Sleep Position Classification from a Depth Camera using Bed Aligned Maps

Timo Grimm*, Manuel Martinez*, Andreas Benz† and Rainer Stiefelhagen*

*Institute of Anthropomatics and Robotics

Karlsruhe Institute of Technology, Karlsruhe, Germany

†Thoraxklinik-Heidelberg GmbH, Heidelberg, Germany

***Abstract*—Sleep position is an important feature used to assess the quality and quantity of an individual's sleep. Furthermore, it is related to sleep disorders like sleep apnoea and snoring, and needs to be tracked in nursery homes to avoid pressure ulcers. Therefore, a gravity sensor attached to the chest is generally used to register body position during sleep studies. We suggest a non-intrusive and cost-efficient approach to detect the sleep position based on a single depth camera. Compared to alternative state-of-the-art approaches, ours require no calibration, and has been evaluated on a real setting comprising 78 patients from a sleep laboratory. We use the Bed Aligned Maps to extract a low resolution descriptor from a depth map which is aligned to the bed position, We perform classification using Convolutional Neural Networks, achieving an accuracy of 94.0%, thus outperforming current state-of-the-art algorithms and even the contact sensor from the sleep laboratory which achieves an accuracy of 91.9%.**

I. Introduction

Having sufficient quality sleep is essential for the physical and mental well-being. Due to population aging, sleep disorders such as obstructive apnoea, insomnia and restless leg syndrome are becoming common [1], [2].

Sleep position is related to several sleep disorders. For instance, sleep apnoea is observed with much higher frequency in supine position [3]–[5]. Hence, according to [6], shifting the body position during sleep is an effective medical treatment for sleep apnoea. Nakano *et al.* [7] note that sleeping in supine position is most likely to cause snoring. Finally, sleep position is monitored in Intensive Care Units and nursery homes to prevent pressure ulcers [8].

Sleep disorders are diagnosed through a polysomnogram, a sleep study performed in specialized sleep laboratories. On polysomnograms, the sleep position is commonly registered using a contact sensor (gravity based) attached to the chest. This approach is intrusive, valid for a one-time study but not suitable for long-term monitoring.

Therefore, we suggest an alternative approach using a single depth camera. Camera-based systems have the advantage that the same sensor can be used for a wide range of tasks such as detecting accidents [9], breathing patterns [10]–[12], awakeness [13], agitation [14]–[16], action recognition [17], [18], etc. Compared to contact sensors, camera-based approaches are inexpensive, easy to install even by non-experts, non-intrusive and portable. Hence, such computer vision monitoring systems are well suited for assisted living and elderly care.

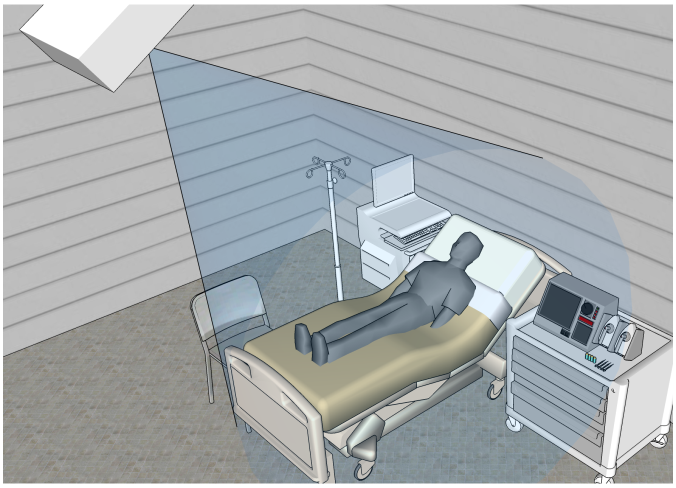

Fig. 1: We classify sleep position from an overhanging camera system attached to the ceiling. The bed position is detected automatically from the depth camera and used for alignment.

We build upon our previous work on Bed Aligned Maps (BAMs) [16], which are low resolution depth based descriptors that use the bed position as an anchor to provide alignment. BAMs were previously used together with Large Margin Nearest Neighbors [19] to predict sleep position in a simulated scenario with great success, however the approach does not generalize well to real scenarios.

