2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) | 979-8-3503-2447-1/23/$31.00 ©2023 IEEE | DOI: 10.1109/EMBC40787.2023.10340645

Unobtrusive Sleep Position Classification Using a Novel Optical Tactile Sensor

Alexander Breuss*[1], Carmelo Sferrazza*[2], Jonas Pleisch[1], Raffaello D'Andrea[2], Robert Riener[1,3]

***Abstract*—Unobtrusive sleep position classification is essential for sleep monitoring and closed-loop intervention systems that initiate position changes. In this paper, we present a novel unobtrusive under-mattress optical tactile sensor for sleep position classification. The sensor uses a camera to track particles embedded in a soft silicone layer, inferring the deformation of the silicone and therefore providing information about the pressure and shear distributions applied to its surface.**

We characterized the sensitivity of the sensor after placing it under a conventional mattress and applying different weights ($258\,\mathrm{g}$, $500\,\mathrm{g}$, $5000\,\mathrm{g}$) on top of the mattress in various predefined locations. Moreover, we collected multiple recordings from a person lying in supine, lateral left, lateral right, and prone positions. As a proof-of-concept, we trained a neural network based on convolutional layers and residual blocks that classified the lying positions based on the images from the tactile sensor.

We observed a high sensitivity of the optical tactile sensor: Even after placing the sensor below a conventional mattress, we were able to detect our lowest test weight of $258\,\mathrm{g}$. Using the neural network, we were able to classify the four sleep positions, lateral left, lateral right, prone, and supine with a classification accuracy of $91.2\,\%$.

The high sensitivity of the sensor, as well as the good performance in the classification task, demonstrate the feasibility of using such a sensor in a robotic bed setup.

***Clinical Relevance*— Positional Obstructive Sleep Apnea is highly prevalent across the general population. Today's gold standard treatment of using CPAP ventilation is often not accepted, leading to unwanted treatment cessations. Alternative treatments, such as positional interventions through robotic beds are highly promising. However, these beds require reliable detection of the lying position. In this paper, we present a novel, scalable and completely unobtrusive sensor that is concealed under the mattress while classifying sleeping positions with high accuracy.**

I. Introduction

Obstructive Sleep Apnea (OSA) is a condition characterized by a recurring collapse of the upper airways during sleep leading to a reduced or complete blockage of airflow. The severity of OSA is measured through the Apnea-Hyponpea Index (AHI) which is calculated by dividing the number of apnea and hypopnea events by the total number of hours of sleep. A normal AHI is below 5 events per hour, while an AHI of 15 or higher is considered indicative of sleep apnea. Sleep apnea is often underdiagnosed. According to estimates, over $93\,\%$ of women and $82\,\%$ of men in the adult employed population with moderate to severe sleep apnea have not been clinically diagnosed [1]. This is especially problematic as sleep apnea is associated with an increased risk of hypertension, type two diabetes mellitus, heart failure, stroke, and death [2], [3], [4], [5]. The gold-standard treatment for patients with OSA is the use of a CPAP mask through which a positive pressure during the end-expiratory phase is applied to keep the airways open. However, correct usage of this mask requires a tight fit and continuous wearing during the entire night. This leads to an overall low acceptance [6], [7].

Fig. 1: Intelligent Sleep Apnea Bed (ISABel): The bed has a motorized backrest and can gently lift the trunk of the user to reduce sleep apneas.

In patients with Positional Obstructive Sleep Apnea (POSA), the majority of apnea events occur in supine position [8]. For these patients, modern treatment approaches include positional interventions that prevent the users from sleeping in supine position. The ISABel I (Sensory-Motor Systems Lab, ETH Zurich, Switzerland, Figure 1) is a trunk-elevating bed that can change the position of the user during sleep [9]. For such a bed to work closed-loop, it is essential to detect the position of the user without influencing sleep quality or raising privacy concerns. Current sensors that classify the sleeping position are often based on accelerometers that are physically attached to the user and might induce discomfort. Existing unobtrusive pressure sensors integrated into the bed often have limited resolution, suffer from crosstalk, provide limited accuracy, or induce electromagnetic fields that can interfere with measurement equipment. [10].

