

SleepVST: Sleep Staging from Near-Infrared Video Signals using Pre-Trained Transformers

Jonathan F. Carter^{*1,2} João Jorge² Oliver Gibson² Lionel Tarassenko¹

¹Institute of Biomedical Engineering, University of Oxford

²Oxehealth Ltd., Oxford

Abstract

Advances in camera-based physiological monitoring have enabled the robust, non-contact measurement of respiration and the cardiac pulse, which are known to be indicative of the sleep stage. This has led to research into camera-based sleep monitoring as a promising alternative to “gold-standard” polysomnography, which is cumbersome, expensive to administer, and hence unsuitable for longer-term clinical studies. In this paper, we introduce SleepVST, a transformer model which enables state-of-the-art performance in camera-based sleep stage classification (sleep staging). After pre-training on contact sensor data, SleepVST outperforms existing methods for cardio-respiratory sleep staging on the SHHS and MESA datasets, achieving total Cohen’s kappa scores of 0.75 and 0.77 respectively. We then show that SleepVST can be successfully transferred to cardio-respiratory waveforms extracted from video, enabling fully contact-free sleep staging. Using a video dataset of 50 nights, we achieve a total accuracy of 78.8% and a Cohen’s κ of 0.71 in four-class video-based sleep staging, setting a new state-of-the-art in the domain.

1. Introduction

Accurate sleep monitoring is critical for the diagnosis of sleep disorders and the discovery of novel treatments and biomarkers. Poor sleep is broadly linked with a number of health conditions, including cardiovascular diseases such as diabetes and hypertension [15]. Additionally, there are links between specific sleep stages and neurodegenerative conditions, such as a decrease in non-rapid-eye-movement stage 3 (N3) sleep with Alzheimer’s disease [23].

Overnight video polysomnography (vPSG), the “gold-standard” [52] technique for sleep monitoring, requires the

use of a large number of contact sensors including electrodes placed on the scalp (EEG), near the eyes (EOG), and under the chin (EMG) of the patient. Human experts (polysomnographers) must then review the recorded data, classifying the patient’s sleep into five stages at 30-second intervals (epochs) and annotating other important events such as leg movements and apnoeas. This manual process often takes multiple hours to complete.

Experts primarily rely on characteristic patterns in brain activity measured from the EEG to classify sleep stages [20]. However, sleep stage information is also encoded in measures of autonomic activity, including cardiac [48] and respiratory signals [19], and body movements [68]. This has led to the development of methods for automatically classifying sleep stages from sensors that measure these parameters, such as smartwatches [63].

Prior work has shown that these physiological parameters can also be measured using video cameras [18, 65, 66], leading to the development of methods for classifying sleep stages entirely from video input [6, 35, 60]. Camera-based methods have particular promise as part of a multi-purpose sleep monitoring system. For example, within elderly care settings, in addition to classifying sleep stages, they could also be used to detect specific sleep movement disorders [18], which are linked with conditions such as Parkinson’s, and which are typically distinguished via manual video review [51]. They could also be used to detect falls [12], one of the leading causes of injury and mortality amongst elderly individuals.

State-of-the-art video-based sleep staging [6] has used the heart rate (HR) and breathing rate (BR) derived from video to perform sleep staging. However, these *rates* are typically measured as averages over time, losing higher frequency inter-pulse and inter-breath interval information. EEG- [40] and wearable-based [22] sleep staging methods have both been improved by using raw waveforms as inputs, rather than derived quantities e.g. signal features [37, 45].

*jcarter@robots.ox.ac.uk

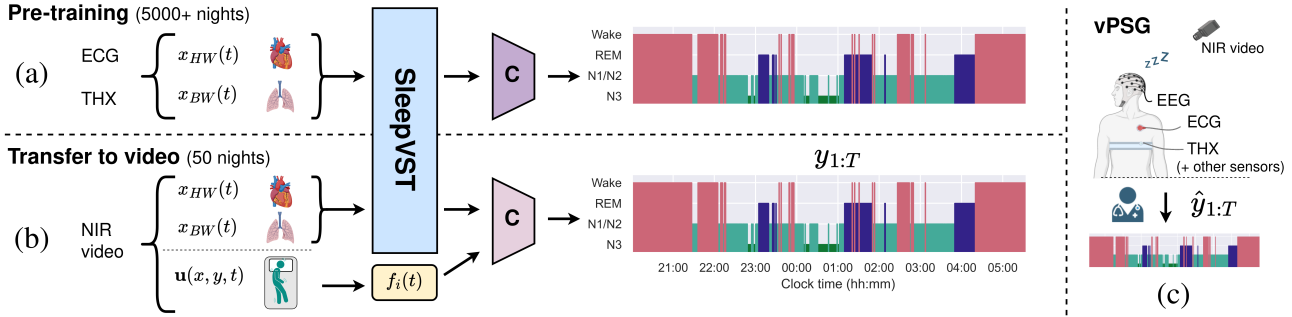


Figure 1. **Sleep staging from near-infrared video signals using SleepVST.** (a) We first pre-train the model end-to-end on cardiac (heart) and respiratory (breathing) waveforms $x_{HW}(t)$ and $x_{BW}(t)$ derived from the electrocardiogram (ECG) and a thoracic respiratory belt (THX) respectively. (b) After pre-training, we use the model as a frozen feature extractor, applying it to cardio-respiratory waveforms derived from near-infrared (NIR) video to generate sequences of features. When transferring to video data, we additionally use a set of motion features $f_i(t)$ derived from an optical flow field $u(x, y, t)$ as inputs to the classifier **C**. This approach allows us to utilise much larger contact-sensor datasets to train SleepVST, whilst also enabling the incorporation of motion information when transferring to video data. (c) We evaluate SleepVST using overnight video polysomnography (vPSG) studies, comparing expert-labelled sleep stage sequences $\hat{y}_{1:T}$ against those generated entirely from near-infrared video $y_{1:T}$ using our method.

In this paper, we introduce SleepVST¹, a transformer model for sleep stage classification from cardio-respiratory waveforms. First, we pre-train the model on waveforms derived from contact sensors (Fig. 1a), after which it outperforms existing methods for cardio-respiratory sleep staging on two widely used public datasets, highlighting the effectiveness of our transformer architecture.

