THE PERFORMANCE OF SOME MACHINE LEARNING APPROACHES IN HUMAN MOVEMENT ASSESSMENT

Johan Hagelbäck¹, Pavlo Liapota², Alisa Lincke¹ and Welf Löwe¹

¹*DISA, Linnaeus University, Växjö/Kalmar, Sweden*

²*Softwerk AB, Växjö, Sweden*

ABSTRACT

The advent of commodity 3D sensor technology enabled, amongst other things, the efficient and effective assessment of human movements. Statistical and machine learning approaches map recorded movement instances to expert scores to train models for the automated assessment of new movements. However, there are many variations in selecting the approaches and setting the parameters for achieving high performance, i.e., high accuracy and low response time. The present paper researches the design space and the impact of approaches of statistical and machine learning on accuracy and response time in human movement assessment. Results show that a random forest regression approach outperforms linear regression, support vector regression and neuronal network approaches. Since the results do not rely on the movement specifics, they can help improving the performance of automated human movement assessment, in general.

KEYWORDS

Statistical, Machine Learning, Human Movement Assessment

1. INTRODUCTION

Healthcare is becoming more and more value centered optimizing patient value (best possible health achieved) and system value (effective treatments at efficient costs). Until recently, measuring value, e.g., clinical outcomes, was not a capability of hospital or health IT (Elton and O’Riordan 2016). With inexpensive sensor technologies and data analytics becoming increasingly available, it is nowadays possible to collect data on the clinical activities of healthcare, the health status of a patient and the change of this status after treatment. We contribute with an automated objective method for the assessment of physical health of the human musculoskeletal system to help diagnose, predict or prevent related pain, injuries and long-lasting diseases. We use commodity 3D camera technology, such as the Microsoft Kinect.¹ Unlike similar tools for physical therapists^{2,3} not aiming to support the caring or nursing process, our approach localizes potential issues and quantifies their severity.

Regardless of the actual physical movement, their assessment in existing studies is usually done by comparing a patient’s performance in the movement to the performance of healthy reference participants. Participants repeat the same movement several times recorded by sensors, video cameras or other devices. Then features are defined and extracted from the recordings. Finally, different machine learning (ML) algorithms apply supervised learning using the features and expert scores as labels. This process makes it is quite efficient to add new movements and to assess them accurately (Dressler et al. 2019a).

As in many data driven applications, selecting the right ML approach is still a tedious, partially manual process characterized by prototyping, trial and error. This is because there are quite a few variants to select from spanning a huge search space. At the same time the variant selection has a large impact on the performance of the final application. Variants include: feature selection, the selection of preprocessing steps, movement alignment (cutting and matching), ML model selection, and the selection of training and testing data. The present paper focuses on selecting variants of ML approaches. More specifically, we assess

¹ <https://en.wikipedia.org/wiki/Kinect>

² <https://www.qinematic.com>

³ <https://kinetisense.com>

accuracy and response time of scoring and time for training the following models: linear regression, random forest, support vector regression, and neuronal networks.

The paper contributes with (i) a general human movement assessment framework, (ii) a component of this framework for substituting and assessing ML models, and (iii) an experimental evaluation of accuracy and response and training time of four different ML models.

It is organized as follows: Section 2 discusses related work in human movement assessment. Section 3 gives basic definitions. Section 4 introduces the human movement assessment framework (i). Section 5 explains the component for human movement quality scoring for different ML models (ii). Section 6 experimentally evaluates the approaches wrt. assessment accuracy and response time (iii). Section 7 concludes the paper and mentions future work.

2. RELATED WORK

Features necessary for human movement assessment extracted with commodity 3D sensor technology, such as the Kinect, seem to agree with the same features extracted with high-end motion capturing systems (Fernández-Baena et al. 2012, Müller et al. 2017, Ma et al. 2018) or with other sensor technology (Baldewijns et al. 2014, Lopez et al. 2018). Only few studies found differences in, e.g., foot placement locations and step lengths (Geerse et al. 2017).

Many studies research the *recognition* of human movements using commodity 3D sensors (Wang et al. 2012, Wang and Zheng 2012, Xia et al. 2012, Jiang et al. 2015, Pazhoumand-Dar et al. 2015). While these approaches are similar to ours in feature extraction and preprocessing, their goal is the classification of movements not their quality assessment.

