the calculation, since the system should allow a certain degree of deviation. This need arises, since an ascent is not always performed in the same way. Skipping of holds or different grip positions produce a variation in every ascent, which is compensated by this method. The distance between two symbols is calculated by adding the encoded angle distance of all three axis, e.g. $dist([0,0,0],[0,1,2]) = 0 + 1 + 2 = 3$. Since there are always two possibilities to transition from one angle to another, the shortest difference has to be chosen. As an example, consider a circle from its center two lines arise. Both of these lines have an outer and an inner angle. As a result, the distance between the two symbols $[0,0,0]$ and $[29,0,0]$ is 1 and not 29.

To generate the final result, the incoming dataset has to be compared to the training datasets which are available for each route. For this, the Levenshtein distances between the strings for the left hand and the strings for the right hand are summed up, which results in a score for the pair of the two datasets (incoming and training). The smaller the score is the higher is the probability that the route is recognized correctly. Thus, the pair with the lowest score is chosen as the winner.

ANALYSIS

The recognition performance of the algorithm was evaluated based on the data corpus collected in the user study. For this, the recorded data was used as both, training and testing data using cross validation methods.

Grip Detection

The first step in the route recognition process is the detection of grips. To evaluate the grip detection of the system, five sample recordings were randomly chosen. For each sample the number of grips was assessed manually by inspecting the video recording. This number served as a ground

Ascent	LH_V	RH_V	LH_D	RH_D	S_V	S_D	Δ
1	10	10	11	8	20	19	1
2	10	10	11	11	20	22	2
3	9	9	10	9	18	19	1
4	9	10	10	12	19	22	3
5	9	10	12	12	19	24	5

Table 1. Grip detection. A random sample of 5 ascents was analyzed. Ground truth based on video recordings and detected number of grips. LH: left hand, RH: right hand, $_V$: video, $_D$: detected, S: sum, Δ: delta.

Route	actual no. of holds	min	max	mean	SD
Route 1	16	21	32	27	2.78
Route 2	21	26	39	32	4.30
Route 3	18	21	34	25	4.32
Route 4	16	19	28	22	3.62
Route 5	17	16	31	22	3.60

Table 2. Statistics w.r.t the number of detected grips for all routes used in the study. The table shows that the number of actual holds correlates with the average number of detected holds per climb.

truth and was compared to the number of automatically recognized grips. This measure was chosen since the number of grips during an ascent varies from climber to climber but correlates with the number of holds mounted to the wall. Table 1 shows the result of this assessment. For each ascent the actual number of holds per hand was transcribed from the video and compared to the number of detected grips. The number of actual grips ranged from nine to ten while the recognized number of grips ranged from eight to twelve. Except for one record, the difference between the actual number of grips and the detected number of grips fluctuated between one and three. One record differed in five grips.

Furthermore, the minimum, maximum, and average number of detected holds was calculated for each route. This was done for the whole data corpus. Additionally, the number of holds attached to the wall was included in this analysis. Table 2 shows the results. The actual number of holds per route ranged from 16 to 21, while the average detected number of grips ranged from 16 to 39. Although the average difference between the actual number of holds and the average number of detected grips was 8.0 a correlation could be identified. A computation of the Pearson correlation resulted in 0.788 which shows a significant relation between the detected number of grips and the actual number of holds mounted on the climbing wall. The standard deviations of the average number of detected grips ranged from 2.78 to 4.32.

Route Recognition

To evaluate the route recognition exhaustive and non-exhaustive cross validation methods were applied: *leave-one-out cross-validation* (LOOCV) and *2-fold cross validation.* (2FCV). When applying LOOCV to every recording dataset, the system tries to recognize the route based on all 49 remaining recordings. For this, the evaluation script iterates over all recordings and queries the analytics toolkit, which responds with a route id. Since all recordings, including the testing data, are linked with the corresponding route, the number of correctly recognized routes can be divided by the number of processed routes. This results in the recognition rate. The results of this validation method indicate how good the system works, when a relatively large number of recordings per route are available.

