Phyjama: Physiological Sensing via Fiber-enhanced Pyjamas

ALI KIAGHADI, University of Massachusetts Amherst, USA
SEYEDEH ZOHREH HOMAYOUNFAR, University of Massachusetts Amherst, USA
JEREMY GUMMESON, University of Massachusetts Amherst, USA
TRISHA ANDREW, University of Massachusetts Amherst, USA
DEEPAK GANESAN, University of Massachusetts Amherst, USA

Unobtrusive and continuous monitoring of cardiac and respiratory rhythm, especially during sleeping, can have significant clinical utility. An exciting new possibility for such monitoring is the design of textiles that use all-textile sensors that can be woven or stitched directly into a textile or garment. Our work explores how we can make such monitoring possible by leveraging something that is already familiar, such as pyjama made of cotton/silk fabric, and imperceptibly adapt it to enable sensing of physiological signals to yield natural fitting, comfortable, and less obtrusive smart clothing.

We face several challenges in enabling this vision including requiring new sensor design to measure physiological signals via everyday textiles and new methods to deal with the inherent looseness of normal garments, particularly sleepwear like pyjamas. We design two types of textile-based sensors that obtain a ballistic signal due to cardiac and respiratory rhythm — the first a novel resistive sensor that leverages pressure between the body and various surfaces and the second is a triboelectric sensor that leverages changes in separation between layers to measure ballistics induced by the heart. We then integrate several instances of such sensors on a pyjama and design a signal processing pipeline that fuses information from the different sensors such that we can robustly measure physiological signals across a range of sleep and stationary postures. We show that the sensor and signal processing pipeline has high accuracy by benchmarking performance both under restricted settings with twenty one users as well as more naturalistic settings with seven users.

CCS Concepts: • **Human-centered computing → Ubiquitous and mobile computing design and evaluation methods**; • **Applied computing → Health care information systems**;

Additional Key Words and Phrases: Smart textile, IoT

ACM Reference Format:
Ali Kiaghadi, Seyedeh Zohreh Homayounfar, Jeremy Gummeson, Trisha Andrew, and Deepak Ganesan. 2019. Phyjama: Physiological Sensing via Fiber-enhanced Pyjamas. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 3, Article 89 (September 2019), 29 pages. https://doi.org/10.1145/3351247

Authors' addresses: Ali Kiaghadi, University of Massachusetts Amherst, College of Information and Computer Sciences, 140 Governors Drive, Amherst, MA, 01003, USA, akiaghadi@umass.edu; Seyedeh Zohreh Homayounfar, University of Massachusetts Amherst, Department of Chemistry, 140 Governors Drive, Amherst, MA, 01003, USA; Jeremy Gummeson, University of Massachusetts Amherst, College of Information and Computer Sciences, 140 Governors Drive, Amherst, MA, 01003, USA; Trisha Andrew, University of Massachusetts Amherst, Department of Chemistry, 140 Governors Drive, Amherst, MA, 01003, USA; Deepak Ganesan, University of Massachusetts Amherst, College of Information and Computer Sciences, 140 Governors Drive, Amherst, MA, 01003, USA.

2474-9567/2019/9-ART89 $15.00
https://doi.org/10.1145/3351247

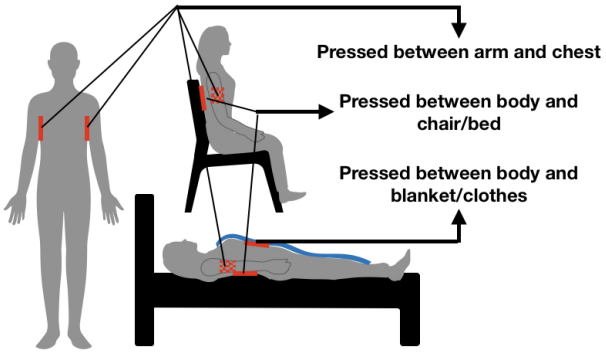

Fig. 1. Many parts of a loose textile are naturally under pressure, for example, between the body and the bed/chair, arm and torso, and clothing under a blanket. These can be leveraged to sense physiological signals in a loose-fitting textile.

1 INTRODUCTION

It is widely thought that electronically active garments are the future of portable, interactive devices. In particular, an exciting new possibility is the design of textiles that use all-textile sensors and actuators that can be woven or stitched directly into a textile or garment. While there are many smart textile-based garments that are already on the market (e.g. Nike's AeroReact [6], Teslasuit [9], and Zephyr Compression shirts [10]), these generally use flexible electronic components that are integrated with textiles. However, enhancing textiles with electronics is demanding because of two reasons: a) they change the aesthetics and tactile perception (or feel) of the textile, and b) the large, varied mechanical stresses to which textiles are subjected to can easily abrade or damage microelectronic components and electronic interconnects.

Our work explores how we can use something that is already familiar, such as cotton/silk thread, fabrics, and imperceptibly adapt it to enable sensing of physiological signals to yield natural fitting, comfortable, and less obtrusive smart clothing. Specifically, we focus on pyjamas as a representative instance of loosely worn and comfortable clothing that can be worn at home and during sleep. A comfortable, loosely worn sleepwear that can measure a variety of physiological signals continuously during sleep and other everyday situations can pave the way towards smart clothing that looks and feels more like normal clothing.

While the ability to instrument everyday textiles opens up exciting new possibilities, a big challenge that we face is designing methods to measure physiological signals using loosely worn clothing. Existing solutions for sensing respiratory and cardiac signals all rely on tightly worn bands or electrodes that are placed at specific locations on the skin. Similarly, many of the ECG-sensing shirts need a tight fit at several locations on the body to obtain the cardiac signal. In contrast, our objective is to enable physiological sensing with a wearable at the other end of the spectrum in terms of looseness i.e. an extremely loose daily-use textile like a pajama that is designed *solely with comfort* in mind.

1.1 Leveraging Pressured Surfaces

While looseness may appear to present a problem, we observe that even when we consider "loose clothing", there are several parts of such a textile that are pressed against the body due to our posture and contact with external surfaces. In fact, once we start carefully observing all the different locations where the textile is naturally pressed, we find that we can classify them in several groups as shown in Figure 1. The first group is locations where there is a force exerted by the body on an external surface, for example, between our torso and a chair or bed. The second group is where different limbs of the body put pressure on the torso. For example, when the arm rests on

its side, it puts pressure on the textile between the arm and torso (i.e. below the armpit). The third group is very light pressure due to a blanket or even pressure due to the weight of the textile on the chest when an individual is lying down.

Often, many such pressured surfaces are present concurrently. When sitting, there is pressure between the body and the chair surface, between the arm and torso, and between the chest and the textile. When sleeping, we have the above set of pressure points but also additional pressure due to a blanket or the pyjama itself pressing against the chest. More surprisingly, pressure between the arm and torso, and between the chest and clothing are even present when standing and there is no contact with an external surface.

In conjunction, these present myriad sensing opportunities but how do we leverage them to measure cardiac and respiratory signals? One option is to use discrete electronic components like ECG electrodes or pressure sensors but we lose the comfortable feel of the textile if we use discrete electronics. A second option is to use textile-based ECG electrodes but this requires tightly worn clothing that is in direct contact with the skin and raises significant robustness issues due to motion artifacts with dry electrodes.

1.2 The Phyjama Approach

The limitations of existing methods led us to explore ways to sense ballistic movements i.e. pressure changes in the textile due to breathing and heartbeats, and measure these changes to extract physiological variables. Our approach seeks to design a novel method that leverages the numerous contact opportunities to measure ballistic movements while relying solely on comfortable textile-based sensing solutions.

But we face several challenges that make it non-trivial to design such a solution. First, there is no existing fabric-based method to sense continuous and dynamic changes in pressure. Existing pressure sensing methods using textiles are binary detectors i.e. they detect high pressure versus low pressure, but they do not measure the amount of pressure in a continuous manner. Second, the dynamic range of pressure at different opportunistic sensing points is many orders of magnitude apart. At one end of the spectrum, a substantial amount of body weight is placed on the textile while sleeping and at the other end, there is a minuscule amount of pressure from the chest on to the textile during inhalation. Third, we need to measure the signal at multiple locations and fuse the information since no single location may have a sufficiently good signal for robustly estimating physiological parameters, and the best location changes depending on the user posture.

Phyjama addresses these challenges using several unique approaches. For locations where there is moderate to large amounts of pressure, we design a novel all-textile pressure sensor that leverages impedance changes to measure pressure changes due to respiration and heartbeats. For locations where there is a tiny amount of pressure but where the fabric is dynamic, we design a triboelectric textile sensor that leverages small amounts of fabric compression to extract the dynamics of the textile. We show that these patches can be combined in typical loose-fitting textiles and their signals fused using a combination of signal processing and machine learning to enable holistic textile-based sensing of physiological variables without sacrificing comfort.

In summary, our contributions are:

- We design a novel distributed multi-modal textile-based sensor that can be integrated with loosely-worn clothing such as pyjamas to measure physiological signals. Our design relies exclusively on textile elements in sensed regions, while using discrete electronic components only in expected locations such as buttons.

- Our design combines a novel fabric-based pressure sensor and a triboelectric sensor, and fuses signals from a distributed set of sensors to extract ballistic signals from multiple locations. We show that this combination of sensors allows us to detect physiological signals across diverse postures and leverage all forms of opportunistic contact between a loose fabric and the body.

- We develop a signal processing pipeline to fuse information from multiple vantage points, and fuse them while taking into account signal quality from each patch. This allows us to extract precise information about heart rate, respiration rate, and sleep posture.
- We implement and evaluate a full version of the Phyjama in two user studies. The first is a benchmarking study across 21 users where we show that we can detect BCG peaks with 97% F1-score, breathing rate with 0.64 resp/min median error and heart rate with 0.5 bpm median error. The second is a one-hour nap study across seven users, four of whom are elderly participants, where we show that we can detect breathing rate with 0.75 resp/min median error and heart rate with 2.5 bpm median error.

2 RELATED WORK

The goal of this work is the design of a comfortable and unobtrusive vital sign monitoring system that can be worn continuously during long duration of wear without impacting sleep. To achieve these aims with loosely fitting textiles, the sensing substrate must be able to simultaneously capture posture information in addition to signals that contain respiration and heart rate information. Existing sensing systems fall short of these aims.

Flexible and discrete sensors in smart textiles A variety of prior work has looked at using flexible but non-textile based sensors that are embedded in textiles. For example, one solution to measure vital signs uses electromechanical film (EMFi) to measure ballistic heart rate [16]. Another solution also senses ballistics using pressure sensors printed on a polymer substrate [50]. Several such approaches have also been presented for posture detection using smart textiles. Sardini et al. [48] weave a serpentine shaped copper wire in the back of a shirt to form a varying impedance due to bending of spine. Dunne et al. [22] use a plastic optical fiber to monitor spinal posture. Lorussi et al. [38] use an array of piezoelectric sensors to find human posture. While the sensors are flexible, they are still made of stiff non-textile components that lack the feel of an everyday textile. In addition, several of these assume tight contact between sensors and skin, which in turn, requires tight clothing.

Several other researchers have integrated discrete sensors like IMUs and pressure switches in textile elements, primarily to obtain postural parameters. Normally, three IMUs are used to capture spinal angle, placed on thoracic, thoraco-lumber, and lumber parts [25, 26, 57]. Since any movement would be sensed by the IMU, the garments are often tight-fitting to avoid unwanted rotation of the IMUs which would substantially increase motion artifact noise. Our work has no discrete sensing elements and directly measures the ballistic signals.

Fabric-based sensors Much of the prior work on physiological sensing with fabric-based sensors are based on tight-fitting garments typically by relying on conductive fabric electrodes (existing methods and requirements of smart textiles are surveyed in [18]). While these electrodes are widely available, they are designed for tight contact with the skin and unsuitable for loosely worn clothing. There has been some work on measuring impedance changes for physiological measurements — for example, [45] integrates piezoelectric elements in a smart textile and tracks changes in impedance using a sinusoid injected across two fabric layers. The work also relies on tightly worn clothing and close skin contact.