We collected a real dataset from a sleep laboratory using an overhanging recording system with depth and infrared cameras (see Fig. 1). We recorded 78 different patients, comprising a total of 94 entire nights. Our recording system was installed in three different rooms with different bed sizes. An important aspect is that no calibration or nurse interaction was required.

This dataset contains people with significant sleep disorders, and a wide variety of ages and body types. To assess the difficulty of the task, it suffices to say that the gravity sensor attached to the chest achieves an accuracy of only 91.9% for sleep position classification.

We use BAMs [16] with a cell size of 5x5cm to generate a depth descriptor of size 40x26. The small size of the descriptor allows us to classify it efficiently using a small Convolutional Neural Network, achieving an accuracy of 94.0%, outperforming LMNN [16] (70.8%) and HoG+SVM [20] (88.9%).

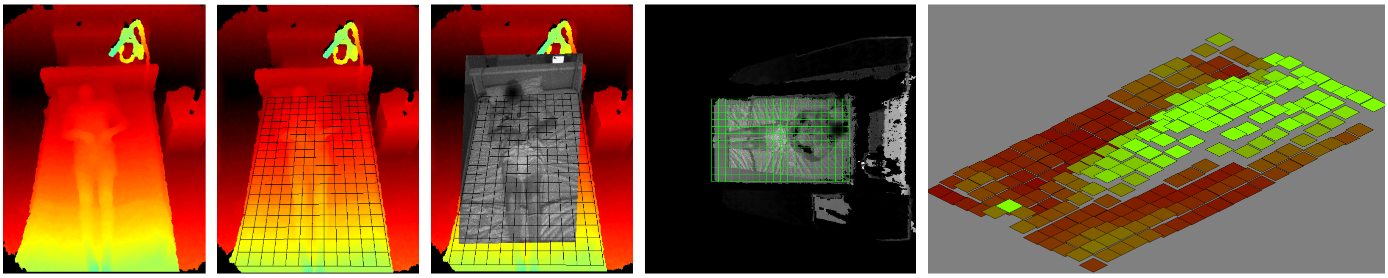

(a) (b) (c) (d) (e)

Fig. 2: The Bed Aligned Map (BAM) is a compact depth based descriptor. From a stream of depth images (a) the bed position is detected (b). An infrared image, used only for labeling and evaluation, is projected over the depth map (c) and a top-down virtual view centered on the bed is generated (d). The bed surface is split into cells and the mean height above the mattress within a cell is used to generate the BAM descriptor (e). To improve visualization, the cell size pictured here is 10x10cm while the cell size used in the paper is 5x5cm. Best viewed in color.

II. Related Work

The problem of sleep position is not uniquely defined. Several approaches (*e.g.* [12], [21], [22]) are based on the United Kingdom Sleep Assessment and Advisory Service (SAAS) who clustered sleep positions in six classes: foetus, log, yearner, soldier, freefaller and starfish [23]. However, sleep laboratories use a simpler classification: supine, left, right and prone [24]. Few papers (*e.g.* [16]) consider a the empty bed class, which is important for unattended applications. Inspired by the sleep laboratory we use as a reference, we use four classes: empty, left, supine, right. Prone position is not considered as it is rare in real conditions and we did not have enough samples in our 600-hour data collection for it to be statistically significant.

Smart beds are a common alternative method to monitor sleep position non intrusively. Motion sensors are placed inside the pillow [25] or onto the bed itself [26], [27]. Hoque *et al.* [27] uses RFID-based sensors equipped with accelerometers that are attached to the bed mattress.

There are camera based approaches that use color [22], infrared [24], thermal, depth cameras [12], [16], or a combination of them [20]–[22]. Yu *et al.* [12] propose an approach for the two class problem (supine and side-lying) based on a depth camera attached to the bed. First head and torso are detected using ellipse fitting, then the position is classified as supine if the topmost pixel on the head is above the topmost pixel of the chest, and side-lying otherwise. Their algorithm is evaluated on 8 volunteers in a simulated experiment. Lee *et al.* [21] describes a system using an overhanging Kinect 2.0 sensor over the bed. They classify between SAAS positions. First they extract body joint positions using Kinect v2 own libraries [28]. They use the relative position of hands and knees with respect to the spine for classification using a parametric approach. The approach requires the patient to not use a blanket. Evaluation and results are not provided.

