
Learning Sleep Stages from Radio Signals: A Conditional Adversarial Architecture

Mingmin Zhao¹ Shichao Yue¹ Dina Katabi¹ Tommi S. Jaakkola¹ Matt T. Bianchi²

Abstract

We focus on predicting sleep stages from radio measurements without any attached sensors on subjects. We introduce a new predictive model that combines convolutional and recurrent neural networks to extract sleep-specific subject-invariant features from RF signals and capture the temporal progression of sleep. A key innovation underlying our approach is a modified adversarial training regime that discards extraneous information specific to individuals or measurement conditions, while retaining all information relevant to the predictive task. We analyze our game theoretic setup and empirically demonstrate that our model achieves significant improvements over state-of-the-art solutions.

1. Introduction

Sleep plays a vital role in an individual’s health and well-being. Sleep progresses in cycles that involve multiple sleep stages: Awake, Light sleep, Deep sleep and REM (Rapid Eye Movement). Different stages are associated with different physiological functions. For example, deep sleep is essential for tissue growth, muscle repair, and memory consolidation, while REM helps procedural memory and emotional health. At least, 40 million Americans each year suffer from chronic sleep disorders ([National Institute of Health](#)). Most sleep disorders can be managed once they are correctly diagnosed ([National Institute of Health](#)). Monitoring sleep stages is beneficial for diagnosing sleep disorders, and tracking the response to treatment ([Carskadon & Rechtschaffen, 2000](#)).

Prevailing approaches for monitoring sleep stages are inconvenient and intrusive. The medical gold standard relies on Polysomnography (PSG), which is typically con-

ducted in a hospital or sleep lab, and requires the subject to wear a plethora of sensors, such as EEG-scalp electrodes, an ECG monitor, multiple chest bands, and nasal probes. As a result, patients can experience sleeping difficulties, which renders the measurements unrepresentative ([Herbst, 2010](#)). Furthermore, the cost and discomfort of PSG limit the potential for long term sleep studies.

Recent advances in wireless systems have demonstrated that radio technologies can capture physiological signals without body contact ([Kaltiokallio et al., 2014](#); [Adib et al., 2015](#); [Zhao et al., 2016](#)). These technologies transmit a low power radio signal (i.e., 1000 times lower power than a cell phone transmission) and analyze its reflections. They extract a person’s breathing and heart beats from the radio frequency (RF) signal reflected off her body. Since the cardio-respiratory signals are correlated with sleep stages, in principle, one could hope to learn a subject’s sleep stages by analyzing the RF signal reflected off her body. Such a system would significantly reduce the cost and discomfort of today’s sleep staging, and allow for long term sleep stage monitoring.

There are multiple challenges in realizing the potential of RF measurements for sleep staging. In particular, we must learn RF signal features that capture the sleep stages and their temporal progression, and the features should be transferable to new subjects and different environments. The problem is that RF signals carry much information that is irrelevant to sleep staging, and are highly dependent on the individuals and the measurement conditions. Specifically, they reflect off all objects in the environment including walls and furniture, and are affected by the subject’s position and distance from the radio device. These challenges were not addressed in past work which used hand-crafted signal features to train a classifier ([Zaffaroni et al., 2014](#); [Tataraidze et al., 2016b](#)). The accuracy was relatively low (~64%) and the model did not generalize beyond the environment where the measurements were collected.

This paper presents a new model that delivers a significantly higher accuracy and generalizes well to new environments and subjects. The model adapts a convolutional neural network (CNN) to extract stage-specific features from RF spectrograms, and couples it with a recurrent

¹MIT CSAIL, Cambridge, MA, USA ²Massachusetts General Hospital, Boston, MA, USA. Correspondence to: Mingmin Zhao <mingmin@mit.edu>.

neural network (RNN) to capture the temporal dynamics of sleep stages.

However, a CNN-RNN combination alone would remain liable to distracting features pertaining to specific individuals or measurement conditions (i.e., the source domains), and hence would not generalize well. To address this issue, we introduce a new adversarial training regime that discards extraneous information specific to individuals or measurement conditions, while retaining all information relevant to the predictive task –i.e., the adversary ensures conditional independence between the learned representation and the source domains.

Our training regime involves 3 players: the feature encoder (CNN-RNN), the sleep stage predictor, and the source discriminator. The encoder plays a cooperative game with the predictor to predict sleep stages, and a minimax game against the source discriminator. Our source discriminator deviates from the standard domain-adversarial discriminator in that it takes as input also the predicted distribution of sleep stages in addition to the encoded features. This dependence facilitates accounting for inherent correlations between stages and individuals, which cannot be removed without degrading the performance of the predictive task.

We analyze this game and demonstrate that at equilibrium, the encoded features discard all extraneous information that is specific to the individuals or measurement conditions, while preserving all information relevant to predicting the sleep stages. We also evaluate our model on a dataset of RF measurements and corresponding sleep stages¹. Experimental results show that our model significantly improves the prediction accuracy of sleep stages. In particular, our model has a prediction accuracy of 79.8% and a Cohen’s Kappa of 0.70, whereas the best prior result for predicting sleep stages from RF signals (Tataraidze et al., 2016b) has an accuracy of 64% and a Cohen’s Kappa of 0.49.

2. Related Work

(a) Sleep Staging: The gold standard in sleep staging is based on Polysomnography (PSG) conducted overnight in a hospital or sleep lab. The subject has to sleep while wearing multiple sensors including an EEG monitor, an EMG monitor, an EOG monitor, nasal probes, etc. A sleep technologist visually inspects the output of the sensors and assigns to each 30-second window a stage label (Rechtschaffen & Kales, 1968).

A few past proposals have tried to automate the process and reduce the number of sensors. These solutions can be classified into four categories according to their source

Table 1. Automated Sleep Staging Systems

Signal Source	Accuracy (acc/κ) ¹	Comfort
EEG	High (83%/0.76) ²	Low
Cardiorespiratory	Medium (71%/0.56)	Medium
Actigraphy	Low (65%/-) ³	High
RF	State-of-the-art	Low (64%/0.49)
	Ours	High (79.8%/0.70)

¹ Four-class subject-independent classification accuracy on every 30-second segment.

² Some studies achieve accuracy over 90% (da Silveira et al., 2016) but they discard artifacts and use segments from the same night to train and test.

