

Stroke-Associated Hemiparesis Detection Using Body Joints and Support Vector Machines

Vishwajith Ramesh
vramesh@ucsd.edu

Kunal Agrawal
kuagrawal@ucsd.edu

Brett Meyer
bcmeyer@ucsd.edu

Gert Cauwenberghs
gert@ucsd.edu

Nadir Weibel
weibel@ucsd.edu

University of California, San Diego
La Jolla, CA 92093, USA

ABSTRACT

Hemiparesis, the weakness of one side of the body, affects the ability of stroke survivors to move and walk. It is generally diagnosed through motor tests performed as part of neurological examinations such as the NIH Stroke Scale (NIHSS), a subjective evaluation that requires the presence of an experienced neurologist. Here we report on an alternative way for computationally identifying hemiparesis that leverages body joint position data captured by the Microsoft Kinect. We employed support vector machines with 14 stroke subjects and 21 controls to characterize hemiparesis based on 4 core body angles recorded while the participants were simply sitting at rest, waiting for their neurologist. When comparing our results to neurologists' NIHSS scores, we were able to always identify right-side hemiparesis, left-side hemiparesis, or no hemiparesis using a leave-one-subject-out analysis. With additional data, our ultimate aim is to include the hemiparesis detection system presented here in a larger, multimodal tool that characterizes stroke based on several stroke-associated deficits. We envision deploying this tool in emergency settings for faster and more precise stroke severity assessments done in real-time.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

Stroke, Hemiparesis, Machine Learning, Support Vector Machines, Kinect, Body-Tracking, Posture Detection

INTRODUCTION AND MOTIVATION

Hemiparesis refers to the partial paralysis of either the left side or the right side of the body. The most common cause of hemiparesis is stroke, with roughly 80% of stroke victims experiencing some type of weakness of one side of the body [14]. In particular, pure motor hemiparesis is the most common form

of hemiparesis [14]. Subjects who suffer from this symptom have weakness in their right or left leg or arm, and therefore have difficulty moving the corresponding limb. In order to diagnose hemiparesis, neurologists conduct NIHSS motor tests to determine whether a stroke subject suffers from the symptom and the severity of the weakness [1]. While the NIHSS motor tests have been found to have strong inter-rater reliability [6], they require an experienced neurologist to conduct them. This limits its application to specific clinical situations (when a neurologist can physically be present). In fact, during emergencies in the field, diagnosis by emergency medical services (EMS) of stroke is often no better than 50% accurate [11]. Patients have to therefore wait to be evaluated by an experienced neurologist at the hospital before being administered Recombinant Tissue Plasminogen Activator (rt-PA), the only FDA-approved therapy for stroke. It is in part for this reason that rt-PA is utilized in less than 5.2% of acute strokes; rt-PA has a narrow therapeutic time window for intervention, within 3 hours of the onset of symptoms [7].

Efforts have been made to more quantitatively and reliably analyze hemiparesis without the need of a neurologist, particularly by analyzing gait. Computational methods to analyze ambulatory activity using wearable three-axis accelerometers, reflective markers, and motion cameras [5] and ankle-mounted step watch activity monitors [3] have all been shown to repeatedly and reliably assess hemiparetic gait. These techniques were designed to monitor hemiparetic stroke subjects outdoors or at their residence, and therefore do not require the presence of a neurologist. Unfortunately, these approaches focus on rehabilitation in the long-term, with the goal of monitoring improvements in movement and gait in recovering stroke subjects, rather than diagnosis in the short term. Furthermore, they require a device to be worn by the subject, which can be obtrusive. These wearables have spatial constraints and need to be fitted to each individual, making them challenging to deploy in the emergency settings in which we are interested.