Most related to our approach, Torres *et al.* [20] use a combination of depth and infrared cameras together with a pressure mattress to classify between SAAS positions. Only one scenario with a fixed camera above the bed is used, so alignment problems are not considered. They use HoG [29] and modified geometric moments [30] as features. The descriptors are combined using coupled-constrained Least Squares to assign weights to each modality, and then multiclass classifiers based on Support Vector Machines and Linear Discriminant Analysis are used. Evaluation is performed on only 5 people in a simulated scenario.

III. Methodology

A. System Setup

We use a multi-camera system fixed to the ceiling above the bed (see Fig. 1). To capture depth we use an Asus Xtion camera based on the same PS1080 platform that powers the Microsoft Kinect. This depth sensor projects a structured infrared pattern over the scene and uses it to extract disparity information at 30 frames per second and a resolution of 640x480 pixels. In addition, an infrared camera captures images calibrated to the depth view which are used only for labeling (see Fig. 2c).

To minimize obstructions to the face and chest areas, the sensor is not installed directly above the bed but above the feet of the patient, with an inclination of 30%. This enables a clear view of the patient even when the bed is articulated (not flat), and avoids the bed trapeze (the holding triangle used to help patients get in and out of the bed).

B. Bed Aligned Maps

The depth camera provides disparity maps, which can be easily converted to depth maps to form a 2.5D scene representation. This representation is robust and light-invariant.

However, as the camera is fixed to the ceiling and the bed has wheels, the relative position of the patient with respect to the camera is different for each recording. This causes an alignment problem, which is worsened by the high dimensionality of the depth map.

We use Bed Aligned Maps (BAMs) [16] to manage the alignment problem by using the bed position as a 3D anchor. We obtain the BAM as described in [16]. Each BAM is created from a single depth image. The bed is localized in 3D space

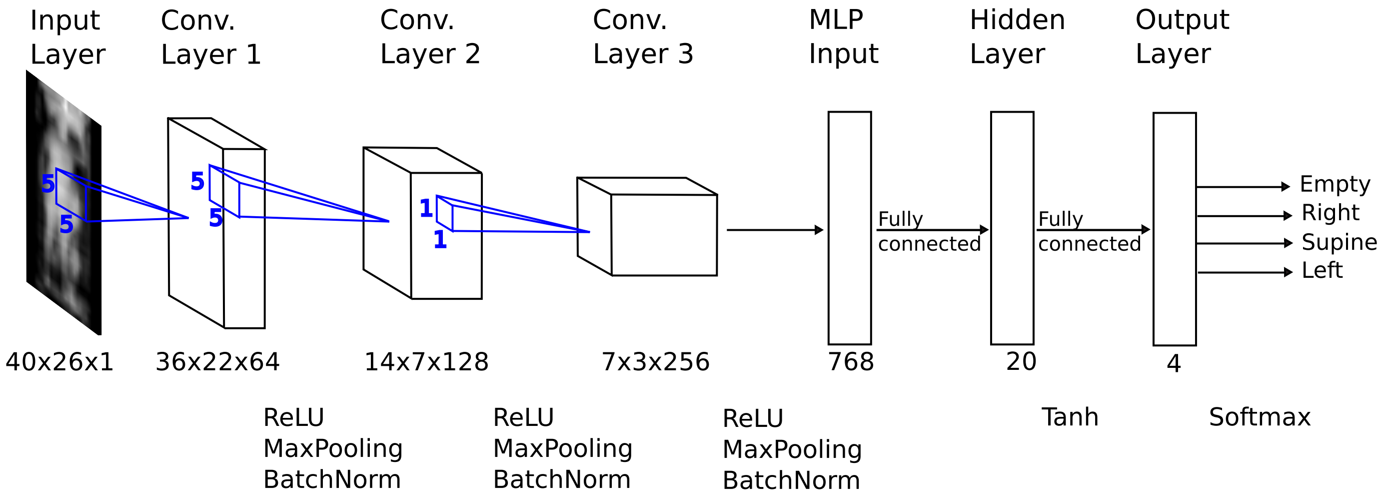

Fig. 3: The CNN architecture used consisting of three convolutional layers and two fully connected layers. Each convolutional layer uses ReLU, 2x2-max-pooling and batch normalization. We use tahn activation for the fully connected layers. The output layer uses softmax with one neuron per class.

and its surface is divided in equally sized cells (5x5cm in this case). A top-down view is generated using the bed as reference, and the point cloud obtained from the depth image is vertically projected over the bed surface with each point being assigned to its corresponding cell. The Bed Aligned Map is the matrix formed by the average cell height above the bed surface (see Fig. 2).