³ Three-class classification based on 5-minute segment.

signal: EEG-based, Cardiorespiratory-based, Actigraphy-based, or RF-based. Table 1 summarizes the state of the art performance in each category. The table shows both the classification accuracy and the Cohen’s Kappa coefficient, κ . The most accurate methods rely on EEG signals (Ebrahimi et al., 2008; Fraiwan et al., 2012; Popovic et al., 2014; Shambroom et al., 2012). However, EEG monitors are also the most intrusive because they require the subject to sleep with a skullcap or a head-band equipped with multiple electrodes, which is uncomfortable and can cause headaches and skin irritation. The second category requires the subject to wear a chest-band and analyzes the resulting cardiorespiratory signals. It is more comfortable than the prior method but also less accurate (Tataraidze et al., 2016a; Long et al., 2014). The third approach is based on actigraphy; it leverages accelerometers in FitBit or smart phones (Hao et al., 2013; Gu et al., 2014) to monitor body movements and infer sleep quality. Yet, motion is known to be a poor metric for measuring sleep stages and quality (Pollak et al., 2001). The last approach relies on RF signals reflected off the subject body during her sleep. It allows the subject to sleep comfortably without any on-body sensors. Yet past approaches in this category have the worst performance in comparison to other solutions.

This paper builds on the above literature but delivers significant new contributions. In comparison to methods that use sources other than RF signals, the paper enables accurate monitoring of sleep stages while allowing the subject to sleep comfortably in her own bed without sensors on her body. Furthermore, due to differences between RF signals and other signal sources, our model has to eliminate extraneous information that are specific to the environment and irrelevant to sleep stage classification. In comparison to past work on learning sleep stages from RF signals (Rahman et al., 2015; Tataraidze et al., 2016b; Liu et al., 2014), our approach significantly improves the prediction accuracy as shown in Table 1. This improvement is due to intrinsic differences between past models and the model in this paper, which avoids hand-crafted features, and learns

¹Dataset is available at:

<http://sleep.csail.mit.edu/>

features that capture the temporal dependencies and transfer well to new subjects and different environments.

(b) Representation Learning: We build on a rich body of literature on CNNs and RNNs which have been successfully used to model spatial patterns (Szegedy et al., 2015; He et al., 2016) and temporal dynamics (Sutskever et al., 2014), including combinations of the two (Pigou et al., 2015). Our CNN differs slightly in terms of convolutions that are adapted to our domain while, architecturally, our RNN is a standard variety LSTM.

Our work also contributes to learning invariant representations in deep adversarial networks. Adversarial networks were introduced to effectively train complex generative models of images (Goodfellow et al., 2014; Radford et al., 2015; Chen et al., 2016) where the adversary (discriminator) was introduced so as to match generated samples with observed ones. The broader approach has since been adopted as the training paradigm across a number of other tasks as well, from learning representations for semi-supervised learning (Makhzani et al., 2015), and modeling dynamic evolution (Vondrick et al., 2016; Purushotham et al., 2017) to inverse maps for inference (Donahue et al., 2017; Dumoulin et al., 2017), and many others. Substantial work has also gone into improving the stability of adversarial training (Metz et al., 2016; Arjovsky et al., 2017; Arjovsky & Bottou, 2017).

On a technical level, our work is most related to adversarial architectures for domain adaptation (Ganin & Lempitsky, 2015; Ganin et al., 2016; Tzeng et al., 2015; 2016). Yet, there are key differences between our approach and the above references, beyond the main application of sleep staging that we introduce. First, our goal is to remove conditional dependencies rather than making the representation domain independent. Thus, unlike the above references which do not involve conditioning in the adversary, our adversary takes the representation but is also conditioned on the predicted label distribution. Second, our game theoretic setup controls the information flow differently, ensuring that only the representation encoder is modified based on the adversary performance. Specifically, the predicted distribution over stages is strategically decoupled from the adversary (conditioning is uni-directional). Third, we show that this new conditioning guarantees an equilibrium solution that fully preserves the ability to predict sleep staging while removing, conditionally, extraneous information specific to the individuals or measurement conditions. Guarantees of this kind are particularly important for healthcare data where the measurements are noisy with a variety of dependencies that need to be controlled.

Finally, our work is naturally also related to other non-adversarial literature on multi-source domain adaptation (Crammer et al., 2008; Long et al., 2015), and work on

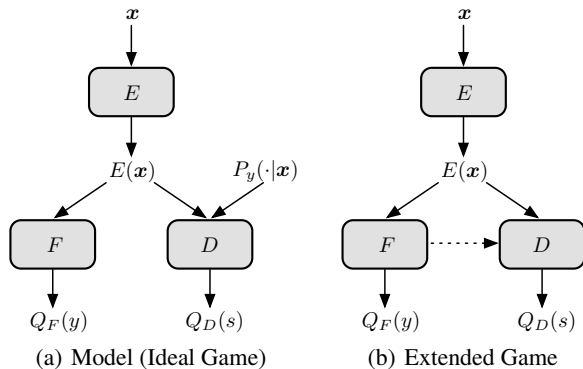


Figure 1. Model and Extended Game. Dotted arrow indicates that the information does not propagate back on this link.

metrics for measuring distance between distributions (Ben-David et al., 2010; Fernando et al., 2013).

3. Model

Let $x \in \Omega_x$ be an input sample, and $y \in \{1, 2, \dots, n_y\}$ an output label. Let $s \in \{1, 2, \dots, n_s\}$ denote an auxiliary label that refers to the source of a specific input sample. We define $\mathbf{x} = [x_1, x_2, \dots, x_t] \in \Omega_x$ as the sequence of input samples from the beginning of time until the current time t .

In the context of our application, the above notation translates into the following: The input sample x is a 30-second RF spectrogram, and the output label y is a sleep stage that takes one of four values: Awake, Light Sleep, Deep Sleep, or REM. The vector \mathbf{x} refers to the sequence of RF spectrograms from the beginning of the night until the current time. Since RF signals carry information about the subject and the measurement environment, we assign each input x an auxiliary label s which identifies the subject-environment pair, hereafter referred to as the source.

Our goal is to learn a latent representation (i.e., an encoder) that can be used to predict label y ; yet, we want this representation to generalize well to predict sleep stages for new subjects without having labeled data from them. Simply making the representation invariant to the source domains could hamper the accuracy of the predictive task. Instead we would like to *remove conditional dependencies between the representation and the source domains*.