After pre-training, we then show that SleepVST can be transferred to cardiac and respiratory waveforms measured entirely from near-infrared video, where our method additionally enables us to incorporate motion features derived from the camera data (Fig. 1b). Using this novel combination of inputs available from a near-infrared camera, we obtain state-of-the-art results for video-based sleep staging. Through a series of ablation studies we demonstrate the effectiveness of our approach. Our work further highlights the potential of camera-based sleep monitoring for important clinical applications.

2. Background and Related Work

Automated sleep staging. To reduce the manual effort of sleep assessment, deep learning methods have been developed for the automatic classification of sleep stages from polysomnography (PSG) recordings, achieving near-human accuracy [39]. Accuracy is commonly measured using Cohen’s kappa statistic (κ , [9]), which measures inter-rater agreement between two sources. Expert annotation of sleep stages is a subjective process, making 100% agreement unachievable [59]. Prior work has shown that sleep scorers agree around 80% of the time ($\kappa = 0.71$ - 0.81 , [11, 24]) when classifying sleep stages according to American Academy of

¹Sleep Video-Signal Transformer.

Sleep Medicine (AASM, [20]) scoring rules. These rules divide sleep into five classes: Wake, N1 (light), N2 (intermediate), N3 (deep) and rapid-eye-movement (REM) sleep. However, due to low inter-scorer agreement in identifying N1 [11], prior work has often treated sleep classification as a four-class problem, by merging N1 and N2 into a single class [22, 50].

Given the drawbacks of conventional polysomnography, prior work has investigated the potential of automatic sleep staging from a reduced number of contact sensors, including the electrocardiogram ([50], $\kappa = 0.66$), and the photoplethysmogram ([44], $\kappa = 0.65$).

Camera-based physiological measurement. Respiratory measurement is possible from video by measuring the induced movements of the chest and abdomen [65]. This has been achieved through methods such as frame-differencing [2, 54], using dense optical flow magnitudes [32, 33] and by tracking the movement of sparse optical flow vectors [25] across video frames.

Camera-based measurement of the pulse is possible using methods which measure specific cardiac phenomena [30]. These include remote photoplethysmography (rPPG, [42, 66]), which measures skin colour changes caused by variations in blood volume over the cardiac cycle, and ballistocardiography [4], which measures head movements induced as blood is pumped through the carotid arteries. More recently, deep learning methods have been developed which learn to extract a pulse signal directly from input video frames [7, 26, 70] rather than targeting specific underlying cardiac phenomena, including unsupervised approaches [49, 69].

Camera-based vital-sign measurement has been validated within clinical settings [55, 62], during overnight

polysomnography recordings [58, 64], and has led to products with regulatory clearance for camera-based measurement of the heart and breathing rate, including the Gili Pro Biosensor [57] and Oxehealth Vital Signs software [56].

Camera-based sleep monitoring. Prior work has investigated the binary problem of distinguishing sleep from wake using video cameras, using measures of activity derived from the input frames [10, 33, 46]. Whilst activity information alone can be effective in distinguishing sleep from wake, cardio-respiratory signals carry important additional information for classifying specific sleep stages [19, 48]. The current state-of-the-art approach for video-based sleep staging ([6], $\kappa = 0.61$) combines multiple modalities derived from video: the heart rate, breathing rate and activity information, to perform sleep stage classification. This approach outperforms prior methods which had used only a single video-derived modality: motion [21, 35] or the cardiac pulse [60].

3. Datasets and Preprocessing

3.1. Contact Sensor Data

To train the SleepVST model, we used the SHHS [43, 71] and MESA [8] contact sensor datasets. These datasets contain PSG recordings and sleep stages annotated according to AASM scoring rules. The cardio-respiratory signals available in each recording include the electrocardiogram (ECG), which measures electrical activity of the heart, and respiratory signals, measured using respiratory bands placed around the abdomen and thorax.

From the available signals, we seek to produce generic cardiac and respiratory waveforms to train the SleepVST model, preventing the model from learning to use information which may be present in signals derived from contact sensors but not present in those derived from video. This is to mitigate the problem of domain shift [28, 47], where the performance of a model is significantly reduced when applied to a different data distribution.

To produce a cardiac (heart) waveform, we first applied the same filtering steps used by the Pan-Tompkins algorithm [36] to the ECG signal. We resampled the output from the moving-window integration step of the algorithm to 10 Hz, then applied a Butterworth bandpass filter with a pass-band of 0.66–2.8 Hz (40–168 beats per minute or BPM) and Gaussian smoothing. To produce a respiratory (breathing) waveform, we used the thoracic respiratory signal as the input source, as previously used in [3]. This was down-sampled to 5 Hz and smoothed using a median filter with a kernel of length 5, to remove artefacts.

3.2. Video Data

We used the dataset previously introduced in [6] for our video-based sleep staging experiments. This dataset con-

sists of overnight vPSG recordings from 50 volunteers, with sleep stages annotated according to the AASM guidelines at 30-second intervals (epochs) by a sleep physiologist. Throughout this work, we refer to this as the Oxford Sleep Volunteers (OSV) dataset. Each recording includes 850 nm near-infrared (NIR) video data captured at 20 FPS. An example camera view can be seen in Figure 5. More information on the dataset is provided in §B.

To derive cardiac and respiratory waveforms from the video data, we used components of FDA-cleared Oxehealth Vital Signs software [56]. For each overnight video recording, we used the software to obtain cardiac and respiratory waveforms sampled at 10 Hz and 5 Hz respectively. We then applied the same filtering steps applied to the contact sensor waveforms: Butterworth filtering and Gaussian smoothing to the heart waveforms, and median filtering to the breathing waveforms.

Figure 2 shows examples of processed cardio-respiratory waveforms measured from contact sensors and video for a short section of a recording from the OSV dataset. These have additionally been normalised to zero mean and unit variance.




Figure 2. Example (normalized) cardiac and respiratory waveforms, $x_{HW}(t)$ and $x_{BW}(t)$, derived from contact sensors (blue) and video (orange) from the OSV dataset.

4. Methods

4.1. SleepVST Architecture

Figure 3 illustrates the SleepVST model architecture, which turns cardiac and respiratory waveforms into sequences of feature vectors using a transformer encoder. These feature vectors are then used to classify each 30-second sleep epoch. During pre-training, we used a linear layer for classification. When transferring to video data, we trained a new classifier that additionally used motion features derived from video as inputs, discussed further in Section 4.2.

