



OPEN

Automated scoring of pre-REM sleep in mice with deep learning

Niklas Grieger¹, Justus T. C. Schwabedal², Stefanie Wendel³, Yvonne Ritze³ & Stephan Bialonski¹✉

Reliable automation of the labor-intensive manual task of scoring animal sleep can facilitate the analysis of long-term sleep studies. In recent years, deep-learning-based systems, which learn optimal features from the data, increased scoring accuracies for the classical sleep stages of Wake, REM, and Non-REM. Meanwhile, it has been recognized that the statistics of transitional stages such as pre-REM, found between Non-REM and REM, may hold additional insight into the physiology of sleep and are now under vivid investigation. We propose a classification system based on a simple neural network architecture that scores the classical stages as well as pre-REM sleep in mice. When restricted to the classical stages, the optimized network showed state-of-the-art classification performance with an out-of-sample F1 score of 0.95 in male C57BL/6J mice. When unrestricted, the network showed lower F1 scores on pre-REM (0.5) compared to the classical stages. The result is comparable to previous attempts to score transitional stages in other species such as transition sleep in rats or N1 sleep in humans. Nevertheless, we observed that the sequence of predictions including pre-REM typically transitioned from Non-REM to REM reflecting sleep dynamics observed by human scorers. Our findings provide further evidence for the difficulty of scoring transitional sleep stages, likely because such stages of sleep are under-represented in typical data sets or show large inter-scorer variability. We further provide our source code and an online platform to run predictions with our trained network.

Sleep is one of the most fundamental and not yet well understood processes that can be observed in humans and most animals. Since the discovery of sleep-wake correlates in electroencephalographic (EEG) signals in the early 20th century, the sleep architecture (i.e., the temporal succession of different states of consciousness) has been under detailed investigation, and it is now known to affect critical processes such as autonomic function, mood, cognitive function, and memory consolidation^{1,2}. Essential to most sleep studies is the manual scoring of EEG recordings, which is known as a time-consuming, labor intensive task. Automated scoring systems have been developed since the late 1960s^{3,4} to facilitate sleep scoring in humans^{5–7} and in animal models such as rodents^{8,9}. While manual scoring remains the gold standard to evaluate sleep studies, such systems could potentially reduce problems such as inter-rater variability and analyst's fatigue introducing systematic errors into the derived sleep architecture.

In rodents, and in particular in mice, the development of systems to automatically score sleep focused on distinguishing between the main sleep stages Wake (wakefulness), REM (rapid eye movement), and NREM (non-rapid eye movement)^{8,9}. While approaches to automatically identify these stages have become increasingly successful, sleep research has uncovered finer structures including sub-stages in mice^{10,11} and *pre-REM* sleep¹², also known as *intermediate stage*^{13,14} or *transition sleep*¹⁵ in rats and mice^{16,17}. Efforts to automatically identify such sleep patterns have been limited to rats, where studies have introduced systems to classify four to seven stages such as transition sleep^{12,18–21}, different Wake states^{12,22}, or different NREM stages^{12,19,21}. In mice, however, published work to date has focused on scoring only the three main stages, with very few exceptions (different Wake states²³, cataplectic events²⁴).

Classical approaches towards automating sleep scoring in rodents are based on *engineered features*⁸ which are often inspired by human sleep scorers who use visual representations of EEG signals for their scoring decisions. Power spectral densities (e.g., magnitudes and ratios of power in delta, theta, sigma bands^{22,25–28}), and EEG amplitudes (e.g., moments of EEG amplitude distributions⁸) are among the most frequently used features. While systems based on carefully engineered features have seen progress in scoring performance, deep-learning-based

¹Department of Medical Engineering and Technomathematics, FH Aachen University of Applied Sciences, Jülich 52428, Germany. ²Independent researcher, Kottbusser Damm 35, Berlin 10967, Germany. ³Department of Medical Psychology and Behavioral Neurobiology, University of Tübingen, Tübingen 72076, Germany. ✉email: bialonski@fh-aachen.de

	training set	validation set	test set	all data
number of mice	14	2	2	18
recording days	42	6	4	52
number of segments	366,720	52,140	35,420	454,280
percentage	80.7 %	11.5 %	7.8 %	100 %

Table 1. Characteristics of the dataset. Data from 18 mice were split into three disjunct sets (training, validation, and test set). To allow for robust assessment of generalization performance of our networks out-of-sample, data of each mouse was assigned to only one of these sets. Each recording was further split into segments of 10 seconds for which sleep stages (labels) were predicted.

systems that can create and use *learned features* have repeatedly been demonstrated to achieve superior classification performance in related fields such as speech or image recognition^{29,30}. These successes indicate that learned features may be particularly effective and expressive representations of data with respect to the classification challenges at hand, as well as various EEG-based challenges including brain-computer interfaces, epileptic seizure detection, or sleep scoring³¹.

While research into human sleep scoring has adopted ideas and methods from deep learning^{6,7}, there are, to our knowledge, only few studies proposing deep-learning-based systems to improve animal sleep scoring, particularly in mice. Deep neural networks introduced in two studies^{32,33} were inspired by image recognition systems and used preprocessed spectrograms of EEG and the electromyogram (EMG) as input “images”. These systems yielded state-of-the-art classification performances but needed an additional hidden Markov model³² that constrained the output of the system on physiologically plausible sleep state transitions or used mixture z-scoring³³ which needs an initial sample of human-annotated data for each mouse. Four studies proposed deep neural networks that were trained end-to-end on raw input EEG and EMG signals. One of the earliest studies³⁴ was based on 22 years of EEG and EMG signals from mice and reported an improvement of sleep scoring accuracy with respect to classical approaches. The system contained a bidirectional long short-term memory (LSTM) module in the classifier head to model long-range non-linear correlations³⁵, and the training also included a retraining scheme. The second study³⁶ reported a convolutional neural network to show better classification performance compared to a random forest trained on engineered features but was not based on human-annotated data (the study relied on software annotated data instead). Finally, the third study²⁴ trained a convolutional neural network to predict sleep stages in narcoleptic and wild-type mice where cataplectic events were derived based on rules operating on the predictions (Wake, REM, or NREM) of the network. All discussed neural networks were trained to distinguish between the main sleep stages Wake, REM, and NREM only.

In this contribution, we introduce a neural network to automatically score the main stages Wake, NREM, and REM as well as pre-REM sleep and artifacts in mice with one EEG channel as input. To the best of our knowledge, this is the first contribution to study the capability of deep neural networks to score pre-REM in mice. We discussed challenges that arise from predicting pre-REM sleep including class imbalance and absence of consensus⁹ for scoring rodent sleep. When constraining our network to score the main stages only, our network showed state-of-the-art performance compared with previous work. We consider network architectures as proposed here and in other contributions^{24,36} to be particularly promising for promoting further development of sleep scoring systems. Finally, source code³⁷ and an online platform³⁸ are provided to allow for adapting and testing our trained network.

Materials and methods

Data acquisition. The dataset consists of polysomnographic recordings of 18 mice with a total recording duration of 52 days (corresponding to 454,301 windows (segments) of 10 seconds each, see Table 1). Each mouse was recorded for 3 days, except for one mouse whose recording spanned 1 day because it unexpectedly passed away. The data were acquired during a study at the University of Tübingen which tested the influence of dietary variations on sleep.

Animal nurturing and treatment. All mice (age about 10 weeks, male, C57BL/6J) were kept at the animal facility FORS at the University of Tübingen, Germany. All applicable local and federal regulations of animal welfare in research (Directive 86/609, 1986 European Community) and ARRIVE guidelines were followed. Experiments were approved by the ethics committee of the Regional Council Tübingen, Germany, permit number MPV 3/15. Mice were housed under SPV conditions in IVC-cages with cedar bedding and cellulose cloth and experiments were conducted at controlled temperature ($20 \pm 2 \text{ C}^\circ$), humidity, and 12h light/dark cycle at all times (light on at 06:00 or 7:00 summer/winter time). Out of the 18 mice, 9 were fed with a different diet (sucrose solution instead of water) to test an experimental paradigm to be published elsewhere. We did not consider the dietary groups as relevant for the present study and did not attempt a statistical comparison thereof. The mouse that passed away unexpectedly was part of the altered-diet group.

Polysomnographic recordings. The mice were chronically implanted with a synchronous recording device of one EMG and four EEG electrodes seven days before the first recording. Surgeries were conducted in the morning between 8am and 12pm. Four stainless steel screws were placed on top of the cortex using the following coordinates: frontal electrode (AP: +1.5 mm, L: +1.0 mm, relative to Bregma), two electrodes parietal left and right

label	Wake	REM	NREM	pre-REM	artifact
all data					
# segments	250,206	23,065	171,270	8,536	1,203
perc.	55.08 %	5.08 %	37.70 %	1.88 %	0.26 %
train. set					
# segments	203,747	18,606	136,586	7,181	600
perc.	55.56 %	5.07 %	37.25 %	1.96 %	0.16 %
reb. train. set 1					
# segments	110,017	66,010	88,013	51,341	51,341
perc.	30 %	18 %	24 %	14 %	14 %
reb. train. set 2					
# segments	150,110	91,531	124,481	-	-
perc.	41 %	25 %	34 %	-	-
validation set					
# segments	28,382	2,544	19,952	752	510
perc.	54.43 %	4.88 %	38.27 %	1.44 %	0.98 %
test set					
# segments	18,077	1,915	14,732	603	93
perc.	51.04 %	5.41 %	41.59 %	1.70 %	0.26 %

Table 2. Statistics of labeled segments of the whole dataset (“all data”), of the training set (“train. set”), the rebalanced training sets (“reb. train. set 1 and 2”), the validation set (“validation set”), and the test set (“test set”). NREM: Non REM sleep, pre-REM: Pre REM sleep, Wake: Wakefulness. perc: percentage, # segments: number of segments.

