

# BMAR: Barometric and Motion-based Alignment and Refinement for Offline Signal Synchronization across Devices

MANUEL MEIER, Department of Computer Science, ETH Zürich, Switzerland

CHRISTIAN HOLZ, Department of Computer Science, ETH Zürich, Switzerland.

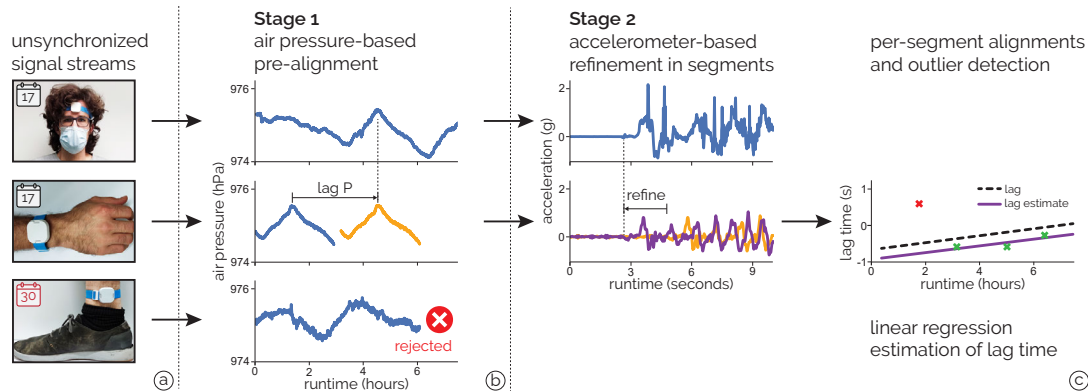


Fig. 1. *BMAR* is a novel method to synchronize and align signals across devices without the need for specific user input, action, or explicit synchronization through wired or wireless communication (e.g., WiFi or BLE). *BMAR* is capable of synchronizing (a) independently recorded signals *after the fact* by (b) first pre-aligning recordings using air pressure as an inexpensive sensing modality that simultaneously allows us to reject non-simultaneous recordings. (c) In a second step, *BMAR* produces a refined signal alignment across sensor devices by cross-correlating accelerometer observations.

A requirement of cross-modal signal processing is accurate signal alignment. Though simple on a single device, accurate signal synchronization becomes challenging as soon as multiple devices are involved, such as during activity monitoring, health tracking, or motion capture—particularly outside controlled scenarios where data collection must be standalone, low-power, and support long runtimes. In this paper, we present *BMAR*, a novel synchronization method that operates purely based on recorded signals and is thus suitable for offline processing. *BMAR* needs no wireless communication between devices during runtime and does not require any specific user input, action, or behavior. *BMAR* operates on the data from devices worn by the same person that record barometric pressure and acceleration—inexpensive, low-power, and thus commonly included sensors in today’s wearable devices. In its first stage, *BMAR* verifies that two recordings were acquired simultaneously and pre-aligns all data traces. In a second stage, *BMAR* refines the alignment using acceleration measurements while accounting for clock skew between devices. In our evaluation, three to five body-worn devices recorded signals from the wearer for up to ten hours during a series of activities. *BMAR* synchronized all signal recordings with a median error of 33.4 ms and reliably rejected non-overlapping signal traces. The worst-case activity was sleeping, where *BMAR*’s second stage could not exploit motion for refinement and, thus, aligned traces with a median error of 3.06 s.

Authors’ addresses: Manuel Meier, manuel.meier@inf.ethz.ch, Department of Computer Science, ETH Zürich, Switzerland; Christian Holz, christian.holz@inf.ethz.ch, Department of Computer Science, ETH Zürich, Switzerland.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, or post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2474-9567/2023/6-ART69 \$15.00

<https://doi.org/10.1145/3596268>

CCS Concepts: • **Hardware** → **Sensor devices and platforms; Sensor applications and deployments; Digital signal processing**; • **Computer systems organization** → *Embedded systems*.

Additional Key Words and Phrases: Synchronization, signal processing, wearable devices, embedded systems, monitoring, data analysis.

#### ACM Reference Format:

Manuel Meier and Christian Holz. 2023. BMAR: Barometric and Motion-based Alignment and Refinement for Offline Signal Synchronization across Devices. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 2, Article 69 (June 2023), 21 pages. <https://doi.org/10.1145/3596268>

## 1 INTRODUCTION

Technological advances have led to a fast increase in the miniaturization of technology, giving rise to many devices that can be worn throughout the day [8]. Among the main uses of wearable devices are activity tracking [1, 13, 15, 27, 31, 35, 55] and health monitoring [3, 9, 31, 38, 55, 62]. Both these types of devices incorporate multiple sensors that generate large amounts of time-dependent data that they must process in real-time, store, or transmit to another device. Oftentimes, the data is processed online to reduce storage or transmission requirements.

When raw recordings obtained on wearable devices are important, however, continuous data streaming places a large toll on battery power and storing all data for offline processing is preferable. This is especially true of data collection experiments that enable later research efforts [11]. Beyond the sensor data collected by one device, data collections commonly include multiple wearable devices, either worn by one individual or across individuals, such that the collected records can later be combined for joint analysis [42]. To analyze data series stored across multiple devices, accurate temporal data synchronization is an essential requirement. Wearable sensing devices usually associate measurements with timestamps, for example, derived from an on-board real-time clock (RTC). Since timekeeping typically relies on a crystal oscillator with an accuracy of 10–100 ppm [57], the clocks between two devices may deviate by several seconds after hours of runtime, even if they matched at one point. Repeated synchronization is therefore essential for accurate data processing at a later point.

A common approach to cross-device time synchronization is using a central reference for timekeeping. The Global Positioning System (GPS) is a popular choice because it synchronizes time at an accuracy of typically 10 ns [16] which exceeds that of an RTC and is only limited by hardware capabilities [34]. However, GPS signals are not reliable in locations with no clear view of the sky and the receiver's power consumption of several milliwatts renders them little suitable for repeated synchronization in low-power applications. Therefore, many wearables rely on a wireless sensor network (WSN) that implements an online synchronization algorithm. While there are no established standards in WSN synchronization, multiple commonly used wireless synchronization protocols exist, including the timing-sync protocol for sensor networks (TPSNs) [21], the flood time synchronization protocol (FTSP) [39], and reference broadcast synchronization (RBS) [20]. However, all these approaches are designed and optimized for accuracy and are not suitable for very low-power applications [37].