Patchification. The continuous heart and breathing waveforms (HW and BW) are divided into N windows i.e. patches, drawing on ideas from PatchTST [34]. We used non-overlapping 30-second windows since this is the inter-

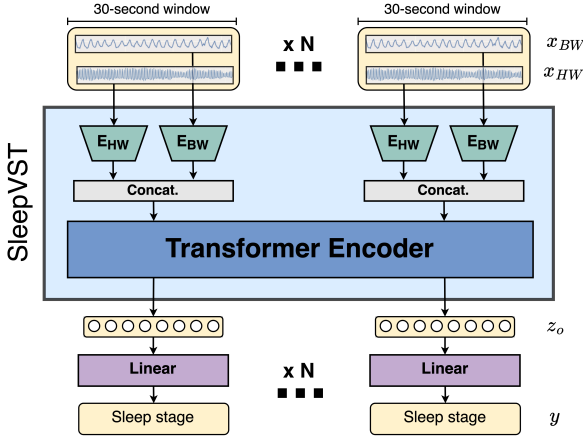


Figure 3. **SleepVST architecture.** Each 30-second window of heart (x_{HW}) and breathing waveforms (x_{BW}) is passed to a **patch encoder**, which turns them into patch-level features. These features are concatenated and passed to a **transformer encoder**. During pre-training, a **linear layer** turns the output feature sequence z_o of length N from SleepVST into sleep stage classifications.

val at which sleep stages are classified according to AASM guidelines. We denote the patchified inputs as $x_{HW} \in \mathbb{R}^{300 \times N}$ and $x_{BW} \in \mathbb{R}^{150 \times N}$ respectively, where N is the sequence length.

Encodings. Each waveform patch is normalised to zero mean and unit variance before being passed to encoding layers E_{HW} and E_{BW} , which transform them into feature vectors $z_{HW} \in \mathbb{R}^{D_{HW} \times N}$ and $z_{BW} \in \mathbb{R}^{D_{BW} \times N}$. The feature vectors from each patch are then concatenated, producing a sequence of feature vectors $z_i \in \mathbb{R}^{(D_{HW}+D_{BW}) \times N}$ which are passed to a transformer encoder.

We used identical, shallow 1D ResNet-style [17] models for E_{HW} and E_{BW} , illustrated in Figure 4. This convolutional encoder design [1] compresses the waveform patches into lower-dimensional feature vectors. This allowed us to keep D_{HW} and D_{BW} small, reducing the computational and memory complexity of the downstream transformer and the likelihood of over-fitting, particularly when transferring the model to the smaller video dataset.

This encoder design additionally provides a high degree of flexibility at the input level, such as the ability to use input signals of differing frequencies. In later ablation studies, this enables us to compare the performance of the SleepVST model with simpler variants which use 30-second patches of 1 Hz heart rate $x_{HR} \in \mathbb{R}^{30 \times N}$ and/or breathing rate $x_{BR} \in \mathbb{R}^{30 \times N}$ time-series as inputs, rather than underlying waveforms. When using either of these as inputs, we used linear projections [13, 34] for the input encoders, because of the lower dimensionality of x_{HR} and x_{BR} . Throughout our experiments we use the same feature dimension i.e. $D_{HW} = D_{BW} = D_{HR} = D_{BR} = 64$.

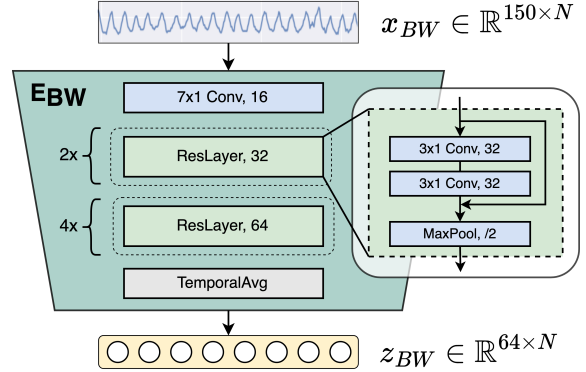


Figure 4. **Waveform encoder design.** Using a series of convolutional layers, the encoder turns sequences of signal patches into sequences of lower-dimensional feature vectors.

Transformer layer. We used a Transformer encoder [61] with sinusoidal position encodings to turn the sequence of input features z_i into an output sequence of features z_o with an identical feature dimension. Normal sleep exhibits long-range structures, known as sleep cycles [38], which typically last around 90–120 minutes and typically occur four to five times per night. From its superior performance in many sequence modelling tasks, including EEG-based sleep staging [41], we expect that the inductive biases of the transformer model (i.e. the attention mechanism) can enable it to learn to use relative position within a sleep cycle to classify sleep stages.

Classification. When training and testing on contact sensor data, we used a linear classification layer $W_C \in \mathbb{R}^{C \times (D_{HW}+D_{BW})}$ to transform the sequence of output feature vectors z_o into a sequence of sleep stage probabilities $y \in \mathbb{R}^{C \times N}$, where the number of classes C depends on the sleep classification strategy used. In line with prior work, we classified sleep stages into four classes during pre-training: Wake–N1/N2–N3–REM, i.e. $C = 4$.

Model training and inference. In all experiments, we used a sequence length of $N = 240$ windows i.e. two hours. During training, we created batches by sampling sequences of length N with a step of 5 minutes from within each recording. The model was trained to minimise the multi-class cross-entropy loss using the AdamW [27] optimiser with a learning rate of 0.0005 and weight decay of 0.001. Other hyper-parameters, such as the no. encoder layers, are detailed in §D. To apply the model to arbitrary length input sequences during evaluation, we re-applied it at regular intervals. More details on this process are provided in §E.

4.2. Motion features from Video

Following prior work [35], we seek to produce motion features from video which encode for information which is pre-

dictive of the sleep stage, such as the rate and distribution of body movements [68]. To do so, we define a set of parameterised time-series features that collectively quantify motion over a range of spatial and temporal scales.

Optical flow estimation. We first estimated a 2D optical flow field $\mathbf{u}(t, x, y) \in \mathbb{R}^2$ over the bed region. Inspired by sleep pose detection work [31], we applied a homography transformation [16] to the individual video frames before estimating the flow field, to achieve camera viewpoint invariance (shown in Figure 5).




Figure 5. **Example processing of an NIR video frame from the OSV dataset.** (a) Real viewpoint. (b) Virtual viewpoint. (c) Head (H), body (B), and outer (O) bed regions. The distance from the camera to the head region is around 1.5 m.