We introduce a multi-domain adversarial model that achieves the above goal. Our model is shown in Fig. 1(a). It has three components: An encoder E , a label predictor F , and a source discriminator D . Our model is set up as a game, where the representation encoder plays a cooperative game with the label predictor to allow it to predict the correct labels using the encoded representation. The encoder also plays a minimax game against the source discrimina-

tor to prevent it from decoding the source label from the encoded representation.

A key characteristic of our model is the conditioning of the source discriminator on the label distribution, $P_y(\cdot|\mathbf{x})$ (see Fig. 1(a)). This conditioning of the adversary allows the learned representation to correlate with the domains, but only via the label distribution –i.e., removes conditional dependencies between the representation and the sources.

The rest of this section is organized as follows. We first formally define three players E , F , and D and the representation invariance they are trained to achieve. In Sec. 3.1, we analyze the game and prove that at equilibrium the encoder discards all extraneous information about the source that is not beneficial for label prediction (i.e., predicting y). Training the ideal model in Fig. 1(a) is challenging because it requires access to the label distribution $P_y(\cdot|\mathbf{x})$. To drive an efficient training algorithm, we define in Sec. 3.2 an extended game where the source discriminator uses the output of the label predictor as an approximation of the posterior probabilities, as shown in Fig. 1(b). We prove that the equilibriums of the original game are also equilibriums in the extended one.

Encoder E : An encoder $E(\cdot) : \Omega_{\mathbf{x}} \rightarrow \Omega_{\mathbf{z}}$ is a function that takes a sequence of input samples \mathbf{x} , and returns a vector summary of \mathbf{x} as $\mathbf{z} = E(\mathbf{x})$.

Label Predictor F : A label predictor $F(\cdot) : \Omega_{\mathbf{z}} \rightarrow [0, 1]^{n_y}$ takes a latent representation $E(\mathbf{x})$ as input and predicts the probability of each label y associated with input \mathbf{x} as $Q_F(y|E(\mathbf{x}))$. The goal of an ideal predictor F is to approximate $P_y(\cdot|\mathbf{x})$ with $Q_F(\cdot|E(\mathbf{x}))$.

The loss of the label predictor, F , given the encoder E , is defined as the cross-entropy between the label distribution $P_y(\cdot|\mathbf{x})$ and $Q_F(\cdot|E(\mathbf{x}))$:

$$\mathcal{L}_f(F; E) = \mathbb{E}_{\mathbf{x}, y}[-\log Q_F(y|E(\mathbf{x}))] \quad (1)$$

During training, the encoder E and predictor F play a co-operative game to minimize the label prediction loss.

Source Discriminator D : We define a source discriminator as $D(\cdot, \cdot) : \Omega_{\mathbf{z}} \times [0, 1]^{n_y} \rightarrow [0, 1]^{n_s}$. It takes the latent representation $E(\mathbf{x})$ and the label distribution $P_y(\cdot|\mathbf{x})$ as inputs, and predicts which source domain (i.e., subject and environment) they are sampled from as $Q_D(\cdot|E(\mathbf{x}), P_y(\cdot|\mathbf{x}))$.

Next, we define the desired representation invariance.

Definition 1 (Representation invariance). *We say that representation E is invariant if $E(\mathbf{x})$ contains no information about s beyond what is already contained in $P_y(\cdot|\mathbf{x})$; that is, $Q_D(\cdot|E(\mathbf{x}), P_y(\cdot|\mathbf{x})) = Q_D(\cdot|P_y(\cdot|\mathbf{x}))$ for the optimal D .*

To measure the invariance of an encoder E , we define the

loss of the source discriminator D as the cross-entropy between $P_s(\cdot|\mathbf{x})$ and $Q_D(\cdot|E(\mathbf{x}), P_y(\cdot|\mathbf{x}))$:

$$\mathcal{L}_d(D; E) = \mathbb{E}_{\mathbf{x}, s}[-\log Q_D(s|E(\mathbf{x}), P_y(\cdot|\mathbf{x}))] \quad (2)$$

During training, encoder E and discriminator D play a minimax game: while D is trained to minimize the source prediction loss, encoder E is trained to maximize it in order to achieve the above invariance.

3.1. Ideal Game

During training, encoder E plays a co-operative game with predictor F , and a minimax game with discriminator D . We define a value function of E , F and D with $\lambda > 0$:

$$\mathcal{V}(E, F, D) = \mathcal{L}_f(F; E) - \lambda \cdot \mathcal{L}_d(D; E) \quad (3)$$

The training procedure can be viewed as a three-player minimax game of E , F and D :

$$\min_E \min_F \max_D \mathcal{V}(E, F, D) = \min_{E, F} \max_D \mathcal{V}(E, F, D) \quad (4)$$

Proposition 2 (Optimal predictor). *Given encoder E ,*

$$\mathcal{L}_f(E) \triangleq \min_F \mathcal{L}_f(F; E) \geq H(y|E(\mathbf{x})), \quad (5)$$

where $H(\cdot)$ is entropy.

The optimal predictor F^* that achieves equality is:

$$Q_{F^*}(y|E(\mathbf{x})) = p(y|E(\mathbf{x})) \quad (6)$$

Proof.

$$\begin{aligned} & \mathcal{L}_f(F; E) \\ &= \mathbb{E}_{\mathbf{x}, y}[-\log Q_F(y|E(\mathbf{x}))] \\ &= \mathbb{E}_{E(\mathbf{x}), y}[-\log Q_F(y|E(\mathbf{x}))] \\ &= \mathbb{E}_{\mathbf{z} \sim P(E(\mathbf{x}))} \mathbb{E}_{y \sim P(y|\mathbf{z})}[-\log Q_F(y|\mathbf{z})] \\ &= \mathbb{E}_{\mathbf{z} \sim P(E(\mathbf{x}))} [H(y|\mathbf{z}) + D_{\text{KL}}(P(y|\mathbf{z}) \| Q_F(y|\mathbf{z}))] \\ &\geq \mathbb{E}_{\mathbf{z} \sim P(E(\mathbf{x}))} [H(y|\mathbf{z})] \\ &= H(y|E(\mathbf{x})) \end{aligned}$$

The equality holds when $D_{\text{KL}}(P(y|E(\mathbf{x})) \| Q_F(y|E(\mathbf{x}))) = 0$ for almost every $\mathbf{x} \in \text{Supp}(P_{\mathbf{x}})$. That is $Q_{F^*}(y|E(\mathbf{x})) = p(y|E(\mathbf{x}))$ for almost every y and $\mathbf{x} \in \text{Supp}(P_{\mathbf{x}})$. \square

Similarly we can prove the following Proposition.