(AP: 2.0 mm, L: ± 2.5 mm), and the reference electrode occipital (AP: -1.0 mm, L: 0 mm, relative to Lambda) according to the atlas of Franklin and Paxinos³⁹. This setup was used to sample the electrical activity of the frontal, parietal-left, and parietal-right lobes. Two flexible stainless steel wires were inserted into the neck muscles to measure the EMG signal. During recordings, mice were plugged with a cable attached to a swiveling commutator, which was connected to the amplifier (Model 15A54, Grass Technologies, USA). The amplified (x1000) signals (EEG and EMG) were digitalized, filtered (EEG: 0.01–300 Hz; EMG: 30–300 Hz), and sampled at a rate of 992.06 Hz. The commutator units were built in-house. The software Spike2, version 6.07, was used for data acquisition and subsequent annotation.

Manual sleep-stage classification. Our expert (SW) manually scored sleep by visual inspection of EMG and EEG time series. The parietal right EEG was chosen for scoring decisions because it had the best signal-to-noise ratio across all mice. Recordings were divided into non-overlapping consecutive windows (segments) of 10 seconds. Three consecutive segments of the EMG and the parietal right EEG signal were visually presented to the expert who scored the middle segment. The expert was not required to score the recordings in sequential order but could jump to other parts of the recordings at all times. Epochs containing features of multiple sleep stages were assigned the stage that made up the majority of the epoch.

Segments were scored as stage “Wake”, “NREM”, “pre-REM”, or “REM” sleep (see Fig. 1 and Table 2). Artifact contaminated segments were scored as “artifact”. The Wake state is characterized by elevated activity in the EMG signals when compared with NREM, pre-REM, and REM sleep. NREM sleep is characterized by slow waves between 0.5–4 Hz (delta band) while REM sleep is associated with rhythmic activity between 7–8 Hz (theta band). The pre-REM stage is characterized by successive short bouts (< 10 s) of high amplitudes similar to spindles and the appearance of a theta rhythm^{13,17}. Wake was the most frequent stage in our dataset (55.08 %, see Table 2), followed by NREM (37.7 %) and REM sleep (5.08 %). The intermediate stage pre-REM occurred only rarely (1.88 %) and only 0.26 % of all segments were contaminated with artifacts.

Dataset preparation for machine learning. *Train-, validation-, and test-set splits.* The scored time series were prepared for training and evaluation of the neural network for automated sleep scoring as follows: To mimic the manual scoring process and make our automated scoring system more flexible with respect to different experimental setups (in which only one EEG derivation may be available), we chose to base the training of the neural network on the parietal right EEG only. The parietal right EEG time series were low-pass filtered using a 4-th order Butterworth low-pass filter (critical frequency: 25.6 Hz) in order to prevent aliasing effects due to subsequent downsampling. The low-pass filter was applied to each time series with a forward- and backward pass to avoid any phase shifts. Subsequently, the low-pass filtered time series were downsampled to 64 Hz by linear interpolation.

The recordings were assigned to a training, a validation, and a test set (see Tables 1 and 2). Subjects were distributed such that each set consisted of recordings from mice of both dietary groups. Validation and test set each consisted of recordings from two mice one of which was from the altered-diet group. The mouse that passed

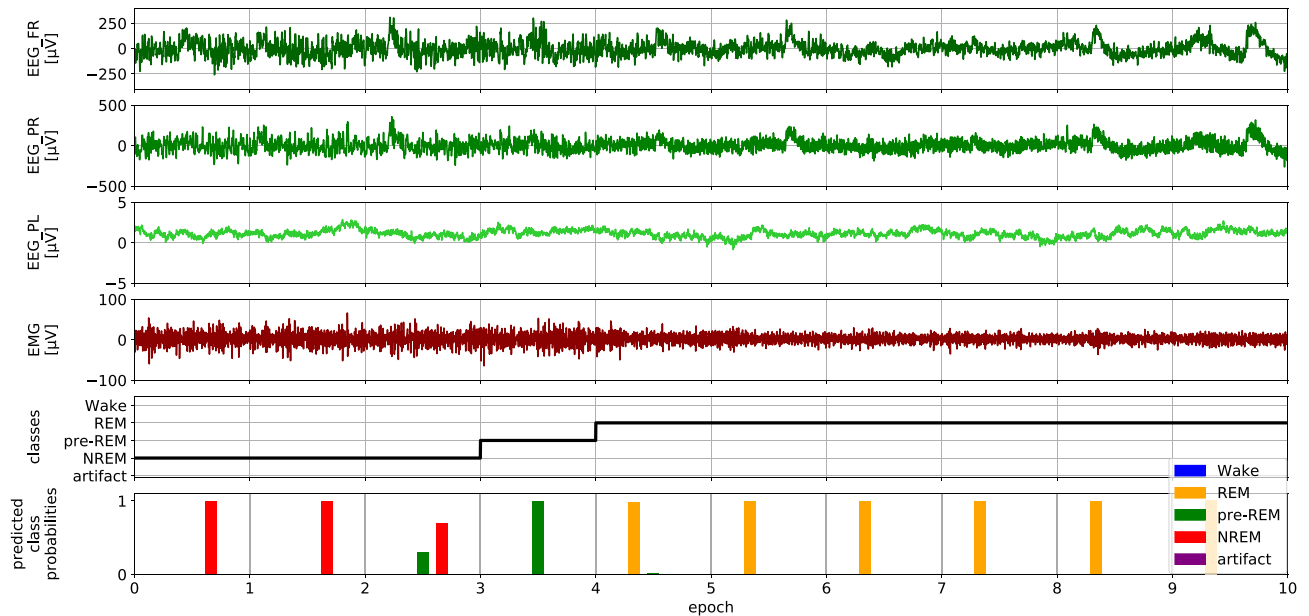


Figure 1. Example EEG time series (panels 1–3: frontal, parietal right, parietal left) and EMG time series (panel 4) of a mouse measured by chronically implanted electrodes (data from the test set). Sleep scores are shown in panel 5, whereas the last panel shows score probabilities predicted by the network. Time series were split into segments of 10 seconds each (marked by vertical lines; see text). For each segment, the sleep score was determined by an expert, and score probabilities were predicted by our network.

away early was assigned to the test set so that later evaluation would allow us to test the trained network on the most novel subject in our dataset.

The network parameters were trained on the training set, whereas the validation set allowed us to optimize hyperparameters of our network, of the training process and of regularization as well as data augmentation strategies. Finally, the test set was used to assess the performance of our network to predict sleep stages out-of-sample (i.e., on data from mice that were not available during training nor validation). We stress the importance of out-of-sample testing since it mimics the situation in sleep laboratories where a trained network predicts sleep stages on data from mice that are unknown to the network. To simulate this situation, we made sure that all data from each mouse was assigned to one and only one of the three sets (see Table 1).

We note that, when the network was trained to predict only the standard set of target stages (Wake, NREM, REM), the training, validation, and test set were created by reinterpreting pre-REM stages as NREM, while segments scored as artifact were removed.

Class rebalancing. The sleep stages identified by the expert occur with widely differing frequencies (e.g., Wake: 55.08 % vs pre-REM sleep: 1.88 %, see first row of Table 2). Such differing frequencies of classes (*class imbalance*) can pose non-trivial challenges when creating data-driven systems due to the tendency of systems to over-classify majority classes and misclassify infrequent classes^{40,41}. To address this issue, we created rebalanced training sets that differed in their class frequencies (see rows 3 and 4 of Table 2) while maintaining the total number of segments in the set. The rebalanced training sets were created before each training epoch by random sampling with replacement from the respective classes of the original training set so that predefined class frequencies were reached. In the rebalanced training set #1, class frequencies were heuristically chosen such that infrequent classes were oversampled and frequent classes undersampled while class frequencies still reflected some of the class imbalance of the original training set. The rebalanced training set #2 was created to investigate the performance of our network to distinguish between the three main stages (Wake, NREM, REM) only.

Data augmentation. Creating additional synthetic training data is known as *data augmentation* and has repeatedly been demonstrated to reduce overfitting and to tackle class imbalance^{42–44}. Moreover, if prior knowledge about the data cannot be easily incorporated into the design of a network architecture, the knowledge can often be utilized to create synthetic training data. This way prior knowledge can indirectly be induced into the network during training. A prime example is the creation of rotation invariance in image recognition systems by supplying rotated images (derived from the original data) during the training process³⁰.

