BodyCompass: Monitoring Sleep Posture with Wireless Signals

SHICHAO YUE, Massachusetts Institute of Technology
YUZHE YANG, Massachusetts Institute of Technology
HAO WANG, Massachusetts Institute of Technology
HARIHARAN RAHUL, Massachusetts Institute of Technology
DINA KATABI, Massachusetts Institute of Technology

Monitoring sleep posture is important for avoiding bedsores after surgery, reducing apnea events, tracking the progression of Parkinson's disease, and even alerting epilepsy patients to potentially fatal sleep postures. Today, there is no easy way to track sleep postures. Past work has proposed installing cameras in the bedroom, mounting accelerometers on the subject's chest, or embedding pressure sensors in their bedsheets. Unfortunately, such solutions jeopardize either the privacy of the user or their sleep comfort.

In this paper, we introduce BodyCompass, the first RF-based system that provides accurate sleep posture monitoring overnight in the user's own home. BodyCompass works by studying the RF reflections in the environment. It disentangles RF signals that bounced off the subject's body from other multipath signals. It then analyzes those signals via a custom machine learning algorithm to infer the subject's sleep posture. BodyCompass is easily transferable and can apply to new homes and users with minimal effort. We empirically evaluate BodyCompass using over 200 nights of sleep data from 26 subjects in their own homes. Our results show that, given one week, one night, or 16 minutes of labeled data from the subject, BodyCompass's corresponding accuracy is 94%, 87%, and 84%, respectively.

CCS Concepts: • **Human-centered computing → Ubiquitous and mobile computing systems and tools**.

Additional Key Words and Phrases: Sleep, Sleep Posture, Wireless Sensing, Contactless Sensing, Healthcare, Domain Adaptation, Transfer Learning, Machine Learning, Deep Learning

ACM Reference Format:
Shichao Yue, Yuzhe Yang, Hao Wang, Hariharan Rahul, and Dina Katabi. 2020. BodyCompass: Monitoring Sleep Posture with Wireless Signals. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 2, Article 66 (June 2020), 25 pages. https://doi.org/10.1145/3397311

1 INTRODUCTION

Each of us has our favorite sleep postures: sleeping on the right side, left side, facing up, or facing down. Significant clinical research has shown that sleep posture is a valuable marker of disease progression, and has a significant impact on health. For instance, patients with Parkinson's disease often suffer from loss of axial movement; and less frequent nocturnal turnovers and longer periods spent recumbent or supine (*i.e.*, facing up) are associated

Authors' addresses: Shichao Yue, Massachusetts Institute of Technology, 32 Vassar St. Cambridge, MA, 02139, scyue@mit.edu; Yuzhe Yang, Massachusetts Institute of Technology, 32 Vassar St. Cambridge, MA, 02139, yuzhe@mit.edu; Hao Wang, Massachusetts Institute of Technology, 32 Vassar St. Cambridge, MA, 02139, hwang87@mit.edu; Hariharan Rahul, Massachusetts Institute of Technology, 32 Vassar St. Cambridge, MA, 02139, rahul@csail.mit.edu; Dina Katabi, Massachusetts Institute of Technology, 32 Vassar St. Cambridge, MA, 02139, dina@csail.mit.edu.

2474-9567/2020/6-ART66 $15.00
https://doi.org/10.1145/3397311

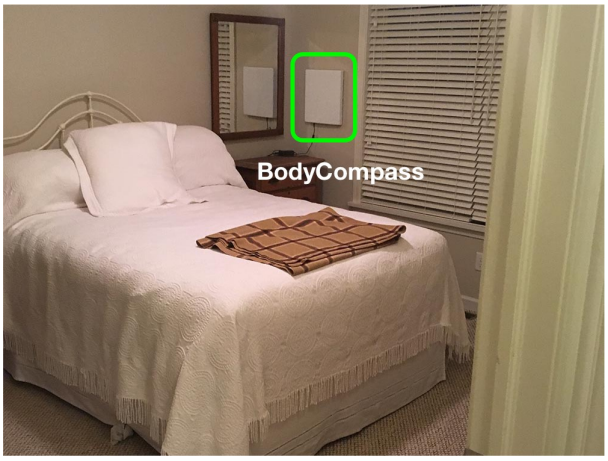

Fig. 1. BodyCompass in one of our deployments. The white box mounted on the wall is the radio. It uploads the RF signals to the cloud where the model processes them to extract sleep posture.

with deterioration in the condition of Parkinson's patients [38]. Similarly, infrequent changes in sleep posture can lead to pressure ulcers in the elderly and post-surgery patients [10]. Studies have also demonstrated that sleeping in a supine position can reduce back pain since it is the position in which the muscles have the least amount of work to do to maintain one's posture against the force of gravity [8]. In contrast, if one has obstructive sleep apnea (OSA), the supine position becomes the worst posture because it imposes unfavorable airway geometry and reduces lung volume. Studies have shown that more than half of all OSA cases can be classified as supine related [28, 29]. Improper sleep posture can even be fatal – sleeping on the stomach can boost the risk of sudden infant death syndrome (SIDS) [9] and sudden death in epilepsy patients [17, 20]. These examples highlight the importance of continuous and fully automatic sleep posture monitoring. Such monitoring can provide doctors with information to better manage patient conditions; it can also provide people themselves information to adjust their posture and reduce their health risks.

Unfortunately, today, there is no good way to provide such sleep posture monitoring. Doctors typically resort to asking patients about their sleep posture, an error-prone mechanism since people routinely, and unknowingly, change their postures while sleeping. Automated monitoring systems primarily fall into two categories. The first category is vision-based. These methods use a camera to monitor the user's sleep, then extract postures from recorded videos with a machine learning system. Deploying cameras in people's bedrooms, however, is privacy-intrusive. Furthermore, cameras have difficulties tracking body posture if the person is covered or lighting is bad, both of which are typical scenarios when sleeping. The second category uses various kinds of on-bed sensors. Such methods require the user to fix the sensor to the surface of the mattress, which can affect sleep comfort.

Ideally, one desires a system that is non-contact, non-intrusive, and works even in dark scenarios typical of sleeping conditions. In this paper, we present *BodyCompass*, an RF-based system for sleep posture monitoring. BodyCompass analyzes the reflections of RF signals to infer subjects' sleep postures. It does so without requiring users to wear or be in contact with any sensors. It is not invasive of privacy, and can also work in the dark. Unlike much previous work, BodyCompass has also been demonstrated to work in the wild - with real subjects sleeping in their own homes, and can generalize to new environments with minimal additional training. Fig. 1 shows BodyCompass deployed in the home of one of our users.

But how can one extract the sleep posture from radio signals? Our idea is to use the multipath effect, a known phenomenon in RF communication systems that refers to the fact that RF signals bounce off different objects and obstacles in the environment and reach the receiver through multiple paths. Past work has shown that the human body acts as a reflector in the low-GHz frequencies, commonly used in commodity radios [52]. As the RF signal is incident upon the human body, it reflects from the body based on the body orientation and bounces off the surrounding objects and walls creating a multipath signature indicative of the body posture. Our objective is to learn an inverse map that receives the reflected multipath profile and attempts to infer the body posture. A key challenge in delivering this idea is that the RF signal bounces off many objects in the environment, not just the human body. Only a subset of the signal path involves reflections from the human body, and hence is relevant to the sleep posture. Thus, one has to extract only the RF reflections that bounced off the human body either directly or indirectly in order to determine the sleep posture.

To address this challenge, we leverage past work that shows how to extract a person's breathing from RF signals. Our intuition is that all paths that bounce off a person's trunk (e.g. chest and belly) during their sleep are modulated by the person's breathing, and hence we can use this property to disentangle these reflections from the rest of the reflections. Specifically, we use standard techniques to separate signals along different paths (FMCW and angle of arrival [1]), and correlate these separated signals individually with the subject's breathing signal to identify the specific signals corresponding to the person in bed. We further design a neural network model that takes this breathing filtered multipath profile, and predicts the sleep posture of the person.

A key question with such a system is how well the neural network model works with different people and in different homes. While RF reflections and the multipath effect naturally depend on the environment, one would hope that with proper design, the model would be able to transfer some of the knowledge across environments. Such a model would learn the underlying features that identify each sleep posture, and tune them to a new environment with a small amount of additional labeled data from that environment. To address this issue, we design our model to be easily transferable. Specifically, given a set of source domains *i.e.,* a number of people and their sleep postures in the training set, and a target domain *i.e.,* a new person in his own home, the model can use a small amount of labeled data (16 minutes to one night) from the new home to optimize its performance for this new environment.

Our model delivers high accuracy. Specifically, our basic sleep posture model using multipath, when trained and tested on the same person and home, achieves an accuracy of 94.1%. The transfer learning model to a new person and a new home has an accuracy of 86.7% with one night of labeled data, and 83.7% with a labeled dataset comprising 8 examples, where in each example, the person lies down in one of his typical sleep postures for a duration of 2 minutes.

To summarize, this paper makes the following contributions:

(1) We present BodyCompass, the first RF-based system that provides accurate sleep posture monitoring overnight in users' own homes. It achieves high accuracy without sacrificing privacy and sleep comfort.
(2) BodyCompass can transfer its model to new homes and users with very little additional training data.
(3) We implement and evaluate BodyCompass extensively in real world settings using data from 26 homes with 26 different subjects and more than 200 nights of sleep.

2 RELATED WORK

Past work on sleep posture monitoring can be divided into two major categories: 1) systems with on-body sensors, and 2) non-contact monitoring systems.

(a) On-body Solutions: On-body sensors can monitor sleep postures accurately [5, 19, 48]. For example, one may attach an accelerometer to the person's chest to monitor their sleep posture. Since gravity always points downwards, the accelerometer's orientation can be calculated by combining the acceleration along three different

axes [5, 48]. However this method is cumbersome and uncomfortable since the accelerometer needs to be fixed on the user's body during their sleep.

(b) Non-Contact Solutions: Contactless systems are more comfortable for the user compared to on-body sensors. Work in this class falls in the following categories. First, vision-based systems [3, 11, 24] deploy RGB or infra-red cameras to record videos of the user's sleep, then process those videos using convolutional neural networks to predict sleep postures. However, cameras, particularly in people's bedrooms, are privacy-intrusive. Further, the accuracy of camera systems decreases significantly in dark settings and when people are covered with a blanket or comforter [3, 11, 24].

Fig. 2. Pressure sensitive bedsheet from [23].

Second, on-bed sensors cover the mattress with an array of pressure sensors [13, 23, 31, 32, 47] or RFID Tags [12, 21]. These solutions are more privacy-preserving than camera systems. However, on-bed sensors, shown in Fig. 2, change the feel of the bed and thus affect the sleep comfort of the subject. Further, most of these systems are evaluated in the lab, as opposed to overnight testing in people's own homes [12, 13, 21, 31, 32, 47].