A recent trend in low-power systems has been in time synchronization that optimizes for power consumption and thus limits communication and the amount of online processing [11]. For example, Makara et al.'s synchronization protocol tolerates average time discrepancies of 10 ms, enabling two individual wearable devices to regularly exchange timestamps through Bluetooth Low Energy (BLE) at an average power consumption of just 133  $\mu W$  [37]. In contrast, many other efforts have opted to perform data synchronization *entirely offline* and based on measurement data alone, which is useful for longitudinal in-the-wild collections. Spilz and Munz and Gilbert et al.'s methods build on magnetometer recordings to detect a predefined event for clock synchronization [26, 52], leveraging electromagnets built into a tray on which the devices are placed. In many applications, such as wearable health sensing or collections of activity data, synchronization is particularly needed when devices are worn by the same person to obtain a coherent dataset for each wearer. For this use case, researchers have also

synchronized data series passively, detecting pre-defined events in the signals from inertial measurement units (IMUs) for synchronization, such as clapping [49] or footsteps [5].

Following from previous efforts in this domain, several key challenges exist for passively synchronizing data series offline: Synchronization must be able to detect whether or not two data series were recorded at the same time as a prerequisite. Synchronization cannot rely on pre-defined activities or events that are required to be performed by a person one or more times. Synchronization cannot assume unique events to be present inside all data series that passively afford accurate alignment. Synchronization must ensure that data series are aligned with a global optimum, particularly in the presence of events that manifest in periodic patterns, such as walking, running, or cycling, and that can lead to locally optimal but overall inaccurate alignment.

To address the challenges above, we propose *BMAR*, a novel offline synchronization method for devices worn by one person. *BMAR* requires no explicit input from users and accurately aligns cross-sensor signal traces without requiring users to perform particular motions or activities.

### 1.1 Offline Signal Synchronization across Body-worn Devices

As shown in [Figure 1](#), *BMAR* synchronizes signal traces individually recorded by wearable devices by aligning them in two stages: a pre-alignment based on air pressure observations and a refinement step that processes accelerometer observations.

In Stage 1, *BMAR* derives a pre-alignment of signal traces based on air-pressure observations. We simultaneously leverage these measurements for the rejection of non-concurrent recordings. Both are possible because air pressure recordings are unique over time but near-identical across all devices worn on the same body. Barometers are inexpensive, small, very low-power, and not affected by motion artifacts, ideally suited for this scenario on wearable devices.

In Stage 2, *BMAR* refines alignment based on accelerometer measurements. Building on our pre-alignment to ensure globally optimal alignment, *BMAR* produces a refined alignment by cross-correlating accelerometer recordings. We compute these for patches of data, which makes our method more robust to short-term misalignments and accounts for clock skew between devices. Accelerometers are ubiquitous in wearable devices thanks to the versatility of their recordings, their low cost, and low power consumption. The distinct high-frequency events in motion across all on-body locations that these sensors observe are ideal for optimizing *BMAR*'s pre-alignment to reduce the alignment error.

To evaluate the effectiveness of our method, we conducted 10 sessions of data collection on 7 separate days. In each session, between 3 and 5 wearable devices recorded data from the same participant for up to 10 hours, resulting in 328 hours of total device runtime. The participant performed one of 3 activities: office work, cycling, or skiing to include a variety of lower and higher-intensity motions. In addition, we included sleep as a 4th activity to verify how well our method performs on recordings with little to no motion.

Our results showed that *BMAR* synchronizes signals with a median error of less than 40 ms during sessions with activity. This is well below the period of human activity and physiological signals [14, 36], setting our method up as a suitable complement to the increasing data monitoring efforts on multi-device platforms.

### 1.2 Contributions

We make the following contributions in this paper:

- (1) *BMAR*, a novel method for signal synchronization that processes recordings offline through a pre-alignment and refinement stage, without relying on assumptions about signal content (e.g., specific actions or user input during recording).
- (2) Particular support for (wearable) low-power platforms, requiring no wireless connectivity or communication during operation, and leveraging only commonplace air pressure and accelerometer measurements.

- (3) An evaluation of our method across 3–5 custom-built wearable devices (with a high-resolution and a low-resolution sensor) across 4 activities against ground-truth measurements. During activities with human motion, our method aligned data traces with a median error of 33.4 ms, whereas this error was 3 s during activities without motion (sleep). Our method also reliably detected and rejected recordings that were not acquired concurrently.

## 2 RELATED WORK

BMAR is related to wireless synchronization methods, data-based synchronization, and wearable applications.

### 2.1 Wireless Synchronization

Over recent years, Wireless sensor networks (WSNs) have attracted more and more interest. Time synchronization performed online or offline between the nodes of a WSN is crucial for the deployment of the Internet of Things (IoT) [4]. Currently, many time-synchronization methods for WSNs use BLE [4, 44]. Especially in the fields of ubiquitous computing, sports science, and health applications, BLE is most commonly used since it is integrated in modern smartphones [44].

With BLE, timestamping errors of less than 50 ns are possible as demonstrated by [Asgarian and Najafi](#) with *BlueSync* [4]. *BlueSync* outperforms previous work by several magnitudes [7, 24, 25, 44, 51, 53]. While [Rheinländer and Wehn](#) achieve an accuracy of 0.9  $\mu\text{s}$  [44], [Sridhar et al.](#) and [Bideaux et al.](#) achieve accuracies of 10  $\mu\text{s}$  [53] and 14.9  $\mu\text{s}$  [7], respectively. [Somaratne et al.](#), [Ghoshdastider et al.](#) and [Ghoshdastider et al.](#) achieve accuracies of 19 ms [51], 37.8 ms [24], and 45.1 ms [25], respectively. While the above works do not have low power consumption as their main goal but achieving a low synchronization error, [Makara et al.](#) aim also to reduce power consumption and achieve a 10 ms error at a power consumption of 133  $\mu\text{W}$  when using two devices [37].