There has been limited work on sensing physiological variables using loose-fitting textiles. One such work is respiration sensing using conductive foam pressure sensors [16]. This is essentially a binary foam-based sensor that moves between an open and short circuit configuration while a person breathes. In contrast, Phyjama provides complete cardiorespiratory rhythm signal while using far more natural fabric elements.

There has also been work on detecting biochemical signals using clothing. For example, prior work has looked at sweat detection — [32] implements a perspiration detection based on fabric sensors placed in the armpit and on the back of a shirt, and [36] detects sweat at joints. While not the focus of our current work, it is possible that some of these methods can be integrated with Phyjama.

Instrumenting furniture and bedding Several prior approaches have explored the use of instrumenting furniture including chairs and beds; approaches in this body of work typically use discrete strain gauges and

custom textiles to sense changes in pressure. eCushin[59] presents e-textiles instrumented in chair's seat cushion to differentiate between multiple sitting postures. Similarly, Tekscan have developed a system to extract pressure heat map between two sheets [8]. Also, Health Chair [28] instruments arm-rest and back of a chair to extract heart rate and respiratory rate of users. Several efforts have also looked at unobtrusively instrumenting beds to measure ballistic heart rate during sleep. One approach leverages highly sensitive geophones to measure the seismic motions induced by individual heart beats and slow moving signals from respiration [33, 34]. Commercial MEMS accelerometer-based units are available that can measure heart rate based on ballistocardiography signals measured via the bed [5].

Wearable devices There are many wearable devices in the market for sleep sensing, most of which use photoplethysmography to measure the pulse wave on the wrist or fingers (e.g. Fitbit [2], Polar Vivofit [3], and Oura Ring[7]). A key distinction is that Phyjama is fully integrated within existing daily wear and does not need additional wearables.

Non-wearable approaches A variety of non-contact methods have recently become popular for measuring respiration and heart rate signals. One body of work is on radar-based sensing of respiration and heart rhythm [11–14, 43, 46]. These methods use FMCW or UWB radars and measure changes in the displacement and the doppler shifts due to respiration and ballistics of the heart. While non-contact sensing is appealing, robustness is a major problem due to occlusions (e.g. blanket), variations in sleep posture, movement artifacts, disaggregation of signals when multiple individuals share the same bed, etc. As a result, these methods typically are more accurate for respiration sensing which causes larger movements than ballistics of the heart. Other non-contact approaches include the use of vision-based and depth camera-based methods such as use of cameras to find physiological variables. These require line-of-sight, proper lighting and a relatively stationary user within an area in front of a camera.

3 FABRIC-BASED SENSOR DESIGN

The central contribution of our work is the design of a distributed all-textile patch architecture that can measure cardiac and respiratory signals. The building blocks of our design are two types of all-textile patches — a resistive patch to measure pressure changes and a triboelectric patch that measures surface charge transfer. The resistive patch is a first-of-its-kind device and we are unaware of similar devices to measure physiological signals; the triboelectric patch is similar to previously published designs but this is the first time it has been shown to detect tiny ballistic signals from the heart. Our overall design is shown in Figure 2 — the Phyjama comprises several patches to enable us to gather physiological signals from multiple vantage points. In the rest of this section, we describe the sensor design challenges involved and how we tackled them in Phyjama.

3.1 The Resistive Patch

The first challenge that we faced was how to design an all-fabric pressure sensor that is sensitive to small changes in pressure due to the ballistics of the heartbeat. At a conceptual level, the design of such a sensor appears quite straightforward – the sensor has two conductive layers with a highly resistive middle layer as shown in Figure 3 (left and middle). The resistivity of the middle layer is inversely proportional to the pressure on the sensor which can be measured.

But the design of the middle fabric is not straightforward since the ballistic signal is extremely weak. On one hand, if the fabric is an insulator like regular cotton, then the resistance is extremely high (teraohms) and it is extremely complex and expensive to design a sensing circuit to measure minute resistance changes at such high electrical impedance. Also, we need high impedance in our circuit to measure changes in a high impedance sensor, but this makes our circuit very sensitive to noise (a small current induced on a high impedance circuit results in higher noise voltage than the same noise on a low impedance circuit). We have many sources of

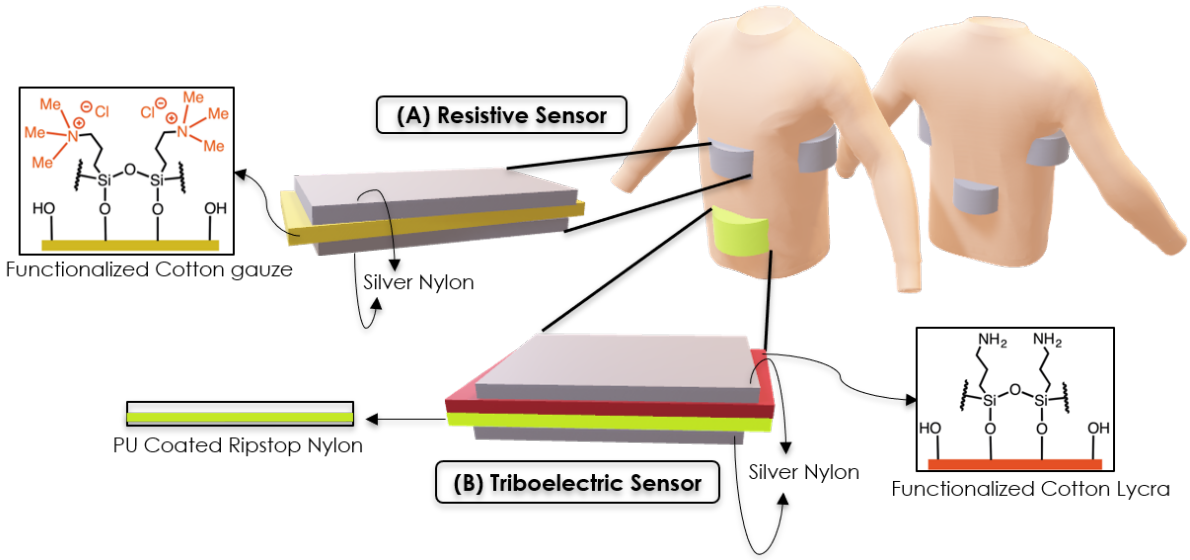

Fig. 2. Phyjama comprises a distributed set of four resistive sensors (A) and a tribolelectric sensor (B).

noise in fabric-based circuits that use large conductive layers including electromagnetic noise, static fields, and motion artifacts, hence we need to operate in a lower impedance regime to minimize the impact of noise on the signal. On the other hand, if the fabric is too conductive, then it can short too quickly after a small amount of pressure is applied and may not be able to cover the range of pressures that are observed in clothing. The pressure between the body and an external surface can vary by more than an order of magnitude depending on whether an individual is seated or lying down; similarly, the pressure between the arm and torso is also much smaller than the pressure between the body and the bed. Thus, we need to operate in a sweet spot where the fabric is optimized with sufficiently high resistance that it does not create a short circuit even under pressure while at the same time being sensitive to small pressure changes due to the ballistics of the heart.

Fabric-level optimizations: To address this issue, we explored a number of textile parameters and surface functionalization reactions to change the surface conductivity of the cotton cloths. First, we explored the impact of weave density on the overall resistivity of the sensor and found that medium-weave cotton gauze minimized shorting events, afforded the most stable pressure-induced electrical signals, and remained comfortable to wear after being incorporated into a garment. Next, a hydrophobic, perfluorinated alkyl acrylate coating was vapor deposited onto cotton cloths using a custom-built vacuum reactor to impart wash stability. Perfluorinated coatings are superhydrophobic and are commonly used to create stain- and sweat-repellant upholstery and active wear. However, this surface coating resulted in fabrics with increased resistivity as compared to pristine samples. Changing the chemical structure of the grafting point to a siloxane moiety did not attenuate the high surface resistivity observed with perfluoroalkyl coatings. We hypothesized that such increases in surface resistivity evolved because the coatings contained saturated alkyl chains without accessible conductive states. As most textile coatings are similarly insulating, we needed to innovate a new surface coating that would impart either electronic or ionic conductivity to the cotton cloths.

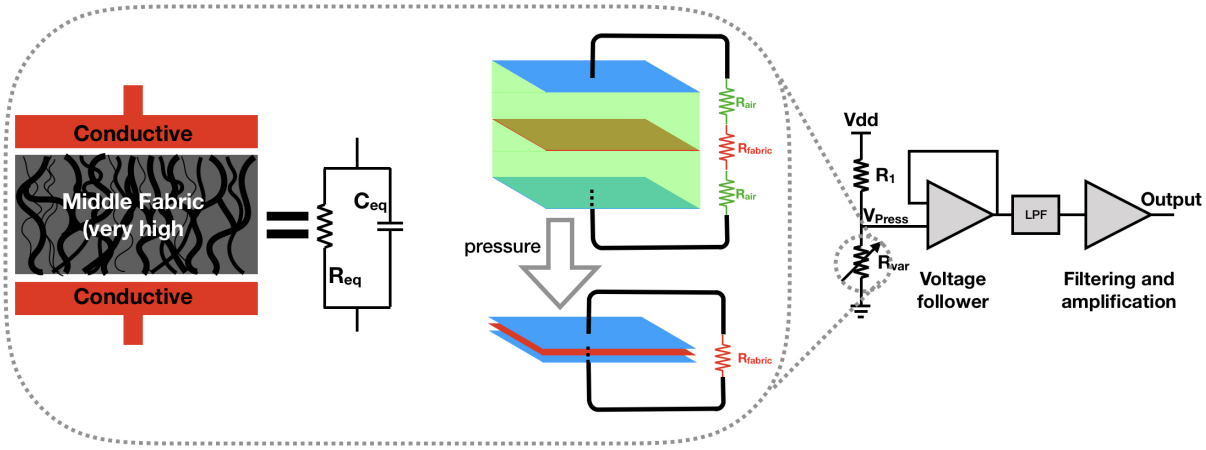

Fig. 3. Fabric structure of resistive pressure-sensing patch and electrical model in shown on the left. The fabric is connected to an analog filtering and amplification circuit as shown on the right.

Finally, we targeted ion conductive coatings because ionic conductors are comparatively more compatible with salt-rich biological systems than electronic materials. We identified a siloxane molecule, N-trimethoxysilylpropyl-N,N,N,-trimethylammonium chloride, containing quaternary ammonium moieties as a potential coating: the siloxane moieties should covalently bond to the free hydroxyl groups present in the repeat unit of cellulose acetate (cotton), while the quaternary ammonium moieties and their chloride counterions would act as ion conductors that should reduce the observed surface resistivity of the fabric. The surface resistivity should also be proportional to the surface concentration of the quaternary ammonium groups, which, in turn, is proportional to the concentration of the siloxane molecule used during the solution-phase functionalization reaction.

Various test sensors of the same size were created by sandwiching a sheet of cotton (either pristine or ion-conductive) between two sheets of silver nylon fabric. As desired, cotton gauze functionalized with N-trimethoxysilylpropyl-N,N,N,-trimethylammonium chloride displayed a more sensitive voltage change with applied pressure, as compared to pristine cotton gauze or cotton lycra. Therefore, three-layer devices containing our ion-conductive cotton gauze proved to be efficient and simple sensor of applied pressure. To impart wash stability to this sensor, we shielded the functionalized surface with an additional hydrophobic, perfluorinated siloxane coating through vapor deposition. The hydrophobic nature of this coating has been shown in prior work to provide the fabric with a strong protective layer against aging processes such as washing or oxidation [56].

We note that this is the first time an ion-conductive cotton cloth was created and incorporated into an all-textile sensor. Commercial textile coatings are aimed at simply imparting hydrophobicity (for stain-repellant fabrics) or creating antimicrobial surfaces. For both functionalities, the necessary coatings are electrically insulating and, therefore, known iterations of functionalized cotton cannot be used in the design of the resistive sensor described in this work.

Sensor model: Having described the sensor chemistry, we now present an electrical model and explain its behavior under pressure. Figure 3 shows the structure of our layered sensor and its electrical equivalent model. The resistance of the functionalized fabric is high enough that we can deal with a range of pressure but low enough that we can use moderate sized resistors in our circuit to minimize noise.