In this work, we present a novel and completely unobtrusive sensor that is based on a rescaled highly sensitive (mN) optical tactile sensor, originally developed for robotic manipulation. ([11], [12], [13], [14], [15]). Next to their high sensitivity, optical sensors can provide a much higher resolution than

*The corresponding authors alexander.breuss@hest.ethz.ch and csferrazza@ethz.ch contributed equally.

[1]A. Breuss, J. Pleisch, and R. Riener are with the Sensory-Motor Systems Lab, Institute of Robotics and Intelligent Systems, ETH Zurich, Switzerland

[2]C. Sferrazza and R. D'Andrea are with the Institute for Dynamic Systems and Control, ETH Zurich, Switzerland

[3]R. Riener is with the Spinal Cord Injury Center, University Hospital Balgrist, Zurich, Switzerland

979-8-3503-2447-1/23/$31.00 ©2023 IEEE

current piezoelectric sensors [10]. Compared to other optical tactile sensors [16], [17], the sensing principle described in [11] and utilized here is particularly suitable for large-scale sensing, because it leverages uncomplicated light and marker placement.

II. Methodology

A. Sensor Design and Material Choice

The new sensor comprised four layers (Figure 3): a soft silicone layer that embedded a spread of clear polyethylene particles (b) was placed on top of a harder clear silicone layer (c), which supported the soft layer and additionally served for uniform light diffusion. These two silicone layers were placed on top of a clear acrylic sheet layer (d) that further increased the structural integrity. On top of the three layers, each of which was $5\,\mathrm{mm}$ thick, a thinner black silicone layer (a) shielded the top surface from ambient light. The purpose of the particles within the soft silicone layer was to enable the tracking of the deformation of the silicone caused by the external forces applied to the top surface. From a pinhole-camera simulation, we determined that approximately $10{,}000$ particles with a diameter of $1.4\,\mathrm{mm}$ were required. To ensure uniform distribution after stirring the particles inside the soft silicone, we selected particles with a density of $0.96\,\mathrm{g/cm^3}$ to $0.98\,\mathrm{g/cm^3}$, similar to the density of the soft silicone layer, such that the particles remained in place after mixing. A black curtain (e) was installed around the camera (ELP USBFHD06H) to block out any remaining ambient light and make the sensor invariant to external light conditions. We placed the camera $28\,\mathrm{cm}$ below the acrylic sheet to capture the entire area and configured the camera to record frames at $30\,\mathrm{fps}$ with a resolution of $640\,\mathrm{px} \times 480\,\mathrm{px}$. To provide constant illumination of the particles for the camera, a band of cold-white LEDs was installed around the soft silicone layers (Figure 2(a)).

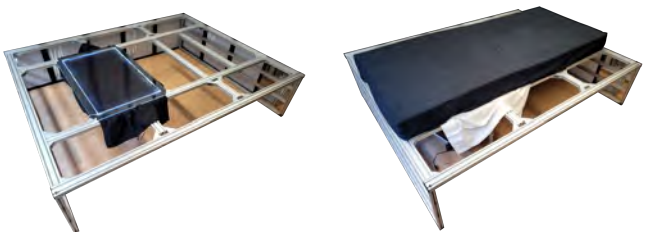

(a) The sensor has dimensions of $50\,\mathrm{cm} \times 80\,\mathrm{cm}$ and is mounted on aluminium profiles.

(b) A conventional standard-sized mattress was placed on top of the sensor.

Fig. 2: Test setup

For manufacturing the sensor, an open-mold approach was used to cast the silicone layers. The plexiglass sheet with added sidewalls served as the mold where the silicone was directly poured into. The LED band was attached to the inside of the sidewalls of the mold. The silicone layers were poured into the mold one after the other, leaving a period of at least 24 hours in between for curing. After all the layers were fully cured, the sidewalls were removed, leaving only the silicone and the LED band on the clear plexiglass sheet.