Third, a few papers have proposed the use of RF signals for monitoring sleep posture [4, 22, 25]. The approach in those papers is intrinsically different from ours; they analyze the signal power as measured by the RSSI (received signal strength indicator) [4] or the power of the frequency sub-channels extracted from the CSI (channel state information) [22, 25]. As studied in [14], they all inherently suffer from interference. That is, they have no ability to separate changes in the signal that are due to the sleeping person from those due to other sources of motion (e.g. a fan, or a person moving in a neighboring room). Such extraneous motion brings randomness and will greatly hamper the robustness of the system in the wild. As a result, all previous papers are evaluated in a single lab environment with one or two subjects consciously performing specified postures.[1] In contrast, we study the spatial pattern of reflections –i.e., the multipath – and ignore the power by re-normalizing the power distribution of each path (see Section 4). Therefore, our system can provide accurate sleep posture monitoring overnight in users' homes and can be easily transferred to new environments.

We also note that past work has demonstrated the feasibility of inferring the human skeleton using only RF reflections [52, 53]. It might seem that one could use such models to infer the skeleton of the person lying in bed and hence their sleep posture. However, due to RF specularity, such models rely on people walking and moving around to achieve good accuracy [52, 53]. Specifically, as described in those papers, a snapshot of RF signal reflections does not capture the full body; Any snapshot captures only a few limbs or body parts that reflect signals directly towards the radio. Hence, their neural networks rely on people moving and walking to expose different body parts in each snapshot so that the network can combine those body parts to create the human skeleton. In contrast, when the person is asleep in bed, the person is mostly static and hence there is not enough motion to allow the neural network to fill in the gaps and combine body parts across different snapshots. To deal with this challenge, our system not only takes the direct reflections towards the radio, but also all the indirect reflections due to multipath. By taking all the multi-path reflections as input, our system estimates the sleep posture accurately even when the person remains static.

Finally, this paper belongs to a growing body of research that focuses on passive monitoring using radio signals. Researchers have demonstrated that by carefully analyzing RF reflections off the human body, they can monitor

[1] We note that while the authors of [22] test their vital sign algorithms outside the lab, the sleep posture is only tested in the lab and in one setting.

people's location [1, 45, 46], gait [15, 41, 50], breathing [30, 43, 49], heart rate [2, 26, 42], falls [37, 39, 44], and sleep quality and stages [14, 33, 36, 54]. Our work builds on this foundation and leverages past work on inferring the breathing signal as a sub-component in our system [49].

3 BODYCOMPASS

BodyCompass is the first RF-based system that provides accurate sleep posture monitoring in the wild, *i.e.*, with subjects sleeping in their own beds in their homes, and it generalizes to new subjects and homes with minimal additional effort. It can be used by healthy individuals interested in monitoring their sleep behavior, or can be provided either to patients to help them modify their sleep posture, or to doctors to assist them in understanding disease prognosis and patient health.

BodyCompass leverages measurements from an FMCW radio equipped with an antenna array [27]. Such radios are commonly used in passive health monitoring using RF signals [14, 33, 49]. They work by transmitting a low power radio signal, and observing its reflections from the surrounding environment. The use of an antenna array combined with FMCW enables the radio to resolve RF reflections from multiple points in space. Specifically, at each instance in time, the radio outputs an array of signal values from various voxels in space, which we refer to as an *RF-snapshot*.

BodyCompass takes a sequence of RF-snapshots from an FMCW radio across a whole night, and produces the sleep postures for the night. A sleep posture is described by an angle between two normal vectors, one of the bed surface and one of the user's anterior trunk surface, as shown in Fig. 3a. For example, 0° represents the user facing upwards and 90° represents the user facing rightwards. Defining sleep posture in terms of angle allows us to differentiate between a slight tilt of the trunk to the right and someone sleeping on their right side. This enables a finer granularity definition of sleep postures that encompasses and expands beyond common posture classes (supine, left side, right side, prone). A fine granularity in posture estimation is important for applications that aim to detect changes in postures, such as tracking the progression of Parkinson's patients by monitoring the frequency of their change of sleep posture.

BodyCompass computes sleep postures using three components:

- A filtered multipath profile feature extractor to estimate the RF reflections that bounced off the person directly or indirectly.
- A source-specific neural network that utilizes the multipath profile features to estimate the sleep posture of a specific person in a specific home.
- A transfer learning model that adapts the source-specific models to estimate the sleep posture of a new person in a new home with minimal additional labeled data.

Below, we describe these components in detail.

4 FILTERED MULTIPATH FEATURE EXTRACTOR

In this section, we describe how BodyCompass extracts filtered multipath features specific to a person from RF-snapshots produced by a FMCW antenna array.

An RF-snapshot consists, for each point of space (RF voxel), of the magnitude of the RF reflection from that point of space. An RF voxel is represented by two coordinates, its distance from the device, and the angle of that position relative to the normal from the device. Specifically, an RF voxel at coordinate (i, j) represents a small cube around the point at traveling distance d_j from the device, and an angle of arrival of α_i, as shown in Fig. 4. We divide the space into N angles, and M distances, and therefore each RF-snapshot is an $N \times M$ matrix. For better visualization, we plot voxels in a standard Cartesian coordinate system, as shown in Fig. 3b, instead of a polar coordinate system.

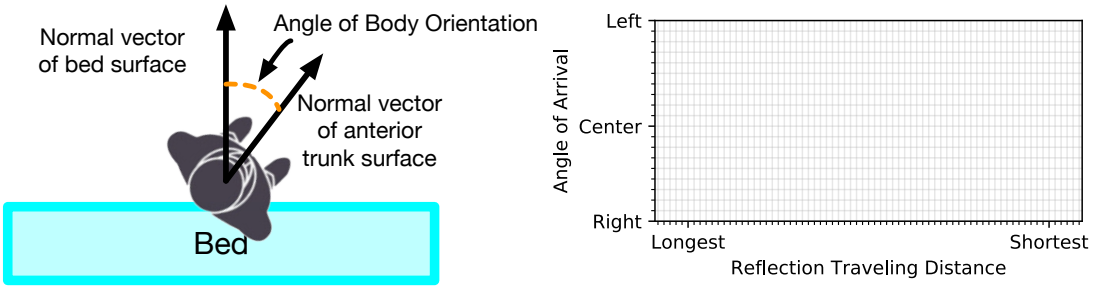

(a) Angle of body orientation

(b) Coordinates of RF-snapshots

Fig. 3. Illustration of body orientation and coordinates of RF-snapshots. In Fig. 3b, we represent all RF voxels in a Cartesian coordinate system. The rightmost pixels are closest to the device, and have the shortest distance. Therefore the direct-path reflections should be to the right of the indirect-path reflections since they travel the shortest path between the user and the device.

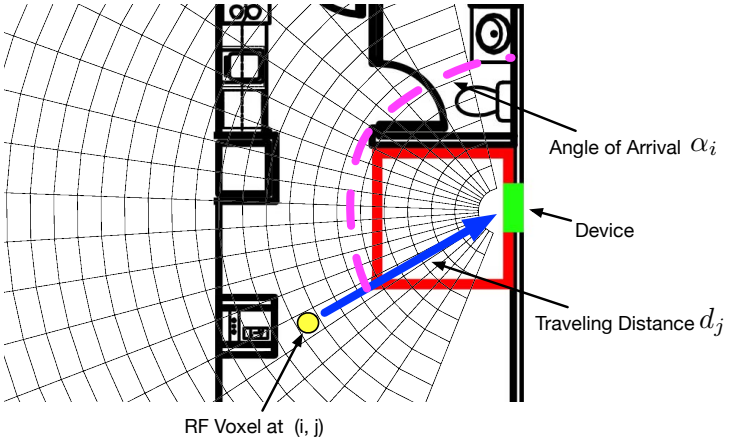

Fig. 4. Illustrative example of an RF voxel. The green rectangle represents the radio location and the red rectangle represents the bed location. Our FMCW antenna array divides the space into small grids. The coordinates of the grid represent the distance from the radio and the angle of arrival.

4.1 Stable Sleep Periods

We first note that sleep postures are not independent over time, since people typically sleep in a posture for some period of time, followed by a movement, after which they settle into a different sleep posture, and so on. BodyCompass therefore first segments the night into a series of stable sleep periods. During each stable period, the orientation of the body is approximately constant, and BodyCompass extracts a single sleep posture from that period. BodyCompass leverages prior work [49] to identify motion events from RF-snapshots, and defines the intervals between such motion events as stable sleep periods.

4.2 Filtered Multipath profile

Next, BodyCompass extracts the multipath profile for each stable sleep period, with the objective of learning the sleep posture from the multipath profile. Recall that the multipath profile captures the pattern of spatial reflections, *i.e.*, how the RF signals bounce around in space before they reach our radio. We represent the multipath profile of a particular stable period with the relative signal power along each path. Thus, the multipath profile of a particular stable period can be computed by taking the RF-snapshots corresponding to the stable sleep period and computing the variance in each voxel.

The multipath profile is affected by the sleep posture of the person, and is therefore informative about their orientation. For instance, when the person is supine, *i.e.*, lying flat on their back, a significant portion of the signal reflects towards the ceiling and bounces off other objects in the environment before reflecting back to the radio, and as a result the multipath profile shows significant dispersion. In contrast, when the person is sleeping on their side, the direct RF reflections will be significantly stronger than the indirect reflections, and the multipath profile will therefore show high concentration.

However, one cannot directly use the overall multipath profile in a stable sleep period to infer sleep posture. This is because such multipath profile contains reflections both from the environment and from the subject. While reflections from static reflectors (*e.g.*, walls, tables) can be removed,[2] reflections from moving objects cannot be easily disentangled, and their contributions can confound the system for two reasons. First, even within a single home, those contributions can change over time even when the sleep posture does not change, for instance, because of movements from a fan, people walking in the environment, or heating, ventilation, and air conditioning (HVAC) systems. Since these changes are not correlated with the sleep posture of the person, they will adversely affect the ability of BodyCompass to infer sleep posture. Furthermore, such reflections are highly specific to each home, and incorporating them into the multipath profile will prevent BodyCompass from generalizing to new homes.

So, how does one filter out environmental contributions while still retaining the multipath contributions from the sleeping subject? Our idea is inspired by the following observation: when breathing, the chest and belly area of the human body move forward and backward. These motions will change the multipath contributions corresponding to the human body in a manner correlated with the breathing signal, while other environment related multipath contributions will not change in a manner correlated with the person's breathing. Fig. 5 shows an illustrative example of this point with a person sleeping facing the device, and a nearby fan.