Leave-one-out cross-validation

This LOOCV is executed two times with a slightly different training set. In the first run the training set contains all records except the one record which is used for testing. To validate if the recognition is influenced by the fact that the training set and also the testing set both contain records of the same person, the following method was applied. Whenever a dataset was used for testing which was recorded by participant A, all records from this participant were removed from the training set. This test should reveal whether the recognition system is influenced by user dependence, or if a *general* corpus of foreign training data is sufficient.

Applying the LOOCV over the complete dataset resulted in a recognition rate of 100%. When using only "foreign" training records, i.e. the training dataset does not contain recordings of the user whose dataset is to be recognized, a recognition rate of 100% was achieved.

Two-fold cross-validation

To investigate how good the system performs when only a small number of recordings per route is available a 2FCV was applied. In a 2FCV the available data is divided into two groups of the same size. One group is used as training data, while the other one is used for testing. In the case of the recognition system, the datasets can not be simply divided into two parts. This is due to the fact that it could not be assured that every route is equally represented in each group. For this, the two sets (*training* and *testing* data) are compiled as follows. The recording datasets of all routes are shuffled and afterwards split into two parts. These two parts are then appended to the *training* group, respectively the *testing* group. The resulting groups are then analyzed with the 2-fold cross-validation described above. This step was performed 100 times to get a more meaningful result.

As for the LOOCV, the 2FCV is also applied in two modes. In the first run the training dataset was not modified and contained a random set of recordings, except the ones that were used for testing. In the second run, the training dataset was modified as follows. At first, all participants which created records used in the training set were identified. Then, all records of these participants were removed from the testing dataset to ensure a user-independent recognition.

The application of the 2FCV resulted in an average recognition rate of 93.11% ($SD = 5.52$). When observing only foreign training records the average recognition rate was 90.19% ($SD = 5.39$).

One Handed Recognition

Motivated by the growing popularity of smartwatches, the validation methods described above were also applied on a modified implementation of the recognition system. The implementation was modified in such a way, that it would only consider either the left or the right hand for training and recognition. This could be easily achieved since both, the data for the left and the right hand were gathered separately. The current implementation of the climb segmentation, that is, identifying the ascent within a recording, uses the sensor data of both hands. Since the grip detection relies on the correct segmentation of the ascent, an evaluation of the grip detection using only one hand was not performed.

Applying the cross-validation methods on only the left and only the right hand resulted in different recognition rates. The leave-one-out cross-validation scored with 84% for the left hand and 94% for the right hand (see Table 3). Using only user independent data sets resulted in recognition rates of 68% and 94%. It can be seen that using only the data from the right hand resulted in a recognition rate almost as good as using both hands.

When applying a 2-fold cross-validation similar observations could be made. Using only data from the left hand resulted in an average recognition rate of 73.15% ($SD = 6.95$) (see Table 4). Applying the 2FCV on only data of the right hand resulted in an average recognition rate of 89.46% ($SD = 4.99$). When considering only user independent records recognition rates of 69.19% ($SD = 6.36$) for the left hand and recognition rates of 87.46% ($SD = 5.4$) could be achieved.

The the chance recognition rate for all validation methods was 20% and lies below the recognition rates our system could achieve.

	LOOCV		
	Left	***Right***	***Both***
complete	84.0%	94.0%	100%
user independent	68.0%	94.0%	100%

Table 3. Summarized evaluation results of the LOOCV. Using only data of the right hand results in higher recognition rates as when only using data of the left hand.

DISCUSSION

Grip Detection

When analyzing the grip detection performance, it can be seen that randomly picked samples (see Table 1) showed good results. The difference between the actual number of attached holds and the number of detected holds can be explained with various possible climbing styles. First, a usual technique is to grip a hold with both hands (also known as matching). This is used whenever the next hold is too far away or difficult to grip. Very fast transitions from one hold to another or the movement away and then again back to hold may also result in detection errors. Extensive use of momentum during an ascent may also influence the efficiency of the detection. Another reason for the varying number of holds is the fact that some climbers do not use every single hold that is available. Additionally, it is also possible to use the wall itself as support, which could be falsely detected as gripping. Also sometimes climbers use a single hold for a really short amount of time, using the momentum to ease the transition to the next hold. All these reasons may lead to a false detection of grips.