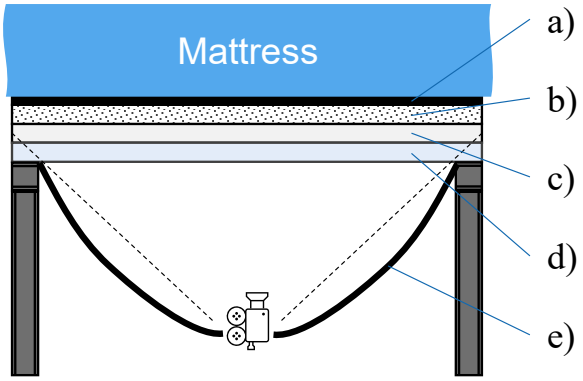

Fig. 3: Cross-sectional view of the sensor setup: a) a thin black silicone layer to block out ambient light; b) a soft silicone layer with embedded particles (Cospheric, United States); c) a harder clear silicone layer; d) an acrylic sheet to support the silicone layers; e) curtain to block out ambient light. The camera is mounted below the acrylic sheet, looking upwards.

We fabricated the sensor for an area of $50\,\mathrm{cm} \times 80\,\mathrm{cm}$ and placed it in the region of the chest and shoulders under the mattress. This placement was chosen as we expected the highest pressure values and therefore the most distinct features from the arms, shoulders, and shoulder blades.

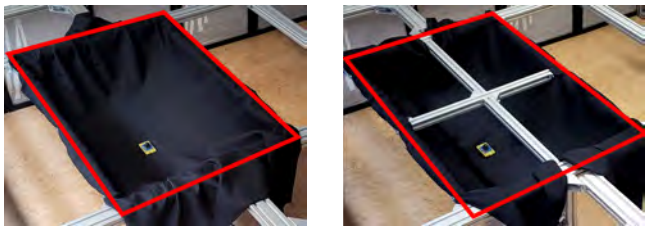

(a) Bottom of the sensor without reinforcement bars.

(b) Bottom of the sensor with reinforcement bars.

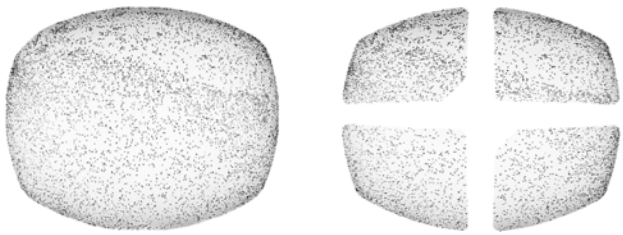

(c) Frame captured in setup without reinforcement bars. Colors inverted for better visibility.

(d) Frame captured in setup with reinforcement bars.

Fig. 4: Overview of the two configurations that were compared. The configuration without reinforcement had a higher particle visibility while the reinforced configuration had a higher stiffness.

B. Sensor Stiffness and Particle Visibility

Due to the large dimension and the flexibility of the silicone and the plexiglass sheet, locally applied pressure led to deformations of the entire sensor base.