Proposition 3 (Optimal discriminator). *Given encoder E ,*

$$\mathcal{L}_d(E) \triangleq \min_D \mathcal{L}_d(D; E) \geq H(s|E(\mathbf{x}), P_y(\cdot|\mathbf{x})) \quad (7)$$

The optimal discriminator D^* that achieves this value is:

$$Q_{D^*}(s|E(\mathbf{x}), P_y(\cdot|\mathbf{x})) = P(s|E(\mathbf{x}), P_y(\cdot|\mathbf{x})) \quad (8)$$

Corollary 3.1. $H(s)$ is an upper bound of the loss of the optimal discriminator D^* for any encoder E .

Next, we state the virtual training criterion of the encoder.

Proposition 4. If predictor F and discriminator D have enough capacity and are trained to achieve their optimal losses, the minimax game (4) can be rewritten as the following training procedure of encoder E :

$$\min_E [H(y|E(\mathbf{x})) - \lambda \cdot H(s|E(\mathbf{x}), P_y(\cdot|\mathbf{x}))] \quad (9)$$

Proof. Based on the losses of the optimal predictor F^* and the optimal discriminator D^* in Proposition 2 and Proposition 3, the minimax game (4) can be rewritten as (9). Thus, encoder E is trained to minimize a virtual training criterion $C(E) = H(y|E(\mathbf{x})) - \lambda \cdot H(s|E(\mathbf{x}), P_y(\cdot|\mathbf{x}))$. \square

Next, we describe the optimal encoder.

Theorem 5 (Optimal encoder). If encoder E , predictor F and discriminator D have enough capacity and are trained to reach optimum, any global optimal encoder E^* has the following properties:

$$H(y|E^*(\mathbf{x})) = H(y|\mathbf{x}) \quad (10a)$$

$$H(s|E^*(\mathbf{x}), P_y(\cdot|\mathbf{x})) = H(s|P_y(\cdot|\mathbf{x})) \quad (10b)$$

Proof. Since $E(\mathbf{x})$ is a function of \mathbf{x} :

$$\mathcal{L}_f(E) = H(y|E(\mathbf{x})) \geq H(y|\mathbf{x}) \quad (11a)$$

$$\mathcal{L}_d(E) = H(s|E(\mathbf{x}), P_y(\cdot|\mathbf{x})) \leq H(s|P_y(\cdot|\mathbf{x})) \quad (11b)$$

Hence, $C(E) = H(y|E(\mathbf{x})) - \lambda \cdot H(s|E(\mathbf{x}), P_y(\cdot|\mathbf{x})) \geq H(y|\mathbf{x}) - \lambda \cdot H(s|P_y(\cdot|\mathbf{x}))$. The equality holds if and only if both (10a) and (10b) are satisfied. Therefore, we only need to prove that the optimal value of $C(E)$ is equal to $H(y|\mathbf{x}) - \lambda \cdot H(s|P_y(\cdot|\mathbf{x}))$ in order to prove that any global encoder E^* satisfies both (10a) and (10b).

We show that $C(E)$ can achieve $H(y|\mathbf{x}) - \lambda \cdot H(s|P_y(\cdot|\mathbf{x}))$ by considering the following encoder E_0 : $E_0(\mathbf{x}) = P_y(\cdot|\mathbf{x})$. It can be examined that $H(y|E_0(\mathbf{x})) = H(y|\mathbf{x})$ and $H(s|E_0(\mathbf{x}), P_y(\cdot|\mathbf{x})) = H(s|P_y(\cdot|\mathbf{x}))$. \square

Adversarial training of D can be viewed as a regularizer, which leads to a common representation space for multiple source domains. From Theorem 5, the optimal encoder E^* using adversarial training satisfies $H(y|E^*(\mathbf{x})) = H(y|\mathbf{x})$, which is the maximal discriminative capability that any encoder E can achieve. Thus, we have the following corollary.

Corollary 5.1. Adversarial training of the discriminator does not reduce the discriminative capability of the representation.

Remark 5.1. During the proof of Theorem 5, we construct an encoder $E_0(\mathbf{x}) = P_y(\cdot|\mathbf{x})$ that can achieve the optimal value of $C(E)$. However, we argue that training will not converge to this trivial encoder in practice. This is because $P_y(\cdot|\mathbf{x})$ is a mapping from the full signal history to the distribution over stages at the current step, therefore itself highly complex. Since we use the RNN state as the encoding $E(\mathbf{x})$, and it feeds into the LSTM gates, distribution over stages at previous step does not represent a sufficient summary of the history until the current one. Therefore, $E(\mathbf{x})$ must be able to anticipate the temporal evolution of the signal and contain a more effective summary than $P_y(\cdot|\mathbf{x})$ would be.

Corollary 5.2. If encoder E and predictor F have enough capacity and are trained to reach optimum, the output of F is equal to $P_y(\cdot|\mathbf{x})$.

Proof. When predictor F is optimal (Proposition 2), $Q_F(y|E(\mathbf{x})) = p(y|E(\mathbf{x}))$. When E is optimal (Theorem 5), $H(y|E(\mathbf{x})) = H(y|\mathbf{x})$, that is $p(y|E(\mathbf{x})) = p(y|\mathbf{x})$. Therefore, $Q_F(y|E(\mathbf{x})) = p(y|\mathbf{x})$. \square

3.2. Extended Game

In practice, estimating the posterior label distribution $P_y(\cdot|\mathbf{x})$ from labeled data is a non-trivial task. Fortunately however our predictor F and encoder E are playing a cooperative game to approximate this posterior label distribution $P_y(\cdot|\mathbf{x})$ with $Q_F(\cdot|E(\mathbf{x}))$. Therefore, we use $Q_F(\cdot|E(\mathbf{x}))$, the output of predictor F , as a proxy of $P_y(\cdot|\mathbf{x})$ and feed it as input to discriminator D (Fig. 1(b)).