The same technology has also been used in movement *quality assessment*, e.g., for identifying anterior cruciate ligament (ACL), for injury risks (Gray et al. 2014), for human abduction-adduction motion analysis (Kim et al. 2016), for ergonomic assessment in real workplace conditions (Plantard et al. 2017), for the measurement of lower extremity jump landing and squatting kinematics, for sit-to-stand analysis (Cippitelli et al. 2015), and for static and dynamic balance assessment (Eltoukhy et al. 2018). Similar studies include (Wang et al. 2013, Paiement et al. 2014, Parmar and Morris 2016, Tao et al. 2016). However, these studies aim at *qualitative* assessments—checking, e.g., whether movements are correctly or incorrectly, well or badly performed—or at identifying different abnormal movement patterns rather than *quantitatively* scoring the quality of a movement and localizing impairments. Some of the studies (Paiement et al. 2014, Tao et al. 2016) use preprocessing transformations such as dimensionality reduction based on manifold learning (Elgammal and Lee 2008), which enables a fast online assessment but even prohibits impairment localization.

Pirsiavash et al. (2014) suggest a learning framework for training models able to quantitatively assess the quality of human movements from 2D videos. Their approach trains a regression model from spatiotemporal pose features to scores obtained from expert judges. Features are extracted using unsupervised feature learning directly from 2D video data. Therefore, a localization of issues is not possible.

Dressler et al. (2019b), (2019a) suggest an automated, quantitative assessment of human movements based on the Kinect camera that scores the overall movement quality and localizes and scores issue. It provides the baseline of our assessment framework and performance comparisons. They use linear regression mapping movement features to expert scores.

Other ML approaches for estimating movement quality classes and scores include support vector machines (Taylor et al. 2010, Parmar and Morris 2016, Bao et al. 2019), Adaboost (Taylor et al. 2010, Parmar and Morris 2016), random forest (Leightley et al. 2017), multi-layer neural network (Parmar and Morris 2016), and distance functions (Houmanfar et al. 2016, Parmar and Morris 2016, Bao et al. 2019).

Liao et al. (2019) use deep neural networks, such as CNN, RNN, and hierarchical neural networks, for assessing the quality of ten different movements with high accuracy. Movements include deep squat, hurdle step, and sit to stand. Our study assesses linear regression, random forest, support vector regression, and neuronal networks to using the deep squat movement.

3. BASIC DEFINITIONS

The Kinect identifies 24 “joints”. Their x , y , and z coordinates are given in a 3D coordinate system with the camera in its origin and a horizontal x , z plane with z the depth coordinate. Because of low reliability of the other joints, we only use the following: head, neck, spine shoulder, spine mid, spine base, left/right shoulder, left/right elbow, left/right wrist, left/right hip, left/right knee, and left/right ankle.

A *sequence* is a list of frames from a recording of a human movement, e.g., a deep squat. It is a *master sequence* if the movement is assessed and scored by a human expert and used for building a (statistical) model of the movement. Any sequence that is to be scored automatically is called a *user sequence*. A *frame* is a record of features. It is a part of a sequence and describes the body posture at a specific point in time during the recorded movement. A *feature* is a part of a frame and describes an aspect of the body posture at specific point in time. A feature is called *direct* if it is directly measured by the 3D camera or *indirect* if it is computed from the direct or other indirect features. The direct features include the x , y , and z coordinates of skeleton joints. Indirect features include the angles between limbs and angles between limbs and the axes of the 3D coordinate system. An *aggregated sequence* is a list of aggregated frames aggregating one or more sequences into one. An *aggregated frame* aggregates two or more frames into one. It is a vector of sample distributions of the feature values of each feature. For scoring, we compare a user sequence with an aggregated master sequence.

4. A FRAMEWORK FOR MOVEMENT ASSESSMENT

All variants of our movement assessment algorithm perform five steps: (1) *Building an aggregated master sequence* is only performed once while the following steps are performed for each user sequence. (2) *Preparing the sequences* mitigates noisy feature values caused by random camera and skeleton recognition errors. (3) *Cutting and Matching* cuts the leading and trailing frames that do not belong to the movement and aligns a user sequence with the aggregated master sequence. (4) *Normalizing* applies linear transformations to the sequences and adds indirect features. (5) *Scoring* computes indicators for individual features and an overall movement score. If not explicitly mentioned otherwise, all steps are performed automatically. Details are introduced in the following paragraphs and precisely defined in (Dressler et al. 2019b).