According to Equation 1, the resistance of a transmission medium is inversely proportional to its thickness.

$$R_{eq} = \rho \frac{l}{A} \tag{1}$$

where ρ is electrical resistivity l is the length, and A is the cross-sectional area of the medium. In our design, l represents the thickness of the middle fabric layer – R_{fabric} in Figure 3.

Let us now see how the sensor works under pressure and what aspects of pressure we can measure. Upon applying inward pressure on two outer fabric layers, we see the two simultaneous phenomena. First, the number of resistive routes between two conductive patches is increased because the air gap reduces between the layers. At the same time, the thickness of the fabric is reduced and the capacitance of the device also changes. Both these factors contribute to reduction in impedance of the fabric as a result of increase in pressure.

From a measurement perspective, it is much simpler to design a circuit to measure resistance changes than capacitance changes, therefore we focus on the resistance changes to measure the ballistic signal. To follow the pressure applied on the fabric, we use a voltage divider to produce a voltage that follows the changes in resistance of the fabric. This voltage contains information about the pressure applied to the fabric; however, it is too coarse grained to be useful for extracting vital signs. This signal is then filtered and amplified in the analog domain before being used for respiration and heartbeat detection.

The circuit schematic is shown in Figure 3 (right). Due to the very small signal generated by heartbeats, we need to increase sensitivity from the source. In other words, the whole design needs to be tuned in such a way that changes in R_{var} can cause maximum possible impact on output voltage. This means we need to increase $\partial V_{press}/\partial R_{var}$.

$$\begin{aligned} V_{press} &= V_{dd} \times \frac{R_{var}}{R_{var} + R_1} \\ \longrightarrow \frac{\partial V_{press}}{\partial R_{var}} &= V_{dd} \times \frac{R_1}{(R_{var} + R_1)^2} \end{aligned} \tag{2}$$

These equations show that sensitivity decreases as R_{var} increases. Maximum sensitivity is achieved when $R_{var} << R_1$. Naively, this can be achieved by choosing an extremely large R_1, however, very large output resistance of the sensor can result in a substantial amount of noise to be injected into the electronics circuit. The more sensible approach is to decrease the resistance of the fabric layer, so we carefully tuned the resistance of the textile to the desired regime.

3.2 The Triboelectric Patch

The second sensor is designed to measure ballistics under very low pressure situations, such as when a fabric rests on the chest when a user is standing or sleeping. While breathing is far too slow to induce sufficiently large changes to the textile to be be detectable, ballistics due to heartbeats induce rapid impulses. While the magnitude of this change is quite small and imperceptible to the naked eye, the dynamics are quite large due to the rapid changes in flow resulting in a strong ballistic force on the chest wall. These facts motivated us to use a triboelectric patch as an additional fabric based sensor to capture dynamics of the body.

The triboelectric patch that we use in this work is constructed using the technology that was proposed by prior works on the triboelectric textiles [36, 55, 61]. Triboelectric textiles measure motion via charge transfer — the tribo patch comprises two dielectric layers which transfer charge between them as the distance between them changes during various movements. In our case, the voltage generated by a triboelectric patch is related to the speed of contact and separation between two fabric layers which allows us to extract the ballistic changes due to heart beats. While triboelectric materials have been used for sensing the movement of joints [36], we are unaware of prior work on leveraging this technology to detect vibrations caused by ballistics of the human heart. We briefly describe how the triboelectric sensor works in this section.

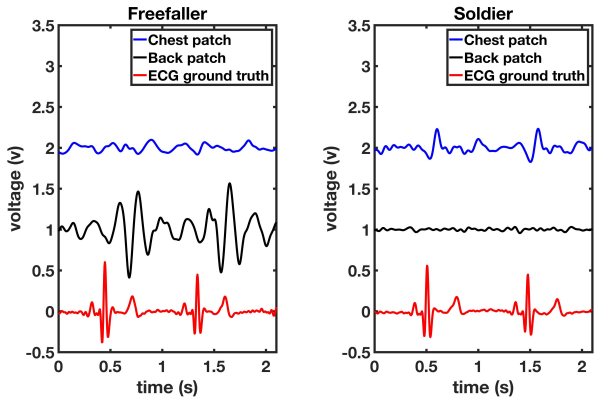

Fig. 4. Figure shows an example of two patches together with ECG ground truth in freefaller posture i.e. lying on front (left panel) and soldier posture i.e. lying on back (right panel). The subject has relatively high body weight, so the patch with higher pressure shows lower sensitivity. We see that each sensor works well in one posture but not the other, demonstrating the need for multi-sensor fusion for robust estimation of physiological measures. (Note that signals are shifted vertically for better presentation.)

The triboelectric textile sensor was created by the face-to-face layering of two different cotton cloths with opposing equilibrium surface charge characteristics [36, 55, 61]. Commercial polyurethane-coated ripstop nylon, which is commonly used for water-repellant outerwear, displays a negative surface charge value, on average, across various ambient environments, due to the presence of the negative triboelectric material, polyurethane. Cotton functionalized with an aminopropyl siloxane, on the other hand, displays a positive surface charge value, on average. When these two fabrics are sandwiched between two silver-nylon cloths, a triboelectric device is formed.

Upon application of pressure, the two oppositely-charged cloth sheets are forced into physical contact, upon which a small amount of surface charge transfer occurs, creating an observable electrical signal. However, this charge transfer event is quickly reversed and the signal quickly decays, even if constant pressure is applied. Due to this behavior, triboelectric devices are perfectly suited for detecting dynamic changes in pressure as a result of ballistics of the heart.

One question that we have not answered is why we need two types of patches — is the resistive patch not sufficient to measure pressure due to ballistics? The difference is that the resistive patch is designed to operate under pressure i.e. it can measure ballistics when sufficient pressure has been exerted on it. The triboelectric patch is designed to operate under very light pressure, for example, due to the textile resting on the body or a thin blanket over the textile. Under higher pressure, there is insufficient change in distance between cloth layers to cause measurable change in charge transfer. Thus, the two types of patches are complementary and cover medium to high pressure situations (resistive) and low pressure situations (tribo).

4 SIGNAL PROCESSING IN PHYJAMA

Having described the design of the textile sensors, we turn to the analysis of the signal from Phyjama. We first provide an overview of design goals and challenges followed by a description of the processing pipeline. An overview of the pipeline is shown in Figure 5.

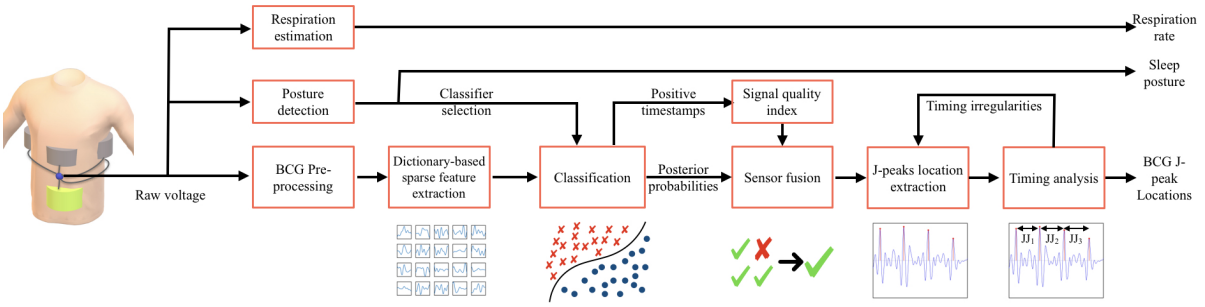

Fig. 5. Phyjama signal processing pipeline.

4.1 Goals and Challenges

Our goal is to provide a comprehensive set of physiological measures of respiratory and cardiac rhythm. These are valuable for many applications — sleep stage classification [58], sleep quality estimation [31], recovery during endurance training [54], stress management [20, 29, 51], and disease prediction [21, 47, 53]. In addition to cardiac and respiratory rhythm, we also get sleep posture as side information from the Phyjama by leveraging the fact that we have several patches on the textile.

The central challenge that we face is that the signals observed by patches depend on several factors including posture, user weight, textile fit, and extent of contact between textile and the body. An example is shown in Figure 4. Two different postures are looked into in this figure, namely, *soldier* and *freefaller*, which refer to lying on back and front. In each case, we see one of the patches performing poorly while the other patch provides a clearer signal. Thus, we see that to obtain robust physiological measures under different real-world situations, we clearly need to fuse information from different sensors.

4.2 Estimating Posture and Respiration

The analog signal from the resistive patches can be directly used to estimate two measures: a) respiration based on baseline variations, and b) posture based on relative pressure across patches.

Estimating posture: The DC baseline directly provides the pressure for each patch which, in turn, gives us information about the contact between different patches and the body. This information can be fused to determine posture. We focus on sleep postures for the Phyjama; in this case, we find that the baseline signals from the patches are highly distinct and a simple decision tree performs near perfectly in distinguishing between postures. We note that posture is useful in two ways in this pipeline — first, sleep posture is a useful output measure by itself, and second, posture is useful to develop a posture-specific classifier that performs better than a posture-agnostic one.

Estimating respiration: The DC baseline can also be used to obtain respiration rate in a straightforward manner. To accurately estimate respiratory rate of the user, we perform two steps. First, we find the frequency bin with the highest power resulted from respiration signal. Second, we perform band-pass filtering based around the FFT peak to avoid counting fluctuations of the second harmonic. The result of the band-pass filter is a signal oscillating around zero. We count the number of zero crossings and divide this number by the duration of the signal to find duration of a half cycle. Since we get a respiration measure from each sensor, we take the median across the four resistive patches to obtain an aggregate measure of respiration rate.

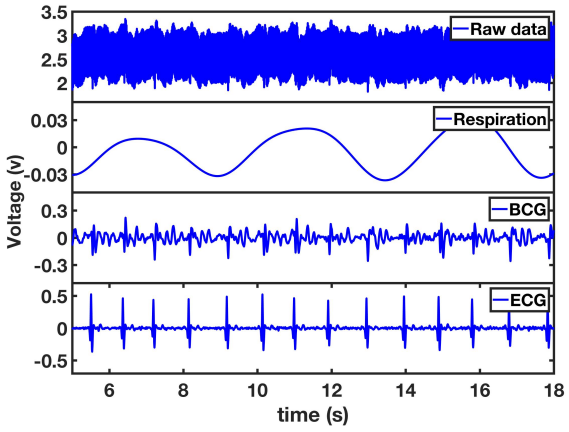

Fig. 6. Top panel shows the raw output signal with all noise sources; second panel shows the respiration signal that is extracted from the low-frequency components of the raw signal; third panel shows the filtered BCG signal; and fourth is ground truth ECG.

4.3 BCG Signal Pre-processing

Unlike posture and respiration, estimating heart rhythm is more challenging. In particular, our objective is to obtain beat-to-beat interval information from the Phyjama signal which can be leveraged to estimate metrics such as heart rate, heart rate variability, sleep stages, and others.

Figure 6 illustrates the challenge in determining the positions of individual BCG peaks. We can see that the respiration signal is quite clear but the BCG signal is more variable and has many peaks that could be misclassified as heart beats. The rest of this section describes our processing pipeline to detect individual heartbeats and peak locations.

Since BCG is a very weak signal, we need to first perform pre-processing to filter out various noise sources (to the extent possible). The output voltage is a combination of DC offset generated by amplifiers, low frequency components corresponding to respiration, higher frequency components corresponding to the BCG signal, and noise in all frequency bins.

4.4 Feature Extraction from Resistive and Tribo Patches

The BCG signal is dependent on which type of sensor we use – the resistive sensors sense pressure changes whereas the tribo sensor measures surface charge transfer. Since these are very different types of signals, we use different feature extraction techniques for these sensors.

Sparse coding features for resistive patches: While ECG feature extraction has been studied for many decades, applying existing techniques to the problem of extracting BCG features from the resistive patches is non-trivial for two reasons. First, the BCG signal varies depending on where the patch touches the body since the ballistic signal is impacted by the skeletal structure, particularly the spine. Second, the types of noise in the patches also differ because motion-induced artifacts like static noise is different across the different locations. This diversity means that traditional detectors can provide sub-optimal performance when subject to these variations. Our work therefore uses unsupervised methods for robust feature extraction to deal with a range of signal variation and noise sources observed in the ballistic signal.