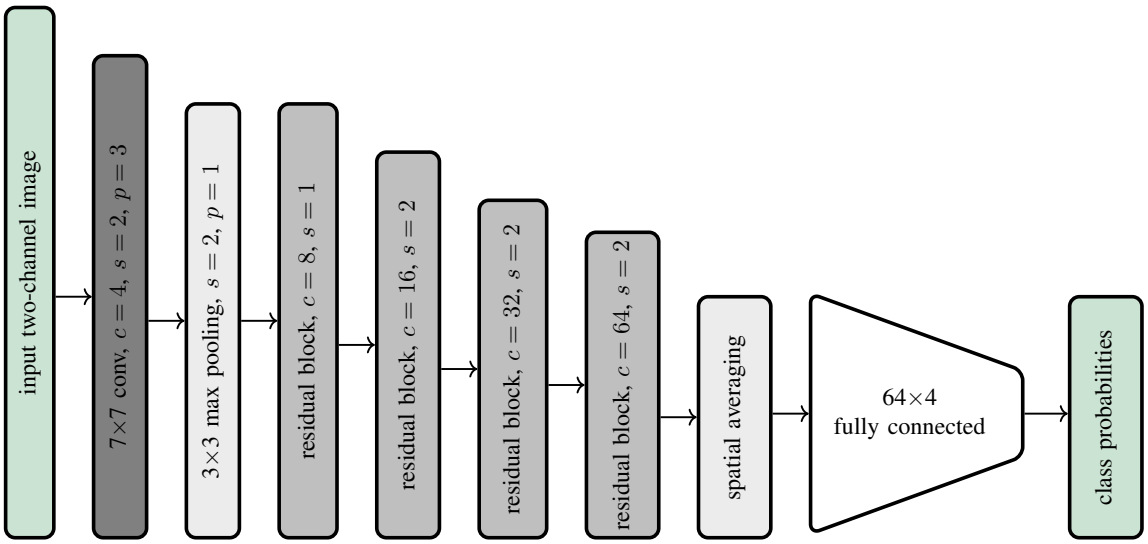

Fig. 5: Learning architecture with the residual blocks depicted in Figure 6. Instance normalization was used after each convolutional layer. This was always followed by rectified linear units (ReLUs), with the exception of the layers undergoing the sum operation in the residual blocks, where the ReLU was only applied after the sum. The spatial averaging layer performs the channel average across all pixels, with the resulting 64 values stacked together before being fed to the fully connected layer. The latter is followed by a softmax activation function that outputs class probabilities.

To reduce this effect, which can hinder proper detection of local deformation, we increased the overall stiffness of the base by adding aluminium reinforcement bars below the plexiglass sheet (Figure 4). However, this reduced the area of visible particles in the camera frame. To assess the tradeoff between increased stiffness and reduced particle visibility, we applied various weights ($258\,\mathrm{g}$, $500\,\mathrm{g}$, $5000\,\mathrm{g}$) to the center of both configurations. For each of these configurations, we used Otsu-thresholding to identify the pixels that showed the most significant difference in intensity between the deformed and the undeformed state [18]. Then we averaged the coordinates of all the pixels that passed this threshold. The intuition is that this operation provided a simple measure of the region where the particles underwent the highest change.

C. Dataset for Supervised Learning

We obtained four 10-minute recordings of one of the authors (male, $73\,\mathrm{kg}$, $187\,\mathrm{cm}$), complying with all the relevant national regulations, institutional policies, and following the ethical principles outlined in the Helsinki Declaration of 1975, as revised in 2000. In every recording, the positions (supine, prone, lateral left, lateral right) were kept for 25 seconds before repositioning. These four sleep positions represent the labels of the supervised learning setup discussed in Section II-D. During the entire recording, we used an ambulatory polysomnography (Nox A1, Nox Medical), to record ground-truth position labels. After removing the repositioning phases from the recorded data, we obtained a total of $44{,}136$ labeled RGB frames. In this dataset, equal amounts of datapoints were divided among the four classes. All the frames were then processed, by converting them to grayscale images and resizing them to a resolution of $160\,\mathrm{px} \times 120\,\mathrm{px}$ for the sake of efficient inference.

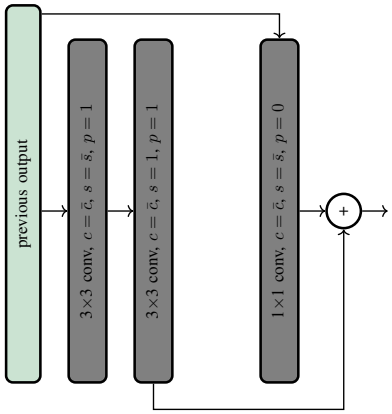

Fig. 6: Residual blocks of the learning architecture shown in Figure 5. All the blocks in green serve as placeholders. "3×3 conv" indicates a convolutional layer with a 3×3 filter size, with c, s, and p denoting the number of output channels, stride, and size of zero-padding, respectively.