An extended three-player game arises: while encoder E still plays a cooperative game with predictor F and a minimax game with discriminator D , discriminator D depends strategically on predictor F but not vice versa. The dotted line in Fig. 1(b) illustrates this dependency.

The relationship between the ideal minimax game (Sec. 3.1) and the extended one is stated below.

Proposition 6. If encoder E , predictor F and discriminator D have enough capacity, the solution that encompasses the optimal encoder, E^* , predictor, F^* and discriminator, D^* , in the ideal minimax game is also an equilibrium solution of the extended game.

Proof. By Corollary 5.2, when encoder E and predictor F are optimal, $Q_F(\cdot|E(\mathbf{x}))$ is equal to $P_y(\cdot|\mathbf{x})$. Thus, the extended game becomes equivalent to the ideal game, and E^* , F^* and D^* is an equilibrium solution of both games. \square

Algorithm 1 Encoder, predictor and discriminator training

Input: Labeled data $\{(\mathbf{x}_i, y_i, s_i)\}_{i=1}^M$, learning rate η .
 Compute stop criterion for inner loop: $\delta_d \leftarrow H(s)$
for number of training iterations **do**
 Sample a mini-batch of training data $\{(\mathbf{x}_i, y_i, s_i)\}_{i=1}^m$
 $\mathcal{L}_f^i \leftarrow -\log Q_F(y_i|E(\mathbf{x}_i))$
 $\mathbf{w}_i \leftarrow Q_F(\cdot|E(\mathbf{x}_i)) \blacktriangleright$ stop gradient along this link
 $\mathcal{L}_d^i \leftarrow -\log Q_D(s_i|E(\mathbf{x}_i), \mathbf{w}_i)$
 $\mathcal{V}^i = \mathcal{L}_f^i - \lambda \cdot \mathcal{L}_d^i$
 Update encoder E :
 $\theta_e \leftarrow \theta_e - \eta_e \nabla_{\theta_e} \frac{1}{m} \sum_{i=1}^m \mathcal{V}^i$
 Update predictor F :
 $\theta_f \leftarrow \theta_f - \eta_f \nabla_{\theta_f} \frac{1}{m} \sum_{i=1}^m \mathcal{V}^i$
 repeat
 Update discriminator D :
 $\theta_d \leftarrow \theta_d + \eta_d \nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m \mathcal{V}^i$
 until $\frac{1}{m} \sum_{i=1}^m \mathcal{L}_d^i \leq \delta_d$
 end for

3.3. Training Algorithm

We implement the extended three-player game with iterative updates of the players (Algorithm 1). Note that, since the output of the label predictor is a proxy of the underlying posterior, and since the source discriminator depends strategically on the predictor but not vice versa, the gradient does not back-propagate from the discriminator to the predictor (i.e., the dotted link in Fig. 1(b)).

The number of training steps in the inner loop usually needs to be carefully chosen (Goodfellow et al., 2014). A large number of steps is computationally inefficient but a small one will cause the model to collapse. This is because the outer players, E and F , can be over-trained against a non-optimal inner player D , and they will try to maximize \mathcal{L}_d at the cost of increasing \mathcal{L}_f . To prevent the model collapse phenomenon, we use an adaptive number of training steps in the inner loop and adjust it dynamically based on \mathcal{L}_d (Algorithm 1). The idea is to use the upper bound in Corollary 3.1 as the stopping criterion for the inner loop.

3.4. Discussion of the Model Benefits

While we described our model in the context of sleep staging, we believe the model can be applied more broadly. Our model is characterized by the 3-way game and the adversarial conditioning on the label distribution. This combination yields the following benefits: 1) It guarantees an equilibrium solution that fully preserves the ability to perform the predictive task while removing any distracting information specific to the source domains. Guarantees of this kind are particularly important in healthcare where the measurements are noisy and have a variety of dependencies that need to be controlled. 2) It allows to properly leverage the

adversarial feedback even when the target labels are uncertain. For example, in the sleep staging problem, each 30-second window is given one label. Yet, many such windows include transitions between sleep stages, e.g., a transition from light to deep sleep. These transitions are gradual and hence the transition windows can be intrinsically different from both light and deep sleep. It would be desirable to have the learned representation capture the concept of transition and make it invariant to the source (see the results in Sec. 4.5). 3) It allows the conditioning to remain available for additional guiding of representations based on unlabeled data. The model can incorporate unlabeled data for either semi-supervised learning or transductive learning within a unified framework.

4. Experiments

In this section, we empirically evaluate our model.

4.1. RF-Sleep Dataset

RF-Sleep is a dataset of RF measurements during sleep with corresponding sleep stage labels. All studies that involve human subjects were approved by our IRB.

Study setup: The sleep studies are done in the bedroom of each subject. We install a radio device in the bedroom. It transmits RF signals and measure their reflections while the subject is sleeping alone in the bed.

Ground truth: During the study, each subject sleeps with an FDA-approved EEG-based sleep monitor (Popovic et al., 2014), which collects 3-channel frontal EEG. The monitor labels every 30-second of sleep with the subject’s sleep stage. This system has human-level comparable accuracy (Popovic et al., 2014), and has already been used in several sleep studies (Lucey et al., 2016; Shah et al., 2016).

Size of dataset: The dataset collects 100 nights of sleep from 25 young healthy subjects (40% females). It contains over 90k 30-second epochs of RF measurements and their corresponding sleep stages provided by the EEG-based sleep monitor. Each epoch has one of four labels Awake, REM, Light Sleep (N1 or N2) and Deep Sleep (N3).

4.2. Parameterization

We parameterize encoder E , predictor F and discriminator D as neural networks. Encoder E is parameterized by a hybrid CNN-RNN model. We adapt a residual networks architecture (He et al., 2016) with 24 convolutional layers to extract features from each 30-second RF spectrogram, and an RNN with LSTM cell (Hochreiter & Schmidhuber, 1997) that takes sequences of CNN features as input. Both predictor F and discriminator D are parameterized by networks with two fully-connected layers.