However, there are also alternatives to BLE, including the work by [Li et al.](#), [Seijo et al.](#), and [Ikram et al.](#) utilizing GPS [34], Wifi [46], and radio signals [28], respectively. [Li et al.](#) achieved synchronization with an accuracy well below 1 ms [34], while [Seijo et al.](#) show that their method achieves sub-nanosecond time transfer accuracy [46].

### 2.2 Data-based Synchronization

Multiple works have introduced data-based synchronization for distributed sensor systems using acceleration data from IMUs. Most of these works focus either on synchronization between IMUs and motion-capturing systems/ video data from cameras or between only IMUs. The latter is typically done by subjecting all IMUs to a sudden acceleration such as shaking multiple IMUs that are held together [22, 41, 45, 49, 59]. With this method, [Shabani et al.](#) obtained a minimum mean lag time bias of 25.56 ms [49]. Recently, [Spilz and Munz](#) introduced a novel data-based synchronization method for synchronizing multiple devices by subjecting them synchronously to a magnetic field impulse generated by an array of inductors causing an event that is detectable by the devices' IMUs [52] and allowing for synchronization with a maximum offset of below 2.6 ms. To synchronize IMUs and motion capturing systems/ video data sudden movements such as a heelstroke are commonly used. The event is then visible both in the video and the acceleration data [5, 10, 32, 43, 60]. For video-to-IMU synchronization without explicit user action, [Zhang et al.](#) proposed SyncWISE, an algorithm that achieves a mean error of less than half a second [61]. When aligning recordings based on data and without distinct events, a measure for signal synchrony is required. [Gashi et al.](#) evaluated several such measures when they explored the synchrony of signals of electrodermal activity between presenters and audience members to determine engagement [23].

### 2.3 Applications of Wearable Devices

Activity tracking and health monitoring are among the main use cases for wearable devices. [Leutheuser et al.](#) combine data from four accelerometers in different body locations for human activity recognition [33]. At the




Fig. 2. Example of aligning a short signal (top, orange) with a longer one (top, gray) based on cross-correlation which results in an objective function (bottom, orange) that is at its maximum where the longer signal is maximal. Similarly, the cross-covariance results in an alignment with the section of the longer signal that is the steepest (red). For 'slow' signals such as ambient pressure, alignment based on sample-by-sample signal delta is better suited (green).

intersection of health monitoring and activity recognition, [Strain et al.](#) perform mortality prediction based on data collected from a wrist-worn wearable [55]. Similarly, [Bari et al.](#) combined synchronized data from separate devices recording audio, physiological, and inertial signals to detect stressful conversations [6].

### 3 PROBLEM DEFINITION AND BACKGROUND ON SIGNAL ALIGNMENT

We now discuss preliminary considerations for our proposed method for synchronizing data series as well as the challenges involved.

#### 3.1 Temporal Alignment of Discrete Finite Signals

The problem we are addressing is specified as follows: Given are discrete sensor recordings from two independent devices  $D_1$  and  $D_2$  which were acquired with temporal overlap but may be of varying duration. A recording always comprises traces of barometer, accelerometer, and optionally other sensors. The goal is to temporally align the two recordings based on the similarities in measurements contained in the sensor traces. We utilize a mixture of methods based on cross-correlation, cross-covariance, and signal delta.

**3.1.1 Cross-correlation and Cross-covariance to Align Discrete Finite Signals.** The approximation of the cross-correlation  $R_{fg}$  of two signals  $f$  and  $g$  of finite lengths  $N_f$  and  $N_g$  can be used to determine the best alignment of the two functions by finding the lag  $m_0$  of  $g$  that maximizes  $R_{fg}$ :

$$R_{fg}[m] = \frac{1}{l[m]} \sum_{n=\max(0,-m)}^{\min(N_f-1, N_g-m-1)} f[n]g[n+m].$$

Here, the result is scaled relative to the length  $l[m]$  of the overlap of the two signals:

$$l[m] = \min(N_f, N_g - m) - \max(0, -m).$$

The quality of this approximation decreases at the edges with shorter window lengths  $l$  when  $m$  is close to  $N_g$  or  $-N_f$ . The cross-correlation is commonly used to align data series such as gyroscope measurements [49] or accelerometer measurements [17], but it does not work well for all signal types. Most notably, signals with a characteristic timescale that is longer than the duration of the computation window (i.e., contain few distinct swings within the computation window) are prone to misalignment [58]. An example of misalignment in such a scenario is shown in [Figure 2](#).

In some cases, this issue can be resolved by subtracting the mean values  $\mu_{f'}$  and  $\mu_{g'}$  from the respective signals prior to multiplication. The resulting formula is an approximation of the cross-covariance of the two signals:

$$R_{fg}[m] = \frac{1}{l[m]} \sum_{n=\max(0,-m)}^{\min(N_f-1, N_g-m-1)} (f[n] - \mu_f)(g[n+m] - \mu_g)$$

However, there are still scenarios in the realm of slow sensing modalities that are misaligned when using this formula for trace alignment as shown in [Figure 2](#).