D. Sleep Position Classifier

We trained a neural network to classify the sleep position of the person on the mattress according to the four classes defined above. The inputs to the neural network were two-channel images, one being the processed grayscale frame of the collected dataset, and the other being the grayscale frame captured at rest, before the person started to deform the sensor under the mattress. The purpose of the image at rest was to provide the network with information about the particle distribution in the undeformed state, as described in [15].

The network architecture is shown in Figure 5, and was based on feature extraction through convolutional layers and residual blocks. The residual blocks are a type of building block that allowed the network to learn deeper representations by retaining information learned in early layers. These were

then followed by a fully connected layer with its outputs processed through a softmax function that returned normalized class probabilities. The dataset was divided into four equal parts, each containing data corresponding to each of the four recordings (and balanced across the four classes). Therefore, 4-fold leave-one-recording-out cross validation was employed to evaluate the performance of the network. The network was then evaluated based on the average test accuracy, as reported in section III. During each run, the network was trained for 50 epochs by minimizing the log-likelihood loss with the Adam optimizer [19], using a batch size of 128 samples. A learning rate of $5e{-}5$ and weight decay of $1e{-}5$ were employed.

III. Results and Discussion

A. Sensor Characterization

The pixel intensity differences between the deformed and undeformed states for one weight are shown in Figure 7. Despite the reduced particle visibility in the reinforced setup, the stiffer structure resulted in a greater distance between the image center and the region with the highest change, allowing for better class separability across all weights. Furthermore, from Figure 8, we intuitively ascertain the high sensitivity of the sensor: For all sensor configurations, we observed a nonzero distance between the image center and the region with the highest change. This shows how the sensor could register even the lowest test weight of only 258 g placed on top of a conventional mattress. Moreover, the sensitivity increased after the reinforcement as the overall increased stiffness led to a greater particle displacement inside the soft silicone.

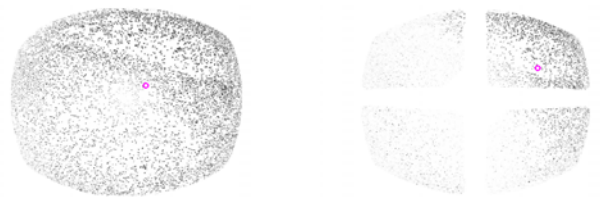

(a) 5000 g, without reinforcement (b) 5000 g, with reinforcement

Fig. 7: Intensity difference between the deformed and undeformed state for both configurations, with black denoting the pixels with the highest difference. The pink circle indicates the region of highest change, computed using pixel coordinate averaging for all pixels that passed the Otsu threshold [18].

B. Classifier Performance

We report an overall classification accuracy of 91.2 % after 4-fold cross validation (Figure 9). Note that data about each of the classes were collected at different overlapping locations across the mattress, which increases the difficulty of the task (e.g., the dataset included examples of the person lying in lateral right and lateral left in the same mattress regions). The tendency of the network to overclassify the prone label may be explained by a reduced class separability with the supine position, which could be improved by controlling the filtering effect of a dedicated mattress. The network inference takes about 1 ms on the single-core of a standard laptop CPU,