To overcome this problem, in this paper we present a means of determining hemiparesis in an unobtrusive way, without the need of an experienced neurologist. Our goal is to implement this approach in emergency medical settings for diagnosis soon after the incidence of stroke. This work is based on the hypothesis that body posture, particularly when sitting, is a meaningful metric for diagnosing hemiparesis in stroke patients [10, 2, 8].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PervasiveHealth '18, May 21–24, 2018, New York, NY, USA

© 2018 ACM. ISBN 978-1-4503-6450-8/18/05...\$15.00

DOI: <https://doi.org/10.1145/3240925.3240979>




Figure 1: *Body tracking data collected from a subject in an outpatient clinic with the Microsoft Kinect v2. The subject here is shown sitting at rest prior to the start of the NIHSS examination. The representation of the body skeleton through the Kinect joints is on the left. Note that because the image and skeleton is mirrored, the viewer’s right side corresponds to the subject’s right side.*

METHODS

In collaboration with a team of neurologists from the University of California, San Diego, we recruited 35 subjects that were seen at the university’s stroke center. All participants in the study agreed to be recorded with sensors by signing a consent form approved by the local Human Research Protections Program.

We leveraged the Microsoft Kinect v2’s ability to automatically capture the movement of 25 distinct body joints and their orientations in 3D space to deliver a body skeleton for each recorded individual in the depth camera’s field of view. The Kinect was mounted on the wall opposite to a patient table, where neurologists typically see patients. We placed the camera fully facing the subject to avoid occluding parts of the body. Data was collected using ChronoSense software [13], at a body tracking capture rate of 30 Hz. We recorded participants while they were sitting at rest waiting for their neurologist and while they were undergoing the NIHSS examination. The duration at rest varies from participant to participant, depending on when the neurologist chose to start the NIHSS examination. Figure 1 shows our data collection set-up.

We recorded 35 subjects who either suffered from right-side hemiparesis (7), left-side hemiparesis (7), or were controls with no hemiparesis whatsoever (21). Control subjects included both healthy participants who never had a stroke, subjects who had fully recovered from their stroke and showed no deficits during the neurologists’ examinations, and subjects who suffered from stroke-related deficits but not hemiparesis (typically impairments in visual field).

Data Preprocessing

For each of the 35 participants, we obtained a time series of 25 body joints as they were sitting at rest. We focused our analysis on this “rest period” because we wanted to capture participants in their natural state; we did not ask them to perform any action or activity. Our system can therefore be easily deployed in emergency departments or inpatient clinic rooms where subjects can be monitored as they are waiting for their neurologist; asking them to perform a set of tasks would be difficult since the subjects seen at these locations are likely to be severely debilitated by their recent stroke.

To account for the differences in heights, physicality, and the body skeletons mapped by the Kinect between participants, we used body angles, a metric that is relative to each participant and that can be compared between participants. To extract the value of the core body angles (shown in blue and green

in Figure 2), we first calculated the 3D vectors between 4 core body angles (highlighted in dark red in Figure 2). We then applied simple trigonometry. The body angle θ between 3D unit vectors a and b , each connecting 2 body joints, was calculated (in degrees) as: $\theta = \arccos(a \cdot b)$. We monitored the core body angles instead of limb angles (about the elbow for example) because even during the rest period, subjects tended to move their arms and legs. We did not ask subjects to remain perfectly still to avoid being obtrusive.

To account for the temporal dynamics of motion, i.e. changes in the 4 body angles, we computed the first order derivative (the slope) of each of the 4 time series. The first derivative was calculated by taking the difference between 2 consecutive data points and dividing by $\Delta t = \frac{1}{\text{Frequency of Data Capture (30 Hz)}}$. Slight changes in a body angle over time – “fidgets” – are captured in the time series of the slope of that body angle; if no fidgets occur, the slope is 0. Since fidgets happen randomly, if at all, we cannot predict when during the rest period they occur or their amplitude (the amount of change in body angles). We therefore averaged the root mean square (RMS) value of the slope of the 4 core body angles over the entire rest period.