We apply two modifications to the original BAM algorithm. First, we use a 5x5cm cell size instead of 10x10cm. Second, we modified the point cloud projection algorithm to deal with the bed trapeze. Although the camera placement allows for an unobstructed view of the patient, the trapeze gets projected back to the bed when the top-down view is generated. To avoid this, the point cloud is generated by processing the depth image sequentially bottom-up from feet to head, and checking that the same direction is maintained on the top-down representation.

The dimensions of the Bed Aligned Map depend on the size of the bed, ranging in our experiments from 40x16 for a normal 80cm width bed, to 40x26 for an 130cm bed used by overweight patients. We deal with variable sized maps by zero padding smaller BAMs to 40x26.

C. Convolutional Neural Network

Compared to images, BAMs encode depth instead of light intensity, however they share a similar 2D structure and therefore Convolutional Neural Networks (CNNs) [31] should outperform non-spatially aware classifiers like Support Vector Machines or simple Multilayer Perceptrons.

Choosing the right model size for the CNN is critical. CNNs are data driven methods where both features and classifiers are automatically learnt from training data. The larger the model, the larger is the amount of data required for the training. In cases where enough data is not available, it is common to start with a pre-trained model from a similar task, and refine the model (usually only the last layers) using the available data. However as there is no available pre-trained model for bed analysis from depth, we need to train our model from scratch.

Non simulated sleep position data is expensive to acquire: our 600 hours of video contain only around 1000 significantly different sleep positions. We need to scale our model accordingly, and in this regard BAMs help significantly as they reduce the dimensionality while keeping the 2D structure (unlike PCA).

Our CNN classifies a single 40x26 cell BAM into one of four possible classes: empty, right, supine, and left. We use three convolutional layers with ReLU, 2x2 max pooling, and batch normalization, followed by two fully connected layers and an softmax layer (see Fig. 3).

As Bed Aligned Maps are small, the best results are obtained using reduced receptive fields for the convolutional layers (5x5 for the first two layers and 1x1 for the third one). However, it must be noted that small changes in the number of layers and the parametrization itself do not have a large impact on the performance. Still, the network has 255976 parameters.

D. Learning

Our training data is normalized to a zero mean and unit standard deviation. It is then resampled to balance the amount of samples on each class, and finally data augmentation is used to diminish overfitting. This includes left and right shifting of the bed as well as mirroring across the vertical axis.

Training is performed using stochastic gradient descent with a learning rate of 0.001 and a mini-batch size of 10. After each mini-batch update step t, we anneal the learning rate by a factor of $0.001/(1+t\cdot 10^{-5})$. To compute the loss between the predicted and the target output, we use the negative log-likelihood criterion. Training stops when validation error ceases to improve, between 30 and 61 epochs. The CNN is implemented in Torch7 [32].

IV. Experimental Results and Discussions

A. Sleep Study Dataset

Over a period of six months we recorded 78 different sleep laboratory patients using our integrated monitoring system (see subsection III-A). Our monitoring system is independent to the rest of devices used in the sleep laboratory, but it is time-synchronized with them. Therefore we can compare the results obtained from our system to the chest sleep position sensor used in the sleep laboratory (SleepSense from S.L.P. Ltd.).

As some patients were recorded for more than one night, the complete recording comprised 94 nights for a total of approximately 600 hours corresponding to 65 million images. We extract 20 samples per night as this dataset is too big to be labeled by hand in its entirety.

Considered classes are supine, left, right and empty. The empty class represents when the patient is not lying on the bed, but may contain a patient sitting on the bed. The prone position (also known as abdominal position) is special. It is an uncommon position to sleep, but it is even more rare in the sleep laboratory as it is uncomfortable for the patients carrying a cable box on their chests. Therefore we discard the very few samples we captured in the prone position, as they were not statistically significant.