Using breathing also allows BodyCompass to identify the orientation of the person, *i.e.*, whether the person is facing up/down when supine, and whether the person is facing towards/away from the device when on their side. This is because during breathing, only the front of the human body moves significantly, whereas the back does not, therefore breaking symmetry between the orientations, and changing the *filtered* multipath profiles in the two cases.

BodyCompass uses DeepBreath [49], to extract the breathing signal of the subject in bed from the RF-snapshots in the stable period. The breathing signal is a time series that reflects the scaled chest displacement of the subject over time. Then, for each RF-voxel, BodyCompass correlates this extracted breathing signal with the time series of signal magnitudes for this RF-voxel obtained from the RF-snapshots. Specifically, for each voxel, BodyCompass computes the absolute value of the Pearson correlation coefficient between the person's breathing and the magnitude of the RF signal received from that voxel, as expressed in the sequence of RF-snapshots. This correlation provides a spatial filter that allows us to extract the voxels in the multipath profile whose signal is highly correlated with the person's breathing.

Next, BodyCompass multiplies the multipath profile with the above filter to extract the filtered multipath profile, which focuses on the signals that bounced directly or indirectly off the subject. The filtered multipath

[2]We remove static reflections by subtracting the average RF-snapshot for each stable period before computing the multipath profile [49].

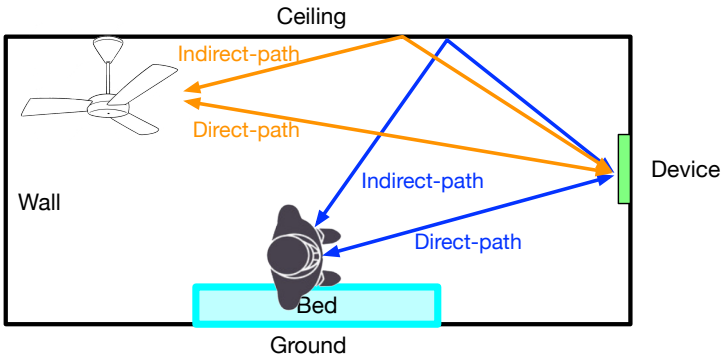

Fig. 5. An illustrative example of signal reflections. In this case the multipath profile contains reflections from the subject (in blue) and a ceiling fan (in orange). Correct processing requires eliminating the fan reflections. Also the figure illustrates how the posture could affect the multipath. When the user is facing the device, the reflections along the direct path have the largest signal variations because the chest movements are most significant in that direction. In contrast, the signal variations along the indirect path are much smaller because the side of the body is not moving significantly.

profile emphasizes pixels with a significant contribution from the subject's breathing while still retaining the relative power contributions from direct and indirect paths corresponding to that breathing.

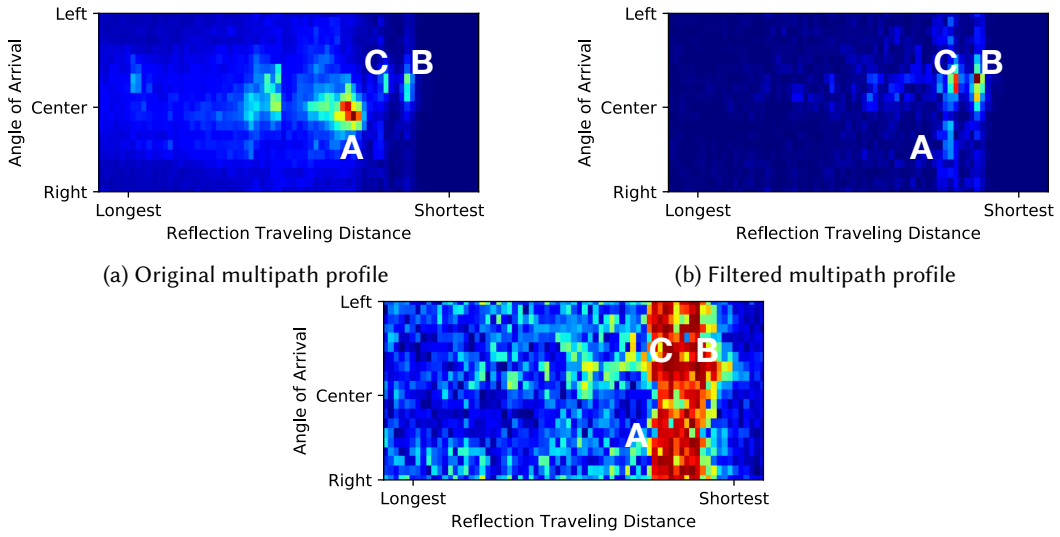

(a) Original multipath profile

(b) Filtered multipath profile

(c) Filter used to extract the multipath associated with the subject. The filter consists of the correlation coefficients with breathing

Fig. 6. Visualization of one stable period. The value of each RF-voxel represents the corresponding attribute of the RF reflection. The visualization is color-coded, the redder the pixel, the higher the relative value of that pixel, and the bluer the pixel, the lower the relative value. Points A, B, C highlight three different kinds of reflections: environmental movement reflections, breathing reflections along direct-path, breathing reflections along indirect-path.

To help understand the process, we plot an example in Fig. 6. Specifically, Fig. 6a shows the relative power in each voxel in space in the original multipath profile. Fig. 6c shows the breathing filter, and Fig. 6b shows the filtered multipath profile. As we can see, pixel A is very bright in the original multipath profile (Fig. 6a), meaning that it has very a high reflection power relative to other pixels. However, since it has a very low correlation coefficient with breathing (Fig. 6c), it is removed in the final filtered multipath profile (Fig. 6b). In this case, pixel A was contributed by environmental movements from a different person walking in the environment. In contrast, while pixels B and C have lower power compared to A, they are emphasized in the filtered multipath profile since they exhibit strong correlation with the breathing signal. It is also worth noting that pixel B is actually the direct-path reflection, and pixel C is one of the indirect-path reflections.

Next we show two typical filtered multipath profiles of two different postures in Fig. 7. When the user is sleeping in a supine position, the filtered multipath profile shows more dispersion because the subject reflects a significant part of the signal towards the ceiling causing more indirect reflections. In contrast, when the user is facing the device most of the signal is directly reflected from the user to the device and hence the power in the filtered multipath profile is concentrated at the user's location, i.e. the direct path.

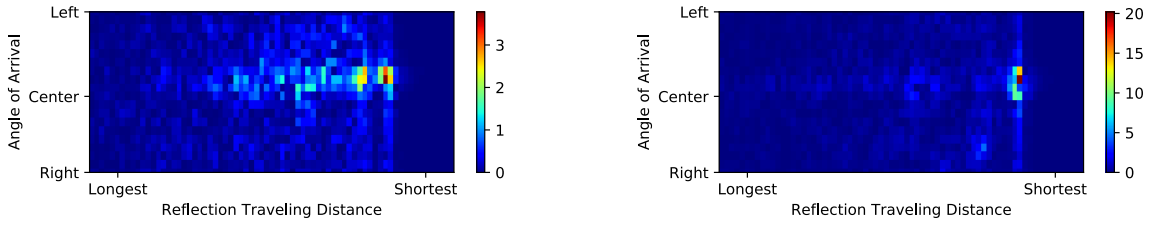

(a) The user facing up

(b) The user facing the device

Fig. 7. Two typical examples of filtered multipath profiles of the user facing up (Fig. 7a) and facing towards the device (Fig. 7b). Compared to Fig. 7b, in Fig. 7a, we can see much higher power in further away pixels. This is because when the user is facing up, he deflects the signal towards the ceiling causing indirect reflections.

5 SOURCE-SPECIFIC SLEEP POSTURE MODEL

Having computed the filtered multipath profile for each stable sleep period corresponding to a source (*i.e.* a user in a specific home), BodyCompass then uses a neural network to predict the sleep posture for that source during each stable sleep period.

Our model uses a multi-layer fully-connected neural network. We deliberately choose a fully-connected neural network instead of the commonly used convolutional neural network (CNN). CNNs are more suitable for natural images because one typically needs to compare each pixel with the pixels in its neighborhood. In contrast, to capture the multipath profile, one needs to compare pixels globally. (See Sec. 8.4 for an empirical comparison of the performance of a fully-connected network and a CNN on this task.)

For training the neural network, we ask subjects to wear accelerometers to collect the ground-truth angular orientation of the body. Detailed ground truth collection process is described in Sec. 7.2. Recall that we express the sleep position in terms of the angle specifying the trunk rotation with respect to the bed. BodyCompass averages the angular values that the accelerometer measures during a stable period to obtain the ground truth sleep posture of the subject in that period.

BodyCompass trains the neural network to predict the sleep angle associated with each filtered multipath profile. To train the network we need to compare the predicted angle with the ground truth angle. Directly

comparing angles however leads to discontinuity since angles wrap around, i.e., 0° and 360° are the same angle, but simply computing their difference will yield a large loss.

Therefore, in order to ensure smoothness of the loss function, BodyCompass's model predicts complex numbers, and uses the phase of the complex number as the angle prediction. Specifically, we define *Circular Loss* as follows:

$$\mathcal{L}_c(\theta) = \mathbb{E}_{x,y\sim p(x,y)}\, arccos\left(\frac{Re(F(x,\theta)\cdot e^{-iy})}{|F(x,\theta)|}\right) \tag{1}$$

where x is the input feature vector (i.e., the filtered multipath profile of a stable segment), y is the ground-truth angle, $F(\cdot,\theta)$ is the model that maps a feature vector into a complex number, θ is the model parameters (the wights of the neural network), arccos denotes the arc cosine function, and $\mathbb{E}$ is the expectation.

The operand of $arccos$: $\frac{Re(F(x,\theta)\cdot e^{-iy})}{|F(x,\theta)|}$ can be interpreted as the cosine similarity between two vectors: one is our prediction $(Re(F(\cdot)), Im(F(\cdot))$ and the other one represent the unit vector of the ground truth angle $(cos(y), sin(y))$. The similarity reaches its maximum when the predicted vector has the same angle as the ground truth (an $arccos$ of 0). And it reaches its minimum when these two vectors are diametrically opposed (an $arccos$ of π radians). This loss function solves the discontinuity problem since it computes the angle difference between our prediction and the ground truth in a differentiable way.

6 TRANSFERRING THE MODEL TO NEW USERS

In the previous section, we explained how to train a model to predict a user's sleep posture accurately given abundant labeled data from the that user. However, data collection is a laborious and time-consuming task for both the user and the operator of the system. Ideally, we would like our system to perform well on new users with minimal effort.

Since the properties of RF signals (power, phase, and multipath) depend on the environment, transfer between different homes is a challenging task. In order to achieve satisfactory performance while reducing the burden on the user, we assume that only limited labeled data from a new user is available. We refer to such labeled examples (where an example is a filtered multipath profile and its correct sleep angle) as *Calibration Points*. As described in Sec. 8.2, the number of calibration points can be as few as 8 examples, each lasting for 2 minutes, for a total of only 16 minutes.