**3.1.2 Signal Delta to Align Discrete Finite Signals.** To align sensor signals with a long characteristic timescale in scenarios where the sensors have recorded the same underlying signal at the same time bar the impact of sensing tolerance (as is the case for air pressure measurements), the point-by-point difference of the signals can be analyzed. Considering the same two signals  $f$  and  $g$  of finite lengths  $N_f$  and  $N_g$  and the overlapping length  $l[m]$  as before, a loss function, say  $L$ , can be formulated as follows:

$$L[m] = \frac{1}{l[m]} \sum_{n=\max(0,-m)}^{\min(N_f-1, N_g-m-1)} L_\delta(f[n] - g[n+m])$$

Where  $L_\delta$  is the Huber loss function which reduces the sensitivity to measured outliers:

$$L_\delta(x) = \begin{cases} \frac{x^2}{2}, & |x| \leq \delta \\ \delta(|x| - \frac{\delta}{2}), & |x| > \delta \end{cases}$$

The resulting loss function  $L[m]$  can be minimized to find the lag  $m_0$  that aligns the two sensor traces the best. Alternatively, the standard deviation of the point-by-point difference of the two signals can be used as a loss function. This makes it invariant to the mean of the signal difference and effectively yields a measure of how well two signals run in parallel:

$$L[m] = \sqrt{\frac{1}{l[m]-1} \sum_{n=\max(0,-m)}^{\min(N_f-1, N_g-m-1)} |f[n] - g[n+m] - \mu|^2}$$

where  $\mu$  is the mean of  $f[n] - g[n+m]$ :

$$\mu = \frac{1}{l[m]} \sum_{n=\max(0,-m)}^{\min(N_f-1, N_g-m-1)} f[n] - g[n+m]$$

### 3.2 Clock Offset, Drift, and Skew

The difference between two RTCs  $\delta(t)$  can be characterized in three components [57]: (1) offset, the constant component of  $\delta(t)$  caused by different starting points, (2) skew, the first derivative  $\delta'(t)$  which indicates that one clock is running faster than the other resulting in a gradual de-synchronization, and (3) drift, the second derivative  $\delta''(t)$  which is the variation in skew. This terminology is consistent with previous definitions as used by Mills [40] and Sundararaman et al. [56]. For RTCs that rely on quartz crystals, clock drift is mainly caused by the temperature dependence of the latter which is typically characterized as  $Err(T) = 0.036ppm \cdot (T - 25^\circ C)$  [57]. This results in only minor clock drift for temperature differences commonly experienced by two wearable devices. Therefore, we assume clock drift to be 0 and de-synchronization effects to be linear between devices in our considerations, similar to related work [50].

### 3.3 Measurement Errors of Air Pressure Sensors

Air pressure measurements are subject to sensor tolerances. The resulting error can be characterized by an absolute accuracy (i.e., the absolute difference to ground truth) and a relative accuracy (i.e., the accuracy of the delta between two measurements by the same sensor). For commercially available barometric pressure sensors, the absolute accuracy is usually about an order of magnitude larger than the relative accuracy [29, 47, 48]. Therefore, BMAR uses absolute measurements to reject non-simultaneous measurements only and relies on relative measurements for trace alignment.

### 3.4 Ensuring within-device Sensor Synchronization

In offline settings, the RTC of a device does not provide absolute time information which could be used for cross-device synchronization but instead keeps track of the runtime since starting the device more accurately than internal sensor clocks would. With typical tolerances in the order of several percent, sensor clocks are not suitable to be used for timekeeping. Hence, a prerequisite for the alignment of multiple sensor traces across devices is the synchronization of traces with the RTC within each device.

*3.4.1 Alignment Offsets Due to Sensor Startup Latency.* Starting each sensor of one device at a known RTC timestamp does not guarantee synchronization. This is due to differences between the latency of each sensor for activation and starting the acquisition of its first sample. Because this latency is not commonly specified in the datasheet, it can usually not be compensated for without extensive calibration measurements. However, assuming that this latency is constant for a specific sensor type and only data from the same sensor type is compared across devices, the latencies cancel out and have no negative impact on the resulting synchronization results.

*3.4.2 Synchronized Sampling.* When sensors support external triggering, the RTC can be used to directly initiate measurements. This keeps sensor data traces perfectly synchronized with the RTC over arbitrarily long periods of runtime. In systems where sensors do not support external triggers from the RTC, sensors must rely on their internal clocks. In this case, the relative error in the sampling rate requires compensation.

## 4 BMAR: METHOD DESIGN FOR OFFLINE SIGNAL SYNCHRONIZATION

Optimal data synchronization across multiple wearable devices worn by a single person requires patterns in the data that are unique over time and do not depend on the body location where the device is worn. The presence of such patterns depends on user behavior and sensing modality.

Therefore, BMAR builds on barometric pressure sensors, which observe changes in the environment that affect all worn devices to a similar extent. This causes the individual on-body location of a sensor to manifest in a negligible amount and does not require our method to rely on specific user actions or activities. Barometric pressure sensors are inexpensive, low-power, and are commonly built into wearable devices. The signal they measure is influenced by changes in the altitude above sea level (e.g., walking up or down stairs, taking elevators, or performing any other motion not in a flat plane), meteorological changes in atmospheric pressure, or changes in ambient pressure caused by the immediate environment (e.g., opening a door of a closed room or car).

Our method aligns the data recorded by any two independent devices  $D_1$  and  $D_2$  or rejects the alignment in case it determines that the devices were not running at the same time or worn by the same person. The goal of this process is to determine the lag time  $t_{lag}(t)$  of  $D_2$  compared to  $D_1$ .  $t_{lag}(t)$  itself is time-dependent due to RTC tolerances in  $D_1$  and  $D_2$  but is assumed to be linear.

### 4.1 Rejection of Non-simultaneous Measurements

The nature of air pressure measurements allows the rejection of two data traces that were not recorded at the same time. The difference of two temporally independent, single-sample air pressure measurements collected




Fig. 3. Approximations of the same  $t_{lag}(t)$  with different  $n_{win}$  for two pairs of recordings. (a) Fewer windows lead to fewer outliers and a more accurate approximation for these recordings. (b) More windows lead to a reduction in alignment error for these recordings.

in the same location are expected to be normally distributed  $X_1 - X_2 \sim \mathcal{N}(0, 2\sigma^2)$  with  $\sigma^2$  being the variance of independent air pressure samples for the specific location. The probability of correct rejection of two such measurements can be calculated as follows:

$$p_{rej} = \frac{1}{2\sigma\sqrt{\pi}} \int_{-2A}^{2A} e^{-\frac{x^2}{4\sigma^2}}$$

Where  $\pm A$  is the range of the absolute accuracy of each barometric pressure sensor. When inserting real-world values such as  $\sigma = 928 Pa$  [12] and  $A = 100 Pa$  [47], the probability of correct rejection is 0.83. Using multiple measurements over time and incorporating a distribution over altitude above sea level, this probability is expected to increase. Here, the sensor accuracy is assumed to have strict lower and upper bounds and thus the rate of false rejection is zero when setting the acceptance tolerance to the same boundaries. Modeling the sensor accuracy as a normal distribution leads to the expected trade-off of minimizing false positives versus false negatives, however, specifications in datasheets usually do not provide enough information for such calculations.