The resulting set is severely unbalanced. Of the 1880 samples extracted and hand labeled, 85 belong to empty beds, 461 to right position, 761 to supine position, and 573 to left position (see Fig. 4). Note that the left position is more common than the right position. This is because in all monitored rooms both the room door and the bathroom are on the left side of the bed. This induces a bias as most patients prefer to sleep facing the door, and also makes the left position more challenging as most patients choose to enter and leave their beds from the left.

B. Evaluation

We perform a 5-fold cross validation, making sure that each fold has a similar size, but comprises a different set of patients. Subsequently, we present our experimental results averaged over all folds.

Chest Sensor: The chest sensor used in the sleep laboratory uses an accelerometer to detect its orientation by means of detecting the gravity vector. Its output is discrete: left, right, supine, prone or up.

The up label means that the patient chest is in vertical position, either sitting on the bed or standing. We associate this label to our empty class.

Among all sleep positions, the sensor shows a significant bias towards the supine position (see Table. I).

Accuracy across all samples is 91.9%, however it must be noted that the sensor tends to oscillate frequently between sleep positions, and often missreports the prone position. It is difficult to estimate the impact of those errors on the accuracy.

Large Margin Nearest Neighbor: Large Margin Nearest Neighbor [19] (LMNN) is a variant of k-Nearest-Neighbor applied to classification where a custom metric is used instead of euclidean. This metric is designed to minimize the leave-one-out error and is closely related to Mahalanobis distance.

LMNNs were used with success to classify sleep position from BAMs on a simulated dataset [16]. The simulated dataset

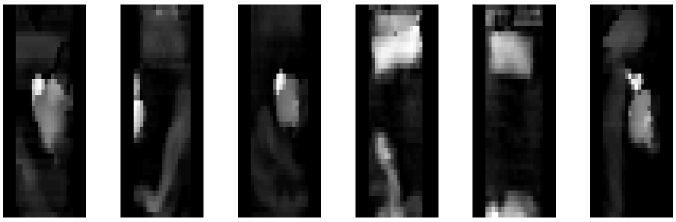

(a) Empty

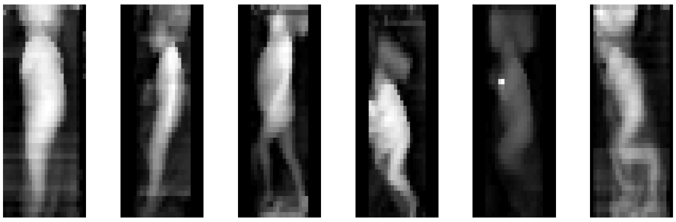

(b) Right

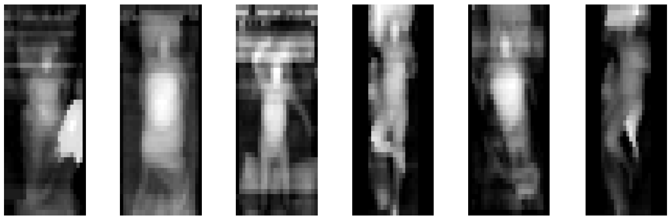

(c) Supine

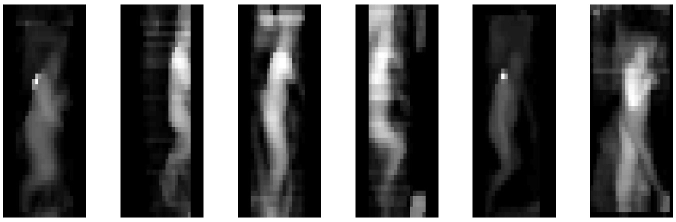

(d) Left

Fig. 4: Sample BAMs. The dataset is diverse with respect to patient appearance (weight, age, gender, corpulence) and sleep attitudes (*e.g.* sleeping with/without blanket). It contains all the common behaviors in a hospital bed, *e.g.* a nurse visible in the leftmost image in (c). The gray scale indicates depth.