Feature definitions. Given an arbitrary region R of the optical flow field, we define the following parametrised signals:

$$v(t; R) = \max_{(x,y) \in R} |\mathbf{u}(t, x, y)| \quad (1)$$

$$s(t; R) = \frac{1}{|R|} \sum_{(x,y) \in R} |\mathbf{u}(t, x, y)| \quad (2)$$

i.e. the maximum (1) and the average (2) of the optical flow magnitude in the region R at time t . Using these signals, we define two time-series features that quantify motion within a time window Δ up to time t :

$$f_1(t; R, \Delta) = \sum_{t-\Delta}^t v(t; R) \quad (3)$$

$$f_2(t; R, \Delta) = \sum_{t-\Delta}^t s(t; R) \quad (4)$$

Plus two additional features which measure the time elapsed since either signal went above a threshold δ :

$$f_3(t; R, \delta) = \sum_{t-\tau}^t 1 \quad (5)$$

where $\tau = \max_{t' \leq t} t'$
s.t. $v(t'; R) > \delta$

$$f_4(t; R, \delta) = \sum_{t-\tau}^t 1 \quad (6)$$

where $\tau = \max_{t' \leq t} t'$
s.t. $s(t'; R) > \delta$

We calculated each feature for all possible combinations of parameters $\Delta \in \{30, 300\}$ seconds, $\delta \in \{0.01, 0.1, 1\}$ and $R \in \{H, B, O\}$, where $\{H, B, O\}$ are approximate head, body and outer bed regions as illustrated in Figure 5. This resulted in $(2+2+3+3)$ time-series features per region, i.e. 30 unique time-series features. Additionally, we provided the video-based classifier (see Section 4.3) with time-shifted features $f_i(t+T)$ for $T \in \{-90, 0, +90\}$ seconds. This gave a total of 90 motion features for each sleep epoch.

4.3. Transferring SleepVST to Video

To transfer the SleepVST model to video using the OSV dataset, we removed the linear classification layer learnt during pre-training and used the model as a frozen feature extractor, applying it to sequences of cardiac and respiratory waveforms measured from video to produce sequences of output features z_o . We then trained a new classifier that used these features, plus the motion features derived in the previous section, to classify individual sleep epochs. This allowed us to leverage the representational capacity of the transformer model, by pre-training on the much larger contact-sensor datasets, whilst also enabling us to incorporate motion information from the video dataset. We followed the design choice of [6] and used a Random Forest [5] for our classifier due to its low computational cost and robustness to overfitting.

5. Experimental Results

Sleep staging performance has commonly been reported in terms of classification accuracy (Acc_T) and total Cohen's kappa (κ_T) over all sleep epochs in the dataset, and/or using per-night mean statistics, Acc_μ and κ_μ , i.e.:

$$\kappa_T = K \left(\sum_{i=1}^N c(\mathbf{y}_i, \hat{\mathbf{y}}_i) \right) \quad (7)$$

$$\kappa_\mu = \frac{1}{N} \sum_{i=1}^N K(c(\mathbf{y}_i, \hat{\mathbf{y}}_i)) \quad (8)$$

where $\mathbf{y}_i, \hat{\mathbf{y}}_i \in \mathbb{R}^{C \times L_i}$ are the one-hot encoded model and scorer labels for recording i of length L_i and $c(\mathbf{y}_i, \hat{\mathbf{y}}_i) \in \mathbb{R}^{C \times C}$ is the confusion matrix between them. $K(\cdot)$ is the function which calculates the Cohen's kappa statistic [9] from a confusion matrix and is defined in §G. We report all of these statistics, to make fairer comparisons with prior work, and to aid future comparisons.




Figure 6. **Example four-class sleep hypnograms from the OSV dataset.** (top) Scored by an expert using signals from the vPSG recording. (bottom) Automatically generated from near-infrared video using SleepVST. (Cohen’s $\kappa = 0.71$ between model and expert sleep stages.)

5.1. Pre-training on Contact Sensors

From the SHHS dataset, we randomly selected 500 participants who participated in both visits, to form a test set of 1000 nights of data. This ensured that no participant appeared in both the training/validation and test partitions. Our test set is listed in §H. From the MESA dataset, we used the same test set of 204 nights as Kotzen *et al.* [22]. The remaining nights from both the SHHS and MESA datasets were pooled together and randomly split into training and validation sets using a 75:25 split.

Table 1 compares the overall performance of SleepVST with prior methods for cardio-respiratory sleep staging from contact sensors. These have used input modalities such as the photoplethysmogram (PPG) waveform and the ECG-derived heart rate. Using the ECG and thoracic respiratory signals, SleepVST outperforms prior methods for sleep staging from cardio-respiratory signals on both datasets.

5.2. Video-based Sleep Staging

In this section, we report the performance of the SleepVST model after transfer to the OSV video dataset. We follow the evaluation procedure of [6], training and evaluating on the 50 nights of data using 10-fold cross-validation with non-overlapping folds each containing five nights of data.

In Table 2, we see that our method achieves significantly better performance in terms of both accuracy and Cohen’s κ statistic when compared with prior video-based methods evaluated on similar study populations.

Figure 7 shows the total test confusion matrix between model and human expert (polysomnographer) classifications on the OSV dataset. We observe particularly high accuracy in distinguishing Wake and REM sleep stages.

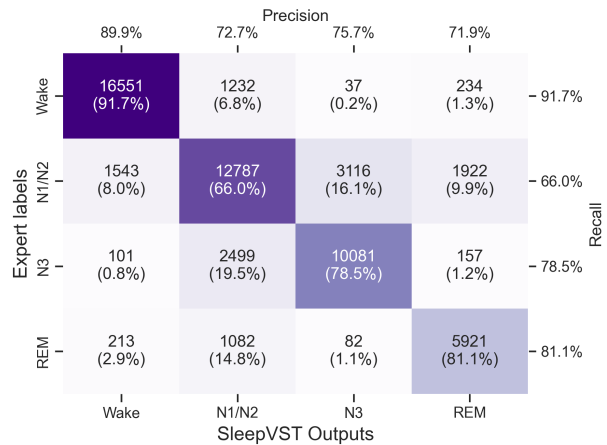


Figure 7. Confusion matrix between SleepVST and expert labels on four-class video based sleep staging.