We expect variability in our data that does not reflect sleep dynamics but can be introduced by the data acquisition, the scoring process, or unrelated physiological processes. Our data augmentation strategies are aimed at inducing invariance in our network with respect to this variability.

(i) **Amplitude variability.** Different acquisition systems can use different signal amplification settings, thereby leading to a variability in the scaling of signal amplitudes. Let \bar{s} denote the part of the time series that is presented to the input layer of the neural network during training. Let s_t denote the sample of this time series at time step

$t \in [1, T]$, where T denotes the length of the time series. During training, instead of \vec{s} we presented the network with the augmented time series

$$\vec{s}^* = \vec{s}(1 + a_a \cdot u) \quad u \in U[-1, 1), \quad (1)$$

where u is a random variable drawn from the continuous uniform distribution in the interval $[-1, 1)$, and where $a_a \in [0, 1]$ is the strength of this type of augmentation. u is newly drawn for each time series entering the network.

(ii) Frequency variability. We expect that the frequency content of EEG signals varies to a certain extent independently of the actual sleep stage, due to different noise sources (e.g., caused by movements of mice or by the measurement device itself). We mimicked such a variability in frequency content by using *window warping*⁴⁵ as an augmentation technique in which a time series is stretched or contracted. Let T denote the length of the time series that enters the network. To create an augmented time series, in a first step, we determined a temporary length

$$T^* = \lfloor T(1 + a_f \cdot u) \rfloor \quad u \in U[-1, 1) \quad (2)$$

of a temporary time series, where $a_f \in [0, 1]$ denotes the strength of the stretching or contracting. The central sample point of the temporary time series is identical to the one of the original time series. In a second step, the temporary time series was resampled by linear interpolation such that its new length equals T . If $T^* > T$, the temporary time series includes sample points prior to and after the original time series. In this case, the resampling operation contracts the time domain and thus frequencies are shifted towards higher values. If $T^* < T$, the resampling operation stretches the time domain and frequencies are shifted towards lower values.

(iii) EEG montage variability (sign flip). Since each EEG signal represents a difference between the electrical potentials at two (or more) measurement sites, the sign of the amplitudes of an EEG signal depends on the chosen *montage* (i.e., the configuration of electrical potential differences between electrodes that is used to represent EEG signals). A change in montages can lead to a flipped sign of the amplitudes of the signal which, however, still reflects the same brain dynamics. The augmented time series \vec{s}^* reads

$$\vec{s}^* = \begin{cases} -\vec{s} & \text{if } u < a_s \\ \vec{s} & \text{else} \end{cases} \quad u \in U[0, 1), \quad (3)$$

where $a_s \in [0, 1]$ is the probability of flipping the sign of the time series amplitudes.

(iv) Time shift variability. Sleep scoring associates a sleep stage to each of the consecutive EEG segments. Thus, the segment length introduces an artificial timescale on which sleep is characterized. Actual sleep stages do not necessarily change at segment boundaries but can change within a segment. To account for this variability, we created augmented data by shifting the time series segments with respect to the sleep scores by a random time shift Δt ,

$$\Delta t = \lfloor T \cdot a_t \cdot u \rfloor \quad u \in U[-1, 1), \quad (4)$$

where $a_t \in [0, 1]$ controls the extent of the shift and thus the strength of this type of augmentation. With $\vec{s} = (s_1, \dots, s_T)^T$ denoting the original time series, the augmented time series reads

$$\vec{s}^* = (s_{1+\Delta t}, \dots, s_{T+\Delta t})^T. \quad (5)$$

Time series were augmented in the order (i), (iii), (ii), (iv) as described above using newly drawn random variables every time the time series entered the neural network for forward propagation. The strength of the augmentations were controlled by the parameters (a_a, a_f, a_s, a_t) .

Machine learning. *Network architecture.* Mimicking the manual scoring process, the input of the neural network consists of three consecutive segments of the parietal right EEG time series, while the outputs of the network are the probabilities of the middle segment to belong to Wake, NREM, pre-REM, REM sleep, or to contain artifacts (class artifact). We used the score with the highest probability as the predicted score by our network.

The network architecture, which draws upon experience gained in previous work by some of the authors⁴⁶, consists of a feature extractor and a classifier (see Fig. 2). Before entering the feature extractor, the time series is batch normalized (“Batchnorm”, see Fig. 2). The feature extractor consists of 8 convolutional layers, each of which is composed of 96 kernels of size $d \times 1 \times 5$, where d denotes the depth of the output volume of the previous layer. Since one time series enters the network, $d = 1$ for the first convolutional layer, while $d = 96$ for all other convolutional layers. To downsample the time series information along the feature extractor, every other convolutional layer uses a stride of 2 with the other layers having a stride of 1. The resulting convolved signals of each layer are nonlinearly transformed by rectified linear units (ReLU) and subsequently batch normalized (Batchnorm). Furthermore, we apply Dropout regularization⁴⁷ to every other convolutional layer. Finally, the output (also called features) of the feature extractor is concatenated into a vector (“Flatten”, see Fig. 2), and Dropout regularization is applied before these features enter the classifier.

The classifier is composed of 2 fully connected (FC) layers with $113 \cdot 96 = 10848$ and 80 neurons, respectively. The first FC layer is equipped with rectified linear units (ReLU) as nonlinearity and Dropout regularization, while the second FC layer uses a softmax activation function to determine the score probabilities \vec{p} which represent the activations of the neurons of the output layer. The output layer consisted of 3 output neurons when the network was trained to distinguish between major sleep stages (Wake, REM, NREM). When the network was trained to also score pre-REM sleep and segments containing artifacts, the output layer consisted of 5 neurons.

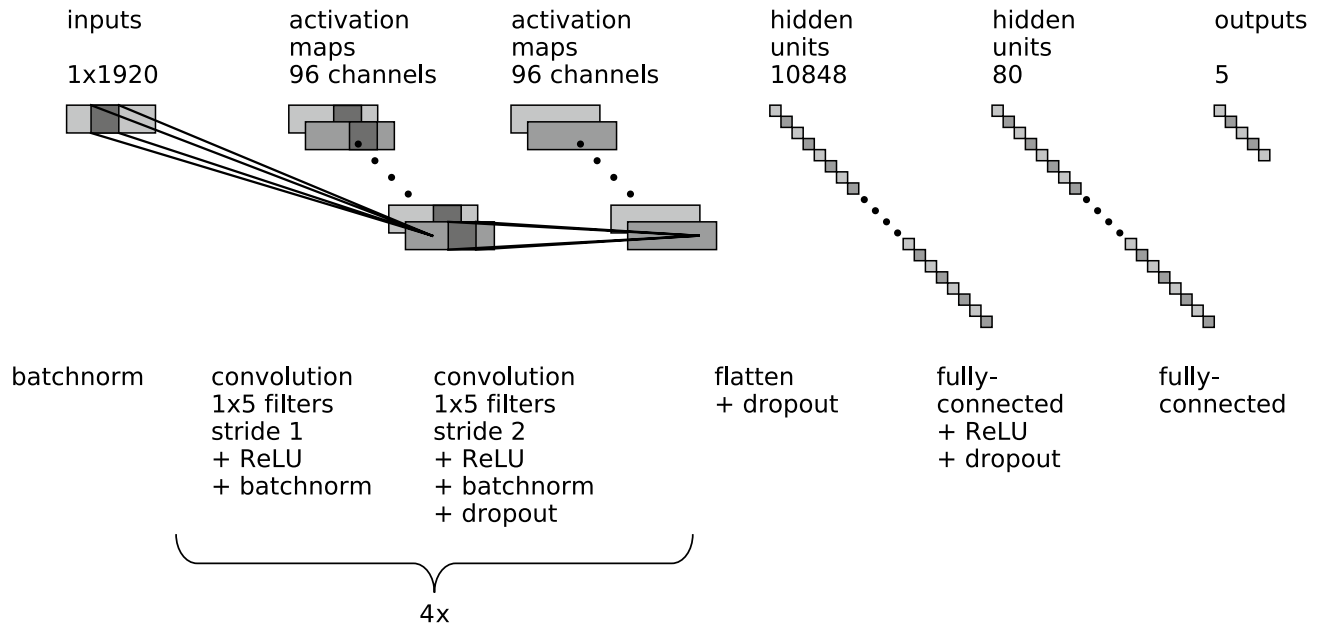


Figure 2. Network architecture. The *feature extractor* consists of 8 convolutional layers, and the output of the last convolutional layer is flattened (Flatten). EEG time series data is batch normalized (Batchnorm) before the first convolutional layer is applied. The *classifier* consists of two fully connected layers. A softmax layer translates the activations of the previous layer into class probabilities.