To determine whether two devices  $D_1$  and  $D_2$  were running at the same time, we consider the alignment of the air pressure measurement data from both devices that minimizes the mean sample-by-sample difference and classify them as non-simultaneously recorded if this value surpasses a threshold.

#### 4.2 Stage 1: Air Pressure-based Pre-alignment

BMAR first analyzes air pressure measurement data series to determine whether  $D_1$  and  $D_2$  were running at the same time and to find an approximation  $t_{lag}^P$  of  $t_{lag}(t)$ .

Since the error of  $t_{lag}^P$  is generally larger than the RTC clock skew effects between  $D_1$  and  $D_2$ , it is computed as a constant value independent of time. We developed two algorithms to compute air pressure-based trace alignment. The *delta-error* method computes the mean Huber loss based on the point-by-point difference between two traces for all possible alignments (Section 3.1.2). The lag time corresponding to the alignment with the minimum mean Huber loss is selected as  $t_{lag}^P$ . For this method, the data from both devices is low-pass filtered to remove noise and normalized to make the alignment invariant to constant offsets due to absolute sensor accuracy tolerance.

The *delta-std* method computes the standard deviation of the point-by-point difference between two traces for all possible alignments and corresponding lag times. The lag time resulting in the minimum standard deviation is selected as  $t_{lag}^P$ . The raw data is used for this method.






Fig. 4. Top and bottom sides of the electronics of our custom wearable device, which measures 23 mm × 22 mm. Components: (a) ADXL355, high-resolution accelerometer, (b) LIS2DH, low-cost accelerometer, (c) power switch, (d) extension connector, used for synchronization ground truth, (e) CR2032 battery, (f) RTC quartz crystal, (g) DA14695 System-on-a-Chip, (h) TH58CYG3S0HRAIJ 8 Gb flash memory, (i) BME280 barometric pressure sensor.

### 4.3 Stage 2: Accelerometer-based Refinement and Optimization

Next, BMAR leverages accelerometer signals to refine  $t_{lag}^P$  and optimize the alignment. We maximize the cross-correlation of the magnitudes to find the best alignment  $t_{lag}^{Accel}(t)$ . To enable the time-dependent approximation of  $t_{lag}(t)$  and to make the algorithm more robust to time-limited measurements that facilitate misalignment, the available data trace is split into  $n_{win}$  windows for each of which a lag time is calculated separately. The result of this is a series of  $n_{win}$  calculated lag times temporally associated with the center point of the respective window that the original trace was split into. We perform ordinary least squares regression (OLS) on these lag times and remove outliers using the Bonferroni outlier test. This results in a linear approximation  $t_{lag}^{Accel}(t)$ .

The optimal  $n_{win}$  depends on the data. Therefore, we repeat the computation, starting with  $n_{win} = 4$  and increase the number of windows by a factor of 2 until the window size drops below 1 second or  $n_{win} > 512$ . The result of this is a set of linear approximations of  $t_{lag}(t)$ . We use the  $r^2$  adjusted for residual degrees of freedom of the OLS to select the most promising approximation. This is based on the assumption that local misalignments within some of the windows do not appear linear over the whole trace, making  $r^2$  an indicator of alignment quality. Figure 3a shows an example where the linear approximation with a smaller  $n_{win}$  is more accurate while Figure 3b shows an example of an alignment where a larger  $n_{win}$  leads to a better result.

## 5 SYSTEM IMPLEMENTATION OF BMAR

To implement and evaluate BMAR, we devised a miniaturized wearable platform that embeds multiple sensors and runs on a small battery. As shown in Figure 4, our platform is coin-sized and affords wear at multiple locations of the user's body.

### 5.1 Device

The low-power sensor device we developed comprises a 23 mm × 22 mm printed circuit board (PCB), which is powered by a CR2032 coin-cell battery (Figure 4e). It centers around a System-on-a-Chip (SoC) (DA14695, Renesas [19], Figure 4g) and logs sensor data to an 8 Gb serial NAND flash memory (TH58CYG3S0HRAIJ, Kioxia America [2], Figure 4h). The data can later be downloaded to a USB port on a PC through a serial port connection on the extension connector of the PCB (Figure 4d). The RTC of the SoC relies on a 32.768kHz external quartz crystal (ECS-.327-7-16-C-TR, ECS Inc. International [30], Figure 4f) with a tolerance of ±20 ppm. Our platform accommodates a series of sensors that are common on wearable devices for the purpose of evaluating our method. Depending on the activated sensors and their sampling frequencies, the device consumes between 2 and 4.5 mW

resulting in a battery life of 35–80 hours. The device runs autonomously and does not require user interaction or wireless communication for operation. It can be turned on and off by a power switch on the board (Figure 4c).

**5.1.1 Sensors.** For later pre-aligning of all data series, the device houses a digital barometric pressure sensor (BME280, Bosch Sensortec, Figure 4i). The sensor is operated in forced mode to enable synchronization to the RTC of the SoC, which triggers a measurement every 100 ms, leading to a sampling rate of the barometer of 10 Hz. We used no oversampling or internal IIR filter. As per the datasheet, the absolute accuracy of the measurements lies within a range of  $\pm 1.0$  hPa ( $\pm 3\sigma$ ) and the relative accuracy is  $\pm 0.12$  hPa ( $\pm 3\sigma$ ) [47] which are low-end characteristics compared to pressure sensors currently available commercially.