Population performance. Figure 8 shows the distribution of Cohen’s κ values across participants from the OSV dataset plotted against age, sex and Fitzpatrick skin type² [14]. We observe a decrease in performance with age, as observed in prior sleep staging work [6, 53]. In older adults, measures of cardio-respiratory activity (e.g. heart rate variability) are known to decline with age [72]. This likely reduces the distance between sleep stages in cardio-respiratory input space and thus makes them harder to distinguish. Additional video data from older subjects would likely help to improve this, which is discussed further in §C.

²Plotted using representative skin tones from [67].

Table 1. Comparison of four-class cardio-respiratory sleep staging methods using contact sensors.

Dataset	N	Method	Modalities	Cohen's κ		Accuracy / %	
				κ_{μ}	κ_T	Acc_{μ}	Acc_T
SHHS [43]	296	Bakker et al. [3]	HR+Thor*+Air [†]	-	0.64	-	76.7
	800	Sridhar et al. [50]	HR (ECG)	0.65	0.66	77.3	77.0
	1000	SleepVST	ECG+Thor*	0.73	0.75	82.8	83.0
MESA [8]	296	Bakker et al. [3]	HR+Thor*+Air [†]	-	0.68	-	79.8
	194	Sridhar et al. [50]	HR (ECG)	-	0.69	-	80.0
	204	Kotzen et al. [22]	PPG	-	0.73 [‡]	-	82.6 [‡]
	204	SleepVST	ECG+Thor*	0.76	0.77	85.1	85.2

*Thoracic respiratory effort. [†]Nasal airflow. [‡]From reported confusion matrices.

Table 2. Comparison of four-class sleep staging methods using video cameras.

Dataset	N	Method	Modalities**	Cohen's κ		Accuracy / %	
				κ_{μ}	κ_T	Acc_{μ}	Acc_T
Healthy infants [21]	8	Kamon et al. [21]	Motion	0.26	-	48.0	-
Healthy adults [35]	6	Nochino et al. [35]	Motion	0.19	-	40.5	-
HealthBed Study [60]	46	van Meulen et al. [60]	HW*	0.49	0.49 [‡]	67.9	67.9 [‡]
Oxford Sleep Volunteers	50	Carter et al. [6]	HR+BR+Motion	0.61	0.64 [‡]	73.4	74.3 [‡]
	50	SleepVST	HW*+BW [†] +Motion	0.68	0.71	77.7	78.8

**Camera-derived. *Cardiac pulse (heart) waveforms. [†]Respiration (breathing) waveforms. [‡]From reported confusion matrices.




Figure 8. Scatter and box plot of Cohen's κ against age, sex and Fitzpatrick skin type² for individuals from the OSV dataset.

Example SleepVST output. Figure 6 shows an example hypnogram generated entirely from near-infrared video using our method. This example corresponds to the median Cohen's κ between model and expert from the OSV dataset. The model correctly identifies all three REM cycles, periods of deep sleep, and brief awakenings throughout the night. Additional hypnogram examples can be found in §C.

5.3. Ablation Studies

Video transfer strategy. In Table 3, we compare the performance of three methods for transferring the SleepVST model to video data:

1. Directly applying the pre-trained SleepVST model to cardiac and respiratory waveforms from video data.
2. Using SleepVST as a feature extractor and:
 - (a) training a new classification layer without using motion features.
 - (b) training a new classification layer that also uses motion features i.e. our best-performing approach.

Here we see that training a new classifier for video data is essential to the effectiveness of our approach. This step likely allows the classifier to ignore parts of the learnt feature space that do not transfer between contact-sensor and video-derived waveforms i.e. sensor-specific features.

Table 3. SleepVST video-based sleep staging performance for different model transfer strategies.

	κ_{μ}	κ_T	Acc_{μ} / %	Acc_T / %
Direct apply	0.510	0.536	66.7	67.3
Transfer w/o motion	0.660	0.689	76.4	77.4
Transfer w/ motion	0.677	0.708	77.7	78.8




Figure 9. Example four-class hypnograms during fragmented sleep from the OSV dataset. The use of motion features improves accuracy during short awakenings between 22.00 and 00.00.

Motion features. Table 4 reports the effectiveness of the two main components of our motion feature set: the set of feature definitions f_i and the set of regions used R . Additional ablations can be found in §C. Here we define and report an additional metric κ_F which is the Cohen’s κ statistic computed over epochs near transitions to/from Wake (W) i.e. where $\hat{y}_i^t \neq \hat{y}_i^{t+\Delta} = W$ or $\hat{y}_i^t = W \neq \hat{y}_i^{t+\Delta}$ for any $\Delta \in [-2, 2]$ epochs. This metric better quantifies accuracy during fragmented sleep, such as in Figure 9.

As measured by κ_F , we see that motion features aid classification during fragmented sleep. The ability to accurately measure sleep fragmentation is particularly important since it is linked with daytime sleepiness and decreased cognitive performance [29].

Table 4. Ablation study of motion feature parameters.

Ablation	Parameter Set	n_f^*	Cohen’s κ	
			κ_T	κ_F
Feature set	$f_i \in \{\}$	0	0.689	0.397
	$f_i \in \{f_1, f_2\}$	36	0.697	0.467
	$f_i \in \{f_3, f_4\}$	54	0.704	0.465
	$f_i \in \{f_1, f_2, f_3, f_4\}$	90	0.708	0.491
Region(s)	$R \in \{H \cup B\}$	30	0.702	0.461
	$R \in \{H \cup B, O\}$	60	0.704	0.468
	$R \in \{H, B, O\}$	90	0.708	0.491

*No. features.

Cardio-respiratory inputs. In Table 5, we compare the performance of the SleepVST model with variants which use heart rate x_{HR} or breathing rate x_{BR} as inputs instead of the underlying waveforms, as described in Section 4.1. We see that the use of waveforms, rather than derived heart and breathing rates, leads to improvements in sleep staging accuracy across all datasets.

Table 5. Comparison of four-class sleep staging performance for different cardio-respiratory input combinations to SleepVST.

Input signals	Cohen’s κ_T		
	SHHS	MESA	OSV
x_{HR}, x_{BR}	0.683	0.722	0.681
x_{HW}, x_{BR}	0.718	0.736	0.684
x_{HR}, x_{BW}	0.723	0.758	0.702
x_{HW}, x_{BW}	0.749	0.765	0.708

6. Conclusions

We have introduced SleepVST, a transformer model for sleep stage classification from cardio-respiratory waveforms. Using a transfer learning approach, we have shown how SleepVST can be used to perform sleep stage classification entirely from near-infrared video.