Building an Aggregated Master Sequence builds an initial aggregated master sequence and aggregates it with the other master sequences: (1.1) *selects* the best master sequence S manually. It should have a constant movement speed without any delays and stops. (1.2) *cuts* off manually leading and trailing frames of S containing postures that do not belong to the movement. (1.3) *prepares* S , i.e., applies Step (2). (1.4) *aggregates subsequent frames* of S . We aggregate frames into 15 groups each containing the same number of frames (± 1). Each group contributes to a separated aggregated frame. (1.5) *normalizes* S , i.e., applies all but the first sub-steps of Step (4). (1.6) *computes a numerical sample distribution* of the feature values, separately for each group and for each feature.

For any of the remaining master sequences S' , (1.7) *prepares* S' , i.e., performs Step (2). (1.8) *matches* S' with current aggregated master sequence, i.e., performs Step (3). This maps each frame of S' to an aggregated frame of the aggregated master sequence. (1.9) *normalizes* the master sequence, i.e., performs Step (4). (1.10) *aggregates subsequent frames* of S' , i.e., for each frame and each feature of S' , it adds the feature value to the sample distribution of the respective feature of the mapped aggregated frame.

Preparing the Sequences: (2.1) *aligns the floor clip plane*, i.e., rotates the joint position vectors for each frame such that the floor clip plane is parallel to the x , z plane. (2.2) *smoothing* computes a sliding average of feature values for all *direct* features. (2.3) *interpolates* positions if a joint was not visible for $< k$ consecutive frames. If a joint was not visible for $\geq k$ consecutive frames, the joint is considered not tracked and an error is reported.

Cutting and Matching: (3.1) *cutting* removes leading and trailing frames of postures not belonging to the movement. (3.2) *matching* aligns a sequence with the aggregated master sequence.

Normalizing: (4.1) *computes a scaling* transformation that moves each joint of the first frame of a sequence to the corresponding mean joint position of the first aggregated frame of the aggregated frame sequence. It applies this scaling to all frames of the sequence. (4.2) *computes a hip rotation* transformation for the first frame of a sequence that lets skeletons “look” towards the camera. It applies this rotation to all frames of the sequence. (4.3) *computes the indirect features* from the direct features for each frame.

Scoring: (5.1) *calculates frame weights* for the aggregated master sequence frames based on the extremeness of the poses in these frames, i.e., the deviation of the features from a normal null pose. This is done automatically but offline, once for the aggregated master sequence. (5.2) *calculates the deviations* of user sequences feature values v as the z -score of v in the corresponding sample distributions from sub-step (1.10) and multiplies them with the respective weight calculated in the previous sub-step. (5.3) *computes for each feature* $v_{avg-max}$ as the average of the three largest weighted feature deviation from the previous sub-step. (5.4.) *calculates a feature score* for each $v_{avg-max}$ as the sample probability that a lower (better) maximum weighted deviation ever exists for a feature, i.e., that the previous sub-step ever gives a smaller (better) value than $v_{avg-max}$. Feature scores are between 0 (best) and 1 (worst) and locate movement issues and their severity. The next and final step (5.5) uses these feature scores in order to compute an overall movement quality score using an ML approach, cf. Section 5.

5. STATISTICAL AND MACHINE LEARNING COMPONENT

Various ML approaches apply to estimate an overall quality score. We assess linear, random forest, and support vector regression, and feed forward deep neural network.

Figure 1. ML Component

The ML component depicted in Figure 1 uses recursive feature elimination for the random forest and the support vector regression approaches to select the subset of features with the best estimation results (Guyon and Elisseeff 2003). For linear regression and neural network, recursive feature elimination is not applied, because the importance of features is inherently regarded. Parameter tuning is applied in all models. The dataset contains the following 40 features scores calculated in step 5: 13 angles deviation between limbs adjacent to the same joint, e.g., elbows, hips, head, neck, spine, 24 weak links features, e.g., heels rise, knees more inwards or outwards, arms fall forward, etc., two time-deviation features relative to the start and end frames, and the mean of these scores. Time-coherence is not considered relevant. Hence, we use aggregated feature scores, not sequences of features. The output is an estimated overall quality movement score. The estimated scores are passed to a model evaluation step where they are compared with human experts scores.

To avoid overfitting the training data, regression is constrained such that pairs of symmetric features, e.g., deviations in the left/right knees, get the same weights. This is achieved by regressing on the *sum* of these instead of the individual feature scores.