For refining synchronization during our optimization stage, we included two accelerometers for comparison. First is an ADXL355 (Analog Devices, Figure 4a), a high-resolution 3-axis digital accelerometer. It operates in FIFO mode, receiving triggers for measurements from the RTC of the SoC. This results in a sampling frequency of 128 Hz. We set the range to  $\pm 2$  g at a resolution of 20 bit ( $4 \mu\text{g}/\text{LSB}$ ) [18]. Second is a LIS2DH (STMicroelectronics, Figure 4b), a low-cost 3-axis digital accelerometer that is found in lower-end wearables such as fitness bands. It also operates in FIFO mode, but does not support external triggers. The accelerometer was configured to sample at 200 Hz using the internal clock source, at a range of  $\pm 2$  g and a resolution of 8 bit ( $4 \text{ mg}/\text{LSB}$ ) [54].

## 5.2 Encasing

We created a 3D-printed case that comprises an insert to house all the electronics and the battery as well as an outer casing with a strap for wear as shown in Figure 5b. The case supports different straps to support wear at different locations on the body.

## 5.3 Synchronization

The firmware of the SoC was implemented in C++, building on freeRTOS to drive all sensors and repeatedly download and store the contents of their FIFO buffers. We implemented all synchronization algorithms in Python 3.11 to process all recordings offline after download.

**5.3.1 Buffered Sampling With Compensation.** Since the low-cost accelerometer does not support external measurement triggering from the RTC but relies on its internal clock source, its measurements require compensation to preserve the synchronization of all sensors within the device. When the SoC retrieves buffers from a sensor, it additionally logs an RTC-based timestamp to memory. During offline processing, the expected total data length after each FIFO readout can be calculated based on the logged RTC timestamps of the readout and the sensor activation. To account for the latency between sensor activation and data availability in the FIFO, we calculate the expected length relative to the first FIFO readout:

$$l_{exp}(t_n) = (t_n - t_0) \cdot f_{sensor} + l(t_0)$$

With the sampling frequency of the sensor  $f_{sensor}$ , the expected length  $l_{exp}$  at the RTC timestamp  $t_n$  of the  $n^{th}$  FIFO readout after the initial readout at RTC timestamp  $t_0$  after which a total of  $l(t_0)$  samples were present. Because FIFO readouts are usually not performed on a perfectly regular basis, a lack of samples after a FIFO read-out at  $t_n$  (i.e.,  $l(t_n) < l_{exp}(t_n)$ ) may automatically be compensated for by a sample which narrowly slipped into the next FIFO readout. We account for this by only adding or removing samples to match the expected length when no inverse operation is necessary in close temporal proximity. This results in a sensor data trace which is synchronized to the device's RTC with a tolerance of  $1/f_{sensor}$ .

**5.3.2 Hyperparameters.** For the rejection of non-simultaneous traces (Section 4.1), we set the threshold of the mean delta between two traces to the absolute tolerance of the used barometric pressure sensor of  $100 \text{ Pa}$  [47], trying to achieve a recall close to 1. Lowering it increases precision and lowers recall. We use the same value for

Table 1. For our evaluation, the experimenter recorded data during these activities, wearing multiple devices at the respective locations on the body for the given durations. Each activity was repeated two or more times.

activity	#	locations	runtime	rep
sleep	5	head, chest, ankle, wrist	9.9 h	4
work	3	head, chest, ankle	3.8 h	3
cycle	4-5	head, chest, ankle, wrist	1.4 h	2
ski	3	chest, wrist	8.8 h	2

the  $\delta$  of the Huber loss in the delta-error method (Section 4.2) such that it only affects rarely occurring outliers in the measurements. We set the maximum range of the accelerometer-based refinement (Section 4.3) to  $\pm 5$  seconds. If the pre-alignment differs more than that, the refinement cannot fully compensate for it. On the other hand, increasing this range increases the risk of misalignment caused by this stage.

## 6 EVALUATION

To quantify the performance of BMAR, we conducted a data recording experiment during everyday activities, using between 2 and 5 devices at once with several repetitions. Table 1 shows an overview of the configurations in our evaluation. Since there is no impact of individual users' behavior on our method, all recordings were conducted by one experimenter.

To analyze the performance of our method, we consider pairs of devices that either did or did not record data at the same time while worn by the experimenter. This means that if 3 devices were worn at the same time, we considered them as 3 pairs of simultaneously running devices in our evaluation ( $D_1$  and  $D_2$ ,  $D_1$  and  $D_3$ ,  $D_2$  and  $D_3$ ) and also included pairings of each device with a random device from another recording which was not acquired at the same time.

### 6.1 Procedure

**6.1.1 Locations.** As shown in Figure 5c, devices were attached to the following locations for a recording session: strapped onto the forehead, strapped onto the sternum, on one of the wrists like a watch, or above one of the lateral malleoli.




Fig. 5. Apparatus implementation for the evaluation. (a) Board to connect multiple devices and trigger ground-truth synchronization events, (b) 3D printed case for the custom wearable sensing devices, here shown with a wrist strap, (c) the 6 sensor locations used in the evaluation of BMAR.

**6.1.2 Activities.** The experiment included two categories of sessions: active wearer and passive wearer. For active sessions, the experimenter performed activities such as cycling, skiing, or office work over the full duration. For this session type, we collected 31 pairs of device recordings running at the same time, with an average runtime of 3.73 hours.

For passive sessions, the experimenter remained stationary over the full duration. We recorded these data during sleep, turning on the devices after getting into bed and turning them off before exiting. For this session type, we obtained 40 pairs of devices with simultaneous recordings and an average runtime of 9.9 hours. Due to the lack of activities and motion, this session type can be considered a worst-case baseline for our analysis.

## 6.2 Ground-truth Synchronization

We devised an apparatus to establish repeated ground-truth synchronization across all devices. The experimenter manually attached this apparatus to all devices at the beginning and the end of each session. This resulted in two ground-truth synchronization events in our logs for reference. The apparatus is shown in Figure 5a. It sends a trigger from a manual, debounced push-button to a digital pin of the SoC in each device. The tethered connection ensured accurate reference values without any impact of wireless communication.

