

# BPSniff: Continuously Surveilling Private Blood Pressure Information in the Metaverse via Unrestricted Inbuilt Motion Sensors

Zhengkun Ye\*, Ahmed Tanvir Mahdad†, Yan Wang\*, Cong Shi‡, Yingying Chen§, Nitesh Saxena†

\*Temple University, Philadelphia, PA, USA 19122

†Texas A&M University, College Station, TX, USA 77840

‡New Jersey Institute of Technology, Newark, NJ, USA 07102

§Rutgers University, New Brunswick, NJ, USA 08901

**Abstract**—Blood pressure (BP) is one of the most essential biomarkers for various diseases. It is considered protected health information under HIPAA and usually needs the user’s consent for access. In this work, we uncover an insidious privacy breach in metaverse usage: private BP information can be covertly obtained from unrestricted motion sensors in virtual reality (VR) headsets. The insight is that the motion sensors can capture the subtle vibrations induced by the blood waves in the major arteries. Such vibrations are highly correlated with users’ cardiac cycles and BP. As adversaries can continuously obtain motion sensor data from VR headsets without users’ consent, they can derive and collect users’ BP information in metaverse apps or websites, leading to more severe consequences, such as discrimination, exploitation, and targeted harassment.

To demonstrate this severe privacy leakage in the metaverse, we develop a practical attack, BPSniff, which can reconstruct fine-grained blood flow patterns and derive BP based on motion sensor data from users’ VR headsets. BPSniff is the first practical attack revealing the BP leakage in the metaverse without using dedicated equipment. Unlike previous mobile sensing approaches that require user-specific calibration, BPSniff bypasses this constraint, enabling truly stealthy passive BP attacks at scale. Our attack first employs a variational autoencoder to reconstruct high-fidelity blood flow patterns from VR headset motion sensor data. We then develop an Adam-optimized long short-term memory (LSTM) regression model that leverages BP-related fiducial features from successive blood flow patterns to continuously estimate the user’s BP. We evaluate BPSniff through extensive experiments and a longitudinal study of 8 weeks, involving 37 participants and two VR headset models. The results show that BPSniff can achieve low mean errors of 1.75 mmHg for systolic blood pressure (SBP) and 1.34 mmHg for diastolic blood pressure (DBP), which are comparable to commercial BP monitors and satisfy the standard (i.e., mean error  $\leq 5.0$  mmHg) specified by FDA’s AAMI protocol.

## 1. Introduction

Blood pressure (BP) is considered as the most important biomarkers for many diseases, providing valuable

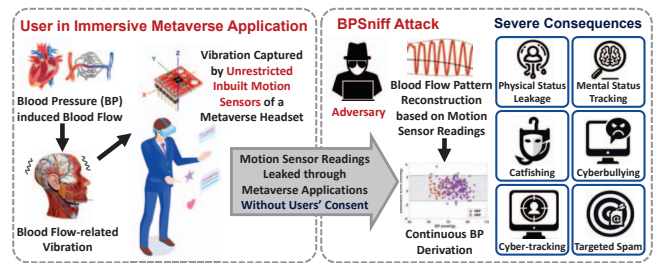


Figure 1: Illustration of the proposed BPSniff Attack.

insights into an individual’s health status. While the metaverse, typically accessed through virtual reality (VR), offers diverse and immersive experiences across numerous applications [1], the severe security and privacy vulnerabilities stemming from its sensor technologies have not been thoroughly examined. With the emerging usage of the metaverse [2], there is an increasing risk that users’ private biomarkers could be leaked from metaverse devices (e.g., VR headsets). The potential leakage of BP information during immersive VR usage remains largely unexplored, but this oversight is consequential. Unauthorized and covert measurements of BP, a biomarker protected by regulations, without user consent represent a significant privacy violation. In addition, BP is strongly associated with mental conditions [3]. BP data can thus reveal not only cardiovascular health [4] but also emotional states [5], stress levels [6], and even reactions to specific virtual stimuli. Therefore, a critical and significant research problem is: “*Is it possible for adversaries to continuously and stealthily derive a user’s private BP information in the metaverse using the unrestricted built-in motion sensors of VR headsets?*”

In this work, we demonstrate that such a BP inference attack based on unrestricted motion sensors is indeed possible. It results in severe privacy leakage and can cause unimaginable consequences: 1) Traditionally, BP measurements are collected in clinical settings using medical instruments and are categorized as protected health information (PHI) under Health Insurance Portability and Accountability Act (HIPAA) [7], requiring strict safeguards and rigorous approval for access. It is a significant privacy concern if unauthorized individuals could stealthily obtain this infor-

mation from unrestricted sensors in the metaverse without using dedicated BP monitoring devices. Adversaries could continuously collect BP information in the background without users' consent over extended periods. The longitudinal BP data could reveal more about users' privacy [8] (e.g., mental states and emotional patterns). 2) Moreover, the consequences could be dreadful if such BP leakage occurs widespread, allowing enterprises or malicious organizations to acquire large-scale BP profiles of a vast number of metaverse users. 3) Considering identities can be derived from multiple channels, such as user profiles on websites or apps, and other fingerprinting methods, linking the users' identities with their BP information could lead to more severe consequences from malicious attacks. For instance, attackers could launch catfishing attacks [9] by leveraging BP data to create more convincing fake personas, manipulating victims' emotions and trust. Moreover, attackers might exploit BP information to identify vulnerable moments for cyberbullying [10] or make biased decisions based on perceived health status, thereby intensifying psychological harm and promoting unfair practices. Additionally, attackers could manipulate virtual reality experiences based on users' BP, potentially inducing or exacerbating cybersickness [11], [12].

To explore the severity of the BP leakage via VR headsets, we develop the first-of-its-kind attack, BPSniff, which allows adversaries to continuously derive users' private BP measurements based on the cardiac cycle-induced vibration signals captured by the inbuilt motion sensors of VR headsets. The inbuilt motion sensors of VR headsets have been considered unrestricted, as the sensor data are freely accessible without user permission [13]. Therefore, adversaries can obtain motion sensor readings remotely (i.e., using a disguised website and a malicious app running at the background) and continuously derive users' BP data without requiring any active user engagement (e.g., gesture or speech). For example, in online VR multiplayer games, an adversary can exploit players' motion sensor data shared with servers to launch BPSniff, inferring players' BP information and manipulating their psychological state in the game. The insight is that motion sensors in VR headsets can capture the vibrations (i.e., body's recoil movements) induced by cardiac movements, known as ballistocardiogram (BCG) [14] waveforms. Researchers have shown that the ascending and descending BP gradients are closely related to BCG waveforms [15], [16]. We exploit this intimate BP-BCG link to craft BPSniff, which reconstructs the complete morphology of successive blood flow patterns and continuously derives BP measurements without using dedicated devices (e.g., sphygmomanometer). The flow and consequences of BPSniff are illustrated in Figure 1. To the best of our knowledge, this is the first work that achieves obtaining users' private BP information in the metaverse continuously, indicating severe privacy vulnerabilities of biomarker leakage.

Realizing BPSniff in practical VR scenarios face several challenges: a) Large body motions like head movements and hand gestures can significantly distort cardiac cycle-

induced vibration signals. And small involuntary motions (e.g., respiratory patterns, blinking, and micro-expressions) also introduce artifacts that overwhelm the vibration signals. To overcome this challenge, we design a data-adaptive multiresolution scheme using empirical mode decomposition (EMD). This approach separates the motion artifacts from the cardiac cycle frequencies by analyzing the data at multiple resolutions, allowing adversaries to accurately isolate and preserve the signals related to blood pressure. b) Traditional BP measurement solutions may require hardware calibration (i.e., adjusts the BP monitoring device against the set standard) and software calibration (i.e., use target users' data to fine-tune the estimation method) to ensure the accuracy on specific devices and users [17], [18], [19]. However, such calibrations are impractical for adversaries as they necessitate timely, covert, individualized data collection, risking detection and thwarting large-scale exploitation.

We address this by developing a calibration-free BP estimation scheme using a pre-trained deep learning model, that enables robust, continuous BP estimation. c) In addition, it's essential yet challenge to model the relationship between blood flow patterns derived from motion sensor readings and systolic/diastolic blood pressure (SBP/DBP). This is crucial because the fiducial features required for estimating SBP and DBP are not readily available from the blood flow patterns. We tackle this challenge by devising a variational autoencoder to reconstruct fine-grained successive blood flow patterns with the help of reference photoplethysmogram (PPG) and extract the essential BP-related fiducial features. We summarize our main contributions as follows:

- We uncover a novel passive privacy attack, BPSniff, leveraging the built-in motion sensors of consumer VR headsets to effectively disclose metaverse users' blood pressure in an unobtrusive manner. This new attack poses a significant threat to the emerging metaverse, as applications can be covertly exploited to unveil users' privacy maliciously.
- We systematically investigate large and small body motions in metaverse applications and their impacts on blood flow-induced signals. We develop data-adaptive multiresolution method using empirical mode decomposition (EMD) to mitigate motion artifacts while preserving the integrity of cardiac cycle frequencies.
- We design and implement a variational autoencoder that allows adversaries to reconstruct successive blood flow patterns from motion sensor readings captured by metaverse headsets, using PPG sensors as a reference. This scheme eliminates the need for calibration process, unlike traditional BP estimation methods. Additionally, we advance feature engineering to extract BP-related fiducial features from the reconstructed blood flow patterns and develop an Adam-Optimized LSTM regression model for robust continuous SBP and DBP estimation.
- Our experimental results with 37 subjects and a longitudinal study of 8 weeks demonstrate that the BPSniff can accurately estimate SBP and DBP with 1.75 mmHg and 1.34 mmHg mean estimation error, respectively. The




Figure 2: Blood pressure pushes blood flow through the blood vessels and squeezes their walls to produce the vibration that metaverse headsets can capture.

results are comparable to commercial BP monitors and meet the gold standard (i.e., mean error  $\leq 5.0$  mmHg) set by the Association for the Advancement of Medical Instrumentation (AAMI) protocol of the Food and Drug Administration (FDA).

## 2. Preliminaries of Stealthy Blood Pressure Measurement in the Metaverse

### 2.1. Sensory Data in VR

**VR Headset-enabled Fine-grained Motion Sensing.** VR headsets, designed for immersive interactions, integrate various sensors, including motion sensors. And BCG signals can be captured at the head by motion sensors [20]. Since VR headsets directly contact the head, these signals (cardiac cycle-induced vibrations) propagate through the headset and are passively detected by built-in motion sensors. We find that when users wear VR headsets on their heads with the headsets closely contacting the bridge of the nose and browbones, the built-in motion sensors can detect the BCG generated by blood volume changes. Due to the high fidelity and sampling rate of these VR built-in motion sensors (e.g., 1k Hz for Meta Quest/Meta Quest 2), they can capture detailed BCG signals, which is not feasible for mobile devices (e.g., smartphones and smartwatches) equipped with motion sensors with the sampling rates below 500 Hz and low-sensitivity sensors for monitoring coarse-grained motions (e.g., step counting) [21], [22]. Note that such low-sampling-rate motion sensors in mobile devices can approximate heart rates but not restore blood flow or heartbeat waveform. As a whole, the high sampling rate and direct contact with the face make VR headsets a promising platform for capturing users' cardiac-cycle induced vibration signals without using dedicated devices.

**Physical Foundation: BP-BCG Relationship.** The human body moves subtly with each heartbeat [23]. These subtle movements, representing the body's recoil to the cardiac expulsion of blood into the arteries, can be recorded as BCG signals [15]. BCG signals are closely related to BP, and their

relationship can be defined by the BP gradients within the aorta [16]. Specifically, BP is the lateral force exerted by blood flow on the arterial wall [24]. Systolic blood pressure (SBP) is the maximum force exerted while the heart is contracting (systole), and diastolic blood pressure (DBP) is the minimum force exerted while the heart is relaxing (diastole). Blood pressure is primarily determined by two key factors: the volume of blood flowing through the vessels and the pumping force exerted by the heart. A pressure gradient [25] is maintained by the heart's pumping action, enabling blood to be propelled from the higher pressure arteries through to the lower pressure capillaries and veins [26]. Essentially, the more blood that is pumped by the heart, the higher the resulting BP within the circulatory system. Changes in blood flow within these vessels cause pressure variations on the vascular walls, leading to periodic expansions and contractions [27]. These regular motions propagate through the facial muscles and skull to the browbones and bridge of the nose, generating periodic vibration signals, i.e., BCG [28], as shown in Figure 2.