Our results advance the state-of-the-art in video-based sleep staging, narrowing the gap to expert-level performance and taking a significant step towards important clinical applications. More broadly, our work shows that a model pre-trained on a task using cardio-respiratory waveforms from large contact-sensor datasets can be effectively transferred to waveforms measured from video. This approach can benefit wider applications such as the detection of arrhythmias and sleep apnoeas, from cardiac and respiratory waveforms respectively.

Finally, we believe there is important future work in learning more expressive representations of motion during sleep. In addition to improving the performance of video-based sleep staging, this could also enable the detection of sleep movement disorders, such as REM behaviour disorder, removing the existing need for time-consuming, manual video review [51].

Acknowledgements

This work was supported by the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems [EP/S024050/1] and funded by Oxehhealth Ltd. Figure 1 was created with BioRender.com. We kindly thank the National Sleep Research Resource for granting access to SHHS and MESA.

References

- [1] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. In *Advances in Neural Information Processing Systems*, pages 12449–12460. Curran Associates, Inc., 2020. 4
- [2] Ying-Wen Bai, Wen-Tai Li, and Cheng-Hsiang Yeh. Design and implementation of an embedded monitor system for body breath detection by using image processing methods. In *2010 Digest of Technical Papers International Conference on Consumer Electronics (ICCE)*, pages 193–194, 2010. 2
- [3] Jessie P. Bakker, Marco Ross, Ray Vasko, Andreas Cerny, Pedro Fonseca, Jeff Jasko, Edmund Shaw, David P. White, and Peter Anderer. Estimating sleep stages using cardiorespiratory signals: validation of a novel algorithm across a wide range of sleep-disordered breathing severity. *Journal of Clinical Sleep Medicine*, 17(7):1343–1354, 2021. 3, 7
- [4] Guha Balakrishnan, Fredo Durand, and John Guttag. Detecting Pulse from Head Motions in Video. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3430–3437, Portland, OR, USA, 2013. IEEE. 2
- [5] Leo Breiman. Random Forests. *Machine Learning*, 45(1): 5–32, 2001. 5
- [6] Jonathan Carter, João Jorge, Bindia Venugopal, Oliver Gibson, and Lionel Tarassenko. Deep Learning-Enabled Sleep Staging From Vital Signs and Activity Measured Using a Near-Infrared Video Camera. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 5940–5949, Vancouver, BC, Canada, 2023. IEEE. 1, 3, 5, 6, 7
- [7] Weixuan Chen and Daniel McDuff. Deepphys: Video-based physiological measurement using convolutional attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 349–365, 2018. 2
- [8] Xiaoli Chen, Rui Wang, Phyllis Zee, Pamela L. Lutsey, Sogol Javaheri, Carmela Alcántara, Chandra L. Jackson, Michelle A. Williams, and Susan Redline. Racial/Ethnic Differences in Sleep Disturbances: The Multi-Ethnic Study of Atherosclerosis (MESA). *Sleep*, 38(6):877–888, 2015. 3, 7
- [9] Jacob Cohen. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1): 37–46, 1960. 2, 5
- [10] Kris Cuppens, Lieven Lagae, Bertien Ceulemans, Sabine Van Huffel, and Bart Vanrumste. Automatic video detection of body movement during sleep based on optical flow in pediatric patients with epilepsy. *Medical & Biological Engineering & Computing*, 48(9):923–931, 2010. 3
- [11] Heidi Danker-Hopfe, Peter Anderer, Josef Zeitlhofer, Marion Boeck, Hans Dorn, Georg Gruber, Esther Heller, Erna Loretz, Doris Moser, Silvia Parapatics, Bernd Saletu, Andrea Schmidt, and Georg Dorffner. Interrater reliability for sleep scoring according to the Rechtschaffen & Kales and the new AASM standard. *Journal of Sleep Research*, 18(1):74–84, 2009. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2869.2008.00700.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2869.2008.00700.x). 2
- [12] Koldo De Miguel, Alberto Brunete, Miguel Hernando, and Ernesto Gambao. Home Camera-Based Fall Detection System for the Elderly. *Sensors*, 17(12):2864, 2017. 1
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, 2021. 4
- [14] T. B. Fitzpatrick. The validity and practicality of sun-reactive skin types I through VI. *Archives of Dermatology*, 124(6): 869–871, 1988. 6
- [15] Marco Hafner, Martin Stepanek, Jirka Taylor, Wendy M. Troxel, and Christian van Stolk. Why Sleep Matters—The Economic Costs of Insufficient Sleep. *Rand Health Quarterly*, 6(4):11, 2017. 1
- [16] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 5
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 4
- [18] Adrienne Heinrich, Xavier Aubert, and Gerard de Haan. Body movement analysis during sleep based on video motion estimation. In *2013 IEEE 15th International Conference on e-Health Networking, Applications and Services (Healthcom 2013)*, pages 539–543, 2013. 1
- [19] D. W. Hudgel, R. J. Martin, B. Johnson, and P. Hill. Mechanics of the respiratory system and breathing pattern during sleep in normal humans. *Journal of Applied Physiology: Respiratory, Environmental and Exercise Physiology*, 56(1): 133–137, 1984. 1, 3
- [20] C. Iber. The AASM Manual for the Scoring of Sleep and Associated Events: Rules, Terminology, and Technical Specification. 2007. 1, 2
- [21] Masamitsu Kamon, Shima Okada, Masafumi Furuta, and Koki Yoshida. Development of a non-contact sleep monitoring system for children. *Frontiers in Digital Health*, 4, 2022. 3, 7
- [22] Kevin Kotzen, Peter H. Charlton, Sharon Salabi, Lea Amar, Amir Landesberg, and Joachim A. Behar. SleepPPG-Net: A Deep Learning Algorithm for Robust Sleep Staging From Continuous Photoplethysmography. *IEEE Journal of Biomedical and Health Informatics*, 27(2):924–932, 2023. 1, 2, 6, 7
- [23] Yee Fun Lee, Dmitry Gerashchenko, Igor Timofeev, Brian J. Bacskaï, and Ksenia V. Kastanenka. Slow Wave Sleep Is a Promising Intervention Target for Alzheimer’s Disease. *Frontiers in Neuroscience*, 14, 2020. 1
- [24] Yun Ji Lee, Jae Yong Lee, Jae Hoon Cho, and Ji Ho Choi. Interrater reliability of sleep stage scoring: a meta-analysis. *Journal of Clinical Sleep Medicine*, 18(1):193–202, 2022. 2
- [25] Michael H. Li, Azadeh Yadollahi, and Babak Taati. A non-contact vision-based system for respiratory rate estimation. In *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 2119–2122, 2014. 2