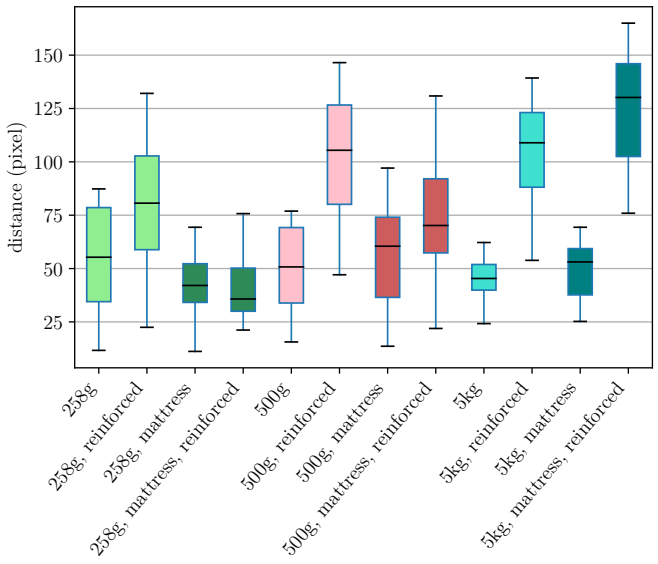

Fig. 8: Distances (px) between the image center and the region with the highest change (pink points in Figure 7) for both cases: weight on top of the mattress and weight directly on top of sensor. In all configurations, the distance was nonzero, indicating that the sensor detected the weight changes.

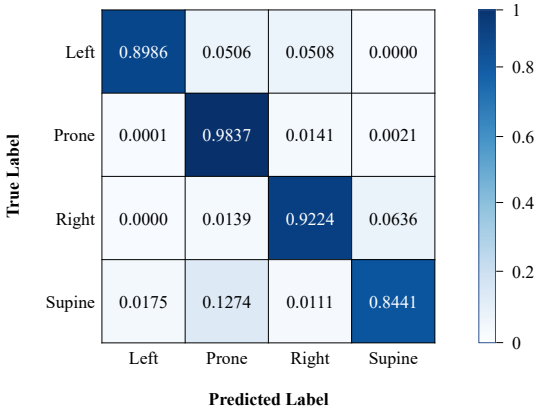

Fig. 9: Confusion matrix of the neural network as described in II-D. Overall classification accuracy was at 91.2 %.

fitting within the camera acquisition frame rate (30 fps). For closed-loop sleep position control, a much lower frequency is likely to be sufficient. However, the fast inference may still be beneficial for improving robustness to image noise by time-averaging over subsequent images.

IV. Conclusion

This work showed how to leverage advances in artificial tactile sensing technology to provide a novel solution for unobtrusive sleep position classification. The sensor prototype described in the paper is of particular interest in sleep robotic applications as it is fully hidden from the user, easily scalable, and highly sensitive, providing very useful information for accurate sleep position classification. Future work includes a more rigorous sensor characterization, the collection of data from multiple participants with different weight, and using the sensor for vital sign estimation and surface contact pressure estimation.

ACKNOWLEDGMENT

Alexander Breuss sincerely thanks the Promedica Foundation for their financial support in conducting this research.

REFERENCES

[1] T. Young, L. Evans, L. Finn, and M. Palta, "Estimation of the clinically diagnosed proportion of sleep apnea syndrome in middle-aged men and women," *Sleep*, vol. 20, no. 9, pp. 705–706, 1997.

[2] J. A. Dempsey, S. C. Veasey, B. J. Morgan, and C. P. O'Donnell, "Pathophysiology of sleep apnea," *Physiological reviews*, vol. 90, no. 1, pp. 47–112, 2010.

[3] V. K. Somers, D. P. White, R. Amin, W. T. Abraham, F. Costa, A. Culebras, S. Daniels, J. S. Floras, C. E. Hunt, L. J. Olson *et al.*, "Sleep apnea and cardiovascular disease: An american heart association/american college of cardiology foundation scientific statement from the american heart association council." *Journal of the American College of Cardiology*, vol. 52, no. 8, pp. 686–717, 2008.

[4] S. Javaheri, F. Barbe, F. Campos-Rodriguez, J. A. Dempsey, R. Khayat, S. Javaheri, A. Malhotra, M. A. Martinez-Garcia, R. Mehra, A. I. Pack *et al.*, "Sleep apnea: types, mechanisms, and clinical cardiovascular consequences," *Journal of the American College of Cardiology*, vol. 69, no. 7, pp. 841–858, 2017.