### 2.2. Proof-of-concept of Stealthy Data Collection

**Calibration-free Blood Pressure Derivation from BCG.** We find that regular changes in blood flow can be captured by photoplethysmogram (PPG), a widely used optical technique to measure blood volume variations [29]. The pulsations of blood volume in tissue recorded by PPG have been connected to arterial pressure [30]. Meanwhile, existing study [31] demonstrates the strong correlation of PPG and BCG when used for blood pressure estimation. Moreover, a prior study [32] has demonstrated that leveraging both BCG and PPG can aid in precise BP estimation, as the method can use the BCG and PPG together to account for changes in vessel radius during mean arterial blood pressure estimation. Existing research [33], [34] has demonstrated that dynamic changes in the consecutive pulse waveforms can be leveraged to provide accurate BP estimation without calibration. In particular, the researchers have shown high correlations between BP estimates derived from PPG morphology information (fiducial features) and BP measurements captured by calibrated BP monitors. This suggests that PPG fiducial features capture user-specific information typically obtained through calibration. By extracting these features from the PPG signal itself, the method may eliminate the need for separate calibration procedures. This offers a promising path towards calibration-free BP estimation that maintains a precision comparable to traditional calibrated techniques. This advantage is particularly essential in the context of adversarial use because attackers usually do not have access to users' devices or data. As such, we propose using PPG signals as a reference to reconstruct successive blood flow patterns from blood vessel vibrations captured by VR headsets.

**Stealthy Script Implementation in the Metaverse.** We validate the viability of stealthily collecting motion sensor data (e.g., accelerometer data) through programming with two VR headsets (Meta Quest [35] and Meta Quest

2 [36]). Our program can be dexterously either run in the background of the VR environment or be embedded in a VR webpage, using the Meta Quest Platform SDK [37] and WebXR Device API [38] (i.e., a JavaScript application programming interface to enable applications to interact with metaverse devices in a web browser). The stolen motion sensor readings can be transmitted to a remote adversarial server. We develop a Meta Quest program in C (programming language) to extract motion sensor data without requiring user permission, running discreetly in the background. According to the Meta Platforms Technologies Privacy Policy [39], it is legitimate for developers to collect this user data. More specifically, we can acquire metaverse users' raw accelerometer data using function `ovr_GetTrackingState()`. The program can also be embedded into various metaverse applications. Our attack is also compatible with OpenVR SDK [40], an open-source VR programming platform that supports more headset models/types and also does not require permissions to access motion sensor readings.

### 3. Threat Model and Attack Overview

#### 3.1. Adversary's Capability

We propose an attack scenario where an adversary aims to infer a victim's private BP information within the metaverse using VR headsets. The adversary can deceive the victim into either installing a malware program developed with mainstream VR SDKs (e.g., Meta Quest Platform SDK) or visiting a malicious VR webpage created with the WebXR Device API. In both cases, the adversary can remotely access the sensor data from the victim's VR headset without permissions. The malware can be disguised as a legitimate VR app or plugin, or embedded into a third-party VR development library commonly used by developers to build VR applications. For example, the adversary might embed malicious code in seemingly benign software, such as a system cleaning utility or a monitoring tool, and post it online to lure victims into installing it. These types of software can run in the background alongside other processes, allowing the malware to stealthily log sensor data while the victim is engaged in VR activities, making detection difficult. In contrast, hosting a malicious VR webpage poses an even greater threat. This method is more accessible to the adversary and more devastating, as it does not require the installation of a third-party app on the victim's device. The adversary can obtain the VR headset's sensor data remotely as soon as the victim visits the malicious webpage.

In our considered attack model, the adversary lacks any prior insider knowledge about the victim's specific VR headset setup or physiological baselines. The adversary is not assumed to have any visibility into the victim's device configurations, such as headset position tracking mechanics determined by external base station placements. Additionally, the adversary also does not require any preliminary training data explicitly labeled with the victim's private BP information. Instead, the adversary can effortlessly calibrate




Figure 3: Illustration of blood flow-induced facial vibrations on accelerometer (z-axis) readings of Meta Quest and PPG signals of subjects.

the attack by independently profiling blood flow patterns from their own commodity VR headset in any real-world setting. This reference mapping between motion sensor readings and blood flow dynamics makes the clandestine real-time monitoring of subsequent victims' private BP a trivial operation. Such assumptions maximize the insidiousness and scalability of our proposed attack - adversaries can indiscriminately violate user privacy without any privileged information. Our proposed attack operates under a highly practical threat model and can be launched surreptitiously and easily in the metaverse.

#### 3.2. Feasibility and Challenges

**Attack Feasibility.** To fully demonstrate the feasibility of using built-in motion sensors to capture blood flow patterns and launch BPSniff, we conduct experiments by simultaneously collecting motion sensor data on a metaverse headset (i.e., Meta Quest) and blood volume data from PPG sensors (300 Hz) with ten participants sitting still. In our experiment, we collect six axes of motion sensor data (i.e., three axes for the accelerometer and gyroscope, respectively) from Meta Quest. We find that the accelerometer on the z-axis has the most significant pulsatile signal, as the facial vibration caused by the blood volume changes in the arteries is perpendicular to the face. As showcased with two participants' datasets in Figure 3, the cardiac cycle-induced vibration (BCG) captured by the motion sensors in the headset produces a waveform similar to that captured by PPG, reflecting the pulsatile signal generated by blood flow in the arteries. We can observe that different participants' facial vibrations both present distinguishable morphological characteristics corresponding to different pulse widths, systolic peak amplitudes, diastolic peak amplitudes, and inflection point amplitudes. Our feasibility study shows strong correlations between the vibrations induced by blood flow in the user's face captured by the metaverse headset and heartbeats, motivating further research on deriving blood pressure from these signals.

**Challenges of Stealthily Deriving Blood Pressure Measurements.** To realize such an attack in the metaverse, we face several practical challenges: 1) *Significant Interference of Motions:* In VR scenarios, users intentionally interact with the metaverse through motions like head rotations and controller movements. These large body motions can significantly distort the BCG induced by blood

flow. Meanwhile, metaverse users exhibit small body motions, including involuntary subtle movements, even when attempting to remain still. These include occasional head nodding, breathing patterns, natural blinking, swallowing, facial muscle activations, and micro-expressions. Both large body motions and small involuntary motions can introduce artifacts that overwhelm the blood pressure-relevant BCG waveforms, which have much smaller magnitudes. 2) *Indirect Relationship between BCG waveforms and Quantitative Blood Pressure Values*: The fiducial features relevant for estimating quantitative BP values (systolic blood pressure (SBP) and diastolic blood pressure (DBP)) are not directly available from the blood flow patterns constructed from BCG waveforms derived from the VR headset’s motion sensor readings. It is critical to model the relationship between the blood flow-induced BCG waveforms and SBP/DBP for effective measurements. 3) *Physiological Characteristic Variations Across Users*: It is desirable to build a calibration-free BP estimation model so that individual users do not need to suffer tedious training processes. However, due to differences in physiological characteristics (e.g., height, weight, age, and sex), it is difficult to train such a general model.

### 3.3. Overview of BPSniff

To continuously derive metaverse users’ private BP information using VR headsets, our newly designed attack, BPSniff, consists of three main modules that work in tandem to accurately estimate SBP and DBP from BCG waveforms, as illustrated in Figure 4. First, the adversary launches BPSniff to collect motion sensor data from VR headsets and pre-processes the data to remove noise and interference in the *Facial Vibration Pre-processing*. This module focuses on removing vibration impacts caused by large-scale and small-scale body movement (e.g., head rotation, head tilting, controller movement, etc.) and respiration. We separate the body motion and cardiac cycle-induced vibration signals using empirical mode decomposition (EMD) and remove the body motion signals.

Next, BPSniff filters the data using a Butterworth bandpass filter (0.6 Hz to 3 Hz) to remove high-frequency thermal noise and impacts from respiration. The filtering thresholds are empirically derived through a comprehensive analysis of motion sensor readings from 10 diverse non-victim participants, varying in age, gender, and race, under different conditions. We perform statistical analyses to identify the optimal thresholds that are most effective in removing the noises while preserving vibration signals related to blood flow patterns. After filtering, BPSniff segments the data based on cardiac cycles. It performs the *Blood Flow Pattern Reconstruction* to reconstruct successive blood flow patterns from the pre-processed data using a variational autoencoder (VAE), which enables the generation of new data by sampling from a learned probability distribution. This allows adversaries to generate data distinct from the input, facilitating a deeper understanding of the underlying blood pressure dynamics. The VAE is offline and trained by

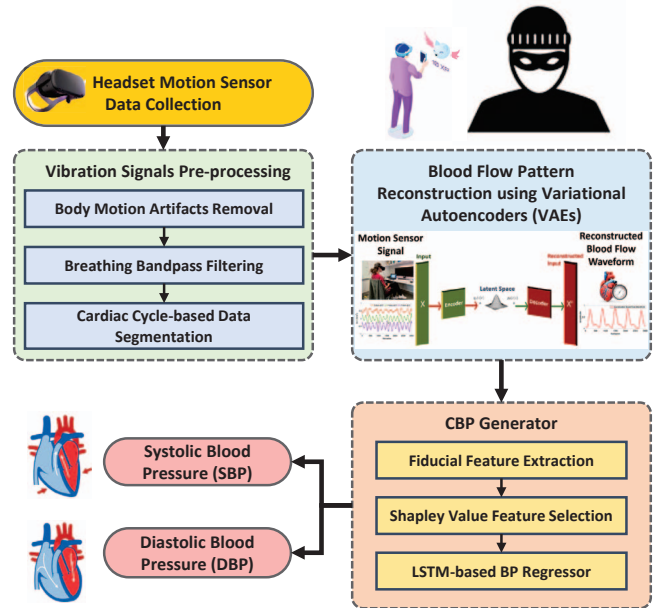


Figure 4: BPSniff: An innovative attack that can continuously derive blood pressure measurements based on passive cardiac cycle-induced vibrations captured by motion sensors of metaverse headsets.

the data labeled by PPG to learn the relationship between successive blood flow patterns and cardiac cycle-induced vibrations. It transforms the pre-processed data into a probability distribution over the space of the blood flow dynamics to reconstruct the successive blood flow patterns.

Last, the *Continuous Blood Pressure (CBP) Generator* extracts BP-related fiducial features from the reconstructed successive blood flow patterns. The features are fed through a Shapley value feature selection algorithm to identify the most relevant features (e.g., pulse interval, peak-to-peak interval) for robust BP estimation. An Adam-optimized LSTM-based BP regressor is trained to map the selected features to SBP and DBP values. LSTMs have shown effectiveness in modeling sequential data and capturing long-term dependencies, making them suitable for analyzing time series data like blood pressure. Our two-stage deep neural network architecture sidesteps the arduous challenge of directly mapping the raw BCG waveforms to quantitative pressure values through an intermediate learned representation of characteristic blood flow patterns.

## 4. Attack Design

### 4.1. Sensor Data Pre-processing

BPSniff takes the acceleration on the z-axis from the headset as input. Our primary signal processing goal is to isolate cardiac cycle-induced vibrations within the 0.6 to 3 Hz frequency range, which is critical for accurate blood pressure estimation. To achieve this adaptively, we employ Empirical Mode Decomposition (EMD). We use




Figure 5: Variational Autoencoder (VAE) for reconstructing blood flow patterns.