Training. Loss function. The training objective is implemented by the loss function \mathcal{L} that becomes small when the classifier assigns a large probability to the correct class. Let \vec{x}_i denote a time series segment i and let k be the index of the middle segment of the input time series for which the network predicts the score probability vector \vec{p}_k . With $c_k \in \{1, \dots, 5\}$ we encoded the correct score for segment k as determined by our human expert. Thus, p_{k,c_k} is the predicted probability of the middle time series segment to belong to the sleep stage that was determined by our human expert. The loss is the negative log-likelihood function with an L2 regularization penalty term L_2 ,

$$\mathcal{L} = -\frac{1}{N_b} \sum_{k=1}^{N_b} \log(p_{k,c_k}) + \underbrace{\frac{\lambda}{2N_b} \sum_w w^2}_{L2 \text{ penalty}}, \quad (6)$$

where N_b is the size of the minibatch, $\lambda \geq 0$ is the L2 regularization strength, and where the sum of the L2 penalty is over all the weights w in the network (except the bias weights).

Gradient descent. Since the neural network is differentiable, the loss function can be minimized by minibatch gradient descent. The network was trained by Adam⁴⁸, a variant of stochastic gradient descent which uses exponential moving averaging of the gradient (equation (8)) and the squared gradient (equation (9)) to allow for adaptive learning rates. Let w_t denote a weight (i.e., a free parameter) of the network and let $g_t = \frac{\partial \mathcal{L}}{\partial w}$ denote the gradient of the loss function with respect to the weight obtained for the minibatch at step t of the training. The weight is updated by

$$w_{t+1} = w_t - \eta \cdot \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon} \quad (7)$$

where η denotes the learning rate, and \hat{m}_t and \hat{v}_t are defined as

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t} \text{ with } m_t = (1 - \beta_1)g_t + \beta_1 m_{t-1} \quad (8)$$

and

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t} \text{ with } v_t = (1 - \beta_2)g_t^2 + \beta_2 v_{t-1}, \quad (9)$$

respectively. We used standard values for the exponential decay rates $\beta_1 = 0.9$ and $\beta_2 = 0.999$ and set $\epsilon = 10^{-8}$ to prevent division by zero in equation (7). Furthermore, to counteract the exploding gradient problem, we employed a common gradient clipping strategy⁴⁹ that ensured by rescaling that the norm of each gradient vector does not exceed a threshold $\theta = 0.1$.

Learning rate protocol. To speed up the learning process, we used large minibatches ($N_b = 256$) and adjusted the learning rate η according to the following protocol: The learning rate was linearly increased⁵⁰ from $10^{-7}N_b$ to $10^{-6}N_b$ over the course of the first 12 *training epochs* (warm-up period). This type of linear scaling of the learning rate with minibatch size in combination with a warm-up period was observed to facilitate the use of large minibatch sizes, thereby shortening training times⁵⁰. A single training epoch was completed when all minibatches of the training set were used once for the gradient updates. The warm-up period was followed by a cool-down period during which the learning rate was exponentially decreased, $\eta_i = 10^{-6}N_b \cdot e^{-\alpha(i-12)}$, where $i > 12$ denotes the training epoch, and where the exponential rate was set to $\alpha = 0.06$. The cool-down period lasted for the remaining training process.

Regularization. We used four mechanisms to prevent our network from overfitting the training data. (i) *L2 regularization*, also known as *weight decay* (see second term in equation (6)), favors neural networks whose weights (parameters) take on low values, thereby effectively restricting the model space of networks where the loss function obtains low values⁵¹. We set the L2 regularization strength to $\lambda = 10^{-4}$. (ii) We used *Dropout regularization* with a dropout probability of $p_{dropout} = 0.2$ in the feature extractor as well as in the classifier of the network (see Fig. 2). Dropout regularization has been reported to prevent complex co-adaptation of neurons and thus drive them to create features on their own⁴⁷. (iii) Limiting the search in parameter space during training (called *early stopping*) has a regularizing effect since it prevents the network from overfitting the training data⁵¹. After each training epoch, the network is evaluated on the validation set by calculating the F1 score (see Sect. "Evaluation measures"). If the F1 score does not improve over 5 consecutive training epochs, the training is stopped. The training is carried out for a minimum of 12 and a maximum of 50 epochs. At the end of the training, the network parameters are returned at that point of the training for which it achieved the best F1 score on the validation set. (iv) Using prior knowledge to create additional training instances is called *data augmentation*. Data augmentation is known to decrease the generalization error of the trained network and thus can be interpreted as a regularizing mechanism. We detail this approach in the section on data augmentation.

We note that we chose regularization parameters for all mechanisms after extensive hyperparameter exploration on the validation set.

Evaluation measures. To investigate and assess the accuracy of our classifiers to score sleep, we used the following evaluation measures.

Confusion matrix. We use confusion matrices to summarize the predictive performance of our networks. The matrix elements describe the number of one of the actual labels (vertical dimension) that have been assigned one of the target labels (horizontal dimension) by the network. Both the actual label frequency (vertical axis) and the predicted label frequency (horizontal axis) can be used to normalize the matrix elements (see Fig. 4 for an example)⁵². When normalized by the actual label frequency, matrix elements show the percentage of segments of a true class being labeled by the network as a predicted class. Each diagonal element corresponds to the network's recall of the respective class. Likewise, when normalized by the predicted label frequency, matrix elements show the percentage of all segments of a class predicted by a network to actually belong to the true class. Each diagonal element corresponds to the network's precision of the respective class.

F1 score. In a binary classification setting (e.g., Wake vs Non Wake; pre-REM vs Non pre-REM) the F1 score, also called F-measure, is defined as the harmonic mean of precision ($p_{precision}$) and recall (p_{recall})⁵³. Estimates of these values ($\hat{p}_{precision}$, \hat{p}_{recall}) are extracted from the normalized confusion matrices:

$$F1 = 2 \frac{\hat{p}_{precision} \cdot \hat{p}_{recall}}{\hat{p}_{precision} + \hat{p}_{recall}} = 2 \left(\frac{1}{\hat{p}_{precision}} + \frac{1}{\hat{p}_{recall}} \right)^{-1} \quad (10)$$

The score takes on values between 0 and 1, with high values indicating better classification performances. Although this metric is not invariant to changes in the data distribution, it is commonly used to compare different algorithms against a consistent dataset⁵³. Note that rebalancing of the training set affects the training score in unpredictable ways; however, we compare our networks based on the F1 scores calculated on the validation set which is not rebalanced. We expanded the definition given in equation (10) for our multiple class classification problem by calculating $\overline{F1}$, the average of the F1 scores of each sleep stage.

Markov transition matrix. Markov transition matrices summarize the empirical transition probabilities of a sequence (e.g., a sequence of sleep stages scored by a human expert or a model) interpreted as a Markov process. They are two-dimensional matrices whose elements describe the number of transitions from one label to another in either absolute numbers or, after normalization, percentages (see Fig. 5 for an example). We created Markov transition matrices for the actual labels and the labels predicted by our networks to measure how well the networks perceived the structure of the sleep architecture.

Results

We investigated how accurate our network predicts sleep stages under various training configurations. We extensively modified the network capacity, the training schedule, class rebalancing weights, and data augmentation. Our experiments were conducted on two sets of target stages; a standard set containing Wake, NREM, and REM stages, and an extended set also including the stages pre-REM and artifact. Our approach was to optimize hyperparameters on the extended set of sleep stages. Afterwards, we translated these parameters to the standard set to investigate if the predictions reproduced the known state-of-the-art performance.

In preliminary experiments, we also considered methods of cross-validation in order to reduce the variance of predictions on the validation set. We discovered, however, that the reduction in variance did not yield additional information that would influence our parameter search. We therefore abandoned this compute-intensive