[5] N. M. Punjabi, B. S. Caffo, J. L. Goodwin, D. J. Gottlieb, A. B. Newman, G. T. O'Connor, D. M. Rapoport, S. Redline, H. E. Resnick, J. A. Robbins *et al.*, "Sleep-disordered breathing and mortality: a prospective cohort study," *PLoS medicine*, vol. 6, no. 8, p. e1000132, 2009.

[6] C. H. K. Lee, L. C. Leow, P. R. Song, H. Li, and T. H. Ong, "Acceptance and adherence to continuous positive airway pressure therapy in patients with obstructive sleep apnea (OSA) in a southeast asian privately funded healthcare system," *Sleep Sci.*, vol. 10, no. 2, pp. 57–63, Apr. 2017.

[7] M.-C. Yang, C.-Y. Lin, C.-C. Lan, C.-Y. Huang, Y.-C. Huang, C.-S. Lim, Y.-C. Liu, and Y.-K. Wu, "Factors affecting cpap acceptance in elderly patients with obstructive sleep apnea in taiwan," *Respiratory Care*, vol. 58, no. 9, pp. 1504–1513, 2013. [Online]. Available: https://rc.rcjournal.com/content/58/9/1504

[8] M. J. Ravesloot, D. White, R. Heinzer, A. Oksenberg, and J.-L. Pepin, "Efficacy of the new generation of devices for positional therapy for patients with positional obstructive sleep apnea: a systematic review of the literature and meta-analysis," *Journal of clinical sleep medicine*, vol. 13, no. 6, pp. 813–824, 2017.

[9] M. Meszaros, A. Breuss, E. Wilhelm, R. Riener, M. Kohler, and E. I. Schwarz, "Robotic beds for the treatment of positional obstructive sleep apnoea – a randomised cross-over pilot trial," *European Respiratory Journal*, vol. 60, no. suppl 66, 2022. [Online]. Available: https://erj.ersjournals.com/content/60/suppl_66/1177

[10] A. Breuss, N. Vonau, C. Ungricht, E. Schwarz, M. Irion, M. Bradicich, F. A. Grewe, S. Liechti, S. Thiel, M. Kohler, R. Riener, and E. Wilhelm, "Sleep position detection for closed-loop treatment of sleep-related breathing disorders," in *2022 International Conference on Rehabilitation Robotics (ICORR)*, 2022, pp. 1–6.

[11] C. Sferrazza and R. D'Andrea, "Design, motivation and evaluation of a full-resolution optical tactile sensor," *Sensors*, vol. 19, no. 4, p. 928, 2019. [Online]. Available: https://www.mdpi.com/1424-8220/19/4/928

[12] C. Sferrazza and R. D'Andrea, "Transfer learning for vision-based tactile sensing," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 7961–7967.

[13] C. Sferrazza, A. Wahlsten, C. Trueeb, and R. D'Andrea, "Ground truth force distribution for learning-based tactile sensing: A finite element approach," *IEEE Access*, vol. 7, pp. 173 438–173 449, 2019.

[14] C. Sferrazza, T. Bi, and R. D'Andrea, "Learning the sense of touch in simulation: a sim-to-real strategy for vision-based tactile sensing," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 4389–4396.

[15] C. Sferrazza and R. D'Andrea, "Sim-to-real for high-resolution optical tactile sensing: From images to three-dimensional contact force distributions," *Soft Robotics*, vol. 9, no. 5, pp. 926–937, 2022, pMID: 34842455. [Online]. Available: https://doi.org/10.1089/soro.2020.0213

[16] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, 2017.

[17] B. Ward-Cherrier, N. Pestell, L. Cramphorn, B. Winstone, M. E. Giannaccini, J. Rossiter, and N. F. Lepora, "The Tactip family: Soft optical tactile sensors with 3d-printed biomimetic morphologies," *Soft robotics*, vol. 5, no. 2, pp. 216–227, 2018.

[18] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.

[19] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.