	2FCV					
	Left		***Right***		***Both***	
	Mean	*SD*	*Mean*	*SD*	*Mean*	*SD*
C	73.15%	6.95	89.46%	4.99	93.11%	5.52
U	69.19%	6.36	87.46%	5.40	90.19%	5.39

Table 4. Summarized evaluation results of the 2FCV. *C*: complete dataset, *U* user independent. As in the LOOCV, only data of the right hand results in higher recognition rates as when only using data of the left hand.

This can be also seen when investigating Table 2. Although the standard deviation of the average number of detected grips is relatively low, there are sometimes outliers. This can be noticed in the maximum number of detected grips in route 2. When investigating the video recording of the corresponding dataset, it could be seen that the participant struggled a lot during her first ascent. This resulted in changing the grip of her hands multiple times before continuing with the climb. As a result, the system detected more grips than a “normal” ascent would require. An implication of that would be, that the recognition performs good at routes which are climbed confidently, but may fail when the climbers struggles during the ascent.

Route Recognition

The evaluation results suggest that automatic route recognition based on wrist worn IMUs is possible and operative. The very good recognition rate of 100% when applying the LOOCV may have several reasons. First, a large number of training records per route are available. Whether this would be the case in a productive environment is discussed later on. Additionally, a relatively small number of routes is used. This needs to be investigated further in the future.

As it can be seen in the difference between the LOOCV and the 2FCV, a larger number of recordings per route are beneficial for the recognition process. When applying the 2FCV on an unfiltered training set, a higher recognition rate than in the other run could be observed. The (filtered) training set used in the second run of the 2FCV only relies on records from participants which were not included in the testing dataset. The cause for the difference between the runs on the filtered and unfiltered dataset could not be definitely identified. One reason could be the resulting smaller number of training data sets. A small training dataset increases the probability of false recognitions. In future work, we plan an extensive study to gain more insights regarding this issue. The second reason would be the difference in climbing styles between the participants which are responsible for the *training* data sets and the participants which are responsible for the *testing* data sets. As it can be seen in the high recognition rate of the LOOCV, a higher number of training datasets would solve this problem.

One Handed Route Recognition

Using only one hand for recognition and training leads to additional promising insights. It could be observed that when using only data from the right hand, higher scores could be achieved as when only data from the left hand was used. This could be due to the fact that most of the participants were right-handed. We conclude that gripping with the dominant hand resulted in more precise movement and in less jitter of the limb. With this, a better recognition of the orientation could be achieved and thus, a more expressive sequence of orientations is used for the recognition process. The fact that the LOOCV for only one limb still performs well may be a result of the large number of records per route. This assumption is encouraged when considering the results of the 2FCV for only one limb. Smaller training sets result in lower recognition rates.

The results suggest that in future work the application of smartwatches for route recognition should be investigated in more detail. Using future sports bands that allow access to the raw data could also be used for the recognition and would make special wrist bands for climbing recognition unnecessary.

Prerequisites

Currently the system relies on the assumption that each recording contains exactly one ascent. In the current prototype, this is achieved by turning the devices on before the climb, and switching them of after lowering. The ultimate goal would be to integrate a climb segmentation which would recognize each climb within a recording of a complete climbing session. This could be established like described in [11].

Another prerequisite is the existence of a sufficient amount of training data. Initial training data could be assessed by the route setter herself. After setting the new route, she could climb the route a couple of times which would suffice for a basic training data set. Another option would be crowdsourcing. A corresponding smartphone application could be used to link recordings of the wristbands to the manual logging of climbed routes. As motivation for the linking of recordings with routes, an achievement system could be deployed. The user could gain achievements like badges or coupons for free admittance in the climbing gym. Considering the scenario described above, the gym operator would also profit from knowing which routes are being climbed. This would justify the expenses of the coupons.