	Empty	Right	Supine	Left
Empty	98.61	0.00	0.69	0.69
Right	1.36	93.49	6.24	0.44
Supine	0.00	1.01	96.98	2.00
Left	0.86	0.45	13.24	85.45

TABLE I: Confusion matrix for the chest sensor from the sleep laboratory. The accuracy is 91.9%.

used was perfectly balanced with one instance per class and person, and clearly distinct sleep positions.

When LMNN is used in the sleep laboratory dataset, results are significantly worse (see Table. II). In this case we used leave-one-patient-out evaluation which is better suited to LMNN, and no rebalancing of the training set, as having repeated samples is harmful. LMNN is unable to overcome the unbalanced training set and shows a strong bias towards the most common class: supine. Overall accuracy drops to 70.8%.

	Empty	Right	Supine	Left
Empty	9.43	11.32	75.47	3.77
Right	0.00	55.88	35.43	8.69
Supine	0.15	5.54	85.83	8.47
Left	0.00	4.42	30.27	65.31

TABLE II: Confusion matrix of the Large Margin Nearest Neighbor, it shows a strong bias towards the supine position and an average accuracy of 70.8%.

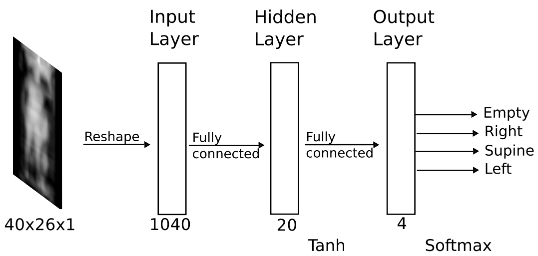

Fig. 5: The architecture of the MLP with 1040 input and 20 hidden neurons. Each 40x26 BAM is reshaped into a 1-dimensional vector before being processed.

Multilayer Perceptron: To quantize the impact of the convolutional layers, we evaluated a simpler neural network approach consisting of a multilayer perceptron (MLP).

The source BAM is flattened to a 1040 dimensional vector and fed to a MLP with 1040 input, 20 hidden and 4 output neurons. Hyperbolic tangent is used as activation function for the hidden neurons and softmax for the output layer (see Fig. 5). Training is performed similarly to the CNN approach, balancing the training set and using data augmentation. We use dropout to prevent overfitting.

As expected, our MLP generalizes better than LMNN, but still shows weak performance with an accuracy of 82.2% (see Table. III).

	Empty	Right	Supine	Left
Empty	82.80	1.60	9.20	6.40
Right	1.60	84.40	8.40	5.60
Supine	1.60	10.40	78.80	9.20
Left	1.20	4.40	9.60	84.80

TABLE III: Confusion matrix of the multilayer perceptron, accuracy of 82.2%.

HoG and Support Vector Machines: Histograms of Oriented Gradients [29] (HoG) and Support Vector Machines [33] (SVM) are known to have a very good synergy to detect and classify human shapes.

It can be compared to a shallow version of a CNN, HoG providing the spatial structure, and SVM classifying on top. Furthermore, it has been suggested recently for sleep position by Torres *et al.* [20].

Our setups used are significantly different and therefore it is not possible to exactly replicate the methods implemented in [20]. Therefore we apply HoG to the BAMs, reducing the feature space from 1040 to 280 dimensions, and apply SVM on top.

As the previous methods, data augmentation and training set balancing was applied, and empirically we found that the RBF-kernel provide the best performance.

The accuracy achieved is 88.9% (see Table. IV). This is better than MLP, as expected due to the HoG descriptor providing structural insight, but not as accurate as the chest sensor.

	Empty	Right	Supine	Left
Empty	96.00	1.20	0.40	2.40
Right	1.60	89.20	6.40	2.80
Supine	0.80	2.80	90.80	5.60
Left	0.40	1.60	12.80	85.20

TABLE IV: Confusion Matrix of the Histogram of Oriented Gradients and Support Vector Machines. The accuracy is 88.9%.

Convolutional Neural Network: Our CNN approach achieves an accuracy of 91.0% without data augmentation, whereas additional data augmentation leads to a further performance boost to 94.0% (see Table. V).