Our data preprocessing yielded a set of 4 features for each of the 35 participants. Each feature was obtained by averaging (over the rest period) the time series of the RMS slope of 1 of the 4 core body angles. Our approach was greatly simplified because we were no longer dealing with time series of dynamically changing body angles. Data normalization was not used; scaling the features to fall in the range $[0, 1]$ ultimately decreased classification performance.

Classification with Support Vector Machines

Each of the 35 participants had a “ground truth” diagnosis of right-side hemiparesis, left-side hemiparesis, or no hemiparesis as determined by experienced neurologists. Neurologists’ diagnoses of hemiparesis through NIHSS motor tests have strong inter-rater reliability [6], and thus we were not concerned with the subjective nature of the NIHSS test for this particular stroke symptom.

In order to predict whether each participant had right-side, left-side, or no hemiparesis based on the 4 average RMS slopes described above, we used support vector machines (SVM)¹ and compared the results with the ground truth label; the SVM classifier was chosen as it is known to perform well for small datasets.

¹We used a C-Support Vector Classification implementation (`svc.svm.SVC`) from the Scikit-learn library [9].




Figure 2: Body geometry recorded by the Microsoft Kinect v2; the orange dots represent the 3D positions of the joints. Highlighted in blue and green are the 4 core body angles analyzed in this study. The vectors from which these angles were calculated are in red.

Because of the low sample size (35 participants) and in order to avoid overfitting, instead of dividing the dataset into a training set and test set, we used leave-one-subject-out cross validation (LOSOXV). Cross validation is an estimate of the generalizability of a model, i.e. the expected fit of a model to a data set independent of the training data set; we used the classification accuracy as a means of summarizing this fit. In LOSOXV, we trained our SVM with 34 participants, while leaving out 1 participant to test the classifier. We repeated this process and rotated this 1 test participant so that every participant in the study was used to *train* the classifier 34 times and was used to *test* the classifier once. For each of the 35 tests, we obtained a predicted label (right-side weakness, left-side weakness, or no weakness), which we then compared to the ground truth label. Finally, to compute a confusion matrix, we tallied the number of true positives, true negatives, false positives, and false negatives over the 35 LOSOXV iterations. Classification accuracy was calculated from the confusion matrix as the number of true positives and true negatives divided by the number of samples.

In order to classify each subject as having right-side hemiparesis, left-side hemiparesis, or no hemiparesis with a binary SVM classifier, we used the *one-against-one* approach for multi-class classification with an ensemble of binary SVM classifiers. In this approach, we used $k(k-1)/2 = 3$ classifiers ($k = 3$ classes), each of which was trained on data from 2 classes. During prediction, each binary classifier was given the same test example and got a vote. The class with the maximum number of votes was predicted.

To account for the imbalance between the sample sizes of the 3 classes (7 participants for right-side hemiparesis, 7 for left-side, and 21 for control), we weighted the cost parameter C of each class i by multiplying it with the inverse of its sample

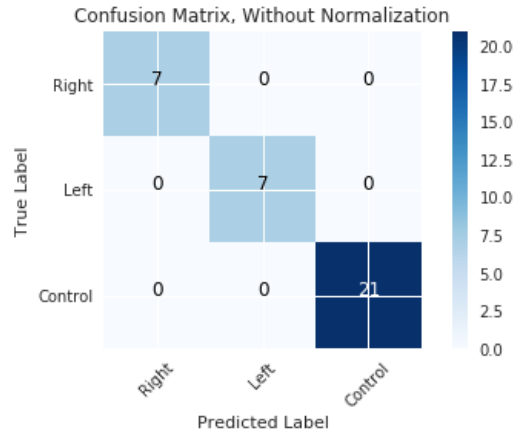


Figure 3: Confusion matrix for the SVM classification results using the average RMS slopes of the core body angles. Columns refer to the label predicted by the SVM analysis, while rows are the real condition of the subjects as determined by experienced neurologists. The dark blue elements of the major diagonal relative to the rest of the matrix indicate that we have all correct classifications.

size; this allowed us to greater penalize mistakes in classes with lower sample size: $\frac{\text{Total Samples in Training Set (34 with LOSOX)}}{\text{Number of Samples of Class } i \text{ in Training Set}}$.