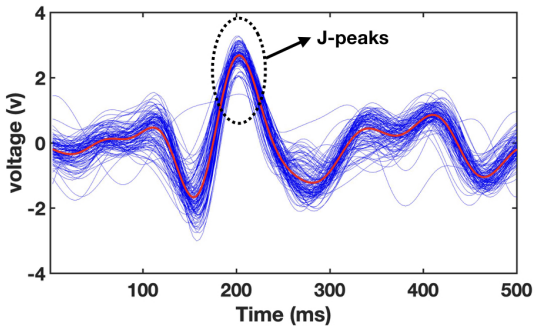

Fig. 7. Multiple traces of BCG signals plotted altogether. Note that amplitude of the signals are normalized in standard deviation.

The approach that we will explore is sparse coding, which has become popular in signal and image processing since it can leverage vast amounts of unlabeled data to generate features [24, 40, 41, 44, 60]. This method has also been applied to a limited extent in the context of ECG signals [15, 39, 42] and BCG signals [35]. The general idea in sparse coding for physiological waveforms is to extract a dictionary of features for detecting the various peaks (e.g. P, Q, R, S, and T in the case of ECG) in a robust manner despite extremely noisy data. In the context of Phyjama, we use sparse coding to learn a sparse dictionary of shapes of the ballistic signals observed at different fabric patches. We provide a brief overview of sparse coding and then describe how we utilize this technique in Phyjama.

Sparse coding is a method for representing a feature vector X in terms of sparse linear combinations $\sum_{k=1}^{K} \alpha_k B_k$ of a set of K basis vectors, B_k. Given a set of basis vectors B_k, the sparse coefficient vector alpha is computed as the solution to the following l_1 regularized optimization problem:

$$\arg\min_{\alpha} \left\| X_n - \sum_{k=1}^{K} \alpha_k B_k \right\|_2^2 + \lambda \left\| \alpha \right\|_1 \tag{3}$$

Given a data set $D = \{X_n\}_{n=1:N}$, the basis is learned to minimize errors between each data case and its reconstruction with the constraint of sparse coefficients. The typical approach to solve this is by using an alternate minimization strategy [23]

Our aim is to recognize the highest BCG peak, also known as J-peak, using sparse coding. Figure 7 shows several instances of such a window overlaid on top of each other for the patch on the users back. We can clearly see that the BCG waveform that we observe via Phyjama is very similar to the pattern presented in literature [19]. In order to find a J-peak, a peak detector with a fairly relaxed threshold is applied over the signal to over-generate candidate peaks.

Note that the sparse coding can be used to learn an over-complete basis in a fully unsupervised manner. This is attractive since we do not need a new user to provide labeled data and can simply expand our dictionary by leveraging raw data from a new user. This can allow us to construct a more representative population-level dictionary without requiring additional labeling overhead for a new user.

Using parameters defined for sparse coding, a dictionary of basis vectors are learned from the time series windows we cropped over candidate peaks. As a result, each window can be represented by a series of weights corresponding coefficients for linear combination of dictionary elements to recreate the window. These weights are used as features for the classification stage.

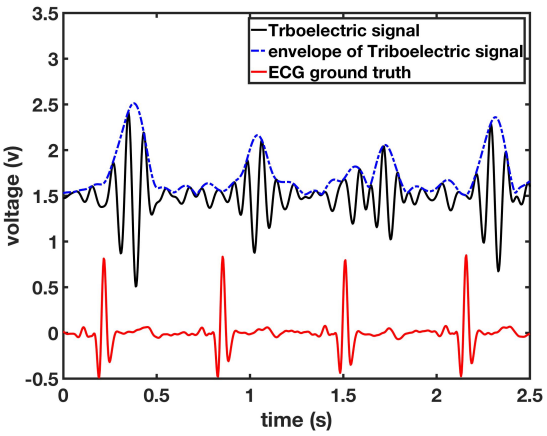

Fig. 8. Triboelectric behavior plotted alongside ECG ground truth. The envelope of the signal has a clear relation with ECG R-peaks. This provide rich information for our classifier model to learn correct J-peak labels.

Feature extraction from tribo patch: The signal obtained from the triboelectric patch is different from the canonical BCG shape that we observe with movement (or pressure) sensors. In the tribo case, we are observing the charge and discharge of the triboelectric material which approximately corresponds to how it compresses and releases as a consequence of the ballistics.

Figure 8 shows an example of the triboelectric waveform — we can clearly see that the ballistics of the heart causes the tribo signal to oscillate much like a spring-mass system with some damping due to the textile properties. The figure also plots the envelope of the triboelectric signal — the amplitude of the envelope roughly correlates with the amount of mechanical energy on skin surface.

One issue with the triboelectric waveform is that the oscillations of the tribo signal does not follow the canonical structure of a BCG waveform. The absence of a clear structure makes it harder to pinpoint which peak corresponds to the J-peak and which corresponds to the other peaks. In addition, the signal peak is also variable and unstable since there is relatively weak contact between the tribo patch and the body (given its location on the stomach).

Instead of using peaks, we use the envelope of the triboelectric signal as the source of features — the envelope loses information about the location of the peaks but is more robust to outliers. After obtaining the envelope of the triboelectric signal, we typically see a correlation between location of the peak of this signal and the expected location of a J-peak. Using this insight, we take 5 samples of the envelope signal with 100 ms interval and use those values as tribo features for classification.

4.5 J-Peak Classification

The next stage classifies the candidate peaks into valid or invalid BCG J-peaks. This stage is executed per-patch i.e. we classify peaks for each patch separately in this stage and then fuse them in subsequent stages.

To perform J-peak classification, the first step is to collect labeled data using an ECG sensor as ground truth. Depending on placement of each fabric patch, the BCG J-peak will have a small delay in regards to its corresponding ECG R-peak. This delay is called the RJ duration and is affected by many factors including an individual's medical condition and patch placement. This duration can reach up to 300 ms [27]. To account for

this delay, we label the largest peak that appears within a 400ms window after an ECG peak as the BCG J-peak. We then manually check a few cases per sensor to ascertain that the labeling is valid.

We use five sets of features for our classifiers: a) the sparse coding feature weights corresponding to our dictionary, b) the posture information coming from the DC baseline, c) the amplitude of the peak, d) five samples from the envelope of the tribo patch centered around the peak, and e) five samples from the envelope of the resistive sensor patch centered around the peak. These features are used to classify each candidate peak.

Once we have the features, the classification model can be any simple machine learning model. We use a linear SVM in our work but other models are equally viable. The classifier is trained based on sparse coding weights and other mentioned time-domain features and the labels provided for each candidate peaks.

At this stage, we also obtain a classification score for the classification of each peak. The classification score is the signed distance from the SVM decision boundary; we use this score in the fusion stage to combine the data from multiple sensor streams and improve the overall results.

4.6 Multi-patch Fusion

The next stage of the processing pipeline fuses the outputs of the individual per-patch classifiers to determine the location of each J-peak in a more accurate manner.

To fuse the outputs, we first need an estimate of the quality of the measurements from each patch. To obtain this, we start by defining a signal quality index that seeks to identify which patches provide the most relevant information so that we can assign more weight to the output from these patches. The signal quality metric that we define is based on the observation that a poor quality sensor generally has high variance in the inter-peak intervals since it has more false positives and false negatives. Thus, we define the Signal Quality Index (SQI) as:

$$SQI_{p,u,s} = 1/std(II_{p,u,s}) \tag{4}$$

where $II_{p,u,s}$ refers to array of inter-beat intervals for each measurement on user u, in position p, and from sensor s. Each element of this array is calculated as the duration between two corresponding consecutive peaks classified as correct J-peaks. :

$$II_{p,u,s}(i) = T^{j}_{p,u,s}(i) - T^{j}_{p,u,s}(i-1) \tag{5}$$

Given the SQI per sensor and classification score for each peak of each sensor s from the SVM classifier, we define the fused score for each peak i as the weighted sum across all sensors. In other words, we simply sum up the scores across the different sensors while considering SQI as weight for each sensor.

$$\text{Fused Score}(i) = \sum_{s=1..4} \text{Score}(i) * \text{SQI}(s) \tag{6}$$

Next, we find the J-peak timestamps by locating positively scored candidate peaks in close proximity. In our implementation, positive labels from different sensors placed closer than 100 ms from one another are considered as an acceptable interval to place a BCG J-peak. To remove false positives and false negatives, we remove the peaks that are too close and re-instate peaks at appropriate locations when we see gaps that are much larger than the average interbeat interval.

The overall process is illustrated in Figure 9. On the left, we see over-generated peaks (red dots), each of which is classified by the per-sensor classifier. In the middle box, we see classification scores for each of the peaks and only a small number have a positive peak. The panel on the right shows the fusion stages using the aggregated scores across sensors — the thin blue rectangles represent the first search intervals with high fused score, and the large green window is the second search stage where the next highest score is selected to fill a missing peak.

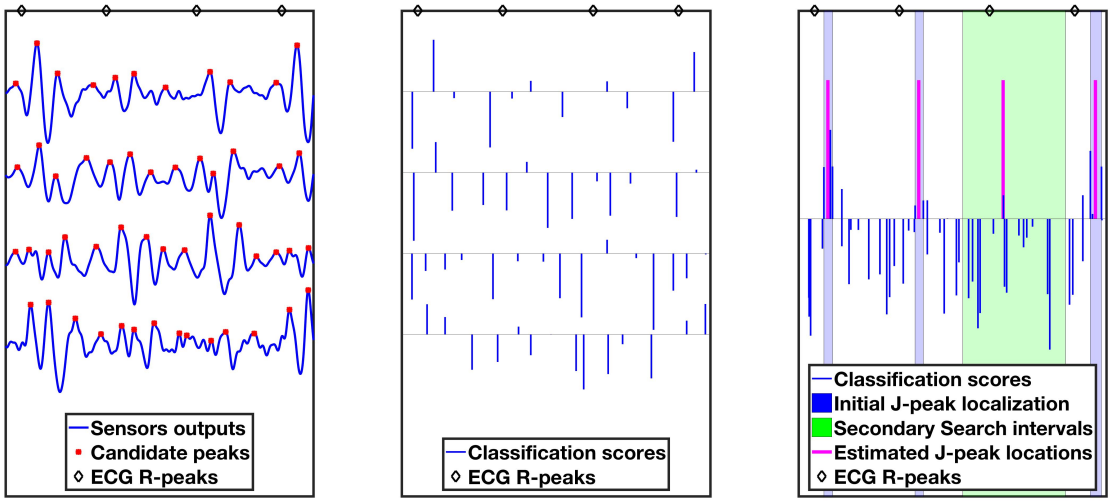

(a) Candidate peaks (b) Classification scores (c) J-peaks location estimation

Fig. 9. Steps taken to estimate locations of BCG J-peaks. a) depicts overgenerated peaks from each sensor and ECG ground truth, b) classification result in form of classification score for each peak, c) J-peak estimation using fused scores, and filling missing J-peaks when gaps are too large. (There are slight timing differences across sensors due to their position, so we fuse scores across a small window.)

5 IMPLEMENTATION

The implementation of Phyjama takes into account several aspects including aesthetics, robustness, and signal quality. From an aesthetics and manufacturing perspective, we wanted to rely solely on textile-based elements for sensing with zero discrete components. This has numerous advantages – the most obvious is that user comfort is maximized if we minimize discrete electronic components at sensitive pressure points[1]. But equally important are the manufacturing advantages since it is much easier to design and fabricate all-textile clothing, and it is much easier to make textile-based elements washable with appropriate hydrophobic coatings.