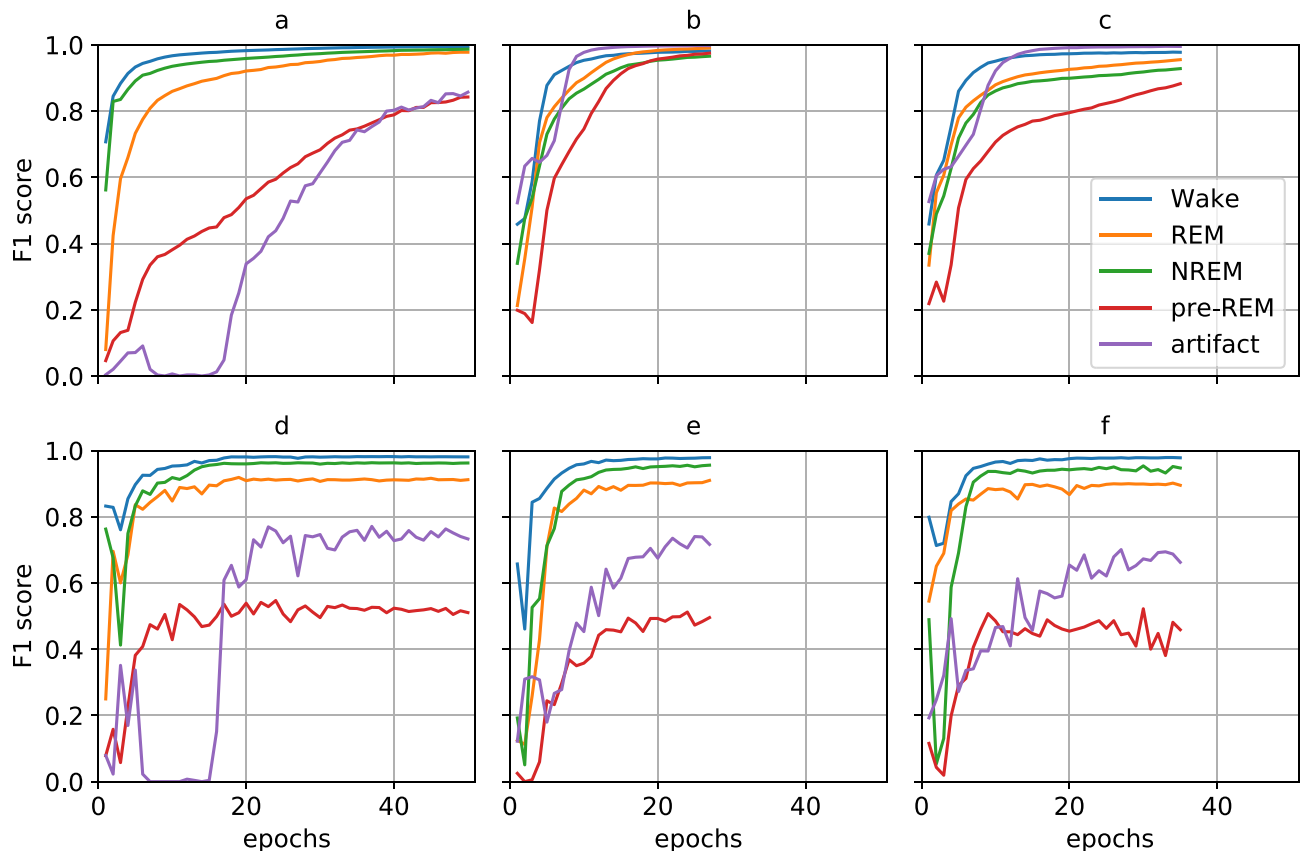


Figure 3. Evolution of F1 scores during training determined on the training set (top row) and validation set (bottom row). The evolution of F1 scores changed depending on whether the network was trained on unbalanced training data (panels **a** and **d**), rebalanced training data #1 (panels **b** and **e**), or rebalanced training data #1 with time shift augmentation (panels **c** and **f**).

approach in favor of exploring a broader set of parameters. Note, that we did not cross-validate across samples in the test set as this would have invalidated the generalization of our final results.

Influence of class rebalancing. To analyze class rebalancing, we compared the epoch-wise improvements in F1 score on the unbalanced and rebalanced dataset (see Fig. 3 panels (a,d) and (b,e) respectively). When trained on the unbalanced training set, F1 scores increased rapidly for the majority classes during the first 10 training epochs, while the F1 scores of the pre-REM and artifact classes only slowly increased. A similar behavior was observed for the F1 scores on the validation set.

Meanwhile, the F1 scores of the pre-REM and artifact classes obtained when training on the rebalanced training set increased more quickly and reached higher optimal values. The optimal F1 scores on the validation set were comparable to the F1 scores obtained on the validation set by the network trained on the unbalanced training set.

While the optimal F1 scores on the validation set were comparable between trainings on the rebalanced and unbalanced training sets, we observed that training on the unbalanced dataset led to *unstable* results: minority-class F1 scores were sometimes not recognized in lapses of several epochs. Such an event is visible for the F1 score of the artifact class in Fig. 3(a, d). Because of this instability, we continued to optimize the training configuration using only rebalanced training sets.

Classification on the extended set of sleep stages. We studied the classification performance under different training configurations when targeting pre-REM and artifact in addition to the standard classes Wake, Non-REM, and REM (see Table 3). Without regularization and augmentation, we observed a mean F1 score of $\bar{F1} = 0.99$ on the training set, and $\bar{F1} = 0.76$ on the validation set. We optimized this F1 score by changing the dropout probability $p_{dropout} \in [0.1, \mathbf{0.2}, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8]$ and L2 regularization strength $\lambda \in [10^{-1}, 10^{-2}, 10^{-3}, \mathbf{10^{-4}}]$ and achieved an improved mean F1 score of $\bar{F1} = 0.81$ on the validation set (optimal values of hyperparameters indicated in bold). To increase the variability of the training data and reduce the possibility of overfitting the minority classes, we continued by applying four data augmentation strategies (see Sect. [Data augmentation](#)) with parameters in the following ranges:

- Amplitude augmentation: $a_a \in [0.5, 0.75, 1.0]$

training configuration	dataset	F1 score	F1 score	F1 score
		(mean)	pre-REM	artifact
no regularization, no data augmentation	train	0.99	0.99	1.00
	valid	0.76	0.44	0.56
best regularization, no data augmentation, <i>Best Network</i>	train	0.98	0.96	1.00
	valid	0.81	0.50	0.74
	test	0.78	0.48	0.59
amplitude ($a_a = 0.5$)	train	0.98	0.96	1.00
	valid	0.81	0.49	0.74
frequency ($a_f = 0.05$)	train	0.97	0.95	1.00
	valid	0.81	0.49	0.72
sign flip ($a_s = 0.5$)	train	0.97	0.94	1.00
	valid	0.80	0.49	0.69
time shift ($a_t = 0.06$)	train	0.94	0.86	1.00
	valid	0.81	0.52	0.67
all data augmentations combined	train	0.92	0.79	0.99
	valid	0.79	0.50	0.64

Table 3. F1 scores for the training, validation, and test set at various training configurations. First, the network was trained without L2 regularization and without dropout (row 1). Next, the network was trained with both regularizations (row 2). Last, the network was trained using different types of data augmentation (rows 3–6; the parameters in column 1 resulted in the best validation F1 scores among the explored parameters), or with all data augmentations combined (row 7). F1 scores averaged over all classes (column 3) characterize how well the networks can predict sleep stages on all five classes. F1 scores in columns 4 and 5 present the capabilities of the networks to predict the two minority classes pre-REM and artifact, respectively. Only the network with the best validation F1 scores (best network) was evaluated on the test set.

- Frequency augmentation: $a_f \in [0.05, 0.1, 0.2, 0.5]$
- Sign-flip augmentation: $a_s \in [0.25, 0.5, 0.75]$
- Time-shift augmentation: $a_t \in [0.01, 0.02, 0.03, 0.06, 0.1]$

Results of the trainings with data augmentation are reported in Table 3. We observed a slight decrease of mean F1 scores on the training set with augmentation (e.g., $\overline{F1} = 0.92$ for all augmentations combined) in comparison to training without augmentation ($\overline{F1} = 0.98$). We interpret this decrease to reflect the regularizing effect of data augmentation that can generally be observed in a network with fixed model capacity when fitting data of increasing variability. Interestingly, this decrease was most pronounced for the pre-REM class (see Fig. 3(c,f)), in contrast to the artifact class that could be fitted perfectly by the network on the training set. Mean F1 scores on the validation set show that none of the data augmentation algorithms could improve on the F1 score achieved by the network trained only with the best regularization parameters. Consequently, we evaluated this network on the test set where we found a mean F1 score of $\overline{F1} = 0.78$.

Comparing the regularized with the un-regularized results indicates that regularization induced increases of the F1 scores in the minority classes. None of the augmentation techniques were able to improve the F1 scores of the individual classes, except time-shift augmentation, which increased the validation F1 score of pre-REM from 0.50 to 0.52.

The test set consisted of one mouse each from both dietary groups. We explored their individual F1 scores and did not find any deviations exceeding 2%. In particular, the pre-REM F1 score for the typical mouse was 0.48; for the altered-diet mouse we found 0.48.

Reduction of the network to the standard stages. Next, we studied the classification performance of our optimized network architecture in distinguishing between the three standard sleep stages Wake, REM, and NREM (see Table 4). With neither regularization nor augmentation, we again observed our network to accurately approximate the training data; on the validation data, we found a mean F1 score of $\overline{F1} = 0.94$. Using the best regularization from the optimization on five stages – $p_{dropout} = 0.2$, and $\lambda = 10^{-4}$ – increased the mean F1 score on the validation set to $\overline{F1} = 0.96$. We accepted this network and found on the test set, reduced to 3 stages, a mean F1 score of $\overline{F1} = 0.95$. A detailed analysis revealed that improvements were made in the REM class, the most under-represented class in the reduced datasets.

These classification performances are comparable with state-of-the-art results obtained in other studies^{32,33,36} and demonstrate that the proposed network architecture is suitable for sleep scoring purposes.

Detailed analysis of the best network. Confusion matrices (see Sect. Evaluation measures) obtained for the network trained without data augmentation (see Fig. 4) show that more than 98% of the Wake segments,

training configuration	dataset	F1 score	F1 score	F1 score	F1 score
		(mean)	Wake	REM	NREM
no regularization, no data augmentation	train	0.99	0.99	1.00	0.99
	valid	0.94	0.98	0.90	0.96
best regularization, no data augmentation, <i>Best Reduced Network</i>	train	0.99	0.99	1.00	0.99
	valid	0.96	0.98	0.92	0.97
	test	0.95	0.98	0.92	0.96

Table 4. F1 scores for the training, validation, and test set for predicting the three main stages Wake, REM, and NREM sleep at various configurations of the hyperparameters. The network was trained on the rebalanced training set #2 (see Table 2). Only the best configuration (best reduced network) was evaluated on the test set.