Given the scarcity of the calibration points, it is not practical to train a model entirely based on those points. Instead, we formulate the task as a semi-supervised domain adaptation problem: we have multiple source users, each with abundant labeled data, and a target user for which we have a few calibration points. We would like the system to achieve high accuracy on the target user given the above information.

6.1 Overview of the Transfer Model

Given a set of source domains i.e. a number of people and their sleep postures in the training set, and a target domain i.e. a new person in their own home, we design a model that learns from the training data of the source domains how to infer sleep posture in the target domain, with a small number of calibration points.

At a high level, our transfer model first preprocesses the training data to ensure that the probability distributions of the source domains look as close as possible to the target domain. Next, given that the amount of data from the target domain is not sufficient to train a model, we try to augment the data from the target by selecting data points from the source domains that look similar to the target data, both in terms of its feature map (*i.e.*, filtered multipath) and the corresponding posture. We use this augmented data to create a virtual target which is similar to the original target but has much more labeled data. Now we can adapt the model from each source domain to work well on the virtual target. The final prediction is then performed by majority voting over all of these adapted models.

In the following sub-sections, we expand on this high-level description, providing the details of the three key components of our transfer model:

(1) **Distribution Alignment**: We explicitly align the data distributions across users, whose differences are caused by different room layouts.

(2) **Data Augmentation**: We generate augmented data by picking data points from source subjects that resemble the calibration points.

(3) **Ensemble Learning**: We use Majority Voting to generate one final prediction that is robust and accurate.

6.2 Distribution Alignment

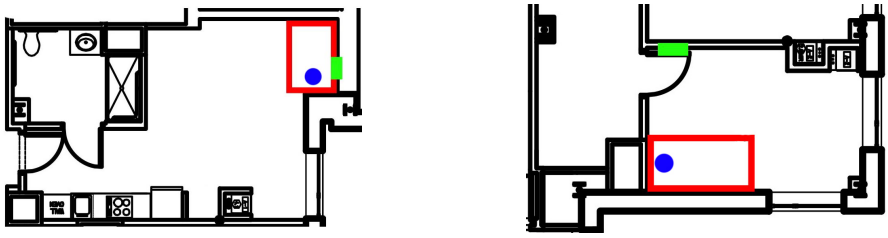

(a) User A's bedroom layout (b) User B's bedroom layout

Fig. 8. Bedroom layouts of two users. Green rectangle shows the location of our device. Red rectangles shows the location of the bed and blue circle shows the position of the pillow.

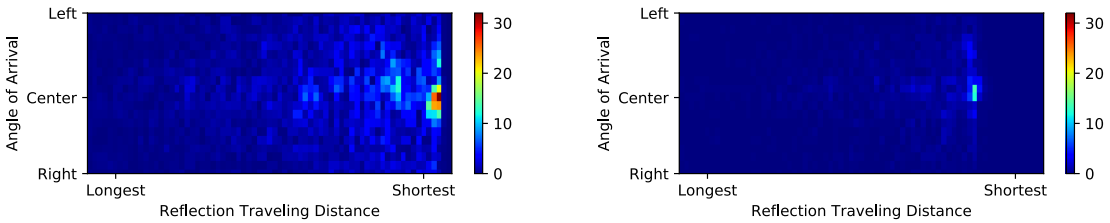

(a) User A's profile when facing to the device (b) User B's profile when facing to the device

Fig. 9. Examples showing how the bed position with respect to the radio affects the signal's strength and location. The figures show that due to differences in the position of the bed with respect to the radio, User A's direct path signal is much stronger and closer to the radio compared to User B's.

Without any alignments, the model's generalization ability will be greatly hampered by the distribution shift between the source user and the target user. One major reason of this distribution shift is caused by different room layouts. For example, we have User A and User B with their floor plan visualized in Fig. 8. User A's bed is very close to the device (the device is right above the bed), and in comparison, User B's bed is far from the device. We show the multipath profiles of the two users in in Fig. 9 (a) and (b). As clear from these figures, the difference in the bed location with respect to the radio impacts the filtered multipath profile in two ways. First, the direct-path reflection of User B needs to travel a longer distance compared to User A. Thus pixels at the same location in the multipath profiles are not directly comparable. Second, the power of RF reflections decreases as

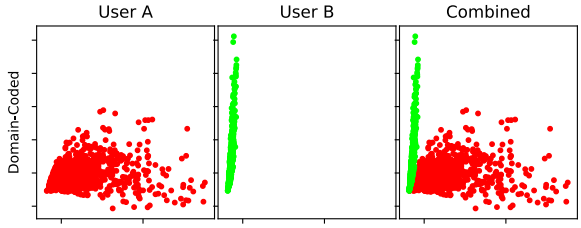

Fig. 10. Visualization of data distribution of User A (red) and User B (green). Since the feature map (i.e., the filtered multipath profile) is high-dimensional, to visualize the data in a two-dimensional space we perform joint Principal Component Analysis (PCA) on all the feature vectors (all filtered multipath profiles) from both A and B using the same set of basis, and plot the data with respect to the two largest principle components. We plot the data of the source and target separately in the left two figures and combined in the right figure. As we can see, the distributions of two users are mismatched significantly.

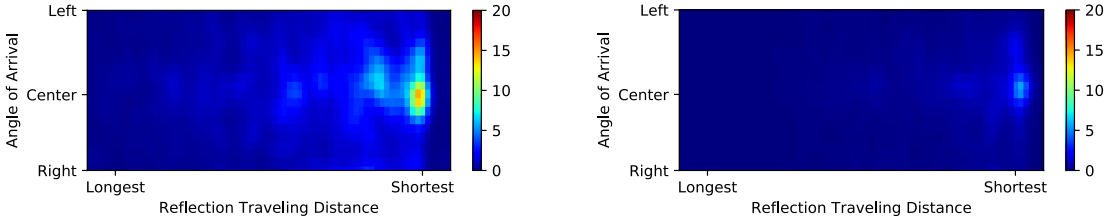

(a) User A's profile after aligning bed location

(b) User B's profile after aligning bed location

Fig. 11. The multipath profiles in Fig. 9 after aligning bed locations. Now the direct path pixels of both Users A and B are at the same location.

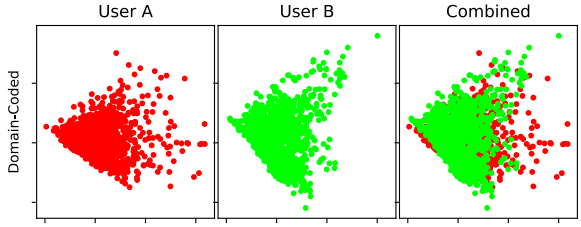

Fig. 12. Visualization of data distribution of User A (red) and User B (green) after bed alignment and power normalization. Compared to Fig. 10, we can see that the two distributions are much better aligned.

the traveling distance increases. Therefore the breathing powers of User B's pixels are much smaller than User A. As a result, and as shown in Fig. 10, the data distributions of User A and User B are significantly mismatched. Below, we explain the two methods we use to align the distributions.

(a) Aligning Bed Locations: We first align the relative location between the bed and radio across all users. While not all differences can be eliminated, this way, we ensure that all direct-path reflections have the same traveling distance. And since we cannot ask every user to move their bed, we can do this alignment virtually by reducing or increasing traveling distances of all the RF reflections in the multipath profile. This method brings us

two advantages: 1) it keeps the original RF reflection pattern, and 2) we only need to know the bed location to perform this alignment.

Identifying the bed position can be done manually, however it is tedious and error-prone. Instead, we propose a robust and accurate way of measuring bed location for aligning. For each user, we do a pixel-wise summation for all of his filtered multipath profiles. Since the direct signal path has the shortest traveling distance, usually it has the highest breathing power in the filtered multipath profiles. Therefore, the pixel with the highest sum will give us an accurate estimation of the bed location. We additionally apply a Gaussian filter with a sigma of 1 to erase small location mismatches.

Looking back to Fig. 8, there is another mismatch. In User A's case, the radio device is on the right-hand side of the user, and for User B, the device is on his left-hand side. This indicates that, when two users are both facing the device, they are actually facing different directions (to the right in Fig. 9a and to the left in Fig. 9b). Therefore we also align directions by flipping the angles. In all of the following discussions and results, -90° (left) represents the direction facing the device (to the right for User A and to the left for User B), and 90° (right) represents the opposite direction.

Fig. 11 shows the filtered multipath profiles of Users A and B above after aligning the bed location. We can see that the direct path pixels of both A and B are now at the same location.

(b) Power Normalization: RF signals attenuate with distance. Hence, the power in a particular pixel in the filtered multipath profile depends on the path length, more so than the sleep posture, as shown in Fig. 11. This dependence can prevent model generalization to a new target user if the target's room layout differs from the source user. Thus we would like to eliminate this dependence. To deal with this issue, we normalize the power distribution in each pixel of the filtered multipath profile (i.e., for each data point, we subtract the mean of the distribution and divide by its standard deviation).

Data distributions for User A and B after both aligning the bed location and normalizing the power are plotted in Fig. 12, and they are aligned much better compared to Fig. 10.

6.3 Target Data Augmentation

Given the small number of calibration points, our information about the target user is limited. One solution is to perform data augmentation. In computer vision tasks, researchers have long been using augmentation techniques such as cropping, rotating, and horizontal flipping to help the model to capture data invariances. After those augmentations, images are still valid images. However, this is not true for our multipath profile. For example, flipping means that the furthest pixel becomes the closest, and the longest indirect reflections becomes the direct-path reflection. Therefore such standard augmentation techniques for images will break the spatial structure of the multipath profile.

Instead, we use data points from the source users that are similar to the calibration points from the target user. Specifically, our augmentation process contains the following steps: First, we align all the data points from all the users (including the calibration points), as described in the previous section. Then given one calibration point (x_0, y_0), we first select all the points (x_i, y_i) that satisfy the condition that the angle difference between y_0 and y_i is smaller than a certain threshold (the default is 20 degrees). Then for those selected points, we further sort them based on the similarity of their multipath profiles to the calibration point as captured by the L2 distance: $\|x_i - x_0\|_2$. Finally, we pick the data points most similar to the calibration point (specifically we pick the 30 most similar source points to each calibration point). We refer to the set of augmented data points as the virtual target.

When adapting the neural network model from a particular source user to the target user, we combine the augmented data points with the labeled data from that source user to train the model and improve the model's performance on the target user. Note that we do not use the calibration points in training any of the adapted

models. We hold the calibration points and use them to select the most effective adapted models as explained in the next section.