6. EXPERIMENTAL ASSESSMENT OF PERFORMANCES

Method: We use an overhead deep squat as the evaluation movement, an exercise standardized by the National Academy of Sports Medicine (NASM)⁴ that comes with movement execution and scoring specifications.

We assess the accuracy of the different approaches as follows. The baseline approach is the digitalization of the NASM overhead deep squat assessment and scoring, i.e., it implements the NASM scoring specification. Manual human-expert-based NASM overhead deep squat assessment and its digitalization show a high agreement (Dressler et al. 2019b) so that the digitalization serves as our ground truth. We assess the agreement of the scores for the four ML learning approaches introduced in Section 5 with the scores of the baseline approach.

A high agreement of is considered a high accuracy. We therefore calculate the Pearson coefficient r for the correlation of scores coming from different approaches with baseline. Altman and Bland argue that a high correlation is necessary but not sufficient for a good agreement between two approaches and suggest an additional analysis of differences (Altman and Bland 1983). Therefore, we additionally assess the differences of the baseline score and score of the approach with the highest correlation to baseline. The accuracy results depend on the sample in training and test data. We therefore repeat the experiments and report the statistics of 100 runs with fair random splitting in training and test data sets each containing 50% of the sample data for a run. We randomly select one run for discussing Bland-Altman statistics.

For time evaluations, preprocessing, cutting, matching, normalization, and feature scoring are excluded from the assessment as these steps are common to all ML variants. Note that we are interested in the relative response time between the variants not in minimizing the absolute response time.

Data: The assessment is based on a sample of 2094 user sequences recorded in 2018, the bulk during so-called company health days in Germany where AIMO⁵ recorded deep squats of health interested employees. The sample data contains scans of female and male persons, 25-65 years old, working in office jobs. No further effort was made to get a representative and fair sample of the full population.

Implementation: All steps of Section 4 use a *Matlab* implementation and so does scoring with linear regression. The machine learning based scoring uses the *Sklearn* (random forest and support vector regression) and the *Keras* libraries based on *TeserFlow* (neuronal network) for *Python*. The baseline approach uses AIMO's production code.

There are 40 features in total. Random forest was eventually trained and tested with 33 features after recursive feature elimination and a maximum depth of 10 (the deeper the tree the slower are training and testing). Support vector regression was trained with 23 features after recursive feature elimination. Linear regression and neuronal network use all features. Different neural network layer architectures were tested. The following layers were found empirically by maximizing the correlation: input layer with 40 neurons activation function *relu*, one hidden layer with 40 neurons and activation function *relu*, and one output layer with 1 neuron and activation function *sigmoid*, optimizer *adam* and loss *binary_crossentropy*.

Response time experiments are run on a MacBook Pro with a 2,3 GHz Intel Core i7 processor.

Results: Figure 2 (left) summarizes the accuracy results. It shows the minimum, first quartile, median, third quartile, and maximum (five-number summary) of the correlation coefficients of the 100 runs assessing the four ML approaches. The mean correlation coefficients (standard errors) are 0.78 (0.9×10^{-3}) for linear regression, 0.83 (1.0×10^{-3}) for random forest, 0.78 (0.9×10^{-3}) for support vector regression, and 0.81 (1.3×10^{-3}) for the neuronal network.

Figure 2 (right) summarizes the response time results. For a movement instance, training (scoring) takes 3 (3) msec for linear regression, 1320 (17) msec for random forest, 30 (8) msec for support vector regression, and 20575 (756) msec for neuronal network.

⁴ https://en.wikipedia.org/wiki/National_Academy_of_Sports_Medicine

⁵ Our industry partner AIMO AB, <https://www.aimo-health.com>

Figure 3 shows the results of Bland-Altman statistics for linear regression (top, left), random forest (top, right), support vector regression (bottom, left) and neuronal network (bottom, right): 95% of the differences between the baseline and the respective approach are in the range of -0.02 ± 0.26 score points for linear regression, -0.01 ± 0.21 score points for random forest, -0.03 ± 0.25 score points for support vector regression, and -0.02 ± 0.23 score points for neuronal network, i.e., between 23% for random forest and neuronal network and 25.5% for linear and support vector regression of the maximum possible range of differences.