To ensure that each device could register the precise moment of a button press, we used a timer block on the SoC to support more finely-grained timing resolution between the RTC ticks. The incoming synchronization signal was logged through the timers event capture functionality, which logs the exact timestamp and triggers an interrupt to process the event. In theory, the accuracy of this mechanism is limited by our timer clock of 32.768 kHz, which corresponds to a tolerance of 0.03 ms. However, the tolerance of the timer event trigger capture is not specified by the SoC manufacturer. Hardware limitations (e.g., stray capacitances and tolerances in logic voltage level detection) may also increase this error.

**6.2.1 Accuracy of Ground-truth Synchronization.** To evaluate the accuracy in practice, we sent multiple synchronization signals to all connected devices in short succession within a few seconds. Assuming no clock skew between the first and last synchronization signal, we compared the alignment of the multiple detected events on the time axis. Across 349 such synchronization sequences with 3 signal triggers each, we found a mean error of  $ME = 1.2 \text{ ms}$  ( $\sigma = 1.1 \text{ ms}$ ) for a single signal trigger. To improve accuracy in the evaluation, we always used the mean across 3 consecutive signal triggers.

## 6.3 Duration of Recordings

For temporal alignment evaluations, the duration of the recording has an impact on the results (i.e., the probability of correct temporal alignment increases with longer recordings). We conducted the evaluations considering various different lengths of recordings by cropping recordings to the respective length. If a recording was multiple times the desired length, we split it and used up to 5 non-overlapping splits separately. The number of traces




Fig. 6. Median alignment error of BMAR's synchronization. The band denotes the range between 10th and 90th percentile.

Table 2. Median alignment error in seconds of the complete synchronization method. The number of pairs of recordings used is stated in brackets.

	length of recording					
	5 min	10 min	15 min	30 min	1 h	≥ 6 h
active	0.442 (155)	0.074 (155)	0.044 (155)	0.041 (107)	0.028 (73)	0.034 (31)
passive	≫60 (200)	≫60 (200)	≫60 (200)	20.2 (200)	5.42 (200)	3.06 (40)

for different lengths is shown in Table 2. When aligning traces, we only cropped or split one of the two traces and preserved the length of the other trace. We considered the length of the shorter trace to be the minimum overlapping time of the two traces to prevent negative effects on the edges. Effectively, this means we measured the accuracy of aligning a trace of a specific length to a longer one that overlaps it.

## 6.4 Results

We first report the results of our complete method spanning both air pressure-based pre-alignment and accelerometer-based refinement. We then report the results of the rejection of non-simultaneously recorded data, of pre-alignment, and of refinement separately.

**6.4.1 Overall Synchronization Accuracy.** As shown in Figure 6 and Table 2, the synchronization of active sessions achieved a median alignment error of 33.4 ms for recording durations of 15 min and longer, 73.6 ms for durations of 10 min, and 443 ms for durations of 5 min. The synchronization of passive sessions (i.e., a worst-case baseline without user motion or activity) resulted in a median alignment error of 3.06 s for recording durations of 6 h and longer, 5.42 s for durations of 2 h, and 20.2 s for durations of 30 min. Figure 7 shows the percentage of synchronizations that our method aligned within a given error tolerance.

We report the errors of unconstrained synchronizations with median and percentiles because the error of alignment for outliers (i.e., failed, effectively random alignments) can be up to four orders of magnitude larger than the median error of all alignments. The impact on the mean error and standard deviation makes them less informative for evaluation.

**6.4.2 Accuracy of Stage 1: Pressure-based Pre-alignment.** In stage 1, air pressure measurements were used to detect non-simultaneous measurements and do a first trace alignment of simultaneously acquired traces.

*Elimination of non-simultaneous measurements.* When using a sensor tolerance of 100 Pa as the threshold for the maximum mean difference between two traces to detect simultaneously recorded air pressure data, we reach a precision and recall of 1.0 for recordings of active sessions that contain at least 35 minutes of data and




Fig. 7. Percentage of traces that BMAR correctly aligned with a given error tolerance.




Fig. 8. Precision and recall rates for BMAR's detection of simultaneously acquired recordings.

recordings of inactive sessions of at least 8 hours of recording time. Figure 8 shows the precision and recall rates for recordings of different lengths.

These results are based on tests with the barometric pressure sensor recordings of each device  $D_n$  both paired with simultaneously and randomly selected, non-simultaneously recorded data. There is a trade-off between precision and recall: Lowering the threshold improves precision but lowers recall. However, such a lowered recall can generally not be compensated for by increasing the duration of the recording.

*Trace alignment based on air pressure.* Aligning active sessions produced a median alignment error of 1.14 s for recording durations of 10 min and longer, 1.53 s for durations of 5 min, and 3.8 s for durations of 3 min, as shown in Figure 9. The synchronization of passive sessions (i.e., our worst-case baseline) resulted in a median alignment error of 3.36 s for recording durations of 3 h and longer, 12.4 s for durations of 1 h, and 25.3 s for durations of 30 min.

Of the two tested alignment methods, the *delta-std* method performed better than the *delta-error* method. Across all recording lengths greater than 0.5 hours (i.e., after initial convergence), the median alignment error of the *delta-std* is on average 48 ms lower for active sessions and 237 ms lower for passive sessions.

**6.4.3 Accuracy of Stage 2: Accelerometer-based Refinement.** Testing the accelerometer-based synchronization refinement in isolation for active sessions and under the assumption of a pre-alignment tolerance of  $\pm 1$  s, produced a median alignment error of 32.0 ms for recording durations of 15 min and longer, 40.6 ms for durations of 5 min, and 218 ms for durations of 1 min, as shown in Figure 10a. The refinement of passive sessions (i.e., worst-case

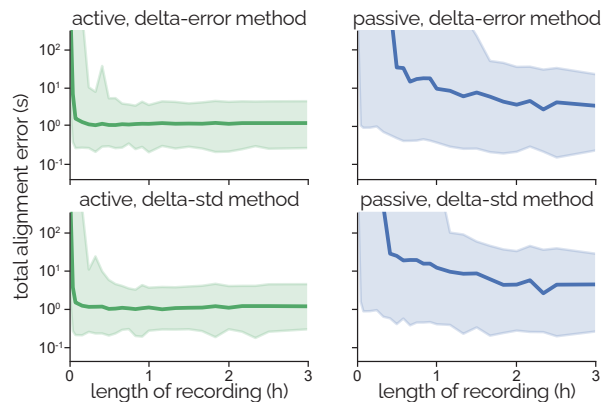


Fig. 9. Median alignment error of BMAR's first stage (air pressure-based pre-alignment). The band denotes the range between the 10th and 90th percentile.