EMD because it offers more adaptive and fine-grained control over frequency separation. EMD is particularly suited for our non-stationary headset acceleration signals because it decomposes the signal into Intrinsic Mode Functions (IMFs) based on its local frequency characteristics, without assuming a fixed frequency content over time. We utilize IMFs to help isolate the blood pressure-relevant signals from both large body motion artifacts as well as subtle involuntary factors (small motions) like occasional head nodding, breathing patterns, and facial muscle activations. Formally, let  $f(t)$  be the non-stationary signal we wish to decompose into  $N$  IMFs  $c_k(t)$ , with  $k = 1, \dots, N$ . We compute the local extrema of  $f(t)$ , and interpolate between the local extrema to obtain the upper and lower envelopes  $u(t)$  and  $l(t)$ . We then get the first IMF  $c_1(t) = f(t) - m_1(t)$ , where  $m_1(t) = (u(t) + l(t))/2$ , repeat above steps with  $f(t) = c_1(t)$  to obtain the second IMF  $c_2(t)$ , and so on until convergence is reached.

The first few IMFs capture the highest frequencies, which in our case include both the desired cardiac cycle vibrations (0.6-3 Hz) and unwanted components like fine motion artifacts or sensor noise. Subsequent IMFs capture progressively lower frequencies, including large motion artifacts (e.g., 3.2-9.6 Hz range [41]), respiration (0.2-0.33 Hz [42]), and other low-frequency physiological processes (below 0.6 Hz [43]) including metabolic, neurogenic, and myogenic activities. To isolate our target 0.6-3 Hz range, we reconstruct the signal using only the IMFs that predominantly contain these frequencies. This is determined empirically by analyzing the frequency content of each IMF. Typically, we find that a combination of the first few IMFs (after discarding the highest-frequency noise) captures our range of interest.

To further enhance the quality of the EMD-processed cardiac cycle-induced vibration signals and optimize the capture of blood flow responses to changes in blood pressure, we implement a strategic filtering approach by applying a Butterworth bandpass filter. The bandpass filter provides a strict 0.6-3 Hz isolation. This is due to EMD can sometimes suffer from 'mode mixing,' where an IMF contains multiple frequency scales. This targeted filtering step eliminates additional noise, ensuring that the final processed signal contains the essential blood flow patterns needed for accurate blood pressure estimation. By safeguarding our frequency range selection, we effectively remove unwanted noise from voluntary and involuntary motions while preserving the critical blood flow signals.

After removing the noise from the cardiac cycle-induced vibration data using EMD, we extracted the cardiac cycle from the data using the peak detection method. We first applied a threshold to the denoised signal to identify the peaks corresponding to each heartbeat. To enhance peak detection accuracy, we used a moving average filter with a window size of 100 ms to smooth the signal. The resulting peaks were then used to extract the consecutive cardiac cycle waveforms, representing blood flow through facial blood vessels.

## 4.2. Blood Flow Pattern Reconstruction

We already showed that the VR headset's motion sensor can capture blood flow patterns. However, we still need to map the sensor data to BP values. Existing studies have shown the relationship between BP and blood flow. Since PPG directly measures blood flow patterns, and its fiducial/morphological features can replace the calibration information. We thus propose using the PPG signal as a reference to reconstruct blood flow patterns for accurate BP measurement. PPG-based techniques don't really "measure" the pressure. Instead, they use the pattern feature analysis and theoretical models to predict the hemodynamics and link them to BP. The PPG signal directly measures the volumetric changes in blood flow. It is also considered the gold standard to capture the pulsatile blood volume changes caused by the cardiac cycle.

In particular, we use a Variational Autoencoder (VAE), as illustrated in Figure 5 to reconstruct blood flow patterns from motion sensor data collected from a VR headset. Our VAE design enforces the latent space to have desirable structural properties, making similar data points in the latent space close to each other. Unlike most Autoencoder (AE) architectures [44], [45] encode a discrete, fixed representation of latent variables, VAEs encode a continuous, probabilistic representation of that latent space. This makes VAE more susceptible to the variation and noise in the input signal when reconstructing the signal since it learns to generate from a continuous distribution. Moreover, the multivariate normal distribution can learn to capture the waveform patterns such as systolic and diastolic cardiac procedure that corresponds to different Gaussian distributions. Meanwhile, the VAE's loss function not only encourages the model to accurately reconstruct data but also ensures the continuity and smoothness of the latent space. By learning a smooth, continuous latent representation, the VAE can interpolate between known data points and extrapolate to new, previously unseen examples, effectively expanding its generative capabilities beyond the training set [46], [47]. As shown in Figure 6, VAE allows for the generation of synthetic blood flow patterns that closely resemble those obtained from medical-grade PPG sensors. This can be useful in cases where PPG sensors are not available or practical to use. Lastly, VAE provides a probabilistic framework for generating blood flow patterns, enabling uncertainty assessment and confident decision-making based on estimated BP values.




Figure 6: Illustration of the comparison of reconstructed blood flow patterns with ground truth PPG signals for two participants.

Specifically, we first collect a dataset of both PPG and motion sensor signals. Then, we use the PPG signals as the ground truth and train VAE to learn the mapping from motion sensor signals to blood flow patterns. We note that PPG sensors are only used by the developer when training the VAE model. Thus, the common issues with PPG sensors like wearable position or skin pigmentation can be avoided under controlled data collection. VAE consists of two parts: an encoder and a decoder. The encoder takes the motion sensor signal and generates a latent code vector. The decoder then takes the latent code vector to produce the reconstructed blood flow pattern. To train VAE, we define a loss function that measures the difference between the reconstructed blood flow pattern and the ground truth PPG signal. The loss function consists of two parts: a reconstruction loss that measures the difference between the reconstructed signal and the ground truth signal, reconstruction loss is defined as:  $L_{rec} = |p - \hat{p}|^2$ , where  $p$  is the ground truth PPG signal,  $\hat{p}$  is the reconstructed blood flow pattern. A regularization loss that encourages the encoder to produce a latent code that follows a prior distribution, specifically a standard multivariate normal distribution. Regularization loss is:

$$L_{reg} = -1/2 \cdot \sum_{i=1}^n (1 + \log(\sigma_i^2) - \mu_i^2 - \sigma_i^2), \quad (1)$$

where  $\mu$  and  $\sigma$  are the mean and standard deviation of the latent code produced by the encoder, and  $n$  is the dimension of the latent code. In this work, we set  $n$  as 22 which is sufficient to capture critical waveform patterns embedded in the signal for reconstruction. The total loss is the sum of the reconstruction loss and the regularization loss:

$$L_{total} = L_{rec} + L_{reg}. \quad (2)$$

During training, we use backpropagation to update the weights of the encoder and decoder to minimize the total loss. After training, we use the decoder to generate blood flow patterns from new motion sensor signals. We note that this is an offline process and users do not need to provide their PPG data.

### 4.3. Cardiac Signal Segmentation

After removing body motion artifacts and respiratory interference, we conduct pulse segmentation to extract the reconstructed blood flow pattern that contains a complete cardiac cycle. We then detect peaks in the constructed blood

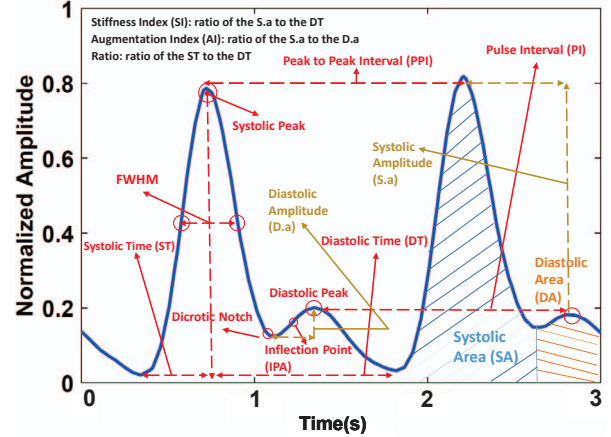


Figure 7: Blood pressure features generated from reconstructed consecutive blood flow patterns.

flow pattern (PPG-like signal). These peaks correspond to the systolic and diastolic phases of the cardiac cycle. Our system uses these peaks to segment the constructed blood flow signal into individual cardiac cycles. Each cardiac cycle is expected to contain a complete heartbeat, including a systolic peak. To determine the starting and ending points of the PPG-like segments, we look for the valley points before and after the systolic peak, respectively. However, the dicrotic notch can sometimes have the lowest amplitude, which can be mistaken for a valley point. To address this issue, we select valley points that are within the typical time ranges before and after each systolic peak, which are in the ranges of 0.15s to 0.26s and 0.44s to 0.74s [48], respectively. By selecting the valleys that fall within these time ranges, the system can accurately extract the PPG-like segments that contain complete cardiac cycles.

### 4.4. Blood Pressure Feature Extraction

Blood flow pattern, by itself, is not directly related to BP, which is a measure of the pressure imposed by blood flow on the wall of vessels; however, subtle variation in the morphology of the blood flow pattern appears to be correlated to BP. The features of the blood flow pattern can be used to encode the complex relationship hidden between the blood flow pattern and BP [49]. We analyze the successive reconstructed blood waveform cycles, rather than the morphology of a single cycle alone. Most previous PPG-based BP estimation methods that extract features from the morphology of a single PPG cycle, which has proven challenging due to inter-person variations. By instead analyzing how the blood flow pattern changes from one cycle to the next, our method can capture the dynamic characteristics related to BP, which may be more robust across individuals. Additionally, prior study [33] proves that calibration-free solution can be achieved by using dynamic changes in the pulse waveform over short intervals and replaces calibration with information from PPG morphology.

In this study, we first extract 11 features [50] from reconstructed blood flow patterns as depicted in Figure 7,

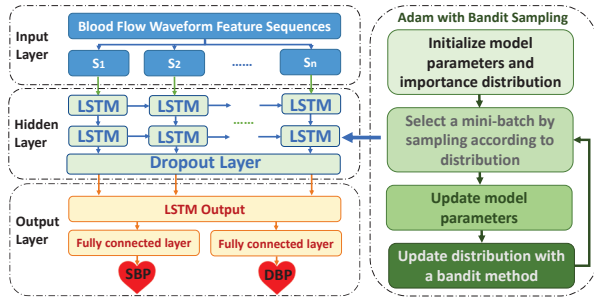


Figure 8: Illustration of the Adam-optimized LSTM neural network architecture.

including Augmentation Index (AI), Stiffness Index (SI), The ratio of the systolic time to the diastolic time, Systolic amplitude (S.a) and diastolic amplitude (D.a), Systolic area (SA) and Diastolic area (DA), Inflection point area (IPA), Systolic Time (ST), Diastolic time (DT), Pulse Interval (PI), Peak-to-peak interval (PPI), and Time to Full-Width at Half-Maximum (FWHM). We then utilize Shapley value as a feature selection tool [51]. The Shapley value estimates the impact of a feature on predictions, assuming that a known random characteristic should have less influence than an informative feature. Specifically, the Shapley value is calculated as follows:

$$i = \frac{1}{n!} \sum_{S \subseteq N \setminus \{i\}} (n - |S| - 1)! |S|! [F(S \cup \{i\}) - F(S)] \quad (3)$$

where  $i$  is the index of the feature being evaluated,  $n$  is the total number of features,  $S$  is assigned as a subset of features that does not include feature  $i$ ,  $F(S)$  is the model's prediction when the features in  $S$  are used,  $F(S \cup \{i\})$  is the model's prediction when the features in  $S$  and feature  $i$  are used. This formula calculates the Shapley value for each feature in the PPG feature set. The resulting values can be used to rank the importance of each feature and identify which features have the greatest impact on the model's predictions. Based on calculations, we select S.a/D.a, SA/DA, ST, DT, PI, PPI and IPA as representative features from consecutive blood flow patterns for building the CBP regression model. We use statistical parameters (e.g., mean/standard deviation of the pattern amplitudes across cycles) from extracted blood pressure features to quantitatively characterize the distribution of the blood flow pattern values and how they vary across successive cardiac cycles, without explicitly modeling the exact morphological shape. We integrate these features into a vector and use the series of feature vectors for BP estimation in the next step.