6.4 Majority Voting

So far, we adapted each of the source models to transfer its knowledge to the virtual target. However, some of the source users in our training set may be very different from the target user, and despite adaptation, their knowledge may not translate well to the target user. Thus we need a mechanism to detect models that are well adapted to the target and combine their predictions.

We define validation accuracy as the model's accuracy on the target's calibration points. This validation accuracy is an estimate of the adapted model's true accuracy on the target data. Thus, we use this validation accuracy to evaluate the source's compatibility with the target. We filter out models that have bad validation accuracy (accuracy worse than 10% compared to the best model), and perform a majority vote among the models with good accuracy.

The majority vote is performed as follows: We create a histogram of the predictions where each angle degree has its own bin. Then we smooth this histogram with a Gaussian filter with a standard deviation of 20. Finally we pick the angle that has the highest value after smoothing as our final predicted angle. It is worth mentioning that, the smoothing is performed in a circular way, *i.e.*, the last bin and the first bin are connected.

7 EXPERIMENT SETUP

7.1 Data Collection

All experiments with human subjects were approved by our IRB, and we have obtained informed consent from every subject. In total, we have collected 224 nights of data from 26 subjects (17 male subjects and 9 female subjects).[3] Subjects' bed sizes cover most common sizes, from twin-size (1m wide) to king-size (2m wide). Each subject sleeps alone in his/her own bedroom. For each subject, we install the radio device on the wall to the side of the bed. The distance between the bed and the radio ranges from 0 meters (right above the bed) to 4 meters (on the other side of the bedroom).

7.2 Ground Truth Collection

To collect the ground truth postures, for each night, we ask the subject to wear two accelerometers, one on the chest and one on the abdomen. Both accelerometers are fixed on the body using sport tapes to prevent sliding during sleep. In Fig. 13, we show the placement of accelerometers on the body.

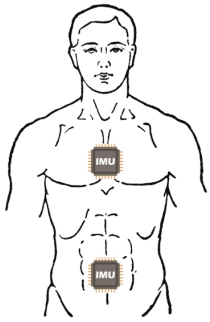

Fig. 13. The placement of accelerometers on the subject's body. The accelerometers are used to collect the ground truth posture.

We align both accelerometers so that the accelerometer's x's positive points towards the subject's head, and z's positive points opposite to the body. Then the angle of body orientation y can be calculated using the following equation [34]:

$$y = atan2(a_y, a_z)$$

where $atan2$ is the 2-argument arctangent function.

We use two accelerometers for mutual validation. There are many factors that can lead to poor data quality. Most commonly, sensors may fall off from the body during sleep, or the subject may fail to align the sensors correctly. By having two accelerometers, we can identify bad data points because readings from those two will no longer be equal. Such bad data points are then excluded from training and testing.

[3] After data cleaning described in Sec. 7.2

7.3 Radio Specification

We use a standard FMCW radio similar to the one used in past work [2, 40, 49, 51]. The radio sweeps the frequencies from 5.4 GHz to 7.2 GHz and transmits at sub-milliwatt power in accordance with the FCC regulations. The radio is equipped with two antenna arrays: a horizontal array and a vertical array. Each array is capable of dividing space into 20 (angle) by 71 (distance) pixels, with an angular resolution of $\sim 8^\circ$ and a distance resolution of $\sim 20cm$. In total we have $2 \times 20 \times 71$ different pixels.

7.4 Model Implementation

We use a fully-connected neural network with 4 layers. After each hidden layer, there is a layer of Batch Normalization [16], a layer of ReLU and a layer of Dropout. Fig. 14 illustrates the neural network architecture.

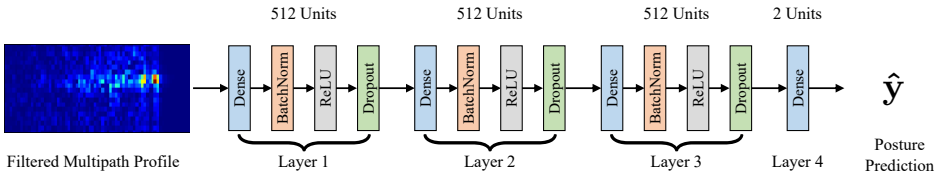

Fig. 14. Architecture of the neural network.

8 EVALUATION

In this section, we empirically evaluate the performance of BodyCompass.

8.1 Metrics

We define *Angle Error* as follows: For a single data point (stable segment) with a ground truth angle of y and a predicted angle of y', its *Angle Error* is the difference between y and y', as defined in Eq. 2. Angle error is always within the range from 0° to 180°.

$$e(y, y') = |(y - y' + 180) \bmod 360 - 180| \tag{2}$$

We use following two metrics that are based on angle error to evaluate the performance of our system.

- **Average Angle Error** (Error): Average Angle Error is the weighted average of all the angle errors for all stable segments in the testing dataset, where the weight of each stable segment is set its time the duration.
- **Threshold Accuracy** (Accuracy): Since most past work computes accuracy with respect to a few key postures (supine, right side, left side, and prone), we similarly estimate accuracy as the percentage of time that the angle error between the prediction and the ground truth is smaller than 45°. This gives us an intuitive understanding of the percentage of time we predict the direction of the user correctly.

8.2 Evaluation Setting

Depending on the amount of data available from the target subject, we present results under three different evaluation settings:

(1) **1-Week**: If we collected enough data from the target subject, we can directly train on the target's data, without transfer learning. Under this setting, for each subject in our dataset, we report the result with leave-one-night-out cross-validation, i.e., using one night for testing and the remaining nights for training, and repeat this process until all the nights have been tested.

(2) **1-Night**: In this setting we limit ourselves to only one night of labeled data from the target subject, and the rest of data from the target is unlabeled. One night is not enough for training. Therefore, we use this one night of data for our calibration points as discussed in Sec. 6.

(3) **16-Minutes**: To further reduce the effort required from users, we present the results with only 8 labeled data points from the target subject. We assume that those 8 points cover most common positions of that user. This can be achieved by asking the user to emulate sleeping in his common sleep postures. In our evaluation, We select those calibration points from the existing sleep dataset by clustering the subject's sleep postures, and picking the center points for each cluster. If the resulting stable segment is longer than 2 minutes we use only a window of two minutes. Thus, collecting these 8 calibration points can be done in 16 minutes.

8.3 Evaluation of BodyCompass's Performance

Table 1. Evaluation results under three different settings with different methods (BodyCompass, k-NN, Random Forest (RF), XGBoost [6] (XGB)). Baseline methods are evaluated under two scenarios: All (A): trained with data from all the subjects; Target (T): trained with data from the target subject only. Note that BodyCompass significantly outperforms all three baselines under all settings.

	BodyCompass	k-NN (A)	k-NN (T)	RF (A)	RF (T)	XGB (A)	XGB (T)
Angle Error (1-week)	15.3° ± 4.4°	NA	31.3° ± 9.7°	NA	33.8° ± 13.0°	NA	33.8° ± 13.3°
Accuracy (1-week)	94.1% ± 4.3%	NA	77.7% ± 9.8%	NA	75.4% ± 12.0%	NA	75.5% ± 12.9%
Angle Error (1-night)	25.6° ± 6.7°	43.1° ± 11.0°	40.6° ± 11.0°	52.5° ± 17.0°	45.4° ± 15.1°	53.9° ± 16.2°	49.2° ± 13.1°
Accuracy (1-night)	86.7% ± 6.7%	65.2% ± 10.5%	67.8% ± 10.2%	54.8% ± 14.5%	62.2% ± 13.8%	53.5% ± 14.2%	59.9% ± 10.5%
Angle Error (16-min)	28.3° ± 8.7°	59.1° ± 19.0°	60.6° ± 19.0°	58.4° ± 20.2°	55.0° ± 18.9°	60.7° ± 20.1°	65.1° ± 13.1°
Accuracy (16-min)	83.7% ± 6.8%	50.3% ± 14.6%	46.4% ± 17.0%	51.0% ± 14.9%	52.2% ± 15.0%	48.7% ± 15.8%	42.8% ± 11.4%

To evaluate the effectiveness of BodyCompass, we compare its performance with three baselines: k-NN (k-Nearest Neighbors), Random Forest and XGBoost [6]. Note that all baselines and BodyCompass take filtered multipath profiles as input. Since the baselines do not have transferability capability like BodyCompass it is not clear how to train them when the available labeled data from the target subject is limited, i.e., in the 1-night and 16-minutes settings. Thus, for these settings, we evaluate the baselines' performance under two different training setups: 1. using data from all the available subjects; 2. using data only from the target subject.

Table 1 compares BodyCompass with the baselines for three different amounts of labeled data from the target subject: 1-week, 1-night, and 16-minutes. The table shows that BodyCompass significantly outperforms all three baselines under all settings. Specifically, BodyCompass and the baselines achieve their best performance when there is sufficient data from the target user. In such setting, BodyCompass's accuracy is 94%, whereas best accuracy across all baseline methods is only 77.7%.

The table also shows that when the amount of labeled data from the target subject is limited (e.g., in the 1-night setting and 16-minutes setting), the accuracies of all baselines are significantly reduced – 27.4% reduction for k-NN, 24.4% reduction for Random Forest, and 26.8% reduction for XGBoost. This is because of the natural variability in sleep postures. For example, even in the same body orientation, a slight change in arm or leg positions can cause a change in the pattern of RF reflections. All three baselines do not have the ability to handle such variability in the absence of a large amount of labeled data from the target. And since differences between

different subjects are large, data from other subjects cannot help the baselines improve their robustness (e.g. under 1-night setting, data from other subjects is detrimental to the final performance). In contrast, BodyCompass aligns the distribution across different users, and performs data augmentation to battle this variability. As a result, BodyCompass can sustain high accuracy of 83.7% even with only 16 minutes of labeled data from the target user.

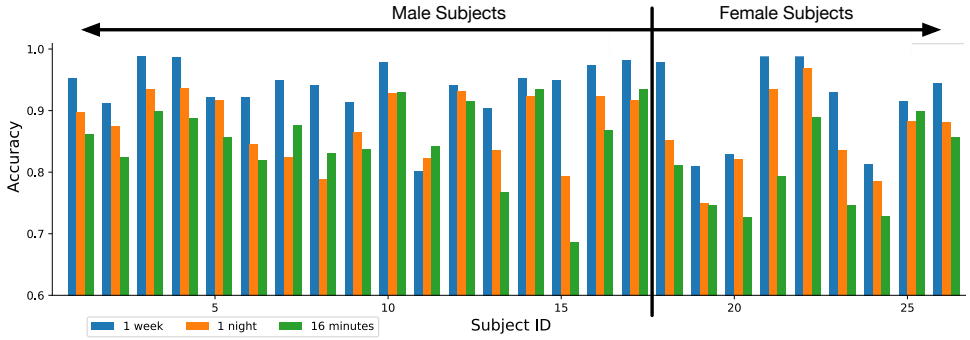

Fig. 15. Accuracy for each of our subjects under three different test settings. Subjects are separated by their gender.