Table 2. Sleep Stage Classification Accuracy and Kappa

Approach	Accuracy	κ
Tataraidze et al. (2016b)	0.635	0.49
Zaffaroni et al. (2014)	0.641	0.45
Ours	0.798	0.70

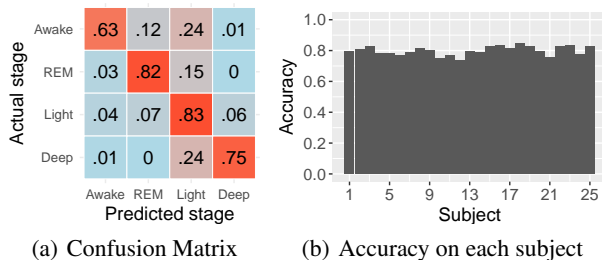


Figure 2. 2(a) shows that our model can distinguish deep and light sleep with high accuracy. And 2(b) illustrates that our model works well for different subjects and environments.

4.3. Classification Results

We evaluate the model on every subject while training on the data collected from the other subjects (i.e., the model is never trained on data from the test subject). The training data is randomly split into a training set and validation set (75%/25%).

We use two metrics commonly used in automated sleep staging, namely Accuracy and Cohen’s Kappa. While accuracy measures the percent agreement with ground truth, Cohen’s Kappa coefficient κ (Cohen, 1960) takes into account the possibility of the agreement occurring by chance and is usually a more robust metric. $\kappa > 0.4$, $\kappa > 0.6$, $\kappa > 0.8$ are considered to be moderate, substantial and almost perfect agreement (Landis & Koch, 1977).

Table 2 shows the accuracy and Cohen’s Kappa of our model compared to the state-of-the-art in classifying sleep stages using RF reflections. Since neither the dataset nor the code used in past papers is publicly available, we compare with their published results. We note however that the Cohen’s Kappa provides some normalization since it accounts for the underlying uncertainty in the data. The table shows that our model has an accuracy of 79.8% and a $\kappa = 0.70$, which significantly outperforms past solutions.

Fig. 2(a) shows the confusion matrix of our model. Fig. 2(b) also shows the accuracy on each subject. It has a standard deviation of 2.9%, suggesting that our model is capable of adapting to different subjects and environments.

Finally, we show in Fig. 3 the full-night predictions along with the ground truth for the average, best, and worst classification accuracy.

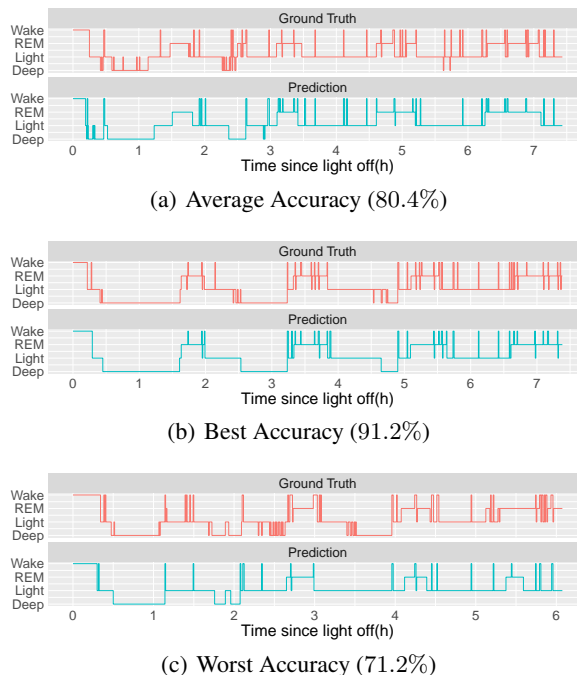


Figure 3. Three examples of full night predictions corresponding to the average, best and worst classification accuracy.

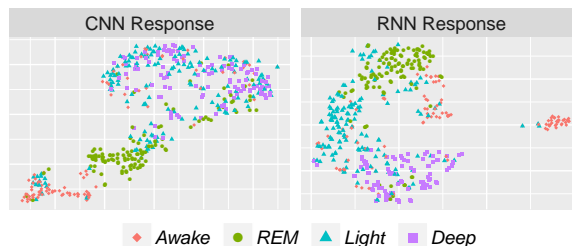


Figure 4. Visualizations of the CNN and RNN responses. CNN can separate Wake REM and from the other stages, yet Deep and Light Sleep can only be distinguished by RNN.

4.4. Understanding the Role of CNN & RNN

We analyze the role of CNN and RNN in predicting sleep stages. To do so, we use t-SNE embedding (Maaten & Hinton, 2008) to visualize the response of our network after CNN and RNN, respectively. Fig. 4 shows the visualization results from one of the subjects. Data points are randomly sub-sampled for better viewing. The result shows that the CNN succeeds at separating the Wake, REM from Light and Deep Sleep. However it fails at separate Light Sleep and Deep Sleep from each other. In contrast, Light Sleep and Deep Sleep form different clusters in the RNN response. These results demonstrate the role of CNN and RNN in our model: CNN learns stage-specific features that can distinguish Wake, REM and from Deep and Light Sleep. RNN captures the dynamics of those features to fur-

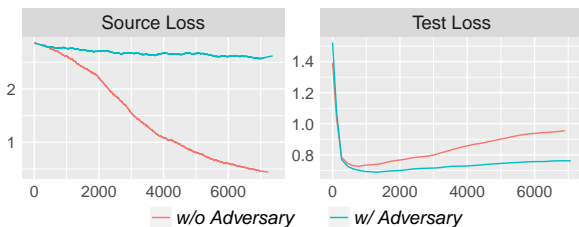


Figure 5. Baseline model and ours are evaluated on same dataset. A higher source loss indicates the removal of source specific information, and a lower test loss shows that the proposed setup can better avoid overfitting.

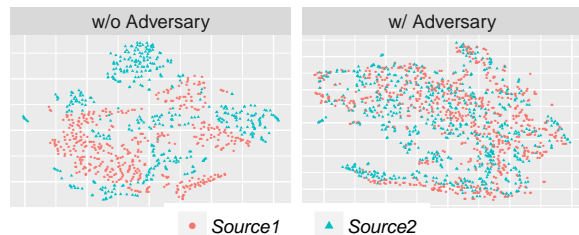


Figure 6. Visualization of learned latent representations from two sources. Data-points are separated when no adversary, yet they are well aligned by proposed setup.

ther determine whether the sleep is light or deep. Note that Light and Deep Sleep are more similar to each other and are typically referred to as NREM, i.e., non-REM.