#### 4.5. Continuous Blood Pressure Estimation

After feature extraction, BPSniff uses a Long Short-Term Memory (LSTM) neural network to learn the temporal dependencies between the representative features and derive hidden representations of the reconstructed blood flow pattern. The LSTM can learn time-dependent relationships in the pattern data that are invariant and unaffected by

intermittent motion distortions. The LSTM contains memory cells that can maintain information over long time intervals, along with gating units that modulate the flow of information. This structure enables the model to learn long-range temporal relationships in the pattern morphology and sequences of features that are predictive of blood pressure changes. In our design, the hidden representations are fed into two separated, fully connected layers, one for SBP and one for DBP estimation. The outputs of the fully connected layers are used to estimate SBP or DBP using the following equation:

$$\hat{y}_{SBP} = \alpha_{SBP} * h_T + \beta_{SBP}, \quad (4)$$

$$\hat{y}_{DBP} = \alpha_{DBP} * h_T + \beta_{DBP}, \quad (5)$$

where  $h_T$  is the hidden state output by the LSTM model at the final time step,  $\alpha_{SBP}$  and  $\alpha_{DBP}$  are the coefficients of the linear regression model for SBP and DBP prediction, respectively, and  $\beta_{SBP}$  and  $\beta_{DBP}$  are the intercept terms of the linear regression model for SBP and DBP prediction, respectively.

The LSTM layers have 128 hidden nodes in each phase. A dropout layer with a rate of 0.2 is applied at the end of the decoding phase to prevent overfitting and improve the generalization performance. The loss function of LSTM is defined as:


$$L = \frac{1}{N} \sum_{i=1}^N (y_{true}^{(i)} - y_{pred}^{(i)})^2, \quad (6)$$

where  $y_{true}$  is the ground truth SBP or DBP,  $y_{pred}$  is the predicted BP, and  $N$  is the number of samples. To optimize the training process for LSTM, we use Adam with Bandit Sampling (ADAMBS) [52], a combination of adaptive moment estimation (Adam) [53] and bandit sampling techniques, as shown in Figure 8. The bandit method is utilized to focus on the most important features of the data and dynamically adjust the distribution to address the challenge of partial feedback. In particular, the model parameters are initialized, and an important distribution is selected. A mini-batch is dynamically formed by sampling from the distribution, and the model parameters are updated using the chosen mini-batch to improve the model's performance. During training, we use the reconstructed blood flow pattern and ground-truth BP values (i.e., averaging three measurements collected from the CONTEC ABPM50 ambulatory BP monitor) to learn the coefficients and intercept terms.

## 5. Attack Evaluation and Results

### 5.1. Experimental Setup

Our experiments exploit the Meta Quest VR headset to collect accelerometer data on the z-axis with a sampling rate of 1k Hz. The reference BP data is collected using the CONTEC ABPM50 ambulatory BP monitor. In addition, we use pulse sensor [54] connected to an Arduino UNO (REV3) board to collect PPG data at 300Hz. Figure 9 (a) showcases the devices. We recruit 37 participants from university students, including 31 males and 6 females aged



(a) Measurement setup.

(b) Three postures used for evaluation.

Figure 9: Illustration of the experimental specifics, including setup devices (e.g., Meta Quest headset and PPG sensors) and the three postures adopted during data collection.

18 to 41, for data collection. Among the ethnicity, there are 25 Asians, 8 Whites, 2 Blacks, and 2 Hispanics/Latinos. All the participants have normal SBP and DBP levels ranging from 86 to 145 mmHg and 50 to 101 mmHg, respectively. During the experiments, participants wear the VR headset, ambulatory BP monitor, and pulse sensor to collect the accelerometer and reference data simultaneously. Each participant collects data for 5 minutes with three typical static postures in VR applications, including standing, sitting, and supine, as shown in Figure 9 (b). The supine position is often associated with longer durations of use, such as binge-watching sessions (typically less active and more relaxed). This scenario provides attackers ample opportunity to collect accurate BP data due to the minimal interference from large body motions. In our settings, we have accounted for user comfort by allowing participants to wear the headset according to their preferences. All participants were instructed to use the VR headset as they normally would during data collection but try to avoid unnecessary major adjustments to the headset itself. The participants were also suggested to minimize interrupting gestures (e.g., pull/push bands, adjust wear position), to ensure the efficiency of the data collection and analysis.

We only use PPG as a reference signal, collected in a controlled manner from the pre-selected area of the temple, as temporal artery in the head is a great site to capture pulse for BP monitoring [55]. During data collection, we ask the participants to use the headsets with minimal voluntary motion. However, we acknowledge that users may exhibit involuntary movements like occasional head nodding, breathing, blinking, swallowing, or facial muscle movements when interacting with the experiment

VR application. These unconscious behaviors may induce extremely minor vibrations that could be captured by the headset sensor. However, as these actions are intermittent and attenuate rapidly through tissue, their influence on the measured physiological vibration signal can be handled by our artifacts removal technique and, thus, is negligible for our purposes. We display gamified immersive elements in our application to enhance the participant engagement.

**Ethical Considerations.** The data collection procedures were approved by our university's Institutional Review Board (IRB). Given the sensitivity of blood pressure data, all participants were required to follow a Minimal Risk Protocol and sign a Research Consent Form. All collected data are securely stored on a local server that is only accessible by the principle instructor and graduate student in this project. Note that researchers have shown that a minimum of 33 participants is sufficient to obtain accurate validation results for BP monitoring [56]. Therefore, our evaluation involving 37 participants yields reasonably precise results.

## 5.2. Attack Evaluation Metrics

We use three widely used metrics including mean error ( $ME$ ), standard deviation of mean error ( $STD$ ), and Pearson's correlation coefficient ( $R$ ) to assess the accuracy and robustness of BPSniff.  $ME$  is the average difference between the predicted and reference BP from the ground truth. We define it as:

$$ME = \frac{\sum_{i=1}^n (X_i - Y_i)}{n}. \quad (7)$$

$STD$  measures the variability of the mean error as defined below, indicating the precision of BPSniff's predictions.

$$STD = \sqrt{\frac{\sum_{i=1}^n (X_i - Y_i - ME)^2}{n}}. \quad (8)$$

$R$  reflects the linear relationship between the predicted and reference BP, which is defined as:

$$R = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2} \sqrt{\sum_{i=1}^n (Y_i - \bar{Y})^2}}. \quad (9)$$

With values closer to 1,  $R$  indicates a stronger correlation. For the formulas of  $ME$ ,  $STD$  and  $R$ , we have  $X_i$  defined as the predicted value of BPSniff,  $Y_i$  represents the reference values of BP from ground truths,  $\bar{X}$  is the mean of predicted BP and  $\bar{Y}$  is the mean of reference BP,  $n$  is the number of samples. We record absolute value for all the results.

## 5.3. Attack Performance

To evaluate the performance of our attack for unseen victims (i.e., the attacker does not have the victims' training data), we conduct leave-one-participant-out validation. In the experiments, we use data from one participant (as a victim) for testing and data from the remaining participants for training. We first show the attack's performance using the mean estimation error of SBP and DBP in a day. As shown in Table 1, our attack can achieve low mean estimation errors




Figure 10: Mean error results of BP estimation for each of the 37 participants.




Figure 11: Standard deviation of BP estimation error for each of the 37 participants.

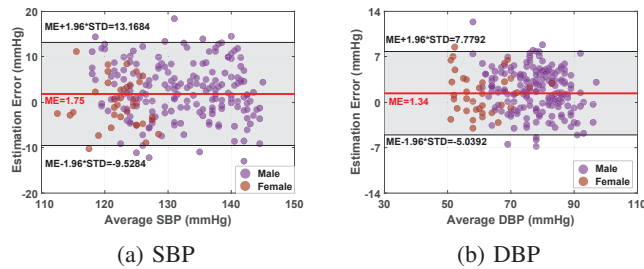


Figure 12: Bland–Altman diagram of SBP and DBP. The plot displays the difference between measurements of BP-Sniff and measurements of ambulatory BP monitor, and includes lines for the mean error and the limits of agreement.

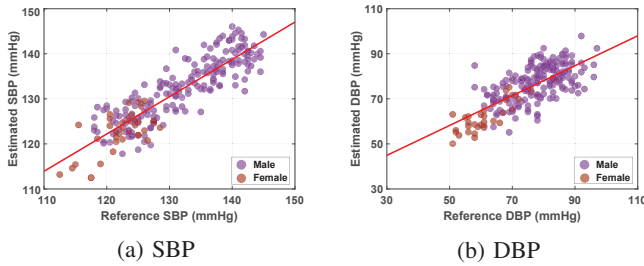


Figure 13: Correlation diagram of SBP and DBP. The plots visualize the relationship between reference BP values and the estimated BP values. The dots reflect BP values for all combinations of training and testing data from 37 subjects.

with SBP and DBP of only 1.75 mmHg and 1.34 mmHg, respectively. These results not only demonstrate the attack’s effectiveness but also comply with the criteria set by the FDA’s AAMI standard [57], which requires a mean error of less than or equal to 5 mmHg and a standard deviation of less than or equal to 8 mmHg for SBP and DBP. Furthermore, Table 2 compares our attack’s performance with the requirement for the Britain Hypertension Society (BHS) standard [58]. We can observe that our attack achieves grade A for both SBP and DBP estimation. These results underscore the severity of our attack, as it allows attackers

TABLE 1: Continuous BP estimation performance.

CBP Estimation (mmHg)	Single Trial	Different Trials, Same Day	Different Days
SBP (ME±STD)	1.75 (±5.76)	3.74 (±6.64)	4.96 (±7.87)
DBP (ME±STD)	1.34 (±4.22)	3.03 (±4.78)	4.72 (±5.91)

TABLE 2: BPSniff performance compared to British Hypertension Society standard (Grade A/B/C).

Method	Type	≤ 5 mmHg	≤ 10 mmHg	≤ 15 mmHg
BHS	A/B/C	60%/50%/40%	85%/75%/65%	95%/90%/85%
BPSniff	SBP/DBP	75.4%/73.4%	89.8%/85.2%	99.5%/96.9%

to obtain highly accurate SBP and DBP readings from unsuspecting victims, compromising their privacy without their consent.

We further use Bland-Altman plots for the estimated SBP and DBP as shown in Figure 12, which contain all combinations of training and testing data. The x-axis represents the average reference BP of all participants, while the y-axis represents the estimation error. Our results show mean errors of 1.75 mmHg and 1.34 mmHg for SBP and DBP, respectively, with standard deviations of 5.76 mmHg and 4.22 mmHg for the mean error of SBP and DBP. These results demonstrate the accuracy of our proposed system for CBP monitoring. The red line represents the mean error, while the black lines represent the limits of agreement ((LOA), defined as  $ME \pm 1.96 \times STD$ ). We observed that more than 95% of the points lie within the limit of agreement in SBP and DBP. In Figure 13, we present a correlation diagram between the estimated results and ground truth. Points that lie on the red line indicate identical values, and the distance from the red line is proportional to the error. The clustered distribution of points around the red line indicates a strong correlation between the estimated and reference values. The Pearson’s correlation coefficients for SBP and DBP are 0.86 and 0.82, respectively, indicating a good correlation between the estimated and reference BP values. For a more detailed comparison, we have plotted Figure 10 and Figure 11 to visualize the ME and STD for each participant.

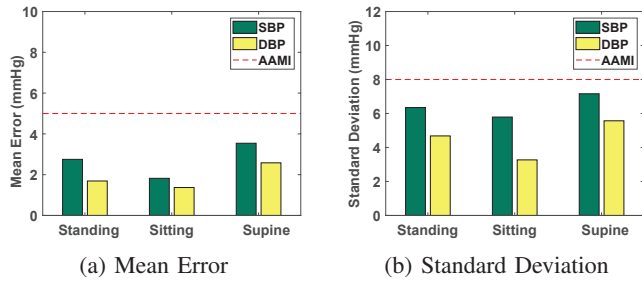


Figure 14: Impact of different postures including standing, sitting, and supine.