Next, we zoom in on BodyCompass and check the accuracy for each target user. In Fig. 15, we plot the average accuracy for each subject with our system under three different settings: 1-week, 1-night, 16-minutes. One can see that while transfer learning gives good accuracy across subjects, not all subjects have equal accuracy. Subjects who are more different from the source subjects will naturally have a lower accuracy when they have only 1-night of labeled data or 8 calibration points. We notice that, in our dataset, on average the accuracy on male subjects is higher than female subjects. Given the small number of subjects, it is not clear whether this accuracy gap is due to gender differences or is specific to the individuals in our dataset.

Another interesting aspect is that we see a slight increase of the accuracy of Subject #11 when moving from 1-week setting to a transfer setting like 1-night or 16-minutes. This is because while collecting data from Subject #11, the accelerometers often fell off during sleep. As a result, a significant part of each night from this subject had bad data with no accurate labels and was ignored. In this case, transfer learning can potentially have higher accuracy because it leverages labeled data from other users.

Finally, Fig. 16a presents the accuracy as a function of body orientation under all three settings. Our system has a high accuracy across all body orientations except for some corner cases where we lack enough training data (See the amount of labeled data for each angle in Fig. 16b). For example, at angle $-180°$, our system has low accuracy due to the fact that the amount of labeled data for this posture is 1/20 amount of labeled data for the supine case, *i.e.,* 0°.

8.4 Evaluating the Components of BodyCompass

We evaluate the contribution of each component in our system by evaluating the performance of the system without that component. Specifically, we evaluate the contribution of using multipath profiles, breathing-filtering multipath profiles, distribution alignment, data augmentation, majority voting, fully-connected neural network, and circular loss. Removing the contribution of the breathing-filtering multipath profiles, data alignment and data augmentation components is straightforward, yet for the rest, we present the following substitutions:

(1) **Substitution for the Multipath Profiles**: Instead of taking the multipath profile as input, we zoom in on the voxels from the bed area, and take only those voxels. This is equivalent to focusing on the direct path

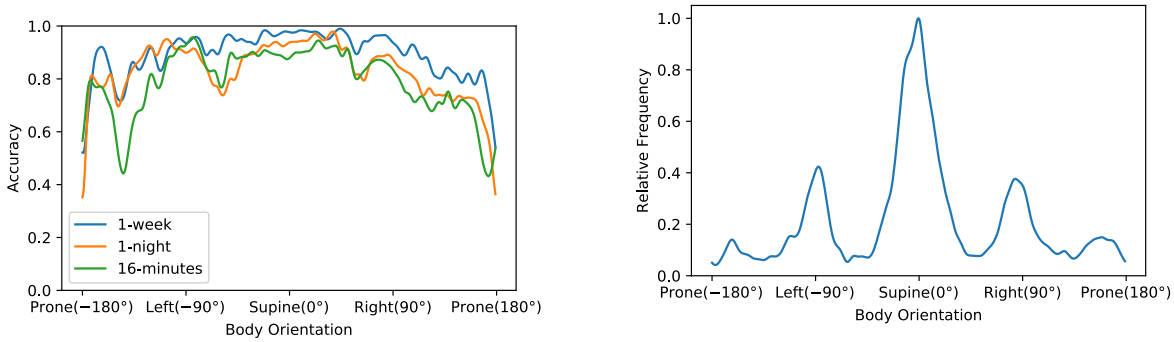

(a) Accuracy w.r.t. Orientation

(b) Amount of labeled data as a function of orientation (normalized by the amount of data for the supine case)

Fig. 16. Accuracy and amount of labeled data for each body orientation.

Table 2. Evaluation of the various components of BodyCompass. The table shows the accuracy under a 1-night setting for the whole system and for the system without a particular component.

	Angle Error	Threshold Accuracy
Full System	25.6° ± 6.7°	86.7% ± 6.7%
CNN instead of Fully-Connected Network	29.5° ± 10.1°	82.5% ± 10.0%
Direct Path instead of Multipath	43.0° ± 11.8°	67.7% ± 9.5%
L2 loss instead of Circular Loss	33.1° ± 12.4°	77.7% ± 13.2%
w/o Breathing Filtering	30.0° ± 10.2°	81.7% ± 9.3%
w/o Distribution Alignment	33.4° ± 11.9°	78.5% ± 10.8%
w/o Data Augmentation	30.7° ± 10.2°	81.5% ± 8.6%
w/o Majority Voting	31.2° ± 11.2°	80.5% ± 9.7%

only, and ignoring any indirect paths that involve signals that bounced off the person and other objects in space.

(2) **Substitution for the Fully-Connected Network**: We substitute the fully-connected neural network with a convolutional neural network (CNN) and evaluate the resulting performance. The CNN model follows the AlexNet model [18].
(3) **Substitution for the Circular Loss**: We can directly regress the angle and use the standard L2 loss.
(4) **Substitution for Majority Voting**: Instead of training on each source subject and performing majority voting, we can combine labeled data from all source subjects and train only one model.

Table 2 provides the evaluation results for BodyCompass's components. All experiments are conducted under 1-night setting. The first row shows BodyCompass's accuracy 1-night setting when all components are active. Comparing the first row with the other rows in the table shows that each of BodyCompass's components offers a considerable improvement to the overall performance, and the removal of any component results in reduced accuracy.

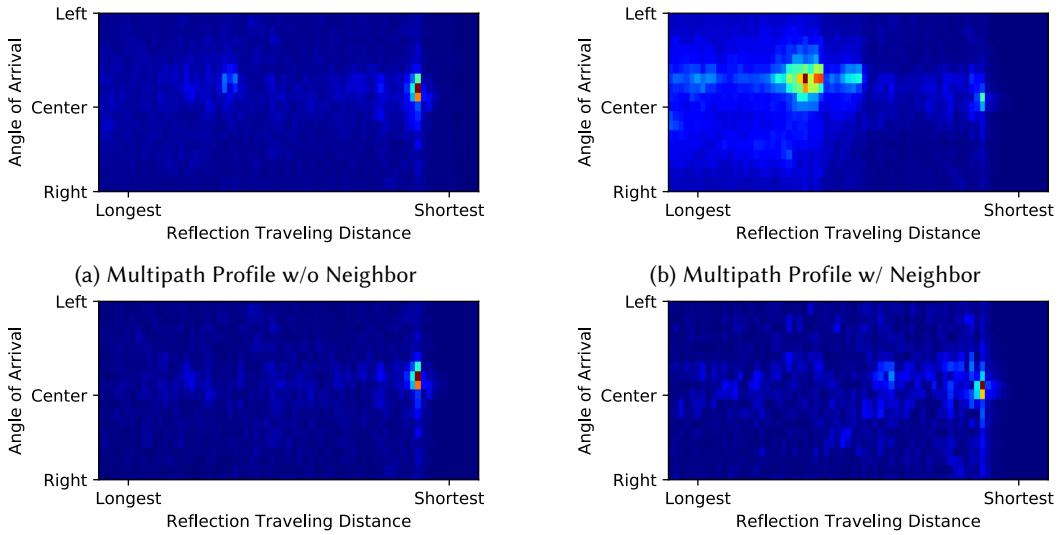

(a) Multipath Profile w/o Neighbor

(b) Multipath Profile w/ Neighbor

(c) Filtered Multipath Profile w/o Neighbor

(d) Filtered Multipath Profile w/ Neighbor

Fig. 17. Robustness to moving neighbors. The figures in the top row show the multipath profile without and with a moving neighbor. Figures in the bottom row *filtered* multipath profile without and with a moving neighbor. The figures show that filtering the multipath profile eliminates extraneous movements from other people, hence boosts the robustness of the system.

8.5 Sensitivity Study

In this section, we test our system's robustness to various factors such as the presence of other people in the environment and their movements, whether the subject has shallow or deep breathing, and the exact location of the radio. We note that when collecting data in the wild, *i.e.*, in people's own homes during their natural overnight sleep, we have no control over the above factors. In fact, in all of the experiments reported in the previous sections, we leave the radio in the home of the subject for about a week to collect the data. We have no control over when the subject goes to bed, where in the bed they sleep, whether their sleep location changes from one night to the next, how they breathe, and who else is at home and how they move while the subject is asleep. Thus, we cannot run sensitivity tests in the wild. We run these tests in a controlled environment where the subject is lying in bed but they are not asleep. In each case, the subject lies in bed at a particular body orientation, while we change the parameter we want to study, and measure BodyCompass's accuracy. The subject then changes his body orientation and we repeat the measurements while varying the parameter of interest.

8.5.1 Sensitivity to Movements by Other People. We evaluate our system's performance when there is a neighbor moving in an adjacent room. The subject sleeps approximately 3 meters away from the radio, whereas the neighbor is about 7 meters away from the radio. We ask the subject to lay down for 2 minutes in each of the following postures: supine, facing left, facing right, and prone. We train the system with examples in which the subject was alone without anyone moving in the neighboring room. We use 6 examples for each sleep posture for training. During testing, we bring a second person to the adjacent room and ask them to move at will, while the subject is lying in bed. We collect 3 examples of each sleep posture for testing.

Fig. 17 shows that our filtering of the multipath profile makes the system robust to extraneous movements such as the presence of a neighbor. The figure plots the multipath profile as well as the filtered multipath profile. As the figure shows, the neighbor's movements have a significant impact on the unfiltered multipath profile (Fig. 17b), yet its impact is removed after filtering the multipath profile using the method in Sec 4.2 (Fig. 17d).

Table 3. Performance w/ neighbor movements.

	Angle Error	Accuracy
Full System	12.7°	100%
w/o Filtering	53.6°	58.3%

Table. 3 presents BodyCompass's accuracy results averaged across the test examples. It shows that our system is robust against extraneous movements by other people.

8.5.2 Sensitivity to Breathing Strength. In this section, we investigate BodyCompass's robustness to variations in breathing strength, *i.e.,* to people having shallow vs. deep breathing. As in the previous section, we ask the subject to lie in bed in various sleep postures and for each posture, vary their breathing depth and rate. We repeat the test for each sleep posture, and with 3 different subjects. During testing, we ask the subject to perform 2 groups of experiments. Each group contains 4 postures (supine, facing left, facing right, and prone) and each posture is repeated 3 times. For the first group, we ask the subject to breathe deeply and slowly, and for the second group, we ask the subject to breathe shallowly and quickly. In Fig. 18, we show the breathing signals as well as their corresponding multipath profiles. Table. 4 reports the average accuracy for different breathing strength. It shows that BodyCompass has high accuracy for both shallow and deep breathing.