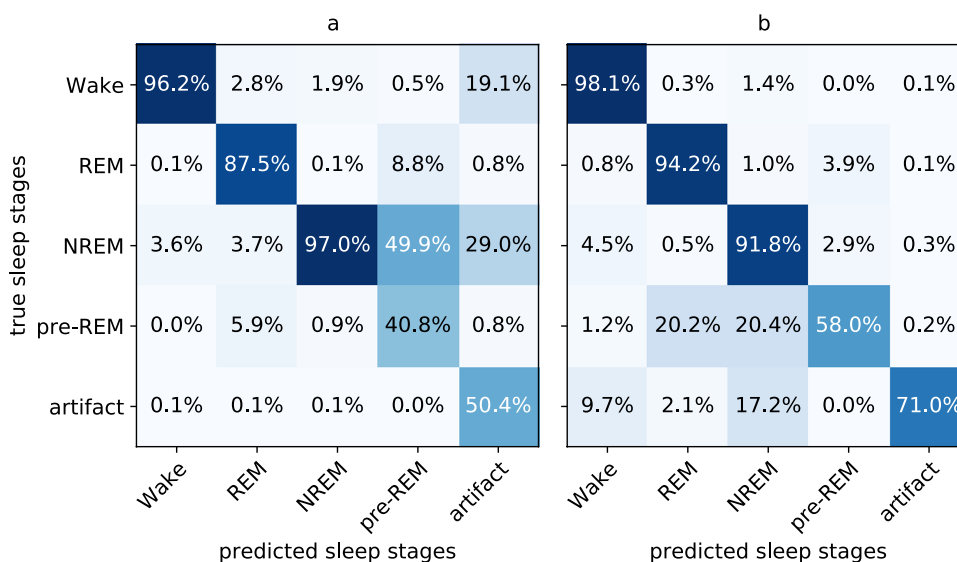


Figure 4. Confusion matrices of the best network evaluated on the test set (see Table 3): **(a)** Confusion matrix averaged by the predicted label frequency; diagonal elements correspond to precision. **(b)** Confusion matrix averaged by true label frequency; diagonal elements correspond to recall. A prediction of pre-REM is about 50% skewed towards NREM being the true stage, whereas a true pre-REM stage has equal probability to be mistaken for REM and NREM by the network.

more than 94% of the REM segments, and close to 92% of the NREM segments in the test set were classified correctly (panel b). The network was also able to score 71% of the artifact segments correctly. The pre-REM stage was scored in 58% of the segments correctly but was confused in about 20% of cases with REM or in 20% of cases with NREM sleep.

The confusion matrices also show that more than 96% of the segments in the test set predicted as Wake, more than 87% of the segments predicted as REM, and 97% of the segments predicted as NREM had been assigned the predicted stage by the human expert (panel a). Furthermore, more than 50% of the segments that were predicted as artifact by our network had been assigned artifact as the true stage. Predictions of pre-REM had pre-REM as the true stage in nearly 41% of cases, while in about 50% of cases, predictions had NREM and about 9% had REM as true stage.

To investigate whether the network was able to capture basic properties of sleep dynamics, we determined Markov transition matrices (see Sect. [Evaluation measures](#)) on the test set based on scores provided by our expert (Fig. 5a) and on scores as predicted by our network (Fig. 5b). The probabilities of transitions between the different stages obtained from the network indeed reflect typical sleep cycles, where Wake is followed by NREM sleep (5.4%), which is followed by either Wake (4.5%), pre-REM (4.1%), or more NREM sleep (91.1%). We highlight that there was only a small transition probability from NREM directly to the REM stage as determined by our network (0.3%), which was close to the transition probability obtained by our expert (0.1%). NREM transitioned via pre-REM to REM stages before the cycle closed with the transition to either Wake (17.7%) or NREM (4.2%). While transition probabilities of our network closely follow those as obtained from the expert, we find notable differences for transitions into and out of the pre-REM stage. pre-REM stages as determined by our network tended to last longer (pre-REM to pre-REM transition probability of 31.6%) than those determined by the scoring expert (17.9%).

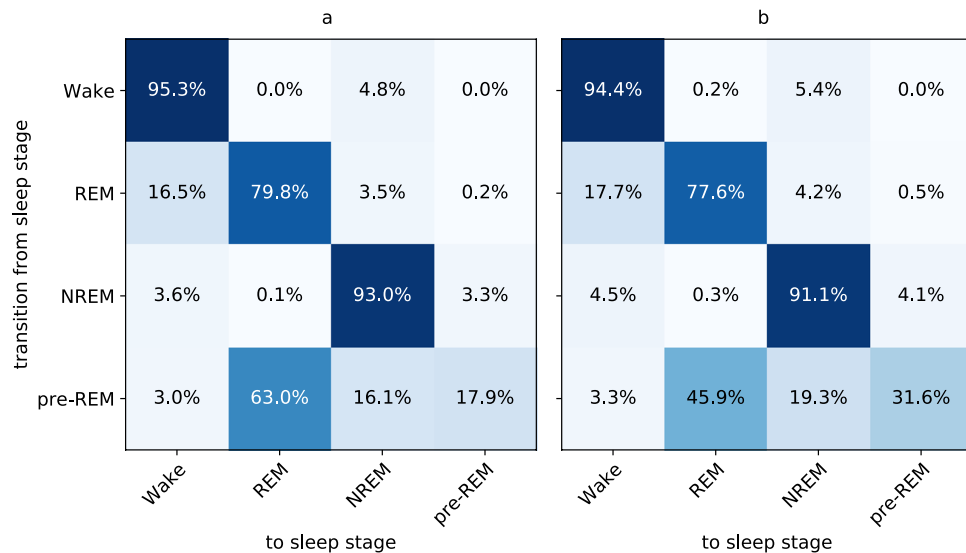


Figure 5. Transition probabilities from one stage (row) to another stage (column) obtained from the validation set with scores provided by our scoring expert (SW) (a) and obtained with scores predicted by the best network (b). The network parameters were trained on the rebalanced training set without augmentation.

Discussion

We proposed a neural network architecture that is able to distinguish between the three main stages Wake, REM and NREM as well as the infrequent stages pre-REM and artifact. When trained on these five stages, we observed our network to classify the majority sleep stages (Wake, REM, NREM) with high precision and recall (see Fig. 4), while the infrequent class pre-REM obtained lower precision and recall and a lower F1 score ($F1 = 0.48$, see Table 3) on out-of-sample data. Transition probabilities between sleep stages predicted by the network were largely consistent with those determined by the human expert. However, pre-REM phases predicted by the network tended to last longer than those by the expert. When the network was trained to distinguish between the three main stages only, classification performance as quantified by the average F1 score on out-of-sample data ($\bar{F1} = 0.95$, see Table 4) was comparable to state-of-the-art classifiers^{32,33}. L2 and Dropout regularization as well as data augmentation increased F1 scores of the minority classes (pre-REM, artifact), while class rebalancing did not. We observed, however, the latter to stabilize the training process.

The F1 score of pre-REM indicates that pre-REM as a transitional sleep stage is particularly challenging to score. This challenge is not confined to scoring sleep in mice but can also be encountered in other species such as rats or humans where transitional sleep stages can also be observed. In rats, systems for sleep scoring obtained consistently lower scoring accuracies for the transition sleep stage compared to the main sleep stages²⁰ with F1 scores ranging from 0.3 to 0.7^{19,21}. In humans, automated scoring systems achieved F1 scores between 0.1 to 0.6 for N1 sleep^{34,55}, a transitional stage between wakefulness and sleep. In mice, we cannot comparatively assess the scoring performance of our network for pre-REM since, to the best of our knowledge, no prior approaches exist at the time of this writing. We note, however, that the obtained F1 score for pre-REM in mice ($F1 = 0.48$) is in the middle of the aforementioned ranges for F1 scores of transitional stages.

We hypothesize that pre-REM sleep is difficult to delineate from NREM and REM sleep for experts and automated systems alike. Indeed, Fig. 1 (bottom row) shows a transition from NREM over pre-REM to REM sleep, where the predicted class probability of pre-REM already increases in the third segment that was labeled as NREM sleep. We observed many transitions like this where the class probabilities of NREM, pre-REM, and REM slowly decreased or increased along many segments. If human scorers had more difficulties in identifying pre-REM compared to other sleep stages, we would expect an increased intra-rater variance (i.e., a lower probability of labeling the same segments as pre-REM when the dataset would be labeled several times by the same person). Moreover, we would also expect larger inter-rater variance for pre-REM than for other stages (i.e., a lower agreement between different human scorers on scoring pre-REM) as has been observed for the transition sleep stage in rats²⁰. Such uncertainties in sleep scores (also called *label noise*) can limit the classification performance achievable by automated scoring systems trained on that data. Besides label noise, another hypothesis that may explain lower classification accuracies for pre-REM is class imbalance⁴⁰. However, since we did not observe increased F1 scores when rebalancing classes (see Figs. 3d and e), we reject this hypothesis for the learning problem and dataset studied here.