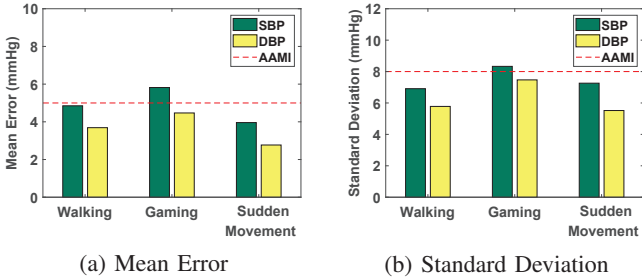


Figure 15: Impact of different moving conditions including walking, gaming, and sudden posture change.

#### 5.4. Impact Factor Studies

**Impact of Subject’s Posture During the BP Measurement.** Human BP can vary for different postures. For example, BP may be higher when standing and tends to decrease when sitting or lying down. Therefore, evaluating BPSniff’s attack performance under the victim’s different postures is crucial for assessing its robustness. As shown in Figure 14, our attack can achieve mean errors of 2.72 mmHg for SBP and 1.64 mmHg for DBP when the victim is standing still, 1.78 mmHg for SBP and 1.57 mmHg for DBP when the victim is sitting still, and 2.54 mmHg for SBP and 1.58 mmHg for DBP when the victim is in a supine position. The standard deviations of SBP/DBP for BPSniff are 6.36/4.69 mmHg, 5.79/3.27 mmHg, and 5.86/3.33 mmHg for standing still, sitting still, and lying supine, respectively. The results indicate our attack’s ability to accurately measure BP across different postures.

**Impact of Subject’s Movements.** Movements can distort the BP readings to be falsely elevated or falsely lowered, leading to inaccurate measurements. To study the impact of subject’s movements on BPSniff, we conduct experiments with participants using Meta Quest with typical movements: walking, gaming, and non-exercise movements. Participants in the experiments are instructed to walk at a natural pace (casual walking scenario that individuals might experience in everyday life), which typically ranges from 3 to 4 miles per hour [59]. The walking environment is designed to mimic typical indoor settings, such as labs, meeting rooms, hallways or open spaces, where participants can move freely without significant obstacles when using VR applications. Meanwhile, the selection of games is intended to simulate a range of user interactions within the metaverse, reflecting how users may experience VR in both casual and active




Figure 16: Impact of different sampling rates in Meta Quest data collection for BP estimation.




Figure 17: Impact of different regression models including Linear Regression, Random Forest, Support Vector Machine, and XGBoost.

gaming scenarios. The games (e.g., PokerVR [60], MiRacle Pool [61], and Chess Club [62]) involve light physical movements, such as reaching or turning. Also, we define non-exercise movements as sudden movements that participants may make during VR interactions. Examples include bending over to pick something up, quickly turning to look at something behind them, or shifting weight from one foot to another quickly. As shown in Figure 15, the BPSniff system has mean errors of 4.86 mmHg for SBP and 3.64 mmHg for DBP when the victim is walking, 5.81 mmHg for SBP and 4.45 mmHg for DBP when the victim is playing game, and 3.89 mmHg for SBP and 2.65 mmHg for DBP when the victim has posture change suddenly. The standard deviations of SBP/DBP for BP are 6.92/5.76 mmHg, 8.34/6.88 mmHg, and 7.26/5.51 mmHg for walking, gaming, and sudden posture change, respectively. Therefore, it is important to minimize subject movement during BP measurement to obtain accurate readings.

**Impact of VR Headset Model.** We validate our BPSniff attack on Meta Quest 2, and conduct a comparative analysis with the Meta Quest. Meta Quest and Meta Quest 2 both use the same series of motion sensor boards developed by Meta. To assess BP estimation-based attack’s efficacy on Meta Quest 2, we enlist a varied group of 10 participants (same 10 engaged in Meta Quest experiments): seven males and three females, spanning ages 17 through 37. The test is implemented in sitting setting. As shown in Figure 18, for Meta Quest, the BPSniff achieves a mean error of 1.85 mmHg and a standard deviation of 5.98 mmHg for SBP estimation, and a mean error of 1.64 mmHg and a standard deviation of 4.94 mmHg for DBP estimation. And for Meta Quest 2, the BPSniff achieves a mean error of 1.97 mmHg and a standard deviation of 6.63 mmHg for SBP estimation,

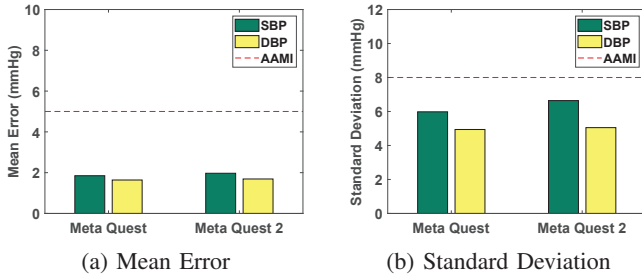


Figure 18: Impact of different VR headset models (Meta Quest and Meta Quest 2).

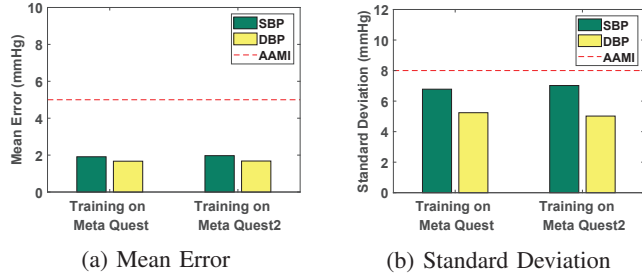


Figure 19: Training on data from different VR headset models to verify the transferability of the attack.

and a mean error of 1.69 mmHg and a standard deviation of 5.05 mmHg for DBP estimation. The results demonstrate similar estimation performance across these two different headset models, confirming the effectiveness of our attack.

**Impact of Sampling Rate.** A higher sampling rate can potentially provide more accurate and detailed information about the victim’s motion. However, most consumer smartphones/smartwatches or low-cost VR headsets are not equipped with high sampling rate IMUs by far like the Meta Quest. We conduct this study to study the impact of the sampling rate. Based on the study results with sitting still subjects using the Meta Quest headset, increasing the motion sensor’s sampling rate in the BPSniff system could enhance BP measurement accuracy. As shown in Figure 16, the sampling rate increased from 100 to 1000 Hz, the mean error in BP measurement decreased steadily, from 6.23 mmHg to 1.75 mmHg, and resulted in a decrease in standard deviation from 12.13 mmHg to 5.76 mmHg, respectively. The results suggest that higher sampling rates can lead to more accurate and consistent BP monitoring for victims. The results also demonstrate that the sampling rate of IMU needs to reach at least 350 Hz to achieve a good attack performance (i.e., meeting the AAMI standard).

**Impact of Regression Models.** We implement different regression models in BPSniff, including linear regression, support vector machine, random forest, and XGboost. We follow standard 7/3 ratio for training/testing and use leave-one-out cross-validation. As shown in Figure 17, XGboost achieves the lowest mean errors for both SBP (1.94 mmHg) and DBP (1.52 mmHg). The mean errors of support vector machine are also low, with SBP (2.23 mmHg) and DBP (1.61 mmHg). Linear regression has mean errors of 2.42 mmHg for SBP and 1.76 mmHg for DBP. Random forest has the highest mean errors, with 4.14 mmHg for SBP and

3.31 mmHg for DBP. These results demonstrate that the performance of our attack still meets the acceptable range criteria defined by the FDA’s AAMI protocol when leveraging four different traditional machine learning models.

## 5.5. Longitudinal Measurements

To assess the long-term viability and efficiency of BP-Sniff, we conduct an extensive longitudinal study simulating real-world scenarios. This study involves 23 participants acting as potential victims, monitored over an 8-week period. The participants are asked to measure their BP using the Meta Quest VR headset and the BP cuff on four different days during the period, i.e., every 2 weeks for 8 weeks. For each day, the participants’ data are collected in three different time periods: 8-10 AM, 2-4 PM, and 7:30-9:30 PM. Table 1 shows the performance of BPSniff when using the testing data collected in different trials of the same day and on different days with the system trained by the data collected on the first day. BPSniff achieves a mean error of 3.74 mmHg and a standard deviation of 6.64 mmHg for SBP estimation and a mean error of 3.03 mmHg and a standard deviation of 4.78 mmHg for DBP across different trials in a day. Additionally, the mean error of SBP and DBP across different days are 4.96 mmHg and 4.72 mmHg, with the standard deviation of 7.87 mmHg and 5.91 mmHg, respectively. Our results demonstrate high accuracy over the longitudinal study without significant performance degradation and manifest that BPSniff consistently meets the FDA’s AAMI standard, even in challenging scenarios across different days.

## 5.6. Transferability of BPSniff

We validate our attack on two models of Meta Quest headsets (Meta Quest and Meta Quest 2), which dominate the consumer VR market [63], [64]. As the most popular and widely used VR headsets, the Quest’s vulnerability to BPSniff presents a large-scale and immediate threat. We conduct experiments with leaving-one-out cross-validation settings for Meta Quest and Meta Quest 2, involving a group of 10 participants. As shown in Figure 19, when using Meta Quest data for training and Meta Quest 2 data for testing, we obtain mean errors of 1.91 mmHg for SBP and 1.67 mmHg for DBP, with the standard deviation of 6.78 mmHg and 5.24 mmHg, respectively. Conversely, with Meta Quest 2 data as the training set and Meta Quest data as the testing set, mean errors are 1.97 mmHg for SBP and 1.68 mmHg for DBP, with the standard deviation of 7.02 mmHg and 5.02 mmHg, respectively. These results demonstrate our attack model’s transferability across different VR headsets/hardware platforms.

## 6. Discussion

**Countermeasures and Mitigations.** One potential defense approach is to develop a privacy-aware sensor management framework. This framework would enhance the

transparency of sensor utilization on metaverse devices, such as VR headsets. Specifically, it could provide real-time statistical assessments of sensor usage, including the duration of sensor use by specific metaverse apps or websites, allowing users to decide if an app or website recording sensor data in the background should be restricted. In addition, an on-device AI sentinel can be built to continuously audit data access patterns, helping users distinguish benign from malicious intent in applications that persistently harvest users' VR motion sensor data. Another potential defense approach is to enhance VR sensor management. For example, developers can implement a permission-based scheme similar to the one used in Android [65]. During installation and runtime, users would be asked to grant permission for metaverse apps or webpages to access specific sensors. This would ensure that users are aware of and can control which sensors are being accessed, enhancing privacy and security.

**Differences in Facial Structures Across Users.** Our current experiments involve participants from different races and diverse facial structures, which may fractionally impact how vibrations from underlying blood vessels propagate through tissue. Facial features like brow ridge protrusion, nasal bone shape, and forehead contour vary among individuals and can influence vibration transmission [66], [67], [68]. Additionally, skin thickness [69], particularly in facial regions where the headset makes contact, can attenuate transmitted vibrations to varying degrees, with thinner skin allowing more effective transmission compared to thicker tissue. Our experiments encompassed a diverse range of participants with varying facial structures to evaluate the robustness of our attack. Despite the heterogeneity in facial features, the standard deviation of error remained relatively small, demonstrating that facial variations do not significantly impact the effectiveness of our attack methodology.

**Generalizability of Practical Attack.** Our attack's scope extends beyond Meta Quest devices. Any augmented reality (AR), VR, or mixed reality (MR) headset featuring always-on, high-sampling-rate, unrestricted motion sensors (i.e., a design philosophy increasingly prevalent in pursuit of better user experiences) is susceptible to our attack. This broad applicability stems from two key factors. First is that our feature selection process ensures the selected features are highly related to BP and less susceptible to variations in different environments or times. Secondly, our LSTM-based BP regressor is designed to capture long-term dependencies in data, allowing it to adapt to subtle changes over time and across settings. The combination of these methods ensures the generalizability of our attack - BPSniff. This means that the BPSniff can be extended to a broader range of AR/VR/MR headset models, including HTC Vive XR Elite [70], Apple Vision Pro [71], and etc.