5.1 Layered Structure of Resistive Sensor

The resistive sensor is comprised of two layers of ion-conductive functionalized cotton gauze, sandwiched between two sheets of silver-plated nylon fabric (purchased from LessEMF). All the textiles were sonicated in water for 15 min, and then rinsed with isopropanol and dried in the air prior to use. To chemically graft the surface of the cotton gauze (purchased from Joann Fabrics Co.), the textile was soaked in N-trimethoxysilylpropyl-N,N,N,-trimethylammonium chloride/isopropanol (15:100 V/V) for 30 min and then cured at 100°C for 2 hours, followed by rinsing with isopropanol and drying in the air. The surface of the functionalized cotton gauze was

[1]In the popular fable "the princess and the pea", a princess is able to feel a pea through twenty feather beds atop twenty mattresses; in the modern context, comfort is highly prized and many of us are highly sensitive to sleep comfort. Even a small discrete sensor padded with textile layers can, like the pea under the bed, cause perceptual discomfort when sleeping in a particular posture for several hours.

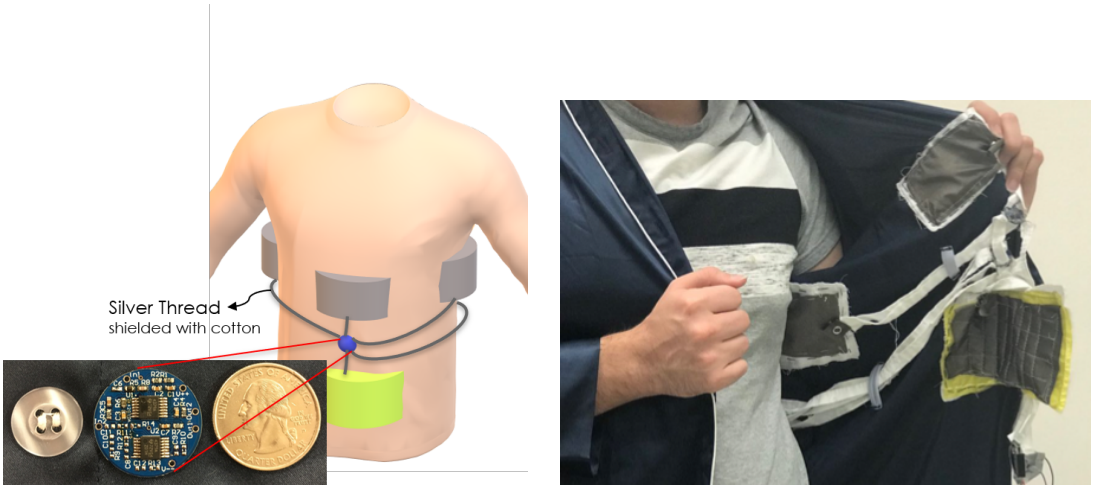

Fig. 10. We illustrate how the patches in Phyjama are interconnected. To avoid using discrete hard components at sensitive pressured locations, the all-textile patches are inter-connected by using silver-plated nylon threads as connectors that are shielded in cotton, The wires coming from each patch end up at a button-size printed circuit board that is placed at the same location of a button on a pyjama. All textile components were placed within the pyjama and were not visible outside.

then modified with a vapor deposition of trichloro (1H,1H,2H,2H-perfluorooctyl)silane, which provides the sensor with washability and durability. The 30-min deposition was conducted in a vacuum custom-built round shaped reactor (290 mm diameter, 70 mm height) at the constant pressure of 1 Torr. The functionalized cotton gauze was then cut into eight 10 cm x 6 cm sheets, each of which was sewn around the perimeter onto a 8 cm x 4 cm sheet of silver fabric. Sewing together each pair of these joined gauze-silver sheets yielded four resistive sensors with a 3-layer structure.

5.2 Layered Structure of Triboelectric Sensor

In the triboelectric sensor, the Polyurethane coated ripstop nylon (purchased from Emma Kites) was used as a negative triboelectric layer. To provide the cotton lycra (purchased from Dharma Trading Co.) with positively-charging surface, the fabric was soaked in (3-aminopropyl) trimethoxysilane/hexane (10:100 V/V) for 30 min, followed by rinsing with isopropanol and drying in the air. The two triboelectric fabrics were then cut into 17 cm x 13 cm sheets and sewn together as they were being placed between two 15 cm x 11 cm sheets of silver nylon fabric. All the chemicals were purchased from Sigma Aldrich Co.

5.3 Assembling the Phyjama

Having designed the individual fabric patches, the next question is how to design and interconnect fabric patches in a way that minimizes the number of discrete hard electronic components. Since the pyjama is designed for maximum comfort, we avoided using wires in our design. instead, we used conductive threads shielded by normal cotton to pass the wires through pajama. Specifically, we used silver-plated nylon threads as connectors (purchased from LessEMF). The threads were shielded in a fabric rod made from cotton (purchased from Dharma Trading Co.) and attached to the silver fabric sheets through snap buttons.

Using these conductive threads, the sensor patches were connected to button-sized PCB boards — two of these boards were responsible for four resistive patches and the last board is connected to triboelectric sensor. While

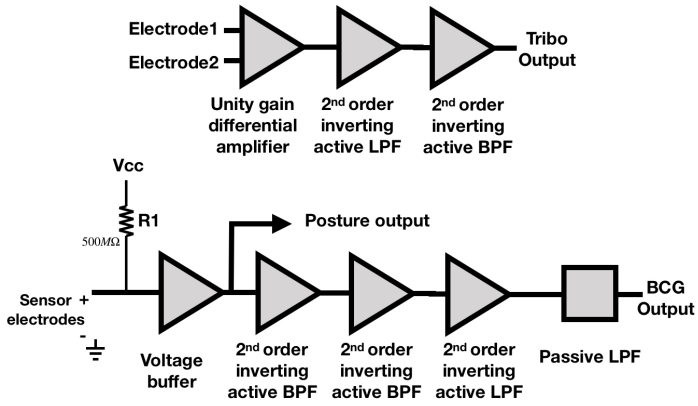

Component model	Description
INA132	Low Power, Single-Supply Differential Amplifier
AD8502	μpower CMOS operational amplifiers
OPA4379	Rail-to-Rail I/O Operational amplifiers
AP7312	3.3 v voltage regulator

Fig. 11. Block diagram of analog circuit boards and components used. Top figure (left) shows the design for triboelectric board and bottom (left) shows design for pressure sensing boards. The right panel lists the components used in analog circuit boards.

these can potentially be combined into a single platform, we designed separate boards for ease of prototyping. We designed the boards to have small form factor, roughly the size of large buttons as shown in Figure 10. We believe the size can be further shrunk down to half the current size and they can be integrated into the buttons of a pyjama. All boards are powered using a single 3V battery.

The PCB boards are optimized for two design goals. First, BCG signals are typically within the 1-10Hz frequency range [49], and the peak power of the BCG signal is in 7-8 Hz frequency bin [27]. We leverage this information to choose a cutoff frequency of 4-10 Hz for faster DC rejection and capturing the strongest BCG frequency component. Second, there is significant power line noise that needs to be rejected to obtain a clean signal. This is complicated by the fact that the noise depends on the proximity of the conductive layer to the body, so an inner layer has more noise than an outer layer, making it difficult to remove noise by differential amplification. The two amplification and noise rejection pipelines that we designed are shown in Figure 11 — the top figure shows the pipeline for the triboelectric patch which comprises an inverting active band-pass filter, an inverting active low-pass filter, and a unity gain differential amplifier. The bottom figure shows the pipeline for the resistive patch which comprises two inverting active band-pass filters, a passive low-pass filters and an inverting active low-pass filter in addition to a voltage buffer.

Our designs went through several fabrication iterations to improve signal-to-noise, reduce form-factor, and reduce power consumption. In its current form, the board for the resistive patch draws about 150μA whereas that for the tribo patch draws 1 mA. In total, the power consumption of the analog boards are around 1.7 mA from a 3.3 v power source, which leads to 5.6 mW of power consumption. The microcontroller/radio board consumes an additional 15 mW when using Bluetooth. Minimizing overall power consumption is possible by further improving the amplifier and improving duty-cycling but we did not exhaustively explore these directions for this paper.

5.4 Optimizing Patch Placement

Optimizing the location of the sensor patches is an important step in our implementation since the signal is sensitive to placement. While this process may eventually be optimized to different body types or even personalized, we optimized patch placement to one individual and used the same setting across all participants.

To find the best placement for the resistive patch on the back and front, we placed a patch at different locations and measured the signal quality while the user is lying down on their front and back, respectively. The

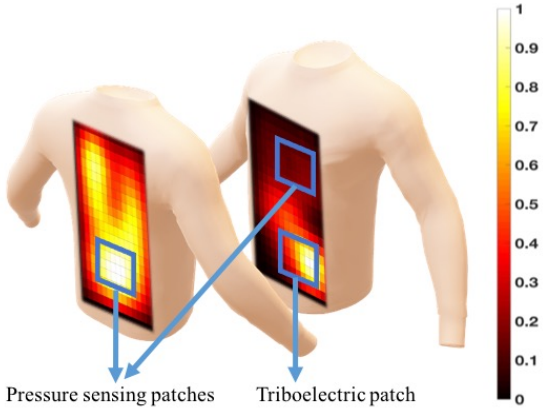

Fig. 12. Ballistics signal power measured across different points on user's body. Signal power is normalized and is plotted as a heat map. Location with strongest signal is allocated to triboelecrtic patch and seconds strongest places are allocated for pressure sensors.

measurement setup was carefully done to minimize folds of the textile and random body movements so that we can isolate the effect of BCG on the output signal.

The patch is placed on 12 different positions on user's back and for each position, 5 measurements are performed each with duration of 30 seconds, resulting in total on 150 seconds of data for each position. Then, J-peaks are manually labeled and the average amplitude across all J-peaks are considered as signal quality factor for each patch, resulting in a 3×4 matrix. The result is then interpolated to achieve higher resolution. A heat map is generated from resulted amplitudes and plotted in Figure 12.

We observe that the front has superior signal strength compared to the back, especially in the stomach area. This is because the spine and rib cage diminish power of heart ballistics. Our decision for where to place the triboelectric patch was also empirically determined. We used only one triboelectric patch to reduce the complexity of dealing with too many patches. While multiple locations may have worked for the tribo patch, we noticed that the worst posture for the resistive patch was when the user was lying on their back, particularly when the individual has high body weight. In this case, the triboelectric patch could compensate for a poor signal from the resistive patch since it can provide an accurate heart rate signal even when the textile is just lying on the subjects chest. Since we could not place both the resistive and tribo patches at the same location, we moved the resistive patch on the chest to its second-best position.

5.5 Data Acquisition

The need to collect raw data from all patches and from a ground truth measurement device (such as ECG or PPG sensor) presented some challenges due to the large number of channels and cumulatively high data rate requirements. For benchmark studies where the subject was stationary, we largely used a tethered setup where the sensors were connected to an eight channel data acquisition unit with 286 *Sample/sec* rate and 16-bit resolution. But this was too limiting for longer-term experimentation.

The received data is processed using MATLAB — we use the SPAMS sparse coding toolbox to extract the sparse dictionary and features, and the Support Vector Machine classifier to classify the peaks using Leave One Subject Out method.

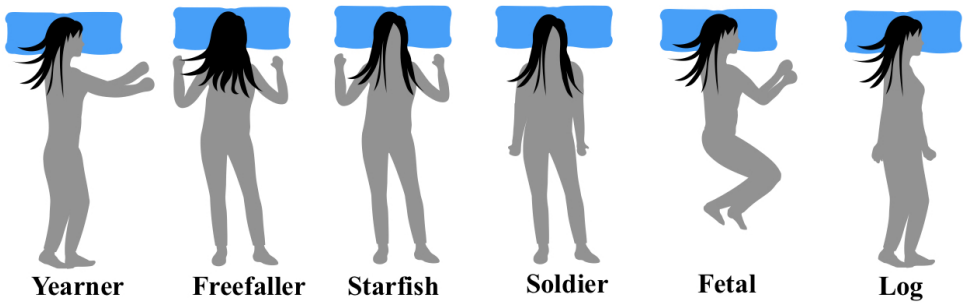

Fig. 13. Most common sleeping postures According to a study conducted by Chris Idzikowski [1].