Figure 2. Pearson coefficient of correlation r of the ML approaches with the baseline approach (left) and time of training and scoring per movement instance (in msec, log scale) of the approaches (right)

Figure 3. Differences of the ML approaches and the baseline approach

Discussion: The *random forest* approach shows the highest correlation ($r=83$) and the lowest range of differences with the scores of the digitalized expert opinions. It responds quite fast in scoring ($t=17\text{ msec.}$) and is sufficiently fast in training new models. It is therefore our champion approach. The neural network performs slightly better than linear and support vector regression. It would require more training data in order to improve further. All approaches show a high correlation and an acceptably low range of differences with the scores of the digitalization of expert opinions. However, the neuronal network approach is quite expensive in training as the *Keras* library runs on CPU and not on a *Tensorflow* GPU.

7. CONCLUSION

The paper presented (i) a general human movement assessment framework allowing the assessments of movements based on data from commodity 3D sensor technology. It reduces the noise of the input sensor and extracts relevant features of movements automatically. These features are input to (ii) a component assessing different ML models that score the quality of movements. We (iii) experimentally evaluated the accuracy and response and training times of four such models. Random forest turned out to be our champion approach.

The selected features, the preprocessing, the matching, and the scoring are agnostic w.r.t. the specific type of movement. None of the steps would be different if another movement, e.g., sit-to-stand, was assessed instead, only the data. Thus, the introduced technology contributes to value-centered healthcare with an improved objective assessment of musculoskeletal insufficiencies of patients. However, for stable conclusions, we need to assess other types of movements on larger and more representative samples of the population, which is matter of future work.

Moreover, the ML approaches outperform the baseline approach (linear regression) only by ca. 6%, which leaves room for improvements. Future work could add new sequential features and use CNN and RNN models trained on a bigger dataset or apply other machine learning models such as probabilistic neural networks.

Finally, we currently develop a mobile client based movement assessment using online algorithms allowing for pipelined preprocessing, cutting, matching, normalization, and scoring with the champion approach. Future work also aims at substituting the 3D with 2D camera technology in the full movement assessment framework allowing for more efficient data collection.

REFERENCES

- Altman, D. G. and J. M. Bland, 1983. Measurement in medicine: the analysis of method comparison studies. *The Statistician* 32: 307-317.
- Baldewijns, G., G. Verheyden, B. Vanrumste and T. Croonenborghs, 2014, Validation of the kinect for gait analysis using the GAITRite walkway. *36th Annual Inter. Conf. of the IEEE Engineering in Medicine and Biology Society, EMBC 2014, Chicago, IL, USA, August 26-30, 2014*.
- Bao, T., B. N. Klatt, S. L. Whitney, K. H. Sienko and J. Wiens, 2019. Automatically evaluating balance: a machine learning approach. *IEEE transactions on neural systems and rehabilitation engineering* 27(2): 179-186.
- Cippitelli, E., S. Gasparrini, S. Spinsante, E. Gambi, F. Verdini, L. Burattini, F. Di Nardo and S. Fioretti, 2015, Validation of an optimized algorithm to use Kinect in a non-structured environment for Sit-to-Stand analysis. *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*.
- Dressler, D., P. Liapota and W. Löwe, 2019a, Data Driven Human Movement Assessment. *11th International Conference on Innovation in Knowledge Based and Intelligent Engineering Systems, (KES), Invited Session on Digital Health, Distance Learning and decision support for eHealth*.
- Dressler, D., P. Liapota and W. Löwe, 2019b, Towards an Automated Assessment of Musculoskeletal Insufficiencies. *11th International Conference on Innovation in Knowledge Based and Intelligent Engineering Systems, (KES), Invited Session on Data Selection in Machine Learning*.
- Elgammal, A. M. and C.-S. Lee, 2008, The Role of Manifold Learning in Human Motion Analysis. *Human Motion – Understanding, Modelling, Capture, and Animation*, Springer.
- Elton, J. and A. O’Riordan, 2016. *Healthcare Disrupted: Next Generation Business Models and Strategies*, Wiley.
- Eltoukhy, M. A., C. Kuenze, J. Oh and J. F. Signorile, 2018. Validation of Static and Dynamic Balance Assessment Using Microsoft Kinect for Young and Elderly Populations. *Jour. of Biomedical and Health Info.* 22(1): 147-153.