**Escalating Threat Leveraging Additional Measurements from VR Controllers.** Metaverse headsets often feature controllers equipped with high-sampling motion sensors, offering an additional vantage point for capturing blood flow-induced vibrations and enhancing blood flow waveform reconstruction. VR controllers thus make the privacy risk even worse for metaverse users. When someone holds a

controller, its sensors track their hand movements. These extra motion sensor readings help adversaries better remove the shaky movements that usually make it hard to read a person's blood pressure during active VR use. This means they can now steal a user's private blood pressure data in many more situations involving many motions. As most VR headsets today come with these controllers, adversaries can target almost any scenario in the metaverse with this enhanced attack.

**Snooping Private BP information from Metaverse Users with Chronic Diseases.** Our current study demonstrates BPSniff's ability to accurately estimate blood pressure in healthy individuals, but its implications for users with chronic diseases are also alarming. We are the first to expose this grave privacy vulnerability: adversaries can surreptitiously harvest sensitive health biomarker-BP through unrestricted VR motion sensors, without user awareness. We posit that BPSniff's threat extends ominously to the metaverse's participants with chronic diseases. Conditions like hypertension, diabetes, and heart disease often manifest unique BP signatures. By training an attack model with data from these user groups, adversaries could not only deduce the presence of such conditions but also track their severity, progression, and even the efficacy of treatments. Moving forward, we aim to explore further along this line in the near future.

**Other Severe Impacts of BPSniff Attack.** BPSniff reveals a troubling new way for adversaries to secretly keep monitoring one of the most personal and visceral bodily signals, continuous blood pressure levels. Adversaries could use leaked BP data to discriminate [72], [73] against individuals with cardiovascular conditions, affecting employment, insurance, and service decisions, thereby reinforcing societal biases. Financial criminals [74] might use BP data to predict stress in traders, optimizing times for phishing or scamming when victims are more vulnerable. Moreover, BP patterns can reveal personal health statuses such as sleep disorders [75], [76], posttraumatic stress disorder [77], [78], or severe anxiety [79], [80], all of which are protected health information. More disturbingly, with continuous BP data, adversaries could reconstruct a user's long-term medical history through simulated metaverse attacks [81]. By identifying those users with heart arrhythmias/valve issues, adversaries could orchestrate 'heart-tuned' audio-visual stimuli in VR, crafting personalized jump-scares or stress-inducing scenarios that risk triggering fatal cardiac events.

## 7. Related Work

### 7.1. Attacks on VR Systems

Most of the early studies on the security and privacy implications of VR systems concentrated on developing new user authentication mechanisms [90], [91], [92], [93], [94]. More recently, studies on VR privacy have been concentrated on leakages in active virtual engagement [95], [96], [97]. Specifically, Dick et al. [95] demonstrate social engineering attacks where users are tricked into revealing details

TABLE 3: Comparisons Between BPSniff and Existing Works.

Research Work	Leakage Info & Target	Device	Sensor Data Usage	No User's Specific Training Needed	Calibration-free	Longitudinal Study
<b>BPSniff (This Work)</b>	<b>Blood Pressure-Malicious</b>	<b>VR</b>	<b>IMU</b>	✓	✓	✓
Face-Mic [82]	Speech Content-Malicious	VR	IMU	✓	×	×
Wu et al. [13]	keystroke Input Info-Malicious	VR	IMU	×	×	×
TyPose [83]	keystroke Input Info-Malicious	VR	IMU	✓	×	×
FaceReader [84]	Respiratory & Heart Rate-Malicious	VR	IMU	✓	×	✓
Crisp-BP [85]	Blood Pressure-Benign	Smartwatch	PPG	✓	×	×
mmBP [86]	Blood Pressure-Benign	Millimetre-Wave Radar	Millimetre-wave	×	×	×
Seismowatch [87]	Blood Pressure-Benign	Smartwatch	ECG	×	×	×
Wang et al. [88]	Blood Pressure-Benign	Smartphone	IMU	✓	×	×
Glabella [89]	Blood Pressure-Benign	Glasses	PPG	×	×	×

about their locations, physical appearances, and movements. Maloney et al. [96] note that users often share personal experiences on VR social media, significantly increasing privacy breach risks. Nair et al. [97] create a game that can expose users' locations and movements that may unveil users' privacy (e.g., height, age, and gender). Moreover, Casey et al. [98] introduce a VR-based attack that can covertly modify VR environments, causing users to collide with physical objects. These attack requires meticulous skill or significant effort to deceive/convince users into disclosing private information.

Recent research works [13], [82], [83], [84] have exposed privacy breaches through unrestricted position and motion sensors on metaverse devices including VR and AR headsets. Among them, Wu et al. [13] have devised a keystroke snooping attack that utilizes the position and orientation of motion sensor readings from VR controllers to deduce users' keystrokes, enabling the inference of passwords or typed content. Shi et al. [82] have introduced Face-Mic, an eavesdropping attack that leverages motion sensors of VR headsets to infer the user's live speech and identity. In addition, Slocum et al. [83] have introduced TyPose, a system capable of inferring the text input of a user on a virtual keyboard by analyzing the user's subtle head movements. Nevertheless, these attacks are effective during user interactions with AR/VR systems, such as engaging in verbal interactions or utilizing gesture-based typing. Furthermore, Zhang et al. [84] have developed an attack, FaceReader, that leverages motion sensor readings captured by VR headsets to reconstruct users' breathing and heartbeat patterns. Such patterns are then utilized for more advanced attacks, including gender recognition, user re-identification, and body fat ratio estimation.

Compared to FaceReader, BPSniff is significantly different. BPSniff unveil a new type of continuous privacy leakage of victim's BP that can lead to more severe consequences (e.g., disclosing emotional states and stress/anxiety levels for unauthorized purposes), whereas FaceReader leaks breathing and heart rates and users' demographic information, which cannot derive BP. Moreover, BPSniff attack is calibration-free and allowing inference of victims' BP

without using their data for training/fine-tuning, which is not achieved in FaceReader. In addition, the longitudinal study demonstrated that BPSniff could robustly estimate BP over an 8-week period, whereas similar validations were not performed in FaceReader.

## 7.2. Benign BP Estimation Using Smart Sensing

Traditional BP monitoring relies on cuff-based BP measurements [99], involving the use of dedicated cuff-based devices around the upper arm. However, this approach only provides one-time measurement, and the rigid setup is considered inconvenient. Researchers have developed cuff-less BP monitoring systems to address the inconvenience of traditional BP measurements. For example, researchers have proposed BP monitoring systems on wrist-worn or in-ear wearable devices using photoplethysmogram (PPG) [85], [89], [100], [101], electrocardiogram (ECG) [87], [102], and impedance plethysmography (IPG) [103]. However, all these studies have indicated the necessity of user-specific training. Xuan et al. [104] lately demonstrate a calibration-free approach requiring the users to wear a plastic clip and keep a specific hand posture, which is not practical in attacks.

Researchers have also explored contactless sensing technologies to allow users monitor BP without wearing any devices. For instance, Sugita *et al.* [105] proposed to leverage light absorption and reflection on different body parts to estimate BP. This method has strict lighting requirements and raises privacy concerns due to the use of video cameras. Radio-frequency signals [86], [106], [107] have also been utilized to capture pulse motions and estimate users' BP. However, the performance of these methods are sensitive to radio interference and body motions.

Recent studies [88], [108] propose using motion sensors on smartphones to capture the minute vibrations induced by cardiac movements in the human chest for BP measurement. However, this method requires the user to attach a phone to their chest (using the phone's accelerometer pressed against the chest), and the vibration signals are subject to changes due to body movements and can vary significantly with

different postures. Therefore, this method is not applicable to metaverse scenarios involving various body motions and postures. Compared to this method, our attack utilizes robust signal processing techniques to isolate BP-related signals. We also designed sensitive feature extraction/selection schemes and an adaptive model to mitigate the body motions and ensure accurate BP information extraction in VR environments.

We summarize the differences between BPSniff and existing studies as illustrated in Table 3.

## 8. Conclusion

We propose BPSniff, the first attack system that continuously estimates systolic and diastolic blood pressure (BP) in the metaverse using cardiac cycle-induced vibrations captured by motion sensors in a VR headset. Our system incorporates a variational autoencoder (VAE) to reconstruct fine-grained cardiac signals resembling blood flow patterns regardless of body motions. A long short-term memory (LSTM) regression model is designed to continuously map the successive blood flow patterns to BP measurements. We conduct extensive experiments with 37 subjects and a longitudinal study of 8 weeks for evaluation. The experiment results show that BPSniff can achieve low mean errors of 1.75 mmHg and 1.34 mmHg and standard deviation errors of 5.76 mmHg and 4.22 mmHg for systolic and diastolic blood pressure, respectively. This result meets the golden standard of the FDA. It attains comparable BP estimation accuracy as the commercial BP monitors, indicating the severe privacy leakage in the metaverse that adversaries can obtain users' protected health information without using dedicated devices.

## Acknowledgment

This work was partially supported by the National Science Foundation Grants CNS2120276, CNS2120396, CNS2145389, CNS2152669, CNS2201465, CNS2329278, CNS2329279, CCF2211163, IIS2311596, IIS2311597, DGE2414365, DGE2414366, DGE2414367, and OAC2139358.

## References

- [1] M. Gupta, "Metaverse applications in the real world," 2023. [Online]. Available: <https://www.forbes.com/sites/forbestechcouncil/2023/01/12/metaverse-applications-in-the-real-world/?sh=36f677a348a2>
- [2] S. Srivastava, "Metaverse use cases and benefits," 2023. [Online]. Available: <https://appinventiv.com/blog/metaverse-use-cases-and-benefits/>
- [3] N. H. Shahimi, R. Lim, S. Mat, C.-H. Goh, M. P. Tan, and E. Lim, "Association between mental illness and blood pressure variability: a systematic review," *Biomedical engineering online*, vol. 21, no. 1, p. 19, 2022.
- [4] B. M. Psaty, C. D. Furberg, L. H. Kuller, M. Cushman, P. J. Savage, D. Levine, D. H. O'Leary, R. N. Bryan, M. Anderson, and T. Lumley, "Association between blood pressure level and the risk of myocardial infarction, stroke, and total mortality: the cardiovascular health study," *Archives of internal medicine*, vol. 161, no. 9, pp. 1183–1192, 2001.
- [5] J. A. McCubbin, M. M. Merritt, J. J. Sollers III, M. K. Evans, A. B. Zonderman, R. D. Lane, and J. F. Thayer, "Cardiovascular-emotional dampening: the relationship between blood pressure and recognition of emotion," *Psychosomatic medicine*, vol. 73, no. 9, pp. 743–750, 2011.
- [6] T. M. Spruill, "Chronic psychosocial stress and hypertension," *Current hypertension reports*, vol. 12, no. 1, pp. 10–16, 2010.
- [7] P. F. Edemekong, P. Annamaraju, and M. J. Haydel, "Health insurance portability and accountability act," 2018.
- [8] I. A. Kretchy, F. T. Owusu-Daaku, and S. A. Danquah, "Mental health in hypertension: assessing symptoms of anxiety, depression and stress on anti-hypertensive medication adherence," *International journal of mental health systems*, vol. 8, pp. 1–6, 2014.
- [9] C. Lauckner, N. Truszczynski, D. Lambert, V. Kottamasu, S. Meherally, A. M. Schipani-McLaughlin, E. Taylor, and N. Hansen, "'catfishing,' cyberbullying, and coercion: An exploration of the risks associated with dating app use among rural sexual minority males," *Journal of Gay & Lesbian Mental Health*, vol. 23, no. 3, pp. 289–306, 2019.
- [10] U. Upadhyay, A. Kumar, G. Sharma, B. B. Gupta, W. A. Alhalabi, V. Arya, and K. T. Chui, "Cyberbullying in the metaverse: A prescriptive perception on global information systems for user protection," *Journal of Global Information Management (JGIM)*, vol. 31, no. 1, pp. 1–25, 2023.
- [11] S. Saharan, S. Singh, A. K. Bhandari, and B. Yadav, "The future of cyber-crimes and cyber war in the metaverse," in *Forecasting Cyber Crimes in the Age of the Metaverse*. IGI Global, 2024, pp. 126–148.
- [12] H. Oh and W. Son, "Cybersickness and its severity arising from virtual reality content: A comprehensive study," *Sensors*, vol. 22, no. 4, p. 1314, 2022.
- [13] Y. Wu, C. Shi, T. Zhang, P. Walker, J. Liu, N. Saxena, and Y. Chen, "Privacy leakage via unrestricted motion-position sensors in the age of virtual reality: A study of snooping typed input on virtual keyboards," in *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2023, pp. 3382–3398.
- [14] I. Sadek, J. Biswas, and B. Abdulrazak, "Ballistocardiogram signal processing: A review," *Health information science and systems*, vol. 7, no. 1, p. 10, 2019.
- [15] C.-S. Kim, S. L. Ober, M. S. McMurtry, B. A. Finegan, O. T. Inan, R. Mukkamala, and J.-O. Hahn, "Ballistocardiogram: Mechanism and potential for unobtrusive cardiovascular health monitoring," *Scientific reports*, vol. 6, no. 1, p. 31297, 2016.
- [16] C.-S. Kim, A. M. Carek, O. T. Inan, R. Mukkamala, and J.-O. Hahn, "Ballistocardiogram-based approach to cuffless blood pressure monitoring: Proof of concept and potential challenges," *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 11, pp. 2384–2391, 2018.
- [17] C. Speechly, N. Bignell, and M. Turner, "Sphygmomanometer calibration: why, how and how often?" *Australian family physician*, vol. 36, no. 10, 2007.
- [18] P. Muntner, D. Shimbo, R. M. Carey, J. B. Charleston, T. Gailard, S. Misra, M. G. Myers, G. Ogedegbe, J. E. Schwartz, R. R. Townsend *et al.*, "Measurement of blood pressure in humans: a scientific statement from the american heart association," *Hypertension*, vol. 73, no. 5, pp. e35–e66, 2019.
- [19] D. B. McCombie, P. A. Shaltis, A. T. Reinsner, and H. H. Asada, "Adaptive hydrostatic blood pressure calibration: Development of a wearable, autonomous pulse wave velocity blood pressure monitor," in *2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2007, pp. 370–373.