We have trained a similar model without the RNN layer on top of CNN. In this case, the overall accuracy decreases by 12.8%, specifically the precision light and deep sleep decreases by 23.5%. This suggests that there are stage-specific information embedded in the temporal dynamics of the RF measurements, and therefore can only be captured and exploited with RNN. Moreover, these temporal dynamics are particularly crucial for distinguishing light and deep sleep. Indeed, there are known temporal patterns that govern the progression of light and deep sleep through the night (Carskadon et al., 2005). For example, the probability of being in deep sleep decreases as sleep progresses. Also, people usually need to go through light sleep before they can get into deep sleep. These temporal dynamics of sleep stages can be captured by RNN and might be exploited to distinguish light and deep sleep.

4.5. Role of Our Adversarial Discriminator

We evaluate the role of our adversarial discriminator in learning transferable features for predicting sleep stages. We first look at the losses on the validation set as training progresses to check whether the extraneous information specific to the individuals and environments can be removed. As a baseline, we compare with a version of our model without the source discriminator. For this baseline,




Figure 7. Visualization of fine-grained alignment on test data. Our model, which conditions the adversary on the posterior distribution, not only aligns deep and light stages, but also aligns the transition periods, which are not directly specified by the labels.

we train a (non-adversarial) discriminator to determine the source of features. Fig. 5 shows that the loss of the source discriminator in the baseline model decreases very quickly while ours stays high (upper bounded by $H(s) = 2.81$ in this case), suggesting that our learned representation is invariant across sources. The figure also shows that adding an adversarial discriminator increases the performance on the test set and can be helpful in reducing over-fitting.

To check that our adversarial model has learned transferable features, we visualize the learned features $E(\mathbf{x})$ on the test data for both models. Color-coding the sources, Fig. 6 shows that our learned features have almost the same distribution on different sources, while the baseline model learns features that are separable.

Next, we illustrate the benefits of conditioning on the posterior distribution, and that it can recover underlying concepts not specified in the labels. We consider the learned features for transition periods between light and deep sleep, which might be a class that is different from both light and deep sleep. We define transition periods as epochs that have both light and deep sleep as neighbors. We visualize it with a different color. Color-coding stages and shape-coding sources, Fig. 7 shows the learned features from transition periods are segregated, as those from light sleep and deep sleep. This indicates that our learned features have recovered the concept of a transition period, which is helpful in understanding and predicting sleep stages.

5. Conclusion

This paper introduce a new predictive model that learns sleep stages from RF signals and achieves a significant improvement over the state-of-the-art. We believe this work marks an important step in sleep monitoring. We also believe that the proposed adversarial setup, which extracts task-specific domain-invariant features, is applicable to other predictive tasks, particularly in health sensing where variations across subjects and measurement conditions could be a major challenge.

Acknowledgments

The authors thank the anonymous reviewers for their helpful comments in revising the paper. We are grateful to the members of the CSAIL for their insightful discussions and to all the human subjects for their participation in our experiments.

References

- Adib, Fadel, Mao, Hongzi, Kabelac, Zachary, Katabi, Dina, and Miller, Robert C. Smart homes that monitor breathing and heart rate. *ACM CHI*, 2015.
- Arjovsky, Martin and Bottou, Léon. Towards principled methods for training generative adversarial networks. *ICLR*, 2017.
- Arjovsky, Martin, Chintala, Soumith, and Bottou, Léon. Wasserstein generative adversarial networks. *ICML*, 2017.
- Ben-David, Shai, Blitzer, John, Crammer, Koby, Kulesza, Alex, Pereira, Fernando, and Vaughan, Jennifer Wortman. A theory of learning from different domains. *Machine learning*, 2010.
- Carskadon, Mary A and Rechtschaffen, Allan. Monitoring and staging human sleep. *Principles and practice of sleep medicine*, 2000.
- Carskadon, Mary A, Dement, William C, et al. Normal human sleep: an overview. *Principles and practice of sleep medicine*, 2005.
- Chen, Xi, Duan, Yan, Houthoof, Rein, Schulman, John, Sutskever, Ilya, and Abbeel, Pieter. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *NIPS*, 2016.
- Cohen, Jacob. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 1960.
- Crammer, Koby, Kearns, Michael, and Wortman, Jennifer. Learning from multiple sources. *JMLR*, 2008.
- da Silveira, Thiago LT, Kozakevicius, Alice J, and Rodrigues, Cesar R. Single-channel eeg sleep stage classification based on a streamlined set of statistical features in wavelet domain. *Medical & biological engineering & computing*, 2016.
- Donahue, Jeff, Krähenbühl, Philipp, and Darrell, Trevor. Adversarial feature learning. *ICLR*, 2017.
- Dumoulin, Vincent, Belghazi, Ishmael, Poole, Ben, Lamb, Alex, Arjovsky, Martin, Mastropietro, Olivier, and Courville, Aaron. Adversarially learned inference. *ICLR*, 2017.
- Ebrahimi, Farideh, Mikaeili, Mohammad, Estrada, Edson, and Nazeran, Homer. Automatic sleep stage classification based on eeg signals by using neural networks and wavelet packet coefficients. *IEEE EMBC*, 2008.
- Fernando, Basura, Habrard, Amaury, Sebban, Marc, and Tuytelaars, Tinne. Unsupervised visual domain adaptation using subspace alignment. *ICCV*, 2013.
- Fraiwani, Luay, Lweesy, Khaldon, Khasawneh, Natheer, Wenz, Heinrich, and Dickhaus, Hartmut. Automated sleep stage identification system based on time-frequency analysis of a single eeg channel and random forest classifier. *Computer methods and programs in biomedicine*, 2012.
- Ganin, Yaroslav and Lempitsky, Victor. Unsupervised domain adaptation by backpropagation. *ICML*, 2015.
- Ganin, Yaroslav, Ustinova, Evgeniya, Ajakan, Hana, Germain, Pascal, Larochelle, Hugo, Laviolette, François, Marchand, Mario, and Lempitsky, Victor. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 2016.
- Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. *NIPS*, 2014.
- Gu, Weixi, Yang, Zheng, Shangguan, Longfei, Sun, Wei, Jin, Kun, and Liu, Yunhao. Intelligent sleep stage mining service with smartphones. *ACM UbiComp*, 2014.
- Hao, Tian, Xing, Guoliang, and Zhou, Gang. isleep: unobtrusive sleep quality monitoring using smartphones. *ACM SenSys*, 2013.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. *CVPR*, 2016.
- Herbst, Ellen et al. Adaptation effects to sleep studies in participants with and without chronic posttraumatic stress disorder. *Psychophysiology*, 2010.
- Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 1997.
- Kaltiokallio, Ossi, Yigitler, Huseyin, Jantti, Riku, and Patwari, Neal. Non-invasive respiration rate monitoring using a single cots tx-rx pair. *IPSN*, 2014.
- Landis, J Richard and Koch, Gary G. The measurement of observer agreement for categorical data. *Biometrics*, 1977.