- Fernández-Baena, A., A. Susín and X. Lligadas, 2012, Biomechanical Validation of Upper-Body and Lower-Body Joint Movements of Kinect Motion Capture Data for Rehabilitation Treatments. *2012 Fourth International Conference on Intelligent Networking and Collaborative Systems*.
- Geerse, D., B. Coolen, D. Koliijn and M. Roerdink, 2017. Validation of Foot Placement Locations from Ankle Data of a Kinect v2 Sensor. *Sensors* 17(10): 2301.
- Gray, A. D., J. M. Marks, E. E. Stone, M. C. Butler, M. Skubic and S. L. Sherman, 2014. Validation of the Microsoft Kinect as a Portable and Inexpensive Screening Tool for Identifying ACL Injury Risk. *Orthopaedic Journal of Sports Medicine* 2.
- Guyon, I. and A. Elisseeff, 2003. An introduction to variable and feature selection. *Journal of machine learning research* 3(Mar): 1157-1182.
- Houmanfar, R., M. Karg and D. Kulić, 2016. Movement analysis of rehabilitation exercises: Distance metrics for measuring patient progress. *IEEE Systems Journal* 10(3): 1014-1025.
- Jiang, M., J. Kong, G. Bebis and H. Huo, 2015. Informative joints based human action recognition using skeleton contexts. *Signal Processing: Image Communication* 33: 29-40.
- Kim, C. Y., J. S. Hong and K. J. Chun, 2016. Validation of feasibility of two depth sensor-based Microsoft Kinect cameras for human abduction-adduction motion analysis. *International Journal of Precision Engineering and Manufacturing* 17(9): 1209-1214.
- Leightley, D., J. S. McPhee and M. H. Yap, 2017. Automated analysis and quantification of human mobility using a depth sensor. *IEEE journal of biomedical and health informatics* 21(4): 939-948.
- Liao, Y., A. Vakanski and M. Xian, 2019. A deep learning framework for assessment of quality of rehabilitation exercises. *arXiv preprint arXiv:1901.10435*.
- Lopez, N., E. Perez, E. Tello, A. Rodrigo and M. E. Valentinuzzi, 2018. Statistical Validation for Clinical Measures: Repeatability and Agreement of Kinect™-Based Software. *BioMed Research International*.
- Ma, M., R. Proffitt and M. Skubic, 2018. Validation of a Kinect V2 based rehabilitation game. *PLOS ONE* 13(8): 1-15.
- Müller, B., W. Ilg, M. A. Giese and N. Ludolph, 2017. Validation of enhanced kinect sensor based motion capturing for gait assessment. *PLOS ONE* 12(4): 1-18.
- Paiement, A., L. Tao, M. Camplani, S. Hannuna, D. Damen and M. Mirmehdi, 2014, Online quality assessment of human motion from skeleton data. *Proceedings of the British Machine Vision Conference, BMVA Press*.
- Parmar, P. and B. T. Morris, 2016, Measuring the quality of exercises. *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), IEEE*.
- Pazhoumand-Dar, H., C.-P. Lam and M. Masek, 2015. Joint movement similarities for robust 3D action recognition using skeletal data. *Journal of Visual Communication and Image Representation* 30: 10-21.
- Pirsiavash, H., C. Vondrick and A. Torralba, 2014, Assessing the Quality of Actions. *Computer Vision – ECCV 2014*, Cham, Springer International Publishing.
- Plantard, P., H. P. H. Shum, A.-S. L. Pierres and F. Multon, 2017. Validation of an ergonomic assessment method using Kinect data in real workplace conditions. *Applied Ergonomics* 65: 562-569.
- Tao, L., A. Paiement, D. Damen, M. Mirmehdi, S. Hannuna, M. Camplani, T. Burghardt and I. Craddock, 2016. A comparative study of pose representation and dynamics modelling for online motion quality assessment. *Computer Vision and Image Understanding* 148: 136-152.
- Taylor, P. E., G. J. M. Almeida, T. Kanade and J. K. Hodgins, 2010, Classifying human motion quality for knee osteoarthritis using accelerometers. *2010 Annual Inter. Conf. of the IEEE engineering in medicine and biology*, IEEE.
- Wang, J., Z. Liu, Y. Wu and J. Yuan, 2012, Mining actionlet ensemble for action recognition with depth cameras. *IEEE Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society.
- Wang, J. and H. Zheng, 2012, View-robust action recognition based on temporal self-similarities and dynamic time warping. *2012 IEEE International Conference on Computer Science and Automation Engineering (CSAE)*.
- Wang, R., G. Medioni, C. J. Winstein and C. Blanco, 2013, Home Monitoring Musculo-skeletal Disorders with a Single 3D Sensor. *2013 IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
- Xia, L., C.-C. Chen and J. K. Aggarwal, 2012, View invariant human action recognition using histograms of 3D joints. *CVPR Workshops*, IEEE Computer Society.