- [20] D. Da He, E. S. Winokur, and C. G. Sodini, "An ear-worn continuous ballistocardiogram (bcg) sensor for cardiovascular monitoring," in *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2012, pp. 5030–5033.
- [21] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, "Activity recognition using cell phone accelerometers," *ACM SigKDD Explorations Newsletter*, vol. 12, no. 2, pp. 74–82, 2011.
- [22] H. Wang, T. T.-T. Lai, and R. Roy Choudhury, "Mole: Motion leaks through smartwatch sensors," in *Proceedings of the 21st annual international conference on mobile computing and networking*, 2015, pp. 155–166.
- [23] J. Gordon, "Certain molar movements of the human body produced by the circulation of the blood," *Journal of anatomy and physiology*, vol. 11, no. Pt 3, p. 533, 1877.
- [24] S. Magder, "The meaning of blood pressure," *Critical Care*, vol. 22, no. 1, p. 257, 2018.
- [25] J. R. Levick, *An introduction to cardiovascular physiology*. Butterworth-Heinemann, 2013.
- [26] A. Baron, G. Brechtel-Hook, A. Johnson, and D. Hardin, "Skeletal muscle blood flow. a possible link between insulin resistance and blood pressure." *Hypertension*, vol. 21, no. 2, pp. 129–135, 1993.
- [27] J. E. Wagenseil and R. P. Mecham, "Vascular extracellular matrix and arterial mechanics," *Physiological reviews*, vol. 89, no. 3, pp. 957–989, 2009.
- [28] T. Ozaki, S. SASAKI, and K. IGARASHI, "Ballistocardiography simplified by recording microvibration of the skin surface of the head," *The Tohoku Journal of Experimental Medicine*, vol. 102, no. 1, pp. 99–100, 1970.
- [29] T. Abay and P. Kyriacou, "Photoplethysmography for blood volumes and oxygenation changes during intermittent vascular occlusions," *Journal of clinical monitoring and computing*, vol. 32, no. 3, pp. 447–455, 2018.
- [30] Y. Liang, M. Elgendi, Z. Chen, and R. Ward, "An optimal filter for short photoplethysmogram signals," *Scientific data*, vol. 5, no. 1, pp. 1–12, 2018.
- [31] 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*, vol. 64, no. 5, pp. 1003–1010, 2016.
- [32] E. S. Winokur, "Single-site, noninvasive, blood pressure measurements at the ear using ballistocardiogram (bcg), and photoplethysmogram (ppg), and a low-power, reflectance-mode ppg soc," Ph.D. dissertation, Massachusetts Institute of Technology, 2014.
- [33] H. Samimi and H. R. Dajani, "A ppg-based calibration-free cuffless blood pressure estimation method using cardiovascular dynamics," *Sensors*, vol. 23, no. 8, p. 4145, 2023.
- [34] M. Kachuee, M. M. Kiani, H. Mohammadzade, and M. Shabany, "Cuff-less high-accuracy calibration-free blood pressure estimation using pulse transit time," in *2015 IEEE international symposium on circuits and systems (ISCAS)*. IEEE, 2015, pp. 1006–1009.
- [35] Meta, "Introducing oculus quest, our first 6dof all-in-one vr system," *Meta Quest Blog*, 2018. [Online]. Available: <https://www.meta.com/blog/quest>
- [36] "Get started with meta quest 2," *Meta*, 2022. [Online]. Available: <https://www.meta.com/quest/products/quest-2/>
- [37] "Platform sdk documentation," *Meta Quest Platform Solutions*, 2024. [Online]. Available: <https://developer.oculus.com/documentation/native/ps-platform-intro/>
- [38] W. C. R. Snapshot, "Webxr device api," 2022. [Online]. Available: <https://immersive-web.github.io/webxr/>
- [39] "Supplemental meta platforms technologies privacy policy," *Meta Privacy Policy*, 2024. [Online]. Available: <https://www.meta.com/legal/privacy-policy/>
- [40] V. Corporation, "Openvr," 2021. [Online]. Available: <https://partner.steamgames.com/doc/features/steamvr/openvr>
- [41] I. O. for Standardization, *Mechanical Vibration and Shock: Evaluation of Human Exposure to Whole-body Vibration*. International Organization for Standardization, 1985.
- [42] F. Benetazzo, A. Freddi, A. Monteriù, and S. Longhi, "Respiratory rate detection algorithm based on rgb-d camera: theoretical background and experimental results," *Healthcare technology letters*, vol. 1, no. 3, pp. 81–86, 2014.
- [43] G. E. Grossman, R. J. Leigh, L. A. Abel, D. J. Lanska, and S. Thurston, "Frequency and velocity of rotational head perturbations during locomotion," *Experimental brain research*, vol. 70, no. 3, pp. 470–476, 1988.
- [44] M. Sewak, S. K. Sahay, and H. Rathore, "An overview of deep learning architecture of deep neural networks and autoencoders," *Journal of Computational and Theoretical Nanoscience*, vol. 17, no. 1, pp. 182–188, 2020.
- [45] K. Pawar and V. Z. Attar, "Assessment of autoencoder architectures for data representation," *Deep learning: concepts and architectures*, pp. 101–132, 2020.
- [46] M. Connor, G. Canal, and C. Rozell, "Variational autoencoder with learned latent structure," in *International conference on artificial intelligence and statistics*. PMLR, 2021, pp. 2359–2367.
- [47] R. Wei, C. Garcia, A. El-Sayed, V. Peterson, and A. Mahmood, "Variations in variational autoencoders—a comparative evaluation," *Ieee Access*, vol. 8, pp. 153 651–153 670, 2020.
- [48] T. Zhao, Y. Wang, J. Liu, J. Cheng, Y. Chen, and J. Yu, "Robust continuous authentication using cardiac biometrics from wrist-worn wearables," *IEEE Internet of Things Journal*, vol. 9, no. 12, pp. 9542–9556, 2021.
- [49] C. Wang, X. Li, H. Hu, L. Zhang, Z. Huang, M. Lin, Z. Zhang, Z. Yin, B. Huang, H. Gong *et al.*, "Monitoring of the central blood pressure waveform via a conformal ultrasonic device," *Nature biomedical engineering*, vol. 2, no. 9, pp. 687–695, 2018.
- [50] A. Dagamseh, Q. Qananwah, H. Al Quran, and K. S. Ibrahim, "Towards a portable-noninvasive blood pressure monitoring system utilizing the photoplethysmogram signal," *Biomedical Optics Express*, vol. 12, no. 12, pp. 7732–7751, 2021.
- [51] D. Fryer, I. Strümke, and H. Nguyen, "Shapley values for feature selection: The good, the bad, and the axioms," *Ieee Access*, vol. 9, pp. 144 352–144 360, 2021.
- [52] R. Liu, T. Wu, and B. Mozafari, "Adam with bandit sampling for deep learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5393–5404, 2020.
- [53] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [54] W. F. E. LLC, "Pulse sensor amped," <https://pulsesensor.com/products/pulse-sensor-amped>, Accessed 2023.
- [55] J. Canning, K. Helbert, G. Iashin, J. Matthews, J. Yang, M. K. Delano, C. G. Sodini, and Q. Zhang, "Noninvasive and continuous blood pressure measurement via superficial temporal artery tonometry," pp. 3382–3385, 2016.
- [56] E. O'Brien, T. Pickering, R. Asmar, M. Myers, G. Parati, J. Staessen, T. Mengden, Y. Imai, B. Waeber, P. Palatini *et al.*, "Working group on blood pressure monitoring of the european society of hypertension international protocol for validation of blood pressure measuring devices in adults," *Blood pressure monitoring*, vol. 7, no. 1, pp. 3–17, 2002.
- [57] G. S. Stergiou, B. Alpert, S. Mieke, R. Asmar, N. Atkins, S. Eckert, G. Frick, B. Friedman, T. Graßl, T. Ichikawa *et al.*, "A universal standard for the validation of blood pressure measuring devices: Association for the advancement of medical instrumentation/european society of hypertension/international organization for standardization (aami/esh/iso) collaboration statement," *Hypertension*, vol. 71, no. 3, pp. 368–374, 2018.