6 DATASET COLLECTION AND GROUND TRUTH LABELING

In this section, we describe two user studies — the first is a benchmarking study to evaluate different building blocks in Phyjama across different physiological and physical parameters of interest, and the second is a longer term study with elderly participants in a more uncontrolled and naturalistic setting. All of these datasets were collected under Institutional Review Board approval.

6.1 Benchmarking Dataset

For this dataset, we asked 21 participants aging from 22 to 38 years old to wear Phyjama and we recorded the output voltage in various stationary conditions. We instrumented a size *XL* pajama shirt with our sensors. However, we did not restrict our recruitment solely to participants of this size since sleepwear is often larger than normal wear and does not always fit exactly to an individual's size (it is also not uncommon to wear larger sleepwear sizes). Participants varied in weight, 107-240 lb, and height, 5'1" to 6'4". 9 out of 21 participants were females.

We collected data in a variety of postures for each individual including six sleep postures and two other stationary postures. Sleep postures are typically classified into six categories as shown in Figure 13 [1], so we collected data from users in all of these postures. In addition to sleep postures, we also look at sitting on a chair and standing as two other postures of interest since they provide a contrast against sleep postures. In particular, standing represents the most difficult scenario since there is no pressure against an external surface to rely on.

The duration of each of these measurements is one minute which leads to total of 8 minutes of recording from single user. Each recording consists of five channels, four of which correspond to pressure sensing patches and one corresponding to the triboelectric patch.

Since the system is designed to capture vital signs, we also need ground truth for the physiological signals. For heart rate, we used a three-channel ECG measurement (2 wrists and an ankle) using the AD8232 evaluation board, [17], and for respiration, we used a PPG sensor to track respiration, Pulsesensor [4].

6.2 Nap Study Dataset

The nap study is designed to evaluate our methods under more realistic conditions. For this, we designed another study where participants are asked to take a nap for one hour while wearing Phyjama. The study was conducted in a sleep study testing center that is specifically designed for naturalistic sleep studies and mimics a realistic environment.

The dataset for this experiment is collected from seven participants consisting of 3 females and 4 males. Four of the seven participants were between 60 - 70, whereas the other three were between 30 - 40, which provides us sufficient data to evaluate performance across different age groups.

One issue we faced was that the wires for ground truth ECG sometimes interfered with the readings for the Phyjama sensor when the user changed posture in their sleep. This was not a problem with our benchmark study since users stayed in a single posture and we could place wires to avoid interference. To address this, we used a finger-worn wireless PPG sensor that wireless transmitted raw data for ground truth rather than ECG electrodes. While PPG is slightly less accurate than ECG, particularly for estimating the respiration signal, this provides a sufficiently good ground truth signal while allowing comfortable sleep.

7 EVALUATION

We present the evaluation in three parts. First, we benchmark the resistive patch that we have designed and show that it is sensitive to the normal range of human weight and sleep activity. Second, we present an analysis of results for the benchmark dataset. We show that the Phyjama provides accurate physiological measures, and breakdown contribution from the different hardware and software building blocks. Third, we analyze performance "in the wild" using the nap study dataset.

7.1 Resistive Patch Benchmarks

In this section, we present benchmarks of the resistive patch — we highlight this sensor since this is a novel device that has not been previously described in research literature.

Sensitivity to pressure: We first validate our claim that the resistive patch is sensitive to typical range of human pressure. In this experiment, we carefully change pressure applied on a $1.5 \times 2.7\text{inch}^2$ patch and recorded the resistance of the fabric. The measurement is repeated 10 times for each pressure point by re-applying the pressure in various rotation and placements to account for probable folds, asymmetry in functionalization and pressure distribution, The response is presented as a box plot in Figure 14(a).

As we can see, the fabric resistance varies monotonically as the amount of pressure is increased. We see that the sensitivity of our pressure sensing patch is inversely related to the amount of pressure applied on fabric surface. To provide a reference, we show roughly the pressure applied by a 240 lb and 107 lb individual when they are lying on their back. We see that our patch is slightly more sensitive to lighter individuals and less sensitive for users who are above 240 lb. Overall, these numbers show that we have good sensitivity in the typical regime of human weight.

Pressure baseline during sleep: We now look at the pressure baseline in a dynamic setting when a user transitions between sleep postures. Figure 14(b) shows the pressure baseline for the different patches (*posture output* in Figure 11). Note that the voltage being measured via the voltage divider circuit in Figure 11 is inversely proportional to the pressure, so lower voltage means higher pressure.

The figure shows that the resistive patch is highly responsive to the range of human pressure. In the soldier position (lying on the back), the back patch has the lowest voltage and the front patch has highest voltage. When the subject transitions to the foetus position, the left patch becomes pressured (since the subject is lying on the left side), whereas the three patches are not under much pressure. Finally, in the freefaller position (subject on stomach), the front patch sees the highest pressure as we would expect. The physiological signals are also visible in the figure — the slower oscillations correspond to the respiration waveform whereas the more rapid but smaller ripples correspond to the heart beat signal.

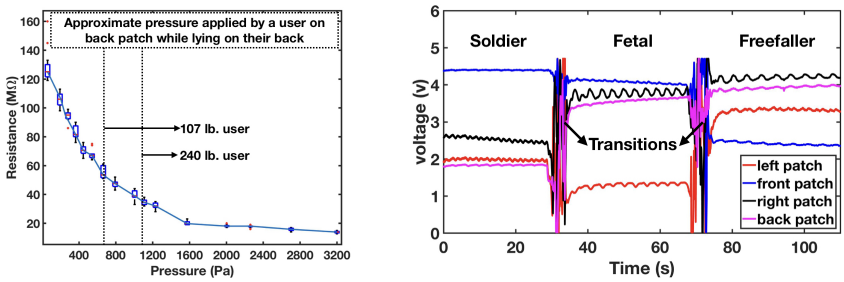

(a) Fabric resistance versus pressure (b) Signal baseline across postures

Fig. 14. Figure benchmarks the resistive patch. On the left, we see that the resistive patch has good sensitivity in the typical regime of human weight. On the right, we see that posture and respiration information is clearly visible from the resistive patch (heart rate extraction requires further processing)

7.2 Measuring Physiological Parameters

In this section, we evaluate the performance of Phyjama in detecting key physiological variables of interest — heart rate and respiration rate. We note that while posture is also an output of the Phyjama, we do not explicitly present results for posture detection since this is trivial to detect from the analog signal across patches. We find that a simple decision tree that looks for the difference between front-and-back patches and left-and-right patches can easily identify posture with 100% accuracy across all subjects. We therefore present results from the other physiological variables of interest.

The users varied across several dimensions including height, body weight, and gender. Among these variables, we found that the most significant impact was due to height which determines where the patch is positioned on the body. Thus, we separated participants into two groups when analyzing the data: the first group consists of participants for whom our Phyjama prototype can fit relatively well and the second group are the ones who are mostly too short to wear Phyjama. For the sake of brevity, these users are called *height matched* and *height unmatched*, respectively. Height matched group includes 11 users whose height vary from 5'7" to 6'3". The rest are height unmatched – this varies quite a bit to include both relatively short and relatively thin individuals (in a couple of instances, the Phyjama shirt reaches just above the knee). Figure 15 shows the results.

Heart rate estimation: Let us first look at heart rhythm metrics. For height-matched users, error in HR estimation is generally less than 1 bpm. The only posture that has high error is *standing* which is to be expected since we do not have any externally pressured surface so we are relying on weaker signals from the pressure of the arm against torso and the tribo patch resting on stomach. But the error is not too high even in the standing case — median HR error is about 2.5bpm. For height un-matched users, the upper quartile and worst-case error is more but the median error is only a little more than then height-matched case (roughly 2bpm HR error).

Respiration rate estimation: The respiration metrics are also very good — median error is generally below 1 resp/minute. In this case, we see that the error is higher for the starfish and soldier positions. This is because the resistive sensor on the back sees a weaker respiration signal due to the spine, and because the tribo sensor on the stomach does not help since it cannot measure slow baseline changes. The signal in this case is primarily due to the resistive sensor on the chest and sensor fusion is less useful in these positions leading to higher error.

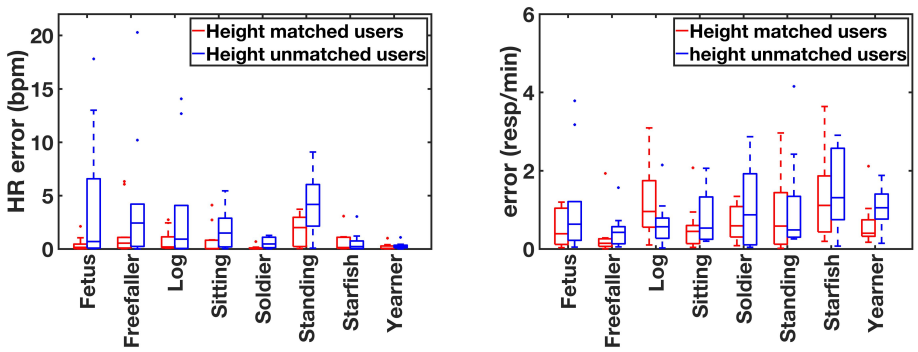

(a) Heartrate error (b) Breathing rate error

Fig. 15. Performance of Phyjama in estimating heart rate and breathing rate in different postures. Left panel shows the error in estimating heart rate and right panel shows breathing rate. In both cases, we observe higher error in standing position and for height unmatched users.

7.3 Breaking Down Contributions

Having discussed the application-level metrics, we now provide a breakdown of how data fusion across different patches benefits overall system performance. Since our system has different hardware components, software building blocks, and application metrics, we provide a few different perspectives on the breakdown to illustrate the advantages of various building blocks. In these results, we do not distinguish between height-matched and height-unmatched users and aggregate results for all users.

Benefits of data fusion on J-peak classification: Here we show the benefits of data fusion in distinguishing J-peaks among all candidate peaks. We provide F1-score as a measure of performance of the classifier. Classification is performed using Leave-One-Subject-out (LOSO).

Figure 16 shows the F1-scores prior to fusion and after fusion. The results on the left show the median F1-score for each posture-patch combination — we see that the results vary quite a bit. For example, the back patch can have poor performance when there is too much pressure on it (log and starfish) or when there is no pressure (fetus), but can offer very good performance in some other positions (freefaller and standing). Similarly, each patch performs better in some scenarios and worse in others. Also, its important to note that no single patch gets an F measure above 90, in fact, in most cases, it hovers between 75-80.

The result on the right shows that F1-score increases dramatically after fusion with median score above 95% in almost all cases. The highest error is for the standing posture for reasons explained earlier. The upper quartiles have somewhat higher error — this is primarily because of the height-unmatched users whose error is higher than the height-matched set.

Benefits of data fusion for HR estimation: We now look at the breakdown from the perspective of an application-level metric, heart rate estimation. We consider three versions of our pipeline. The first version corresponds to the best-case performance when a single sensor is used. We select the best sensor for each user and posture for these numbers; clearly, this is not viable in practice but this gives us an upper bound on single-sensor performance. The second version fuses the posterior probabilities across the sensors without weighting them by the quality index. The third version is the full pipeline with SQI-based weighting.

The result is shown in Figure 17(a). It is clear that sensor fusion greatly reduces the system error (about 4× reduction). The use of a weighted measure using the signal quality index improves results further (about 50%).

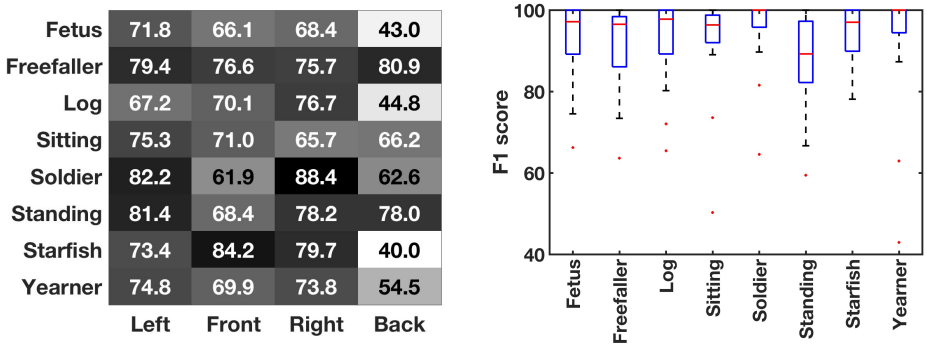

(a) Median F1 score before fusion (b) F1 score after fusion

Fig. 16. F1 scores for classifying J-peaks (aggregated over all participants). On the left, we show the median F1 score from the classification phase prior to fusion. We see that the scores are relatively poor for individual sensors. On the right, we show the F1-score after fusion which is considerably higher; the median F1-score is often close to 1.