One of the limitations of this study is the lack of separate annotations by several experts from which a group consensus for each sleep stage could have been derived. Such annotations would have also allowed to assess the extent of inter-rater variance for each class, thereby allowing to judge the network's performance in classifying pre-REM sleep compared to human scorers. We note, however, that the scoring performance of the network when restricted to Wake, REM, and NREM is comparable to values reported for inter-rater agreement in the literature³³. Another limitation is that we could not investigate whether achievable F1 scores change due to age,

behavior, disorders, dietary constraints, or genetic variations of mice since the dataset was too small. We expect some variability in F1 scores as indicated by previous studies (e.g., for genetic variations^{32,56}) and consider future studies into this direction as desirable.

We consider several strategies as promising to improve upon existing data-driven systems for rodent sleep scoring like the neural network presented in this paper. (i) Statistical approaches such as *confident learning*^{57–59} aim at identifying mislabeled training instances and have been demonstrated to successfully increase classification performance in tasks such as image recognition⁵⁷ by removing or correcting wrong labels, thereby reducing label noise. This approach might prove particularly fruitful for infrequent stages such as pre-REM sleep as we expect such stages to be more affected by label noise. (ii) Constraining the sequence of predicted sleep stages to transition probabilities as observed in the training data³² may improve network predictions for infrequent stages such as pre-REM. We note, however, that such an approach can introduce bias in the frequencies of sleep stages when data was obtained under different experimental conditions. The reliable detection of such changes in sleep stage frequencies is, however, the matter of inquiry in many sleep studies. (iii) When label noise can be reduced or when datasets scored by a group of experts (and thus with known inter-rater variabilities) become available, neural network architectures that could model long temporal successions^{60,61} of sleep stages might be particularly promising candidates to achieve new state-of-the-art classification performances for sleep scoring tasks. (iv) Finally, due to the lack of consensus in rodent sleep scoring⁸, we consider the creation of a public dataset annotated by a committee of experts to be particularly helpful to advance the creation of systems to automate scoring of animal sleep. Such a dataset should include data from different labs (to assess cross-lab variability of system predictions). Labels by a group of experts would establish a gold standard and would allow for an assessment of inter-rater variability, even for pre-REM sleep. Such approaches have been successfully pursued in other areas of sleep research, for instance for the challenge of sleep spindle detection⁶².

We are confident that deep-learning based systems, such as introduced here, will facilitate sleep studies with large amounts of data, including long-term studies and studies with thousands of subjects. The automation of the manual sleep scoring process, a cognitive repetitive task, will likely support the reproducibility of studies and allow researchers and lab personnel to focus on productive tasks.

Data availability

The datasets analyzed during the current study are available from the corresponding author on reasonable request.

Received: 9 December 2020; Accepted: 24 May 2021

Published online: 10 June 2021

References

- Bashan, A., Bartsch, R. P., Kantelhardt, J. W., Havlin, S. & Ivanov, P. C. Network physiology reveals relations between network topology and physiological function. *Nat. Commun.* **3**, 702. <https://doi.org/10.1038/ncomms1705> (2012).
- Rasch, B. & Born, J. About sleep's role in memory. *Physiol. Rev.* **93**, 681–766. <https://doi.org/10.1152/physrev.00032.2012> (2013).
- Drane, D. B., Martin, W. B. & Viglione, S. S. Pattern recognition applied to sleep state classification. *Electroen. Clin. Neuro.* **26**, 238 (1969).
- Smith, J., Negin, M. & Nevis, A. Automatic analysis of sleep electroencephalograms by hybrid computation. *IEEE T. Syst. Sci. Cyb.* **5**, 278–284. <https://doi.org/10.1109/TSSC.1969.300220> (1969).
- Aboalayon, K., Faezipour, M., Almuhammadi, W. & Moslehpour, S. Sleep stage classification using EEG signal analysis: a comprehensive survey and new investigation. *Entropy* **18**, 272. <https://doi.org/10.3390/e18090272> (2016).
- Faust, O., Razaghi, H., Barika, R., Ciaccio, E. J. & Acharya, U. R. A review of automated sleep stage scoring based on physiological signals for the new millennia. *Comput. Meth. Prog. Biol.* **176**, 81–91. <https://doi.org/10.1016/j.cmpb.2019.04.032> (2019).
- Fiorillo, L. *et al.* Automated sleep scoring: a review of the latest approaches. *Sleep Med. Rev.* **48**, 101204. <https://doi.org/10.1016/j.smrv.2019.07.007> (2019).
- Robert, C., Guilpin, C. & Limoge, A. Automated sleep staging systems in rats. *J. Neurosci. Meth.* **88**, 111–122. [https://doi.org/10.1016/S0165-0270\(99\)00027-8](https://doi.org/10.1016/S0165-0270(99)00027-8) (1999).
- Bastianini, S., Berteatti, C. & Gabrielli, A. Recent development in automatic scoring of rodent sleep. *Arch. Ital. Biol.* **153**, 58–66. <https://doi.org/10.12871/000398292015231> (2015).
- Katsageorgiou, V.-M. *et al.* A novel unsupervised analysis of electrophysiological signals reveals new sleep substages in mice. *PLOS Biol.* **16**, e2003663. <https://doi.org/10.1371/journal.pbio.2003663> (2018).
- Lacroix, M. M. *et al.* Improved sleep scoring in mice reveals human-like stages. *bioRxiv*. <https://doi.org/10.1101/489005> (2018).
- Ruigt, G., Proosdij, J. V. & Delft, A. V. A large scale, high resolution, automated system for rat sleep staging. i. methodology and technical aspects. *Electroen. Clin. Neuro.* **73**, 52–63. [https://doi.org/10.1016/0013-4694\(89\)90019-9](https://doi.org/10.1016/0013-4694(89)90019-9) (1989).
- Glin, L. *et al.* The intermediate stage of sleep in mice. *Physiol. Behav.* **50**, 951–953. [https://doi.org/10.1016/0031-9384\(91\)90420-s](https://doi.org/10.1016/0031-9384(91)90420-s) (1991).
- Gottesmann, C. Detection of seven sleep-waking stages in the rat. *Neurosci. Biobehav. R.* **16**, 31–38. [https://doi.org/10.1016/S0149-7634\(05\)80048-X](https://doi.org/10.1016/S0149-7634(05)80048-X) (1992).
- Mandile, P., Vescia, S., Montagnese, P., Romano, F. & Giuditta, A. Characterization of transition sleep episodes in baseline EEG recordings of adult rats. *Physiol. Behav.* **60**, 1435–1439. [https://doi.org/10.1016/S0031-9384\(96\)00301-0](https://doi.org/10.1016/S0031-9384(96)00301-0) (1996).
- Gottesmann, C. The transition from slow-wave sleep to paradoxical sleep: evolving facts and concepts of the neurophysiological processes underlying the intermediate stage of sleep. *Neurosci. Biobehav. R.* **20**, 367–387. [https://doi.org/10.1016/0149-7634\(95\)00055-0](https://doi.org/10.1016/0149-7634(95)00055-0) (1996).
- Amici, R., Jones, C. A., Perez, E. & Zamboni, G. *The Physiologic Nature of Sleep*, chap. A physiological view of REM sleep structure, 161–186 (Imperial College Press, 2005).
- Benington, J. H., Kodali, S. K. & Heller, H. C. Scoring transitions to REM sleep in rats based on the EEG phenomena of Pre-REM sleep: an improved analysis of sleep structure. *Sleep* **17**, 28–36. <https://doi.org/10.1093/sleep/17.1.28> (1994).
- Neckelmann, D., Olsen, Ø. E., Fagerland, S. & Ursin, R. The reliability and functional validity of visual and semiautomatic sleep/wake scoring in the Møll-Wistar rat. *Sleep* **17**, 120–131. <https://doi.org/10.1093/sleep/17.2.120> (1994).
- Gross, B. A. *et al.* Open-source logic-based automated sleep scoring software using electrophysiological recordings in rats. *J. Neurosci. Meth.* **184**, 10–18. <https://doi.org/10.1016/j.jneumeth.2009.07.009> (2009).