Table 4. Performance when subjects breathe at different strengths.

Strength	Angle Error	Accuracy
Deep	13.6° ± 13.5°	97.2% ± 4.8%
Shallow	8.6° ± 4.8°	100% ± 0%

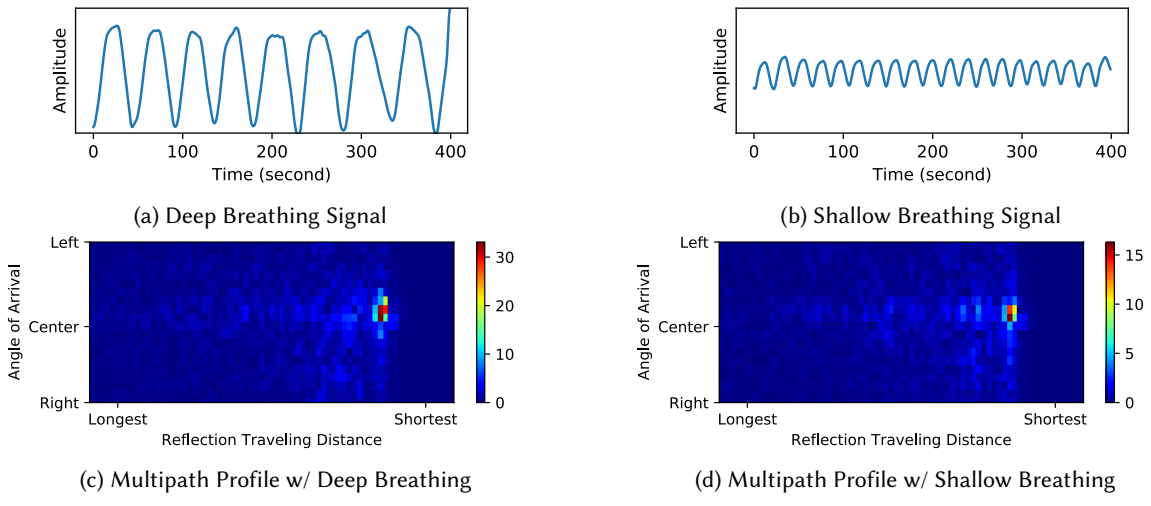

(a) Deep Breathing Signal

(b) Shallow Breathing Signal

(c) Multipath Profile w/ Deep Breathing

(d) Multipath Profile w/ Shallow Breathing

Fig. 18. Breathing signals and their corresponding multipath profiles. Although the amplitude and frequency of the two breathing signals are quite different, their multipath profiles have similar patterns.

We also note that the accuracy results in these sensitivity tests are higher than the numbers presented in Table. 1 for testing in the wild. This is because when the subject is awake, the subject is able to accurately control his posture; In contrast, when the subject is asleep, the limbs can take various positions; Also the subjects may

use pillows to support their bodies, and their use of pillows may change across days. Thus, overall there is much more variability in the wild.

8.5.3 Sensitivity to Device Location. Because the radio has directional antennas and the breathing signal is relatively weak, the device should face the bed to get a good SNR. However, we do not require the device to be exactly facing the chest of the person, nor do we require the device to be at a specific distance from the person.

In this section, we show that BodyCompass is fairly robust to variability in device distance and deviation from facing the chest of the person. In contrast to the previous two sensitivity tests, our in-the-wild deployments exhibit significant diversity in terms of device distance and azimuthal angle with respect to the person's chest. Specifically, in terms of distance between the chest and the device, our deployments cover a range from $0.5m$ to $4m$. In terms of the angle between the device and the person's chest, our deployments cover a range up to plus/minus 35°, where a zero degree means that the device is facing the chest of the person.

Fig. 19 plots the accuracy of BodyCompass as a function of the distance to the person, and the angle to person's chest. The results in the figure show that BodyCompass works reliably for different location settings.

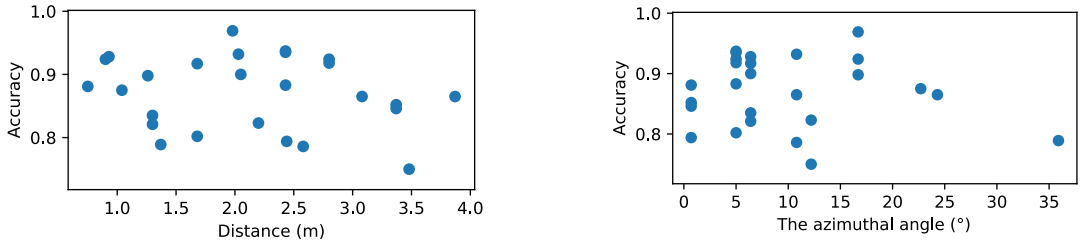

(a) Distance from the radio to the subject

(b) The azimuthal angle from the device to the subject's chest

Fig. 19. Scatter plots of accuracy w.r.t. difference location settings. As two plots show, our system can accommodate a wide range of location settings.

8.6 Example Application: Monitoring the Frequency of Posture Changes

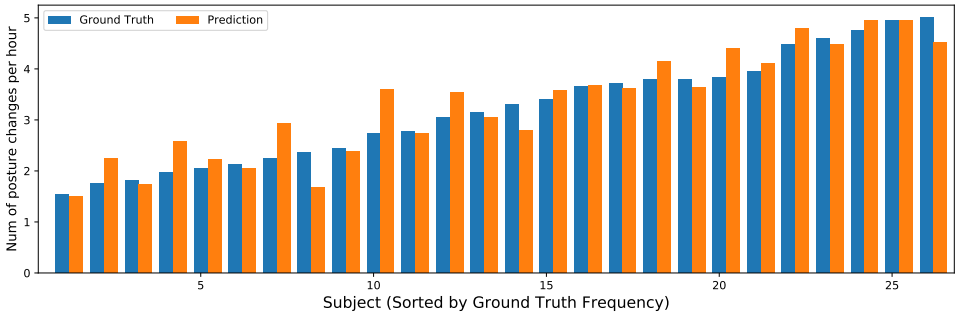

Fig. 20. Ground truth and predictions of posture shift frequency for each subject.

The frequency of posture shift (moving from one posture to another) during sleep is an important sleep-related metric. The literature shows that posture shift frequency is correlated with aging [7, 35] and sleep qualities [7].

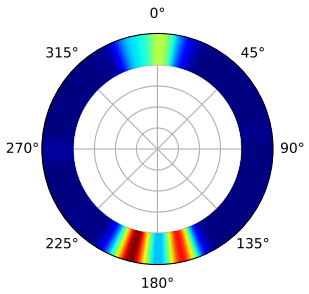

Fig. 21. Angle histogram of Subject #15. Histogram is color-coded into the ring, as a more reddish color represent more occurrence.

Further for patients with Parkinson's disease, less frequent nocturnal turnovers reflects a deterioration in the disease condition [38].

Note that not every motion is a posture change. For example, when the user just changes his arm position, his body angle remains the same. To thoroughly evaluate our system's performance, we adopt the definition of posture change in Skarpsno et al. [35]: to be quantified as a posture change, a body angle change of at least 30° is required.

Fig. 20 plots the frequency of posture shifts for all 26 subjects, ordered in ascending order. The figure compares the ground truth and our model with 1-night of training data from the target user. The figure shows that our model is capable of tracking the frequency of posture shifts accurately, with an average relative error of only 10.3%. These results indicate that our model can be used to track changes in posture shift in Parkinson's patients.

8.7 Failure Case Analysis

Looking back at Fig. 15, we can see that even when the amount of labeled data is scarce, our system is still able to deliver satisfactory accuracies for most subjects. However, there are few exceptions where we can see a significant reduction in performance when the amount of labeled data is reduced. For example, for Subject #15, we have the largest reduction in performance, from 95.0% to 68.7%. To understand the reason, we plot his angle histogram in Fig. 21. We can see that most of his time is spent in prone position (sleep on stomach), and he never sleeps facing left or right. Recall in Sec. 6.2, we explicitly align the distribution to be the same. However this subject's intrinsic distribution is not similar to any other subjects. Therefore the alignment cannot fully succeed, which causes bad transfer performance. We expect to see an increase of our system's performance on this subject if we have more subjects with similar sleep postures.

9 CONCLUSION

We present BodyCompass, a wireless system that provides accurate sleep posture monitoring in the wild. By explicitly extracting RF reflections from the user and designing appropriate machine learning algorithms, our system can accurately capture the user's posture and is able to transfer its knowledge to a new home with minimal additional data. A user study in 26 different homes with 26 subjects and more than 200 nights shows that BodyCompass is highly accurate, with an accuracy of 94% using 1 week of data from the user, and 83.7% using only 16 minutes of data. We believe that this work can serve as a practical sleep posture monitoring system, enabling easy adoption and helping doctors and patients address this unmet need.

ACKNOWLEDGMENTS

The authors thank the members of the NETMIT group for their insightful discussions, anonymous reviewers for their valuable comments and all the volunteers for participating in our study.

REFERENCES

[1] Fadel Adib, Zachary Kabelac, and Dina Katabi. 2015. Multi-person localization via {RF} body reflections. In *12th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 15)*. 279–292.

[2] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C Miller. 2015. Smart homes that monitor breathing and heart rate. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, 837–846.

[3] Sina Akbarian, Ghazaleh Delfi, Kaiyin Zhu, Azadeh Yadollahi, and Babak Taati. 2019. Automated Non-contact Detection of Head and Body Positions during Sleep. *IEEE Access* (2019).

[4] Paolo Barsocchi. 2012. Position recognition to support bedsores prevention. *IEEE journal of biomedical and health informatics* 17, 1 (2012), 53–59.

[5] Kang-Ming Chang and Shin-Hong Liu. 2011. Wireless portable electrocardiogram and a tri-axis accelerometer implementation and application on sleep activity monitoring. *Telemedicine and e-Health* 17, 3 (2011), 177–184.

[6] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.

[7] Joseph De Koninck, Dominique Lorrain, and Pierre Gagnon. 1992. Sleep positions and position shifts in five age groups: an ontogenetic picture. *Sleep* 15, 2 (1992), 143–149.

[8] Gustavo Desouzart, Rui Matos, Filipe Melo, and Ernesto Filgueiras. 2016. Effects of sleeping position on back pain in physically active seniors: A controlled pilot study. *Work* 53, 2 (2016), 235–240.

[9] Terence Dwyer, A-LB Ponsonby, Neville M Newman, and Laura E Gibbons. 1991. Prospective cohort study of prone sleeping position and sudden infant death syndrome. *The Lancet* 337, 8752 (1991), 1244–1247.