- [26] Xin Liu, Brian Hill, Ziheng Jiang, Shwetak Patel, and Daniel McDuff. EfficientPhys: Enabling Simple, Fast and Accurate Camera-Based Cardiac Measurement. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5008–5017, 2023. [2](#)
- [27] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019. [4](#)
- [28] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a Closer Look at Domain Shift: Category-Level Adversaries for Semantics Consistent Domain Adaptation. pages 2507–2516, 2019. [3](#)
- [29] S E Martin, H M Engleman, I J Deary, and N J Douglas. The effect of sleep fragmentation on daytime function. *American Journal of Respiratory and Critical Care Medicine*, 153(4): 1328–1332, 1996. Publisher: American Thoracic Society - AJRCCM. [8](#)
- [30] Daniel McDuff. Camera Measurement of Physiological Vital Signs. *ACM Computing Surveys*, 55(9):1–40, 2023. [2](#)
- [31] Sara Mahvash Mohammadi, Shirin Enshaeifar, Adrian Hilton, Derk-Jan Dijk, and Kevin Wells. Transfer Learning for Clinical Sleep Pose Detection Using a Single 2D IR Camera. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 29:290–299, 2021. [5](#)
- [32] K. Nakajima, A. Osa, and H. Miike. A method for measuring respiration and physical activity in bed by optical flow analysis. In *Proceedings of the 19th Annual International Conference of the IEEE Engineering in Medicine and Biology Society.*, pages 2054–2057 vol.5, 1997. [2](#)
- [33] Kazuki Nakajima, Yoshiaki Matsumoto, and Toshiyo Tamura. Development of real-time image sequence analysis for evaluating posture change and respiratory rate of a subject in bed. *Physiological Measurement*, 22(3):N21–N28, 2001. [2](#), [3](#)
- [34] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *The Eleventh International Conference on Learning Representations*, 2022. [3](#), [4](#)
- [35] Teruaki Nochino, Yuko Ohno, Takafumi Kato, Masako Tanike, and Shima Okada. Sleep stage estimation method using a camera for home use. *Biomedical Engineering Letters*, 9(2):257–265, 2019. [1](#), [3](#), [4](#), [7](#)
- [36] Jiapu Pan and Willis J. Tompkins. A Real-Time QRS Detection Algorithm. *IEEE Transactions on Biomedical Engineering*, BME-32(3):230–236, 1985. [3](#)
- [37] James Pardey, Stephen Roberts, Lionel Tarassenko, and John Stradling. A new approach to the analysis of the human sleep/wakefulness continuum. *Journal of sleep research*, 5(4):201–210, 1996. [1](#)
- [38] Aakash K. Patel, Vamsi Reddy, and John F. Araujo. *Physiology, Sleep Stages*. StatPearls Publishing, 2022. [4](#)
- [39] Huy Phan and Kaare Mikkelsen. Automatic sleep staging of EEG signals: recent development, challenges, and future directions. *Physiological Measurement*, 43(4):04TR01, 2022. [2](#)
- [40] Huy Phan, Oliver Y. Chén, Minh C. Tran, Philipp Koch, Alfred Mertins, and Maarten De Vos. XSleepNet: Multi-View Sequential Model for Automatic Sleep Staging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9): 5903–5915, 2022. [1](#)
- [41] Huy Phan, Kaare Mikkelsen, Oliver Y. Chén, Philipp Koch, Alfred Mertins, and Maarten De Vos. SleepTransformer: Automatic Sleep Staging With Interpretability and Uncertainty Quantification. *IEEE Transactions on Biomedical Engineering*, 69(8):2456–2467, 2022. [4](#)
- [42] Ming-Zher Poh, Daniel J. McDuff, and Rosalind W. Picard. Advancements in Noncontact, Multiparameter Physiological Measurements Using a Webcam. *IEEE Transactions on Biomedical Engineering*, 58(1):7–11, 2011. Conference Name: IEEE Transactions on Biomedical Engineering. [2](#)
- [43] S. F. Quan, B. V. Howard, C. Iber, J. P. Kiley, F. J. Nieto, G. T. O’Connor, D. M. Rapoport, S. Redline, J. Robbins, J. M. Samet, and P. W. Wahl. The Sleep Heart Health Study: design, rationale, and methods. *Sleep*, 20(12):1077–1085, 1997. [3](#), [7](#)
- [44] Mustafa Radha, Pedro Fonseca, Arnaud Moreau, Marco Ross, Andreas Cerny, Peter Anderer, Xi Long, and Ronald M. Aarts. A deep transfer learning approach for wearable sleep stage classification with photoplethysmography. *npj Digital Medicine*, 4(1):1–11, 2021. [2](#)
- [45] S.J. Redmond and C. Heneghan. Cardiorespiratory-based sleep staging in subjects with obstructive sleep apnea. *IEEE Transactions on Biomedical Engineering*, 53(3):485–496, 2006. [1](#)
- [46] A. J. Schwichtenberg, Jeehyun Choe, Ashleigh Kellerman, Emily A. Abel, and Edward J. Delp. Pediatric Videosomnography: Can Signal/Video Processing Distinguish Sleep and Wake States? *Frontiers in Pediatrics*, 6:158, 2018. [3](#)
- [47] Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2):227–244, 2000. [3](#)
- [48] Z. Shinar, A. Baharav, Y. Dagan, and S. Akselrod. Automatic detection of slow-wave-sleep using heart rate variability. In *Computers in Cardiology 2001. Vol.28*, pages 593–596, 2001. [1](#), [3](#)
- [49] Jeremy Speth, Nathan Vance, Patrick Flynn, and Adam Czajka. Non-contrastive unsupervised learning of physiological signals from video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14464–14474, 2023. [2](#)
- [50] Niranjan Sridhar, Ali Shoeb, Philip Stephens, Alaa Kharbouch, David Ben Shimol, Joshua Burkart, Atiyeh Ghoreyshy, and Lance Myers. Deep learning for automated sleep staging using instantaneous heart rate. *npj Digital Medicine*, 3(1):1–10, 2020. [2](#), [7](#)
- [51] Ambra Stefani and Birgit Högl. Sleep in Parkinson’s disease. *Neuropsychopharmacology*, 45(1):121–128, 2020. [1](#), [8](#)
- [52] Ambra Stefani, David Gabelia, Thomas Mitterling, Werner Poewe, Birgit Högl, and Birgit Frauscher. A Prospective Video-Polysomnographic Analysis of Movements during Physiological Sleep in 100 Healthy Sleepers. *Sleep*, 38(9):1479–1487, 2015. [1](#)
- [53] Haoqi Sun, Wolfgang Ganglberger, Ezhil Panneerselvam, Michael J Leone, Syed A Quadri, Balaji Goparaju, Ryan A