21. Wei, T.-Y. *et al.* Development of a rule-based automatic five-sleep-stage scoring method for rats. *Biomed. Eng. Online* **18**, 92. <https://doi.org/10.1186/s12938-019-0712-8> (2019).
22. Kohtoh, S. *et al.* Algorithm for sleep scoring in experimental animals based on fast fourier transform power spectrum analysis of the electroencephalogram. *Sleep Biol. Rhythms* **6**, 163–171. <https://doi.org/10.1111/j.1479-8425.2008.00355.x> (2008).
23. Gelder, R. N. V., Edgar, D. M. & Dement, W. C. Real-time automated sleep scoring: validation of a microcomputer-based system for mice. *Sleep* **14**, 48–55. <https://doi.org/10.1093/sleep/14.1.48> (1991).
24. Exarchos, I. *et al.* Supervised and unsupervised machine learning for automated scoring of sleep-wake and cataplexy in a mouse model of narcolepsy. *Sleep* **43**, zsz272. <https://doi.org/10.1093/sleep/zsz272> (2020).
25. Rytkönen, K. M., Zitting, J. & Porkka-Heiskanen, T. Automated sleep scoring in rats and mice using the naive Bayes classifier. *J. Neurosci. Meth.* **202**, 60–64. <https://doi.org/10.1016/j.jneumeth.2011.08.023> (2011).
26. Sunagawa, G. A., Séi, H., Shimba, S., Urade, Y. & Ueda, H. R. FASTER: an unsupervised fully automated sleep staging method for mice. *Genes Cells* **18**, 502–518. <https://doi.org/10.1111/gtc.12053> (2013).
27. Bastianini, S. *et al.* SCOPRISM: a new algorithm for automatic sleep scoring in mice. *J. Neurosci. Meth.* **235**, 277–284. <https://doi.org/10.1016/j.jneumeth.2014.07.018> (2014).
28. Caldart, C. S. *et al.* Sleep identification enabled by supervised training algorithms (SIESTA): an open-source platform for automatic sleep staging of rodent polysomnographic data. *bioRxiv*. <https://doi.org/10.1101/2020.07.06.186940> (2020).
29. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444. <https://doi.org/10.1038/nature14539> (2015).
30. Goodfellow, I. J., Bengio, Y. & Courville, A. C. *Deep Learning*. Adaptive computation and machine learning (MIT Press, 2016).
31. Roy, Y. *et al.* Deep learning-based electroencephalography analysis: a systematic review. *J. Neural Eng.* **16**, 051001. <https://doi.org/10.1088/1741-2552/ab260c> (2019).
32. Miladinović, D. *et al.* SPINDLE: end-to-end learning from EEG/EMG to extrapolate animal sleep scoring across experimental settings, labs and species. *Plos Comput. Biol.* **15**, e1006968. <https://doi.org/10.1371/journal.pcbi.1006968> (2019).
33. Barger, Z., Frye, C. G., Liu, D., Dan, Y. & Bouchard, K. E. Robust, automated sleep scoring by a compact neural network with distributional shift correction. *PLOS ONE* **14**, e0224642. <https://doi.org/10.1371/journal.pone.0224642> (2019).
34. Yamabe, M. *et al.* MC-SleepNet: large-scale sleep stage scoring in mice by deep neural networks. *Sci. Rep.* **9**, 15793. <https://doi.org/10.1038/s41598-019-51269-8> (2019).
35. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput.* **9**, 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735> (1997).
36. Svetnik, V., Wang, T.-C., Xu, Y., Hansen, B. J. & Fox, S. V. A deep learning approach for automated sleep-wake scoring in pre-clinical animal models. *J. Neurosci. Meth.* **337**, 108668. <https://doi.org/10.1016/j.jneumeth.2020.108668> (2020).
37. Grieger, N. Source code of the model presented in Grieger *et al.*, “Automating sleep scoring in mice with deep learning”. https://github.com/nik-grie/mice_tuebingen (2020).
38. Schwabedal, J. T. C. HypnoX web: a web implementation of deep learning models for classifying sleep stages in mice. <https://jusju.sjus.github.io/hypnox-web> (2020).
39. Paxinos, G. & Franklin, K. B. *Paxinos and Franklin's the Mouse Brain in Stereotaxic Coordinates* (Academic Press, London, 2019).
40. Haixiang, G. *et al.* Learning from class-imbalanced data: review of methods and applications. *Expert Syst. Appl.* **73**, 220–239. <https://doi.org/10.1016/j.eswa.2016.12.035> (2017).
41. Johnson, J. M. & Khoshgoftaar, T. M. Survey on deep learning with class imbalance. *J. Big Data* **6**, 27. <https://doi.org/10.1186/s40537-019-0192-5> (2019).
42. Fawaz, I. H., Forestier, G., Weber, J., Idoumghar, L. & Muller, P. A. Deep learning for time series classification: a review. *Data Min. Knowl. Disc.* **33**, 917–963. <https://doi.org/10.1007/s10618-019-00619-1> (2019).
43. Wen, Q. *et al.* Time series data augmentation for deep learning: a survey. *CoRR*. [arxiv:2002.12478](https://arxiv.org/abs/2002.12478) (2020).
44. Iwana, B. K. & Uchida, S. An empirical survey of data augmentation for time series classification with neural networks. *CoRR*. [arxiv:2007.15951](https://arxiv.org/abs/2007.15951) (2020).
45. Le Guennec, A., Malinowski, S. & Tavenard, R. Data augmentation for time series classification using convolutional neural networks. in *ECML/PKDD Workshop Advanced Analytics and Learning on Temporal Data* (Riva Del Garda, Italy, Sept 2016).
46. Schwabedal, J. T. C., Sippel, D., Brandt, M. D. & Bialonski, S. Automated classification of sleep stages and EEG artifacts in mice with deep learning. *CoRR*. [arxiv:1809.08443](https://arxiv.org/abs/1809.08443) (2018).
47. Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
48. Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. in *3rd International Conference on Learning Representations, ICLR 2015* (San Diego, CA, USA, 7–9 May 2015).
49. Pascanu, R., Mikolov, T. & Bengio, Y. On the difficulty of training recurrent neural networks. *Proceedings of the 30th International Conference on Machine Learning, ICML 2013*, vol. 28 (Atlanta, GA, USA, 16–21 June 2013) 1310–1318 (2013).
50. Goyal, P. *et al.* Accurate, large minibatch SGD: training ImageNet in 1 hour. *CoRR*. [arxiv:1706.02677](https://arxiv.org/abs/1706.02677) (2017).
51. Montavon, G., Orr, G. B. & Müller, K. *Neural Networks: Tricks of the Trade*, vol. 7700 of *Lecture Notes in Computer Science* (Springer, 2012), 2nd edn.
52. Ting, K. M. *Encyclopedia of Machine Learning and Data Mining*, chap. Confusion Matrix, 260 (Springer US, Boston, MA, 2017).
53. Tharwat, A. Classification assessment methods. *Appl. Comput. Inform.* <https://doi.org/10.1016/j.aci.2018.08.003> (2018).
54. Chambon, S., Galtier, M. N., Arnal, P. J., Wainrib, G. & Gramfort, A. A deep learning architecture for temporal sleep stage classification using multivariate and multimodal time series. *IEEE T. Neur. Sys. Reh.* **26**, 758–769. <https://doi.org/10.1109/tnsre.2018.2813138> (2018).
55. Supratak, A., Dong, H., Wu, C. & Guo, Y. DeepSleepNet: a model for automatic sleep stage scoring based on raw single-channel EEG. *IEEE T. Neur. Sys. Reh.* **25**, 1998–2008. <https://doi.org/10.1109/TNSRE.2017.2721116> (2017).
56. Franken, P., Malafosse, A. & Tafti, M. Genetic variation in EEG activity during sleep in inbred mice. *Am. J. Physiol-Reg I* **275**, R1127–R1137. <https://doi.org/10.1152/ajpregu.1998.275.4.r1127> (1998).
57. Northcutt, C. G., Jiang, L. & Chuang, I. L. Confident learning: estimating uncertainty in dataset labels. *J. Artif. Intell. Res.* **70**, 1373–1411. <https://doi.org/10.1613/jair.1.12125> (2021).
58. Lipton, Z. C., Wang, Y. & Smola, A. J. Detecting and correcting for label shift with black box predictors. *Proceedings 35th International Conference on Machine Learning, ICML 2018*, vol. 80, 3128–3136 (Stockholm, Sweden, 10–15 July 2018).
59. Huang, J., Qu, L., Jia, R. & Zhao, B. O2U-Net: a simple noisy label detection approach for deep neural networks. in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, vol. 80, 3325–3333. <https://doi.org/10.1109/ICCV.2019.00342> (IEEE, Seoul, Korea (South), 27 Oct–2 Nov 2019).
60. Vaswani, A. *et al.* Attention is all you need. in *Annual Conference Neural Information Processing Systems 2017*. 5998–6008 (Long Beach, CA, USA, 4–9 Dec 2017).
61. Wu, N., Green, B., Ben, X. & O'Banion, S. Deep transformer models for time series forecasting: the influenza prevalence case. *CoRR*. [arxiv:2001.08317](https://arxiv.org/abs/2001.08317) (2020).
62. Lacourse, K., Yetton, B., Mednick, S. & Warby, S. C. Massive online data annotation, crowdsourcing to generate high quality sleep spindle annotations from EEG data. *Sci. Data* **7**, 190. <https://doi.org/10.1038/s41597-020-0533-4> (2020).

Acknowledgements

The mice experiments were funded by the DFG project GZ: RI 2360/2-1 received by Yvonne Ritze. We are grateful to M. Reißel and M. Grajewski for providing us with computing resources.

Author contributions

Y.R. conceived the mice experiments and collected the data; S.W. annotated the data; J.T.C.S. and S.B. conceived the deep learning experiments; N.G. conducted the deep learning experiments, analyzed the results, and created the figures; S.B. wrote the first draft of the manuscript; all authors reviewed the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to S.B.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021