Limitations

The current implementation of the system limits the recognition to routes which were fully climbed, i.e. climbed from the bottom to the top without falling or resting. Routes which are only climbed partially are not recognized, since the system would not find a match to a partial route. This is due to the fact that the edit distance, which is the similarity measure, is also influenced by the different number of recognized grips during two ascents. A comparison between an incomplete route and the training set would result in a higher edit distance per se. Depending on the intended use of the system, e.g. only as list of completed routes in a climbing gym, this might not be flaw. In fact, a common training method for endurance training is to climb a well known, but hard route multiple times [7].

To recognize the route of a newly recorded ascent, the recording has to be compared to each training dataset. Although this process requires more resources as more routes are available in the system, this is not a big concern. Since a real-time recognition is not necessary, the recognition process can be performed offline in a server structure. Systems like Strava and their segment recognition prove that this is a viable solution.

Shaking the arms for relief, fast switching of holds, chalking (i.e. applying magnesium carbonate to the hands to remove perspiration and thus reduce slipping), or clipping the rope in a belay might evoke irritations of the grip detection. This results in too many or too few orientations represented by an ascent. As a result, the edit distances derivates from the actual value. To overcome this issue, additional machine learning could be applied to recognize and discard non-transition movements like the ones listed above. With this technique grips could be identified better and non-transition movements could be discarded. This would result in a better segmentation, which would give better route recognition rates.

CONCLUSION AND OUTLOOK

This paper introduced ClimbSense, a system that is able to record and automatically recognize the route which a user is climbing. The climber is tracked by wrist-worn IMUs. A corpus of climbing data was collected in a user study with 8 climbers of different skill levels. The features which were extracted from the collected data were used as training data for the recognition system. To verify the recognition system, cross-validation methods were applied to a set of the corpus data. The analysis of the recognition system showed promising results. Eight participants climbed five routes, resulting in 50 recorded data sets. These datasets were used for a combination of two cross-validation methods. The leave-one-out cross-validation resulted in a recognition rate of 100%. A 2-fold cross-validation was performed 100 times and resulted in a recognition rate of 93.11%. Considering the good performance of the recognition process when using the data of only one limb, the use of smartwatches as tracking device should be investigated in more detail. In general, the recognition could be improved in future work by an extended user study. Further machine learning methods could be applied to obtain a better grip detection. A better grip detection might also result in a higher recognition rate.

Since the recognition relies on training data, two concepts were introduced to obtain training records. One option would be, that the setter of the route could climb the route a couple of times, while wearing the wristbands. As an alternative, the training sets could be compiled through crowd sourcing. Climbers could wear the wristbands and link recording with routes with the help of a smartphone application. By improving the one handed recognition, open fitness tracker or smart watches could be used to crowdsource the training data. Furthermore, using future generations of smart watches makes it obsolete to carry a smartphone around. Incentives for the use of the application could be given by virtual achievements or coupons for drinks.

Although the results of the evaluation are satisfying, there are aspects that need more attention in future works. First, a larger number of routes have to be included in the evaluation. This would give insights into how well the system scales. With an increasing number of routes, the detection of movements and grips has to be refined. Currently the system does not distinguish between transitions from one hold to another and other movements like clipping, shaking the arms for relief or chalking. These movements could be detected using machine learning methods. To get closer to the goal of making this system consumer ready, climb segmentation as in [11] would be desirable. Currently only recordings containing one single ascent can be processed. Furthermore the system only detects complete ascents of a route. Falls, resting on the rope or partially climbed routes will certainly result in recognition failures.

Including performance assessment into the system would make it possible to detect user problems that arise at specific routes. These problems could be deducted by the stability and power the user shows during an ascent. With the help of these data and the knowledge of available routes, the system could make suggestions that would tackle the climber's weaknesses and helps her to improve her climbing skills.

Another way of extending the system would be the implementation of a real-time recognition. Currently the user would have to upload the data into the online portal to receive the results of her climbing session. With the immediate knowledge about the already climbed routes a virtual climbing coach could adapt to the user and suggest routes which fit into the current state of training. This would be especially helpful for smartwatches. In fact, smartwatches or wrist-worn activity trackers could replace special sensors like the ones used in this paper. This is shown by the promising results of the one handed route recognition evaluation.

ACKNOWLEDGMENTS

This research project is partially supported by the Deutsche Forschungsgemeinschaft (DFG KR3319/8-1).