The CNN achieves superior performance compared to all other classifiers and even outperforms the chest sensor. The CNN behaves similarly to the HoG+SVM combination. In both cases the best classified class is the empty bed and, most importantly, the errors in sleep position are reasonable. This means that in both cases the left position is more easily confused with the supine position than with the right position, and the same happens in the opposite direction. With such accurate performance, most errors can be attributed to edge cases were the patient sleep position is between both positions.

	Empty	Right	Supine	Left
Empty	98.40	0.00	0.00	1.60
Right	0.40	93.20	4.40	2.00
Supine	0.40	1.60	94.00	4.00
Left	0.00	1.20	4.80	94.00

TABLE V: Confusion Matrix of the CNN. The accuracy is 94.0%. This outperforms all other classifiers and the chest sensor.

V. Conclusions

We have explored the problem of automatically determining the sleep position from a single depth image. We have shown how the problem is effectively treated using a classifier on top of spatial descriptors, as the HoG+SVM and the CNN approaches achieved great results while non-spatial classifiers like LMNN and MLP had mediocre performance.

Our evaluation is performed on 78 patients from a sleep laboratory where we achieved a classification accuracy of 94.0% with the CNN approach, surpassing the chest sensor used in the laboratory (accuracy 91.9%). Compared to other approaches, ours uses a single modality, requires no calibration or user interaction, has been validated on a real scenario, and achieves better performance. Further work is directed towards an holistic sleep quality monitoring system using the same setup for nursery homes, assisted living, and ageing-at-home scenarios.

Acknowledgments

This work is supported by the German Federal Ministry of Education and Research (BMBF) within the SPHERE project.

References

[1] S. Stranges, W. Tigbe, F. X. Gómez-Olivé, M. Thorogood, and N.-B. Kandala, "Sleep problems: an emerging global epidemic? Findings from the INDEPTH WHO-SAGE study among more than 40,000 older adults from 8 countries across africa and asia," *Sleep*, vol. 35, no. 8, pp. 1173–1181, 2012.

[2] W. H. Organization *et al.*, "WHO technical meeting on sleep and health," *WHO regional office for Europe*, 2002.

[3] R. D. Cartwright, "Effect of sleep position on sleep apnea severity." *Sleep: Journal of Sleep Research & Sleep Medicine*, 1984.

[4] T. Penzel, M. Moller, H. F. Becker, L. Knaack, and J.-H. Peter, "Effect of sleep position and sleep stage on the collapsibility of the upper airways in patients with sleep apnea," *SLEEP-NEW YORK-*, vol. 24, no. 1, pp. 90–95, 2001.

[5] W. Richard, D. Kox, C. den Herder, M. Laman, H. van Tinteren, and N. de Vries, "The role of sleep position in obstructive sleep apnea syndrome," *European Archives of Oto-Rhino-Laryngology and Head & Neck*, vol. 263, no. 10, pp. 946–950, 2006.

[6] Y. Itasaka, S. Miyazaki, K. Ishikawa, and K. Togawa, "The influence of sleep position and obesity on sleep apnea," *Psychiatry and clinical neurosciences*, pp. 340–341, 2000.

[7] H. Nakano, T. Ikeda, M. Hayashi, E. Ohshima, and A. Onizuka, "Effects of body position on snoring in apneic and nonapneic snorers," *SLEEP-NEW YORK THEN WESTCHESTER-*, vol. 26, no. 2, pp. 169–176, 2003.

[8] L. Paquay, R. Wouters, T. Defloor, F. Buntinx, R. Debaillie, and L. Geys, "Adherence to pressure ulcer prevention guidelines in home care: a survey of current practice," in *Journal of Clinical Nursing*, 2008.

[9] P. Kittipanya-Ngam, O. Guat, and E. Lung, "Computer vision applications for patients monitoring system," in *Information Fusion (FUSION)*, 2012.

[10] H. Aoki, Y. Takemura, K. Mimura, and M. Nakajima, "Development of non-restrictive sensing system for sleeping person using fiber grating vision sensor," in *Micromechatronics and Human Science*, 2001.

[11] M. Martinez and R. Stiefelhagen, "Breath Rate Monitoring During Sleep using Near-IR Imagery and PCA," in *International Conference on Pattern Recognition (ICPR)*, 2012.