- Liu, Xuefeng, Cao, Jiannong, Tang, Shaojie, and Wen, Jiaqi. Wi-sleep: Contactless sleep monitoring via wifi signals. *RTSS*, 2014.
- Long, Mingsheng, Cao, Yue, Wang, Jianmin, and Jordan, Michael I. Learning transferable features with deep adaptation networks. *ICML*, 2015.
- Long, Xi, Yang, Jie, Weysen, Tim, Haakma, Reinder, Foussier, Jérôme, Fonseca, Pedro, and Aarts, Ronald M. Measuring dissimilarity between respiratory effort signals based on uniform scaling for sleep staging. *Physiological measurement*, 2014.
- Lucey, Brendan P, Mcleland, Jennifer S, Toedebusch, Cristina D, Boyd, Jill, Morris, John C, Landsness, Eric C, Yamada, Kelvin, and Holtzman, David M. Comparison of a single-channel eeg sleep study to polysomnography. *Journal of sleep research*, 2016.
- Maaten, Laurens van der and Hinton, Geoffrey. Visualizing data using t-sne. *JMLR*, 2008.
- Makhzani, Alireza, Shlens, Jonathon, Jaitly, Navdeep, Goodfellow, Ian, and Frey, Brendan. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- Metz, Luke, Poole, Ben, Pfau, David, and Sohl-Dickstein, Jascha. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.
- National Institute of Health. Sleep disorders. http://www.ninds.nih.gov/disorders/brain_basics/understanding_sleep.htm#sleep_disorders.
- Pigou, Lionel, Van Den Oord, Aäron, Dieleman, Sander, Van Herreweghe, Mieke, and Dambre, Joni. Beyond temporal pooling: Recurrence and temporal convolutions for gesture recognition in video. *IJCV*, 2015.
- Pollak, Charles P, Tryon, Warren W, Nagaraja, Haikady, and Dzwonczyk, Roger. How accurately does wrist actigraphy identify the states of sleep and wakefulness? *SLEEP-NEW YORK*, 2001.
- Popovic, Djordje, Khoo, Michael, and Westbrook, Philip. Automatic scoring of sleep stages and cortical arousals using two electrodes on the forehead: validation in healthy adults. *Journal of sleep research*, 2014.
- Purushotham, Sanjay, Carvalho, Wilka, Nilanon, Tanachat, and Liu, Yan. Variational recurrent adversarial deep domain adaptation. *ICLR*, 2017.
- Radford, Alec, Metz, Luke, and Chintala, Soumith. Un-supervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Rahman, Tauhidur, Adams, Alexander T, Ravichandran, Ruth Vinisha, Zhang, Mi, Patel, Shwetak N, Kientz, Julie A, and Choudhury, Tanzeem. Dopplesleep: A contactless unobtrusive sleep sensing system using short-range doppler radar. *ACM UbiComp*, 2015.
- Rechtschaffen, Allan and Kales, Anthony. A manual of standardized terminology, techniques and scoring system for sleep stages of human subjects. *US Government Printing Office, US Public Health Service*, 1968.
- Shah, Purav C, Yudelevich, Eric, Genese, Frank, Martillo, Miguel, Ventura, Iazsmin B, Fuhrmann, Katherine, Mortel, Marie, Levendowski, Daniel, Gibson, Charlisa D, Ochieng, Pius, et al. Can disrupted sleep affect mortality in the mechanically ventilated critically ill? *A state of unrest: Sleep/SDB in the ICU and hospital*, 2016.
- Shambroom, John R, Fábregas, Stephan E, and Johnstone, Jack. Validation of an automated wireless system to monitor sleep in healthy adults. *Journal of sleep research*, 2012.
- Sutskever, Ilya, Vinyals, Oriol, and Le, Quoc V. Sequence to sequence learning with neural networks. *NIPS*, 2014.
- Szegedy, Christian, Liu, Wei, Jia, Yangqing, Sermanet, Pierre, Reed, Scott, Anguelov, Dragomir, Erhan, Dumitru, Vanhoucke, Vincent, and Rabinovich, Andrew. Going deeper with convolutions. *CVPR*, 2015.
- Tataraidze, Alexander, Korostovtseva, Lyudmila, Anishchenko, Lesya, Bochkarev, Mikhail, and Sviryayev, Yurii. Sleep architecture measurement based on cardiorespiratory parameters. *IEEE EMBC*, 2016a.
- Tataraidze, Alexander, Korostovtseva, Lyudmila, Anishchenko, Lesya, Bochkarev, Mikhail, Sviryayev, Yurii, and Ivashov, Sergey. Bioradiolocation-based sleep stage classification. *IEEE EMBC*, 2016b.
- Tzeng, Eric, Hoffman, Judy, Darrell, Trevor, and Saenko, Kate. Simultaneous deep transfer across domains and tasks. *ICCV*, 2015.
- Tzeng, Eric, Hoffman, Judy, Saenko, Kate, and Darrell, Trevor. Adversarial discriminative domain adaptation. *NIPS Workshop on Adversarial Training*, 2016.
- Vondrick, Carl, Pirsaviash, Hamed, and Torralba, Antonio. Generating videos with scene dynamics. *NIPS*, 2016.
- Zaffaroni, A, Gahan, L, Collins, L, O'hare, E, Heneghan, C, Garcia, C, Fietze, I, and Penzel, T. Automated sleep staging classification using a non-contact biomotion sensor. *Journal of Sleep Research*, 2014.
- Zhao, Mingmin, Adib, Fadel, and Katabi, Dina. Emotion recognition using wireless signals. *ACM MobiCom*, 2016.