REFERENCES

1. Aladdin, R., and Kry, P. Static pose reconstruction with an instrumented bouldering wall. In *Proc. VRST* (2012), 177–184.
2. Cormen, T. H., Leiserson, C. E., and Rivest, R. L. *Introduction to Algorithms , Second Edition*, vol. 7. The MIT Press, 2001.
3. Curmi, F., Ferrario, M. A., and Whittle, J. Sharing real-time biometric data across social networks. In *Proc. DIS* (2014), 657–666.
4. Daiber, F., Kosmalla, F., and Krüger, A. BouldAR - Using Augmented Reality to Support Collaborative Boulder Training. In *Proc. CHI EA* (2013), 949–954.
5. Daukantas, S., Marozas, V., and Lukosevicius, A. Inertial sensor for objective evaluation of swimmer performance. In *Proc. BEC* (2008), 321–324.
6. Fuss, F. K., and Niegl, G. Instrumented climbing holds and dynamics of sport climbing. In *The Engineering of Sport 6*, E. F. Moritz and S. Haake, Eds. Springer New York, 2006, 57–62.
7. Horst, E. J. *How to Climb 5.12*, 2nd ed. Falcon, 2003.
8. Kajastila, R., and Hämäläinen, P. Augmented climbing: Interacting with projected graphics on a climbing wall. In *Proc. CHI EA* (2014), 1279–1284.
9. Kooyman, D. J., James, D. A., and Rowlands, D. D. A feedback system for the motor learning of skills in golf. *Procedia Engineering 60*, 0 (2013), 226 – 231.
10. Kranz, M., Möller, A., Hammerla, N., Diewald, S., Plötz, T., Olivier, P., and Roalter, L. The mobile fitness coach: Towards individualized skill assessment using personalized mobile devices. *Pervasive and Mobile Computing 9*, 2 (2013), 203–215.
11. Ladha, C., Hammerla, N. Y., Olivier, P., and Plötz, T. Climbax: Skill assessment for climbing enthusiasts. In *Proc. UbiComp* (2013), 235–244.
12. Liljedahl, M., Lindberg, S., and Berg, J. Digiwall: An interactive climbing wall. In *Proc. ACE* (2005), 225–228.
13. Morris, D., Saponas, T. S., Guillory, A., and Kelner, I. Recofit: Using a wearable sensor to find, recognize, and count repetitive exercises. In *Proc. CHI* (2014), 3225–3234.
14. Ojala, J., and Saarela, J. Understanding social needs and motivations to share data in online sports communities. In *Proc. MindTrek* (2010), 95–102.
15. Ouchi, H., Nishida, Y., Kim, I., Motomura, Y., and Mizoguchi, H. Detecting and modeling play behavior using sensor-embedded rock-climbing equipment. In *Proc. IDC* (2010), 118–127.
16. Pfeil, J., Mitani, J., and Igarashi, T. Interactive climbing route design using a simulated virtual climber. In *SA Sketches* (2011), 2:1–2:2.
17. Premerlani, W., and Bizard, P. Direction cosine matrix imu: Theory. *DIY Drones.[Online] http://goo.gl/cCGNXB* (2009).
18. Rooksby, J., Rost, M., Morrison, A., and Chalmers, M. C. Personal tracking as lived informatics. In *Proc. CHI* (2014), 1163–1172.
19. Stamm, A., James, D., and Thiel, D. Velocity profiling using inertial sensors for freestyle swimming. *Sports Engineering 16*, 1 (2013), 1–11.
20. Stamm, A., Thiel, D. V., Burkett, B., and James, D. A. Towards determining absolute velocity of freestyle swimming using 3-axis accelerometers. *Procedia Engineering 13*, 0 (2011), 120 – 125.
21. Stiefmeier, T., Roggen, D., and Tröster, G. Gestures are strings: efficient online gesture spotting and classification using string matching. *Proc. BodyNets* (2007).
22. Thong, Y., Woolfson, M., Crowe, J., Hayes-Gill, B., and Jones, D. Numerical double integration of acceleration measurements in noise. *Measurement 36*, 1 (July 2004), 73–92.