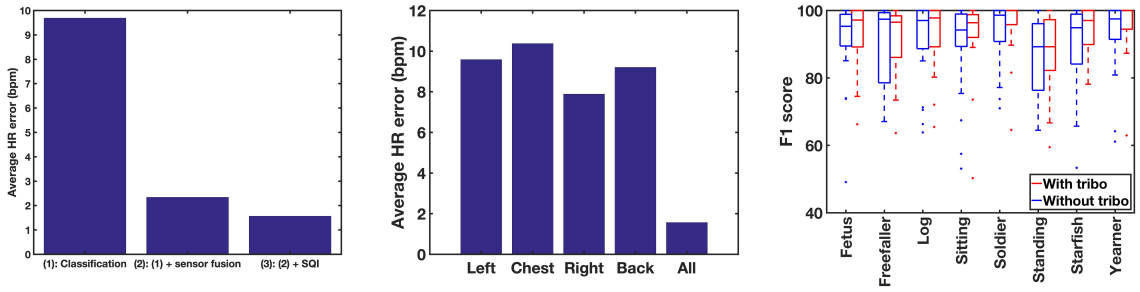

(a) Breakdown by processing blocks (b) Breakdown by sensor patches (c) Breakdown by sensor type

Fig. 17. Breakdown of contribution of different blocks in the signal processing pipeline and sensor hardware. On the left, we show the contribution of different signal processing blocks — we see that sensor fusion plays a huge role in improving Phyjama's performance. In the middle, we show the corresponding performance if only one sensor were used as opposed to fusing information from all sensors together — we see that a distributed set of patches helps improve results. On the right, we see the F1-scores with and without the tribo sensor — we see that information from the tribo sensor helps improve median and reduces outliers.

While not shown, breaking the results down by height-matched versus height-unmatched shows that the numbers are much higher for the height-matched users (6× reduction due to sensor fusion and 2× reduction from using signal quality. Thus, the potential gains can be higher once the textile is matched to the user size.

Sensor contribution: Another interesting question is how much each sensor contributes the overall results and whether there is one sensor that is superior to others in terms of determining physiological measures of interest. To answer this question, we plot the accuracy of Phyjama if only one sensor patch were used and contrast this against the case where the sensor information are fused together. The result is shown in Figure 17(b).

We see that each sensor has high error in its own estimate of HR, however, after sensor fusion, the estimation error drops by 4–5×. This result also highlights the benefits of sensor fusion and shows that any one sensor would not do as well as fusing the readings.

Tribo contribution: Our final set of benchmarks looks at the contribution of the triboelectric patch to overall classification performance. Figure 17(c) shows that the envelope features from the tribo patch is informative and improves overall performance. While average improvement is between 2.5 – 5%, the tribo sensor is particularly helpful with the upper quartile of error cases and outliers which reduce dramatically in many cases.

Overall, these results clearly show the benefits of having a distributed array of sensors on the textile. Unlike traditional wearables like smartwatches that can only measure at a single point on the body, we have five distributed sensors whose information is fused, therefore, the Phyjama can capture a strong signal even if one or two sensors are erroneous due to their positioning.

7.4 Case Study: Nap Monitoring

In this study, we explore the effectiveness of models that we developed from the benchmarking study for a more realistic longer-term case study involving monitoring hour-long naps. We compare heart rate, breathing rate and posture estimated by our hardware and algorithms against ground truth.

We made several efforts to keep the study and evaluation as realistic as possible. We gave the users *no explicit instructions* regarding how to take a nap. They often moved around a bit before lying down to nap. Some of them also used a blanket whereas others did not. We also did not use any method to *personalize* the signal processing pipeline to the users. So, the entire pipeline was trained from the benchmark dataset and directly applied to the new user with no performance tuning or transfer learning.

Posture detection results: In this study, subjects typically changed their posture two to three times during their nap. Given the uncontrolled nature of the study, users often rested in postures that were combinations of multiple base postures. In this experiment, we considered an estimation correct if it was among one of user's contributing postures — for instance, if a user is sleeping on their right with one leg and one arm in fetal position and the other leg and arm in log position, both fetal and log are counted as acceptable postures. We found that the posture classification block detects all postures correctly i.e. it has 100% accuracy.

HR and BR detection: The results for heart rate and respiration rate are shown for each subject are shown in Figure 18. We see that the results are generally quite good, with median error less than 1 cycle for breathing rate and 2.5bpm for heart rate. We note that for subject #3, the resistive sensor on the right failed. Despite this issue, both heart rate and respiration estimates are very good, demonstrating the benefits of sensor fusion.

One aspect that we believe can be improved is the upper quartile errors. The reason for these errors are many — users often moved around before they fell asleep and some had leg movements during their nap whereas our model is trained from the benchmark dataset which only included stationary data. Also, as with our benchmark study, there is sometimes a size mismatch for the Phyjama shirt that we used since we did not restrict height and weight for users we recruited. Despite these differences, the sensor is quite robust and can get a good signal a significant fraction of the time. We expect these results can be substantially improved with more labeled data, as well as personalization strategies to tailor the signal processing pipeline to each user. But the results are very promising and suggest that there is significant clinical utility for the Phyjama.

7.5 User Comfort

Finally, we look at subjective measures of comfort that were obtained from the users. A major advantage of the Phyjama is the comfortable and unobtrusive nature of its design. The sensors are integrated into everyday nightwear with discrete elements placed in expected locations like a button; in addition, users do not need to

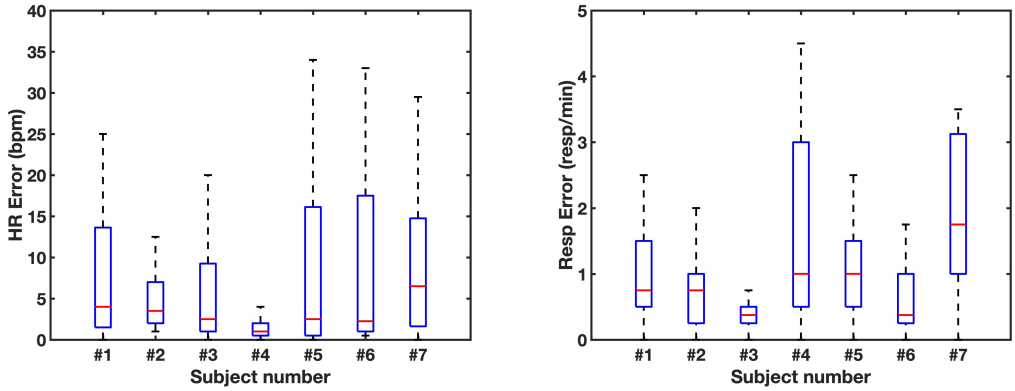

(a) Heart rate error (b) Respiration rate error

Fig. 18. Results for longitudinal measurements for heart rate and respiration rate across all users. On the left, we see the error for HR estimates across the seven users — accuracy is quite good with median error of 2.5 *bpm*. On the right, we see the respiration rate estimates which is also very good and has median error of 2.5 *bpm*. Despite one of the sensors failing for subject #3, estimates are good because there is sufficient signal across other sensors.

Table 1. Summary of subjective assessment report during one hour experiment

Question	**Benchmarking study**	**Nap study**
How comfortable is Phyjama?	Avg = 4.95	Avg = 5
Are you interested in tracking vital signs during sleep?	YES=17 NO=4	YES=6 NO=1
Would you prefer Phyjama or Fitbit?	Phyjama=16 Fitbit=5	Phyjama=5 Fitbit=2
Does Phyjama interrupt your breathing?	NO=21 YES=0	NO=7 YES=0

remember to wear an additional device that would be unusual during sleep like a fitness band. To evaluate comfort, we asked four questions from the participants:

(1) How comfortable is Phyjama? (rate from 1 to 5, 5 being very comfortable).
(2) Do you prefer to track your vital signs during sleep? (This was asked as a general question regarding inclination towards logging vital signs.)
(3) If you were to track your vital signs, would you prefer Phyjama or Fitbit [2]? (Regardless of their interest in vital signs tracking, we asked if they would prefer Phyjama over Fitbit for this purpose.)
(4) Does Phyjama interrupt your respiration or impact its pattern? (As a sleep wear, it is very important that the outcome is completely unobtrusive.)

Results are summarized in Table 1. The results show clearly that users found the Phyjama comfortable, unobtrusive and often preferable to wristworn wearables.

8 DISCUSSION AND FUTURE WORK

The design of Phyjama is complex due to its vertical integration of novel fabric-based sensors, analog hardware, and signal processing/machine learning analytics. Given the vast design space, we could not exhaustively explore

all avenues and restricted ourselves to the most promising directions towards designing a high performance system. There is, therefore, significant scope for follow-on work that digs deeper along the different directions.

Real-time execution of the signal processing pipeline: In this paper, we designed a signal processing pipeline without explicit hardware restrictions since we wanted to understand best-case performance. However, in a real-world system, the computational pipeline will need to be optimized to execute locally on the embedded processor integrated with the Phyjama. We intend to explore the use of low-power processors for efficient execution of machine learning pipelines [52].

Reducing power consumption: Our current prototype consumes roughly 5.6mW for the analog front-end, and an additional 15mW for the microcontroller when using the Bluetooth radio. Thus, the overall power consumption is about 20mW which is relatively high for continuous operation. Some of the directions we plan to explore for reducing hardware power consumption include better duty-cycling strategies across the sensor patches using methods such as sparse and burst sampling [30].

Expanding physiological and physical measures: In this paper, we have focused on the use of the Phyjama to measure heart rate, breathing rate, and posture. But the Phyjama opens up many other possibilities including heart rate variability (by measuring inter-peak separation) and blood pressure (by measuring BCG pulse transmit time between chest and wrist using an additional sleeve patch [37]). Another direction is fine-grained posture sensing since people rarely sleep in exactly one of the canonical sleep postures; rather, they often slept in combinations of postures. We plan to explore these avenues in future work.

Manufacturing cost: Another direction we will explore is to fully understand cost and tradeoffs between manufacturing the entire textile with embedded sensors versus sewing on our sensors to an existing pyjama. We use mass-produced cotton fabrics and a widely-available, cheap siloxane reagent (that is used as an antimicrobial coating for medical devices) to create our pressure sensors, and then simple sew these sensors onto any kind of commercial garment. Therefore, in principle, our prototype does not require customized and novel manufacturing routines which means that multiple units can be made with reasonable speed and low cost. We will further explore these tradeoffs in future work.

Personalized design: In this work, we used a single size of the Phyjama across all users but significant improvements in signal quality may be possible if we personalize the design to a users height and weight. For example, if a user has high BMI that would increase the pressure and reduce signal amplitude but we may be able to compensate by using a larger patch. Thus, there is interesting avenues for future work where we look at custom designs of the Phyjama that is tailored to each individual's physical characteristics.

9 CONCLUSION

In conclusion, we have designed an all-textile sensing system for physiological sensing using everyday clothing like pyjamas. Our work has several important applications in healthcare and potentially in other fields such as entertainment. The sensing techniques that we develop in this paper bring together cutting-edge approaches in disparate areas including textile chemistry and functionalization, textile-electronics co-design, and wearable sensing and analytics. The overall architecture and design principles used to combine these elements promise to be greater than the sum of their parts and fundamentally change the way we think about clothing.

ACKNOWLEDGMENTS

The authors would like to thank all of the reviewers for detailed feedback on the paper. This work was supported by the National Science Foundation (1763524).