[12] M.-C. Yu, H. Wu, J.-L. Liou, M.-S. Lee, and Y.-P. Hung, "Multiparameter Sleep Monitoring Using a Depth Camera," in *Biomedical Engineering Systems and Technologies*. Springer, 2013, pp. 311–325.

[13] M. Naufal Bin Mansor, S. Yaacob, R. Nagarajan, and M. Hariharan, "Patient monitoring in ICU under unstructured lighting condition," in *Industrial Electronics & Applications (ISIEA)*, 2010.

[14] J. Geoffrey Chase, F. Agogue, C. Starfinger, Z. Lam, G. M. Shaw, A. D. Rudge, and H. Sirisena, "Quantifying agitation in sedated ICU patients using digital imaging," in *Computer methods and programs in biomedicine*, 2004.

[15] J. Geoffrey Chase, F. Agogue, C. Starfinger, Z. Lam, G. Shaw, A. Rudge, and H. Sirisena, "Quantifying agitation in sedated ICU patients using digital imaging," in *Computer Methods and Programs in Biomedicine*, 2004.

[16] M. Martinez, B. Schauerte, and R. Stiefelhagen, "BAM! Depth-Based Body Analysis in Critical Care," 2013, pp. 465–472.

[17] J. Wang, Z. Liu, Y. Wu, and J. Yuan, "Mining Actionlet Ensemble for Action Recognition with Depth Cameras," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.

[18] M. Martinez, L. Rybok, and R. Stiefelhagen, "Action recognition in bed using BAMs for assisted living and elderly care," 2015.

[19] L. K. S. Kilian Q Weinberger, John Blitzer, "Distance metric learning for large margin nearest neighbor classification," in *Advances in Neural Information Processing Systems (NIPS)*, 2006.

[20] C. Torres, V. Fragoso, S. Hammond, J. Fried, and B. Manjunath, "Eye-CU: Sleep Pose Classification for Healthcare using Multimodal Multiview Data," in *Winter Conference on Applications of Computer Vision (WACV)*, 2016.

[21] J. Lee, M. Hong, and S. Ryu, "Sleep monitoring system using kinect sensor," *International Journal of Distributed Sensor Networks*, 2015.

[22] W. Huang, A. A. P. Wai, S. F. Foo, J. Biswas, C.-C. Hsia, and K. Liou, "Multimodal sleeping posture classification," in *International Conference on Pattern Recognition (ICPR)*, 2010.

[23] C. Idzikowski, "Sleep position gives personality clue," *BBC News*, 2003.

[24] C.-H. KUO, F.-C. YANG, M.-Y. TSAI, and M.-Y. LEE, "Artificial neural networks based sleep motion recognition using night vision cameras," *Biomedical Engineering: Applications, Basis and Communications*, 2004.

[25] T. Harada, A. Sakata, T. Mori, and T. Sato, "Sensor pillow system: monitoring respiration and body movement in sleep," in *International Conference on Intelligent Robots and Systems (IROS)*, 2000.

[26] K. Malakuti and A. B. Albu, "Towards an intelligent bed sensor: Non-intrusive monitoring of sleep irregularities with computer vision techniques," in *International Conference on Pattern Recognition (ICPR)*, 2010.

[27] E. Hoque, R. F. Dickerson, and J. A. Stankovic, "Monitoring body positions and movements during sleep using wisps," in *Wireless Health*, 2010.

[28] J. Shotton, T. Sharp, A. Kipman, A. Fitzgibbon, M. Finocchio, A. Blake, M. Cook, and R. Moore, "Real-time human pose recognition in parts from single depth images," *Communications of the ACM*, vol. 56, no. 1, pp. 116–124, 2013.

[29] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005.

[30] M.-K. Hu, "Visual pattern recognition by moment invariants," *Transactions on Information Theory*, vol. 8, no. 2, pp. 179–187, 1962.

[31] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient Based Learning Applied to Document Recognition," *Proceedings of the IEEE*, 1998.

[32] R. Collobert, K. Kavukcuoglu, and C. Farabet, "Torch7: A matlab-like environment for machine learning," in *BigLearn, NIPS Workshop*, 2011.

[33] M. A. Hearst, S. T. Dumais, E. Osman, J. Platt, and B. Scholkopf, "Support vector machines," *Intelligent Systems*, vol. 13, pp. 18–28, 1998.