- Tesh, Oluwaseun Akeju, Robert J Thomas, and M Brandon Westover. Sleep staging from electrocardiography and respiration with deep learning. *Sleep*, 43(7):zsz306, 2020. 6
- [54] K. S. Tan, R. Saatchi, H. Elphick, and D. Burke. Real-time vision based respiration monitoring system. In *2010 7th International Symposium on Communication Systems, Networks & Digital Signal Processing (CSNDSP 2010)*, pages 770–774, 2010. 2
- [55] Alexander Trumpp, Johannes Lohr, Daniel Wedekind, Martin Schmidt, Matthias Burghardt, Axel R. Heller, Hagen Malberg, and Sebastian Zauneder. Camera-based photoplethysmography in an intraoperative setting. *BioMedical Engineering OnLine*, 17(1):33, 2018. 2
- [56] U.S. Food and Drug Administration (FDA). De Novo Classification Request for Oxehealth Vital Signs, 2020, https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN200019.pdf. 3
- [57] U.S. Food and Drug Administration (FDA). De Novo Classification Request for Gili Pro Biosensor, 2020, https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN200038.pdf. 3
- [58] Mark van Gastel, Sander Stuijk, Sebastiaan Overeem, Johannes P. van Dijk, Merel M. van Gilst, and Gerard de Haan. Camera-Based Vital Signs Monitoring During Sleep – A Proof of Concept Study. *IEEE Journal of Biomedical and Health Informatics*, 25(5):1409–1418, 2021. 3
- [59] Hans van Gorp, Iris A M Huijben, Pedro Fonseca, Ruud J G van Sloun, Sebastiaan Overeem, and Merel M van Gilst. Certainty about uncertainty in sleep staging: a theoretical framework. *Sleep*, 45(8), 2022. 2
- [60] Fokke B. van Meulen, Angela Grassi, Leonie van den Heuvel, Sebastiaan Overeem, Merel M. van Gilst, Johannes P. van Dijk, Henning Maass, Mark J. H. van Gastel, and Pedro Fonseca. Contactless Camera-Based Sleep Staging: The HealthBed Study. *Bioengineering*, 10(1):109, 2023. 1, 3, 7
- [61] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. *arXiv:1706.03762 [cs]*, 2017. arXiv: 1706.03762. 4
- [62] Mauricio Villarroel, Sitthichok Chaichulee, João Jorge, Sara Davis, Gabrielle Green, Carlos Arteta, Andrew Zisserman, Kenny McCormick, Peter Watkinson, and Lionel Tarassenko. Non-contact physiological monitoring of preterm infants in the Neonatal Intensive Care Unit. *npj Digital Medicine*, 2(1):1–18, 2019. 2
- [63] Olivia Walch, Yitong Huang, Daniel Forger, and Cathy Goldstein. Sleep stage prediction with raw acceleration and photoplethysmography heart rate data derived from a consumer wearable device. *Sleep*, 42(12), 2019. 1
- [64] Qiongyan Wang, Hanrong Cheng, and Wenjin Wang. Feasibility of Exploiting Physiological and Motion Features for Camera-based Sleep Staging: A Clinical Study. In *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 1–5, 2023. ISSN: 2694-0604. 3
- [65] Wenjin Wang and Albertus C. den Brinker. Algorithmic insights of camera-based respiratory motion extraction. *Physiological measurement*, 43(7):075004, 2022. 1, 2
- [66] Wenjin Wang, Albertus C. den Brinker, Sander Stuijk, and Gerard de Haan. Algorithmic Principles of Remote PPG. *IEEE Transactions on Biomedical Engineering*, 64(7):1479–1491, 2017. 1, 2
- [67] William H. Ward, Fernando Lambreton, Neha Goel, Jian Q. Yu, and Jeffrey M. Farma. Clinical Presentation and Staging of Melanoma. In *Cutaneous Melanoma: Etiology and Therapy*. Codon Publications, Brisbane (AU), 2017. 6
- [68] Johanna Wilde-Frenz and Hartmut Schulz. Rate and Distribution of Body Movements during Sleep in Humans. *Perceptual and Motor Skills*, 56(1):275–283, 1983. 1, 5
- [69] Yuzhe Yang, Xin Liu, Jiang Wu, Silviu Borac, Dina Katabi, Ming-Zher Poh, and Daniel McDuff. SimPer: Simple Self-Supervised Learning of Periodic Targets, 2023. 2
- [70] Zitong Yu, Yuming Shen, Jingang Shi, Hengshuang Zhao, Philip Torr, and Guoying Zhao. PhysFormer: Facial Video-based Physiological Measurement with Temporal Difference Transformer. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4176–4186, New Orleans, LA, USA, 2022. IEEE. 2
- [71] Guo-Qiang Zhang, Licong Cui, Remo Mueller, Shiqiang Tao, Matthew Kim, Michael Rueschman, Sara Mariani, Daniel Mobley, and Susan Redline. The National Sleep Research Resource: towards a sleep data commons. *Journal of the American Medical Informatics Association: JAMIA*, 25(10):1351–1358, 2018. 3
- [72] Dan Ziegler, G. Laux, K. Dannehl, M. Spüler, H. Mühlen, P. Mayer, and F.a. Gries. Assessment of Cardiovascular Autonomic Function: Age-related Normal Ranges and Reproducibility of Spectral Analysis, Vector Analysis, and Standard Tests of Heart Rate Variation and Blood Pressure Responses. *Diabetic Medicine*, 9(2):166–175, 1992. 6