[10] Claudia Gorecki, S José Closs, Jane Nixon, and Michelle Briggs. 2011. Patient-reported pressure ulcer pain: a mixed-methods systematic review. *Journal of pain and symptom management* 42, 3 (2011), 443–459.

[11] Timo Grimm, Manuel Martinez, Andreas Benz, and Rainer Stiefelhagen. 2016. Sleep position classification from a depth camera using bed aligned maps. In *2016 23rd International Conference on Pattern Recognition (ICPR)*. IEEE, 319–324.

[12] Enamul Hoque, Robert F Dickerson, and John A Stankovic. 2010. Monitoring body positions and movements during sleep using wisps. In *Wireless Health 2010*. ACM, 44–53.

[13] Chi-Chun Hsia, Yu-Wei Hung, Yu-Hsien Chiu, and Chia-Hao Kang. 2008. Bayesian classification for bed posture detection based on kurtosis and skewness estimation. In *HealthCom 2008-10th International Conference on e-health Networking, Applications and Services*. IEEE, 165–168.

[14] Chen-Yu Hsu, Aayush Ahuja, Shichao Yue, Rumen Hristov, Zachary Kabelac, and Dina Katabi. 2017. Zero-Effort In-Home Sleep and Insomnia Monitoring using Radio Signals. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1 (2017).

[15] Chen-Yu Hsu, Yuchen Liu, Zachary Kabelac, Rumen Hristov, Dina Katabi, and Christine Liu. 2017. Extracting gait velocity and stride length from surrounding radio signals. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2116–2126.

[16] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).

[17] Robert Kloster and Torstein Engelskjøn. 1999. Sudden unexpected death in epilepsy (SUDEP): a clinical perspective and a search for risk factors. *Journal of Neurology, Neurosurgery & Psychiatry* 67, 4 (1999), 439–444.

[18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.

[19] Hong Ji Lee, Su Hwan Hwang, Seung Min Lee, Yong Gyu Lim, and Kwang Suk Park. 2013. Estimation of body postures on bed using unconstrained ECG measurements. *IEEE journal of biomedical and health informatics* 17, 6 (2013), 985–993.

[20] Jennifer A Liebenthal, Shasha Wu, Sandra Rose, John S Ebersole, and James X Tao. 2015. Association of prone position with sudden unexpected death in epilepsy. *Neurology* 84, 7 (2015), 703–709.

[21] Jia Liu, Xingyu Chen, Shigang Chen, Xiulong Liu, Yanyan Wang, and Lijun Chen. 2019. TagSheet: Sleeping Posture Recognition with an Unobtrusive Passive Tag Matrix. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 874–882.

[22] Jian Liu, Yingying Chen, Yan Wang, Xu Chen, Jerry Cheng, and Jie Yang. 2018. Monitoring vital signs and postures during sleep using WiFi signals. *IEEE Internet of Things Journal* 5, 3 (2018), 2071–2084.

[23] Jason J Liu, Wenyao Xu, Ming-Chun Huang, Nabil Alshurafa, Majid Sarrafzadeh, Nitin Raut, and Behrooz Yadegar. 2014. Sleep posture analysis using a dense pressure sensitive bedsheet. *Pervasive and Mobile Computing* 10 (2014), 34–50.

[24] Shuangjun Liu and Sarah Ostadabbas. 2017. A vision-based system for in-bed posture tracking. In *Proceedings of the IEEE International Conference on Computer Vision*. 1373–1382.

[25] Xuefeng Liu, Jiannong Cao, Shaojie Tang, and Jiaqi Wen. 2014. Wi-Sleep: Contactless sleep monitoring via WiFi signals. In *2014 IEEE Real-Time Systems Symposium*. IEEE, 346–355.

[26] Xuefeng Liu, Jiannong Cao, Shaojie Tang, Jiaqi Wen, and Peng Guo. 2015. Contactless respiration monitoring via off-the-shelf WiFi devices. *IEEE Transactions on Mobile Computing* 15, 10 (2015), 2466–2479.

[27] Bassem R Mahafza. 2005. *Radar systems analysis and design using MATLAB*. Chapman and Hall/CRC, USA.

[28] Akshay Menon and Manoj Kumar. 2013. Influence of body position on severity of obstructive sleep apnea: a systematic review. *ISRN otolaryngology* 2013 (2013).

[29] ALISTER McKENZIE Neill, Susan Michelle Angus, Dimitar Sajkov, and RONALD DOUGLAS McEVOY. 1997. Effects of sleep posture on upper airway stability in patients with obstructive sleep apnea. *American journal of respiratory and critical care medicine* 155, 1 (1997), 199–204.

[30] Phuc Nguyen, Xinyu Zhang, Ann Halbower, and Tam Vu. 2016. Continuous and fine-grained breathing volume monitoring from afar using wireless signals. In *IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications*. IEEE, 1–9.

[31] Sarah Ostadabbas, Maziyar Baran Pouyan, Mehrdad Nourani, and Nasser Kehtarnavaz. 2014. In-bed posture classification and limb identification. In *2014 IEEE Biomedical Circuits and Systems Conference (BioCAS) Proceedings*. IEEE, 133–136.

[32] M Baran Pouyan, Sarah Ostadabbas, Masoud Farshbaf, Rasoul Yousefi, Mehrdad Nourani, and MDM Pompeo. 2013. Continuous eight-posture classification for bed-bound patients. In *2013 6th International Conference on Biomedical Engineering and Informatics*. IEEE, 121–126.

[33] Tauhidur Rahman, Alexander T Adams, Ruth Vinisha Ravichandran, Mi Zhang, Shwetak N Patel, Julie A Kientz, and Tanzeem Choudhury. 2015. Dopplesleep: A contactless unobtrusive sleep sensing system using short-range doppler radar. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 39–50.

[34] Grace Gita Redhyka, Dika Setiawan, and Demi Soetraprawata. 2015. Embedded sensor fusion and moving-average filter for Inertial Measurement Unit (IMU) on the microcontroller-based stabilized platform. In *2015 International Conference on Automation, Cognitive Science, Optics, Micro Electro-Mechanical System, and Information Technology (ICACOMIT)*. IEEE, 72–77.

[35] Eivind Schjelderup Skarpsno, Paul Jarle Mork, Tom Ivar Lund Nilsen, and Andreas Holtermann. 2017. Sleep positions and nocturnal body movements based on free-living accelerometer recordings: association with demographics, lifestyle, and insomnia symptoms. *Nature and science of sleep* 9 (2017), 267.

[36] Alexander Tataraidze, Lyudmila Korostovtseva, Lesya Anishchenko, Mikhail Bochkarev, Yurii Sviryaev, and Sergey Ivashov. 2016. Bioradiolocation-based sleep stage classification. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2839–2842.

[37] Yonglong Tian, Guang-He Lee, Hao He, Chen-Yu Hsu, and Dina Katabi. 2018. RF-based fall monitoring using convolutional neural networks. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 137.

[38] Kenji Uchino, Makoto Shiraishi, Keita Tanaka, Masashi Akamatsu, and Yasuhiro Hasegawa. 2017. Impact of inability to turn in bed assessed by a wearable three-axis accelerometer on patients with Parkinson's disease. *PloS one* 12, 11 (2017), e0187616.

[39] Hao Wang, Daqing Zhang, Yasha Wang, Junyi Ma, Yuxiang Wang, and Shengjie Li. 2016. RT-Fall: A real-time and contactless fall detection system with commodity WiFi devices. *IEEE Transactions on Mobile Computing* 16, 2 (2016), 511–526.

[40] Tianben Wang, Daqing Zhang, Yuanqing Zheng, Tao Gu, Xingshe Zhou, and Bernadette Dorizzi. 2018. C-FMCW based contactless respiration detection using acoustic signal. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 4 (2018), 170.

[41] Wei Wang, Alex X Liu, and Muhammad Shahzad. 2016. Gait recognition using wifi signals. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 363–373.

[42] Xuyu Wang, Chao Yang, and Shiwen Mao. 2017. PhaseBeat: Exploiting CSI phase data for vital sign monitoring with commodity WiFi devices. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 1230–1239.

[43] Xuyu Wang, Chao Yang, and Shiwen Mao. 2017. TensorBeat: Tensor decomposition for monitoring multiperson breathing beats with commodity WiFi. *ACM Transactions on Intelligent Systems and Technology (TIST)* 9, 1 (2017), 8.

[44] Yuxi Wang, Kaishun Wu, and Lionel M Ni. 2016. Wifall: Device-free fall detection by wireless networks. *IEEE Transactions on Mobile Computing* 16, 2 (2016), 581–594.

[45] Teng Wei and Xinyu Zhang. 2015. mtrack: High-precision passive tracking using millimeter wave radios. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. ACM, 117–129.

[46] Chenshu Wu, Zheng Yang, Zimu Zhou, Xuefeng Liu, Yunhao Liu, and Jiannong Cao. 2015. Non-invasive detection of moving and stationary human with wifi. *IEEE Journal on Selected Areas in Communications* 33, 11 (2015), 2329–2342.

[47] Xiaowei Xu, Feng Lin, Aosen Wang, Chen Song, Yu Hu, and Wenyao Xu. 2015. On-bed sleep posture recognition based on body-earth mover's distance. In *2015 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 1–4.

[48] Heenam Yoon, Suhwan Hwang, Dawoon Jung, Sangho Choi, Kwangmin Joo, Jaewon Choi, Yujin Lee, Do-Un Jeong, and Kwangsuk Park. 2015. Estimation of sleep posture using a patch-type accelerometer based device. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 4942–4945.

[49] Shichao Yue, Hao He, Hao Wang, Hariharan Rahul, and Dina Katabi. 2018. Extracting multi-person respiration from entangled rf signals. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018), 86.

[50] Feng Zhang, Chen Chen, Beibei Wang, and KJ Ray Liu. 2018. WiSpeed: A statistical electromagnetic approach for device-free indoor speed estimation. *IEEE Internet of Things Journal* 5, 3 (2018), 2163–2177.

[51] Zhenyuan Zhang, Zengshan Tian, and Mu Zhou. 2018. Latern: Dynamic continuous hand gesture recognition using FMCW radar sensor. *IEEE Sensors Journal* 18, 8 (2018), 3278–3289.

[52] Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, and Dina Katabi. 2018. Through-wall human pose estimation using radio signals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7356–7365.

[53] Mingmin Zhao, Yonglong Tian, Hang Zhao, Mohammad Abu Alsheikh, Tianhong Li, Rumen Hristov, Zachary Kabelac, Dina Katabi, and Antonio Torralba. 2018. RF-based 3D skeletons. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*. 267–281.

[54] Mingmin Zhao, Shichao Yue, Dina Katabi, Tommi S Jaakkola, and Matt T Bianchi. 2017. Learning sleep stages from radio signals: A conditional adversarial architecture. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 4100–4109.