We performed a grid search over all possible combinations of the following values of SVM hyper-parameters. The performance metric optimized was the classification accuracy. We used the radial basis function, linear, sigmoid, and polynomial kernels (with 2, 3, or 4 possible degrees for the polynomial kernel). Based on suggestions in literature, we varied the penalty parameter C and gamma by orders of magnitude [4]. To reduce the chance of overfitting, a concern given the small size of the dataset, we capped C at 1.0 to favor a soft-margin classifier. Possible values for C were therefore 10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} , 10^{-1} , and 1.0. Possible values for gamma were $2e-9$, $2e-6$, $2e-3$, 0.25, 1.0, $2e3$, $2e6$, $2e9$. We report the classification accuracy of the best, most optimal classifier given by the grid search in the Results section below.

RESULTS

The grid search revealed that a radial basis function kernel ($C = 10^{-5}$, gamma = $2e-9$) gave the best performance in terms of classification accuracy. By applying an SVM algorithm with a radial basis function kernel under the LOSOXV training/testing methodology, with a feature set of 4 average core body angle RMS slopes per participant as described earlier, we obtained the confusion matrix (without normalization) shown in Figure 3.

For each participant we compared the prediction with the corresponding ground truth label from the neurologist to calculate the overall accuracy. There were no instances of a misclassification. All subjects were classified correctly.

DISCUSSION

With 35 subjects, the SVM was able to classify right-side hemiparesis, left-side hemiparesis, and no hemiparesis based on 4 core body angles with 100% accuracy with the LOSOX approach. We were able to achieve this accuracy when we averaged the RMS slopes of body angles over the rest period for each subject. Subjects only need to sit at rest without being asked to move, waiting for a neurologist to start the NIHSS examination, for our algorithm to detect hemiparesis with high accuracy. This is exciting because it allows us to implement our algorithm in emergency settings where subjects are not likely to move. Gait analysis in ambulances or in the emergency department, for example, is difficult. We believe that our method is less obtrusive than wearable accelerometers, and does not need to be tailored to each subject because body angles are relative to individuals and can be compared across them.

Note that we did not have enough subjects to train the best classifier – a SVM with radial basis function kernel and $C = 10^{-5}$ and $\gamma = 2e^{-9}$ – on a random subset of say 30 subjects (85% of the data), and to subsequently test it on the remaining 5 subjects (15%). When we tried to do so, the accuracy on the test set varied depending on the 5 subjects chosen, even when maintaining the 7:7:21 ratio of right-side, left-side, and control classes when randomly choosing the 5 subjects (1 right-side, 1 left-side, and 3 control test subjects). This was likely due to the small size of the dataset.

We can therefore only conclude that using support vector machines to diagnose hemiparesis based on body posture is a **promising** approach. We are currently developing an alternate methodology in which we segment the time series of each subject into chunks to create additional “virtual” examples; while we have few subjects, we have several seconds of data for each subject. LOSOX can still be applied, but instead of leaving out only a single example per subject, we leave out several examples per subject. And with the addition of virtual examples, we hope to be able to create a test set from the full dataset without compromising the training accuracy. We are also continuing data collection in order to be able to get more subjects to properly test the classifier. While we perform grid search and use LOSOX to minimize overfitting, we still need to properly check whether our algorithm overfits on the training data with a test set. We are also adapting our algorithm to accommodate body skeletons generated by techniques that rely on deep learning and computer vision [12]. Such techniques may do a better job of handling occlusions of the core body than the Kinect depth camera, and may make our system more robust overall.