REFERENCES

[1] Bbc news | health | sleep position gives personality clue. http://news.bbc.co.uk/2/hi/health/3112170.stm.
[2] Fitbit official site for activity trackers & more. https://www.fitbit.com/home.
[3] Garmin vivofit. https://buy.garmin.com/en-US/US/p/582444.
[4] Heartbeats in your project. https://pulsesensor.com/. Accessed: 05/10/2019.
[5] Murata bcg mems sensor. https://www.mouser.com/new/Murata/murata-bcg-mems-sensor/.
[6] Nike men's aeroreact. nike.com. https://www.nike.com/us/en_us/c/running/mens-aeroreact.
[7] Oura sleep sensing ring. https://ouraring.com/.
[8] Pressure mapping, force measurement, & tactile sensors | tekscan. https://www.tekscan.com/. (Accessed on 12/13/2018).
[9] Teslasuit - full body haptic vr suit. https://teslasuit.io/.
[10] Zephyr performance systems. https://www.zephyranywhere.com/.
[11] F. Adib, Z. Kabelac, and D. Katabi. Multi-person localization via rf body reflections. In *NSDI*, pages 279–292, 2015.
[12] F. Adib, Z. Kabelac, D. Katabi, and R. C. Miller. 3d tracking via body radio reflections. In *NSDI*, volume 14, pages 317–329, 2014.
[13] F. Adib and D. Katabi. *See through walls with WiFi!*, volume 43. ACM, 2013.
[14] F. Adib, H. Mao, Z. Kabelac, D. Katabi, and R. C. Miller. Smart homes that monitor breathing and heart rate. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 837–846. ACM, 2015.
[15] A. Adler, M. Elad, Y. Hel-Or, and E. Rivlin. Sparse coding with anomaly detection. *Journal of Signal Processing Systems*, 79(2):179–188, 2015.
[16] J. Alametsä, A. Värri, M. Koivuluoma, and L. Barna. The potential of emfi sensors in heart activity monitoring. In *2nd OpenECG Workshop Integration of the ECG into the EHR and Interoperability of ECG Device Systems, Berlin, Germany*. Citeseer, 2004.
[17] Analog Devices. *EVAL-AD8232 Evaluation Board | Analog Devices*.
[18] G. Andreoni, C. E. Standoli, and P. Perego. Defining requirements and related methods for designing sensorized garments. *Sensors*, 16(6), 2016.
[19] A. O. Bicen, D. Whittingslow, and O. Inan. Template-based statistical modeling and synthesis for noise analysis of ballistocardiogram signals: A cycle-averaged approach. *IEEE Journal of Biomedical and Health Informatics*, pages 1–1, 2018.
[20] J. F. Brosschot, E. Van Dijk, and J. F. Thayer. Daily worry is related to low heart rate variability during waking and the subsequent nocturnal sleep period. *International journal of psychophysiology*, 63(1):39–47, 2007.
[21] Z. Dong, X. Li, and W. Chen. Frequency network analysis of heart rate variability for obstructive apnea patient detection. *IEEE journal of biomedical and health informatics*, 22(6):1895–1905, 2018.
[22] L. E. Dunne, P. Walsh, B. Smyth, and B. Caulfield. Design and evaluation of a wearable optical sensor for monitoring seated spinal posture. In *2006 10th IEEE International Symposium on Wearable Computers*, pages 65–68, Oct 2006.
[23] M. Elad. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer Publishing Company, Incorporated, 1st edition, 2010.
[24] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image Processing*, 15(12):3736–3745, Dec 2006.
[25] J. E. Estrada and L. A. Vea. Real-time human sitting posture detection using mobile devices. In *2016 IEEE Region 10 Symposium (TENSYMP)*, pages 140–144, May 2016.
[26] A. Fathi and K. Curran. Detection of spine curvature using wireless sensors. *Journal of King Saud University - Science*, 29(4):553 – 560, 2017. SI: Smart materials & applications of new materials.
[27] J. Gomez-Clapers, A. Serra-Rocamora, R. Casanella, and R. Pallas-Areny. Towards the standardization of ballistocardiography systems for j-peak timing measurement. *Measurement*, 58:310 – 316, 2014.
[28] E. Griffiths, T. S. Saponas, and A. J. B. Brush. Health chair: Implicitly sensing heart and respiratory rate. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '14, pages 661–671, New York, NY, USA, 2014. ACM.
[29] M. Hall, R. Vasko, D. Buysse, H. Ombao, Q. Chen, J. D. Cashmere, D. Kupfer, and J. F. Thayer. Acute stress affects heart rate variability during sleep. *Psychosomatic medicine*, 66(1):56–62, 2004.
[30] N. Hua, A. Lall, J. Romberg, J. J. Xu, M. al'Absi, E. Ertin, S. Kumar, and S. Suri. Just-in-time sampling and pre-filtering for wearable physiological sensors: going from days to weeks of operation on a single charge. In *Wireless Health 2010*, pages 54–63. ACM, 2010.
[31] S. H. Hwang, H. N. Yoon, Y.-J. G. Lee, D.-U. Jeong, K. S. Park, et al. Nocturnal awakening and sleep efficiency estimation using unobtrusively measured ballistocardiogram. *IEEE Transactions on Biomedical Engineering*, 61(1):131–138, 2014.
[32] J. Jia, C. Xu, S. Pan, S. Xia, P. Wei, H. Y. Noh, P. Zhang, and X. Jiang. Conductive thread-based textile sensor for continuous perspiration level monitoring. *Sensors*, 18(11), 2018.
[33] Z. Jia, M. Alaziz, X. Chi, R. E. Howard, Y. Zhang, P. Zhang, W. Trappe, A. Sivasubramaniam, and N. An. Hb-phone: a bed-mounted geophone-based heartbeat monitoring system. In *Proceedings of the 15th International Conference on Information Processing in Sensor Networks*, page 22. IEEE Press, 2016.

[34] Z. Jia, A. Bonde, S. Li, C. Xu, J. Wang, Y. Zhang, R. E. Howard, and P. Zhang. Monitoring a person's heart rate and respiratory rate on a shared bed using geophones. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*, page 6. ACM, 2017.
[35] C. Jiao, B.-Y. Su, P. Lyons, A. Zare, D. K. Ho, and M. Skubic. Multiple instance dictionary learning for beat-to-beat heart rate monitoring from ballistocardiograms. *IEEE Transactions on Biomedical Engineering*, 2018.
[36] A. Kiaghadi, M. Baima, J. Gummeson, T. Andrew, and D. Ganesan. Fabric as a sensor: Towards unobtrusive sensing of human behavior with triboelectric textiles. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, SenSys '18, pages 199–210, New York, NY, USA, 2018. ACM.
[37] C. Kim, A. M. Carek, R. Mukkamala, O. T. Inan, and J. Hahn. Ballistocardiogram as proximal timing reference for pulse transit time measurement: Potential for cuffless blood pressure monitoring. *IEEE Transactions on Biomedical Engineering*, 62(11):2657–2664, Nov 2015.
[38] F. Lorussi, W. Rocchia, E. P. Scilingo, A. Tognetti, and D. D. Rossi. Wearable, redundant fabric-based sensor arrays for reconstruction of body segment posture. *IEEE Sensors Journal*, 4(6):807–818, Dec 2004.
[39] B. Mailhé, R. Gribonval, F. Bimbot, M. Lemay, P. Vandergheynst, and J.-M. Vesin. Dictionary learning for the sparse modelling of atrial fibrillation in ecg signals. In *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*, pages 465–468. IEEE, 2009.
[40] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary learning for sparse coding. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pages 689–696, New York, NY, USA, 2009. ACM.
[41] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *J. Mach. Learn. Res.*, 11:19–60, Mar. 2010.
[42] A. Natarajan, A. Parate, E. Gaiser, G. Angarita, R. Malison, B. Marlin, and D. Ganesan. Detecting cocaine use with wearable electrocardiogram sensors. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 123–132. ACM, 2013.
[43] K. Niu, F. Zhang, J. Xiong, X. Li, E. Yi, and D. Zhang. Boosting fine-grained activity sensing by embracing wireless multipath effects. In *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*, pages 139–151. ACM, 2018.
[44] B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23):3311 – 3325, 1997.
[45] R. Paradiso, G. Loriga, and N. Taccini. A wearable health care system based on knitted integrated sensors. *IEEE Transactions on Information Technology in Biomedicine*, 9(3):337–344, Sept 2005.
[46] T. Rahman, A. T. Adams, R. V. Ravichandran, M. Zhang, S. N. Patel, J. A. Kientz, and T. Choudhury. 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*, pages 39–50. ACM, 2015.
[47] F. Roche, J.-M. Gaspoz, I. Court-Fortune, P. Minini, V. Pichot, D. Duverney, F. Costes, J.-R. Lacour, and J.-C. BartheÌĄleÌĄmy. Screening of obstructive sleep apnea syndrome by heart rate variability analysis. *Circulation*, 100(13):1411–1415, 1999.
[48] E. Sardini, M. Serpelloni, and V. Pasqui. Wireless wearable t-shirt for posture monitoring during rehabilitation exercises. *IEEE Transactions on Instrumentation and Measurement*, 64(2):439–448, Feb 2015.
[49] D. Shao, F. Tsow, C. Liu, Y. Yang, and N. Tao. Simultaneous monitoring of ballistocardiogram and photoplethysmogram using a camera. *IEEE Transactions on Biomedical Engineering*, 64(5):1003–1010, May 2017.
[50] Y. Shu, C. Li, Z. Wang, W. Mi, Y. Li, and T.-L. Ren. A pressure sensing system for heart rate monitoring with polymer-based pressure sensors and an anti-interference post processing circuit. *Sensors*, 15(2):3224–3235, 2015.
[51] P. K. Stein, M. S. Bosner, R. E. Kleiger, and B. M. Conger. Heart rate variability: a measure of cardiac autonomic tone. *American heart journal*, 127(5):1376–1381, 1994.
[52] https://greenwaves-technologies.com/gap8-product/. GAP8: Ultra-low power, always-on processor for embedded artificial intelligence.
[53] E. Vanoli, P. B. Adamson, Ba-Lin, G. D. Pinna, R. Lazzara, and W. C. Orr. Heart rate variability during specific sleep stages: a comparison of healthy subjects with patients after myocardial infarction. *Circulation*, 91(7):1918–1922, 1995.
[54] V. Vesterinen, K. Häkkinen, E. Hynynen, J. Mikkola, L. Hokka, and A. Nummela. Heart rate variability in prediction of individual adaptation to endurance training in recreational endurance runners. *Scandinavian journal of medicine & science in sports*, 23(2):171–180, 2013.
[55] Z. L. Wang, J. Chen, and L. Lin. Progress in triboelectric nanogenerators as a new energy technology and self-powered sensors. *Energy & Environmental Science*, 8(8):2250–2282, 2015.
[56] S. Wanwong, W. Sangkhun, S. Z. Homayounfar, K.-W. Park, and T. L. Andrew. Wash-stable, oxidation resistant conductive cotton electrodes for wearable electronics. *RSC Adv.*, 9:9198–9203, 2019.
[57] W. Y. Wong and M. S. Wong. Smart garment for trunk posture monitoring: A preliminary study. *Scoliosis*, 3(1):7, May 2008.
[58] M. Xiao, H. Yan, J. Song, Y. Yang, and X. Yang. Sleep stages classification based on heart rate variability and random forest. *Biomedical Signal Processing and Control*, 8(6):624–633, 2013.

[59] W. Xu, M. Huang, N. Amini, L. He, and M. Sarrafzadeh. ecushion: A textile pressure sensor array design and calibration for sitting posture analysis. *IEEE Sensors Journal*, 13(10):3926–3934, Oct 2013.

[60] J. Yang, K. Yu, Y. Gong, and T. Huang. Linear spatial pyramid matching using sparse coding for image classification. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1794–1801, June 2009.

[61] L. Zhang, Y. Yu, G. P. Eyer, G. Suo, L. A. Kozik, M. Fairbanks, X. Wang, and T. L. Andrew. All-textile triboelectric generator compatible with traditional textile process. *Advanced Materials Technologies*, 1(9):1600147–n/a, 2016. 1600147.