- [58] E. O'Brien, J. Petrie, W. Littler, M. De Swiet, P. L. Padfield, D. Altman, M. Bland, A. Coats, N. Atkins *et al.*, "The british hypertension society protocol for the evaluation of blood pressure measuring devices," *J hypertension*, vol. 11, no. Suppl 2, pp. S43–S62, 1993.
- [59] M. Schimpl, C. Moore, C. Lederer, A. Neuhaus, J. Sambrook, J. Danesh, W. Ouwehand, and M. Daumer, "Association between walking speed and age in healthy, free-living individuals using mobile accelerometry—a cross-sectional study," *PloS one*, vol. 6, no. 8, p. e23299, 2011.
- [60] I. Mega Particle, "Pokervr - pure, simple poker," 2019. [Online]. Available: <https://www.meta.com/experiences/pokervr-pure-simple-poker/2257223740990488/>
- [61] P. Works, "Miracle pool has launched on the meta main store!" 2023. [Online]. Available: <https://www.meta.com/experiences/miracle-pool-early-access/5796340207136099/>
- [62] Odders, "The greatest strategy game of all time is now available in mixed reality!" 2021. [Online]. Available: <https://www.meta.com/experiences/chess-club/5353996901307344/>
- [63] "The market for virtual reality headsets in 2024: meta dominates." [Online]. Available: <https://www.zreality.com/the-market-for-virtual-reality-headsets-in-2024-meta-dominates/?lang=en>.
- [64] B. ELAD, "Vr headset statistics by per unit sales, geography, shipments and usages," 2024. [Online]. Available: <https://www.coolest-gadgets.com/vr-headset-statistics>
- [65] Z. Xu and S. Zhu, "Semadroid: A privacy-aware sensor management framework for smartphones," in *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy*, 2015, pp. 61–72.
- [66] A. Gad, M. Laurino, K. R. Maravilla, M. Matsushita, and W. H. Raskind, "Sensorineural deafness, distinctive facial features, and abnormal cranial bones: a new variant of waardenburg syndrome?" *American Journal of Medical Genetics Part A*, vol. 146, no. 14, pp. 1880–1885, 2008.
- [67] Q. Li, X. Zhou, Y. Wang, J. Qian, and Q. Zhang, "Facial paralysis in patients with hemifacial microsomia: frequency, distribution, and association with other omens abnormalities," *Journal of Craniofacial Surgery*, vol. 29, no. 6, pp. 1633–1637, 2018.
- [68] L. Bartel and A. Mosabbir, "Possible mechanisms for the effects of sound vibration on human health," in *Healthcare*, vol. 9, no. 5. MDPI, 2021, p. 597.
- [69] M. Xu and J. Yang, "Human facial soft tissue thickness and mechanical properties: a literature review," in *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 57045. American Society of Mechanical Engineers, 2015, p. V01AT02A045.
- [70] K. Wickens, "Htc vive xr elite," 2024. [Online]. Available: <https://www.pcgamer.com/htc-vive-xr-elite-review/>
- [71] J. Chokkattu, "Review: Apple vision pro," 2024. [Online]. Available: <https://www.wired.com/review/apple-vision-pro/>
- [72] N. Krieger, J. T. Chen, P. D. Waterman, C. Hartman, A. M. Stoddard, M. M. Quinn, G. Sorensen, and E. M. Barbeau, "The inverse hazard law: blood pressure, sexual harassment, racial discrimination, workplace abuse and occupational exposures in us low-income black, white and latino workers," *Social science & medicine*, vol. 67, no. 12, pp. 1970–1981, 2008.
- [73] C. R. Center, "The metaverse: Opportunities, risks, and harms," 2024. [Online]. Available: <https://cyberbullying.org/metaverse>
- [74] I. S. General, "Metaverse - a law enforcement perspective use cases, crime, forensics, investigation, and governance white paper," 2024. [Online]. Available: <https://www.interpol.int/content/download/20828/file/Metaverse%20-%20a%20law%20enforcement%20perspective.pdf>
- [75] D. A. Calhoun and S. M. Harding, "Sleep and hypertension," *Chest*, vol. 138, no. 2, pp. 434–443, 2010.
- [76] J. M. Dopp, K. J. Reichmuth, and B. J. Morgan, "Obstructive sleep apnea and hypertension: mechanisms, evaluation, and management," *Current hypertension reports*, vol. 9, no. 6, pp. 529–534, 2007.
- [77] E. J. Paulus, T. R. Argo, and J. A. Egge, "The impact of posttraumatic stress disorder on blood pressure and heart rate in a veteran population," *Journal of traumatic stress*, vol. 26, no. 1, pp. 169–172, 2013.
- [78] D. Edmondson, J. A. Sumner, I. M. Kronish, M. M. Burg, L. Oye-siku, and J. E. Schwartz, "The association of posttraumatic stress disorder with clinic and ambulatory blood pressure in healthy adults," *Psychosomatic medicine*, vol. 80, no. 1, pp. 55–61, 2018.
- [79] Y. Pan, W. Cai, Q. Cheng, W. Dong, T. An, and J. Yan, "Association between anxiety and hypertension: a systematic review and meta-analysis of epidemiological studies," *Neuropsychiatric disease and treatment*, pp. 1121–1130, 2015.
- [80] T. Qiu, Z. Jiang, X. Chen, Y. Dai, H. Zhao *et al.*, "Comorbidity of anxiety and hypertension: common risk factors and potential mechanisms," *International Journal of Hypertension*, vol. 2023, 2023.
- [81] M. H. page, "Vr headsets can be hacked with an inception-style attack - researchers managed to crack meta's quest vr system, allowing them to steal sensitive information, and manipulate social interactions." 2024. [Online]. Available: <https://www.technologyreview.com/2024/03/11/1089686/hack-vr-headsets-inception/>
- [82] C. Shi, X. Xu, T. Zhang, P. Walker, Y. Wu, J. Liu, N. Saxena, Y. Chen, and J. Yu, "Face-mic: inferring live speech and speaker identity via subtle facial dynamics captured by ar/vr motion sensors," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, 2021, pp. 478–490.
- [83] C. Slocum, Y. Zhang, N. Abu-Ghazaleh, and J. Chen, "Going through the motions: {AR/VR} keylogging from user head motions," in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 159–174.
- [84] T. Zhang, Z. Ye, A. T. Mahdad, M. M. R. R. Akanda, C. Shi, Y. Wang, N. Saxena, and Y. Chen, "Facereader: Unobtrusively mining vital signs and vital sign embedded sensitive info via ar/vr motion sensors," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, pp. 446–459.
- [85] Y. Cao, H. Chen, F. Li, and Y. Wang, "Crisp-bp: Continuous wrist ppg-based blood pressure measurement," in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, 2021, pp. 378–391.
- [86] Z. Shi, T. Gu, Y. Zhang, and X. Zhang, "mmbp: Contact-free millimetre-wave radar based approach to blood pressure measurement," in *Proceedings of the Twentieth ACM Conference on Embedded Networked Sensor Systems*, 2022, pp. 667–681.
- [87] A. M. Carek, J. Conant, A. Joshi, H. Kang, and O. T. Inan, "Seismowatch: wearable cuffless blood pressure monitoring using pulse transit time," *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies*, vol. 1, no. 3, pp. 1–16, 2017.
- [88] L. Wang, X. Wang, Y. Zhang, X. Ma, H. Dai, Y. Zhang, Z. Li, and T. Gu, "Accurate blood pressure measurement using smartphone's built-in accelerometer," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 8, no. 2, pp. 1–28, 2024.
- [89] C. Holz and E. J. Wang, "Glabella: Continuously sensing blood pressure behavior using an unobtrusive wearable device," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 3, pp. 1–23, 2017.

- [90] Y. Chen, Z. Yang, R. Abbou, P. Lopes, B. Y. Zhao, and H. Zheng, "User authentication via electrical muscle stimulation," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–15.
- [91] M. Funk, K. Marky, I. Mizutani, M. Kritzler, S. Mayer, and F. Michahelles, "Lookunlock: Using spatial-targets for user-authentication on hmds," in *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–6.
- [92] C. George, M. Khamis, E. von Zezschwitz, M. Burger, H. Schmidt, F. Alt, and H. Hussmann, "Seamless and secure vr: Adapting and evaluating established authentication systems for virtual reality." *NDSS*, 2017.
- [93] S. Luo, A. Nguyen, C. Song, F. Lin, W. Xu, and Z. Yan, "Oculock: Exploring human visual system for authentication in virtual reality head-mounted display," in *2020 Network and Distributed System Security Symposium (NDSS)*, 2020.
- [94] F. Mathis, H. I. Fawaz, and M. Khamis, "Knowledge-driven biometric authentication in virtual reality," in *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–10.
- [95] E. Dick, "Balancing user privacy and innovation in augmented and virtual reality," Information Technology and Innovation Foundation, Tech. Rep., 2021.
- [96] D. Maloney, S. Zamanifard, and G. Freeman, "Anonymity vs. familiarity: Self-disclosure and privacy in social virtual reality," in *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology*, 2020, pp. 1–9.
- [97] V. Nair, G. M. Garrido, D. Song, and J. O'Brien, "Exploring the privacy risks of adversarial vr game design," *Proceedings on Privacy Enhancing Technologies*, 2023.
- [98] P. Casey, I. Baggili, and A. Yarramreddy, "Immersive virtual reality attacks and the human joystick," *IEEE Transactions on Dependable and Secure Computing*, vol. 18, no. 2, pp. 550–562, 2019.
- [99] J. Penaz, "Photoelectric measurement of blood pressure, volume and flow in the finger," in *Digest of the 10th international conference on medical and biological engineering-Dresden, 1973*, vol. 104, 1973.
- [100] A. Gaurav, M. Maheedhar, V. N. Tiwari, and R. Narayanan, "Cuffless ppg based continuous blood pressure monitoring—a smartphone based approach," in *2016 38th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*. IEEE, 2016, pp. 607–610.
- [101] N. Bui, N. Pham, J. J. Barnitz, Z. Zou, P. Nguyen, H. Truong, T. Kim, N. Farrow, A. Nguyen, J. Xiao *et al.*, "ebp: A wearable system for frequent and comfortable blood pressure monitoring from user's ear," in *The 25th annual international conference on mobile computing and networking*, 2019, pp. 1–17.
- [102] S. S. Thomas, V. Nathan, C. Zong, E. Akinbola, A. L. P. Aroul, L. Philipose, K. Soundarapandian, X. Shi, and R. Jafari, "Biowatch—a wrist watch based signal acquisition system for physiological signals including blood pressure," in *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2014, pp. 2286–2289.
- [103] S.-H. Liu, D.-C. Cheng, and C.-H. Su, "A cuffless blood pressure measurement based on the impedance plethysmography technique," *Sensors*, vol. 17, no. 5, p. 1176, 2017.
- [104] Y. Xuan, C. Barry, J. De Souza, J. H. Wen, N. Antipa, A. A. Moore, and E. J. Wang, "Ultra-low-cost mechanical smartphone attachment for no-calibration blood pressure measurement," *Scientific reports*, vol. 13, no. 1, p. 8105, 2023.
- [105] N. Sugita, M. Yoshizawa, M. Abe, A. Tanaka, N. Homma, and T. Yambe, "Contactless technique for measuring blood-pressure variability from one region in video plethysmography," *Journal of Medical and Biological Engineering*, vol. 39, pp. 76–85, 2019.
- [106] H. Zhao, X. Gu, H. Hong, Y. Li, X. Zhu, and C. Li, "Non-contact beat-to-beat blood pressure measurement using continuous wave doppler radar," in *2018 IEEE/MTT-S International Microwave Symposium-IMS*. IEEE, 2018, pp. 1413–1415.
- [107] I. Kim and Y. A. Bhagat, "Towards development of a mobile rf doppler sensor for continuous heart rate variability and blood pressure monitoring," in *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2016, pp. 3390–3393.
- [108] E. J. Wang, J. Zhu, M. Jain, T.-J. Lee, E. Saba, L. Nachman, and S. N. Patel, "Seismo: Blood pressure monitoring using built-in smartphone accelerometer and camera," in *Proceedings of the 2018 CHI conference on human factors in computing Systems*, 2018, pp. 1–9.

## **Appendix A. Meta-Review**

The following meta-review was prepared by the program committee for the 2025 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

### **A.1. Summary**

This paper presents a practical attack called BPSniff, which can reconstruct fine-grained blood flow patterns and derive blood pressure (BP) based on unrestricted motion sensor data from users' VR headsets. The authors conduct extensive experiments to confirm the efficacy of BPSniff.

### **A.2. Scientific Contributions**

- Independent Confirmation of Important Results with Limited Prior Research
- Identifies an Impactful Vulnerability
- Provides a Valuable Step Forward in an Established Field

### **A.3. Reasons for Acceptance**

It is well-known that unrestricted motion sensors in virtual reality headsets can infer private user information. The authors demonstrate that these sensors can accurately estimate users' blood pressure, achieving results comparable to commercial blood pressure monitors. Notably, this inference does not require specialized equipment, user-specific training, or calibration. This critical privacy vulnerability was previously unrecognized.

### **A.4. Noteworthy Concerns**

- 1) The evaluation is limited to MetaQuest and MetaQuest 2, and the authors do not provide sufficient evidence for the system's transferability to other VR devices or to a broader population beyond the 37 users tested.