CONCLUSION

In this paper, we introduced our approach to detect hemiparesis by using core body angles and support vector machines. We demonstrated that leave-one-subject-out cross validation works well, and confirmed our hypothesis that body posture when sitting – as described by the core body angles – is affected by hemiparesis. The results also support the medical literature that indicate that hemiparetic patients have difficulty maintaining their posture and balance, due to reduced mus-

cle strength on the paretic side [10]. The ability to maintain balance requires postural control, which involves limiting the body’s sway and keeping its center of gravity within its base of support. Studies have used a force platform to measure center-of-pressure displacements in sitting stroke patients and found an increased postural disturbance in hemiparetic patients [2], as well as an inability to recover from leaning either forward or to the paretic side while sitting [8].

All in all, we believe that our research has the potential to provide great advancements to the field of neurology. Developing a computational technique to reliably identify and quantify the degree of different stroke-related deficits will allow for an accurate, per-subject signature that will aid in the prediction of long-term stroke outcomes. Our hemiparesis technique is the first step towards this tool, which will also enable faster, more accurate diagnoses and thus more effective administrations of time-dependent stroke medication.

REFERENCES

1. T. Brott, H. P. Adams, C. P. Olinger, J. R. Marler, W. G. Barsan, J. Biller, J. Spilker, R. Holleran, R. Eberle, and V. Hertzberg. Measurements of Acute Cerebral Infarction: A Clinical Examination Scale. *Stroke*, 20(7):864–870, 1989.
2. N. Genthon, N. Vuillerme, J. Monnet, C. Petit, and P. Rougier. Biomechanical Assessment of the Sitting Posture Maintenance in Patients with Stroke. *Clinical Biomechanics*, 22, 2007.
3. E. Haeuber, M. Shaughnessy, L. W. Forrester, K. L. Coleman, and R. F. Macko. Accelerometer Monitoring of Home- and Community-Based Ambulatory Activity After Stroke. *Archives of Physical Medicine and Rehabilitation*, 85(12):1997–2001, 2004.
4. C. Hsu, C. Chang, and C. Lin. A Practical Guide to Support Vector Classification. www.csie.ntu.edu.tw/~cjlin/papers/guide/guide.pdf, 2016.
5. H. Lee, J. Kim, H. Myoung, J. Lee, and K. Lee. Repeatability of the Accelerometric-Based Method to Detect Step Events for Hemiparetic Stroke Patients. In *EMBC*, 2011.
6. B. C. Meyer, T. M. Hemmen, C. M. Jackson, and P. D. Lyden. Modified National Institutes of Health Stroke Scale for Use in Stroke Clinical Trials. *Stroke*, 33(5):1261–1266, 2002.
7. D. Mozaffarian, E. Benjamin, A. Go, et al. American Heart Association Statistical Update: Heart Disease and Stroke statistics - 2015 Update. *Circulation*, 131(1):188, 2014.
8. D. Nichols, L. Miller, L. Colby, and W. Pease. Sitting Balance: Its Relation to Function in Individuals with Hemiparesis. *Archives of Physical Medicine and Rehabilitation*, 77, 1996.
9. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
10. D. Popovic and T. Sinkjaer. *Control of Movement for the Physically Disabled: Chapter 4 (Restoring Movement: State of the Art)*. Springer Science and Business Media, 2012.
11. P. Ramanujam, K. Z. Guluma, E. M. Castillo, M. Chacon, M. B. Jensen, E. Patel, W. Linnick, and J. V. Dunford. Accuracy of Stroke Recognition by Emergency Medical Dispatchers and Paramedics - San Diego Experience. *Prehospital Emergency Care*, 12(3):307–313, 2008.
12. S. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional Pose Machines. In *CVPR*, 2016.
13. N. Weibel, S. Rick, C. Emmenegger, S. Ashfaq, A. Calvitti, and Z. Agha. LAB-IN-A-BOX: Semi-Automatic Tracking of Activity in the Medical Office. *Personal and Ubiquitous Computing*, 19(2):317–334, 2015.
14. T. C. Weiss. Hemiparesis - Types, Treatment, Facts and Information. www.disabled-world.com/health/neurology/hemiparesis.php, 2017.