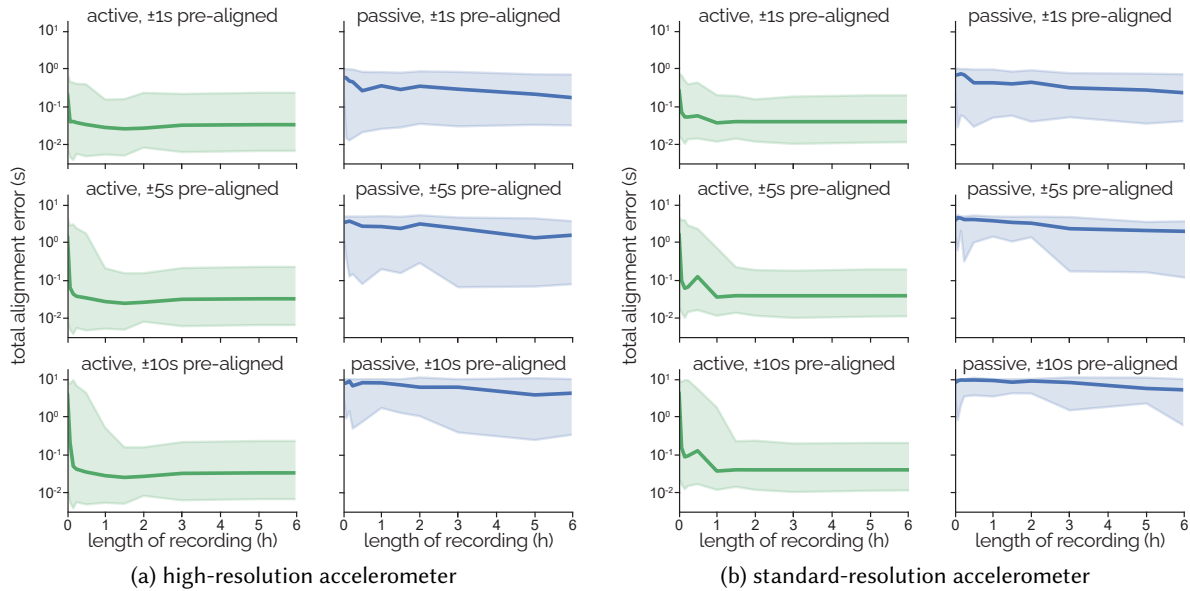


Fig. 10. Median alignment error after BMAR’s second stage (accelerometer-based refinement) for different pre-alignment errors.

baseline) resulted in a median alignment error of 274 ms for recording durations of 15 min and longer, 580 ms for durations of 5 min, and 606 ms for durations of 1 min. As shown in Figure 10a, the median error increases with higher pre-alignment tolerances.

We tested the refinement with data from both a high-resolution accelerometer (ADXL355, Analog Devices) and a standard-resolution accelerometer (LIS2DH, STMicroelectronics). The high-resolution accelerometer produced better results, yielding an average alignment error that was 21.5 ms lower when analyzing all recording durations, which resulted in a median alignment error below 0.1 s.

As shown in Figure 10b, the results we obtained with the standard-resolution accelerometer for active sessions and under the assumption of a pre-alignment tolerance of  $\pm 1$  s, produced a median alignment error of 44.2 ms for recording durations of 15 min and longer, 71.9 ms for durations of 5 min, and 276 ms for durations of 1 min.

The refinement of passive sessions (i.e., worst-case baseline) based on the standard-resolution accelerometer resulted in a median alignment error of 345 ms for recording durations of 15 min and longer, 685 ms for durations of 5 min, and 703 ms for durations of 1 min.

## 7 DISCUSSION, LIMITATIONS, AND FUTURE WORK

Our evaluation confirmed the robustness of our method to detect recordings that were acquired simultaneously to synchronize data stemming from multiple wearable devices. The accuracy our method achieved for recordings during active sessions is high enough to correctly match each occurrence of fast repetitive events across devices. This includes the fastest vital sign (i.e., heartbeats of a child at 3.5 Hz) and fast body motions (i.e., the cadence of a professional sprinter at 4 Hz or the catch rate of the world’s fastest juggler at 9 Hz). Faster motions by humans usually occur with longer intervals in between, making cross-device analysis easier. Unsurprisingly, the results were worse for our baseline recordings, which did not include any activity other than sleeping. However, given

long recordings of multiple hours, non-simultaneous recordings are still rejected reliably and the synchronization tolerance reaches 3.06 seconds which may be sufficient for the analysis of changes in sleep position.

### 7.1 The Relevance of Pre-alignment Using Air Pressure Measurements

Due to differences in the motion characteristics across body parts, IMU data does not lend itself well to aligning data traces without additional constraints. Using cross-correlation, for example, the resulting objective function exhibits many maxima and trace alignment without restrictive constraints on the search window or specifically performed motions by the wearer is challenging. For a more general, rough alignment, the use of a sensing modality with a longer characteristic timescale is more effective to produce a cost function that monotonically converges to a global minimum.

Temperature, perhaps the most commonly logged sensing modality, does include features common across devices worn in different body locations. Events like moving from indoors to outdoors with an abrupt change in ambient temperature could be used for synchronization, but they may be attenuated or delayed due to varying thermal insulation of the devices by clothing.

Air pressure is thus a suitable modality for cross-device synchronization: Air pressure measurements are barely affected by anything other than the geographical location, time of recording, and sensor tolerance. The former two are identical for multiple devices worn simultaneously.

### 7.2 Impact of Sensor Accuracy

As mentioned in Section 6.4.3, our standard-resolution accelerometer performed similarly well as the high-resolution accelerometer during refinement, indicating that our method does not require an expensive accelerometer for accurate signal synchronization. The range of commercially available barometric pressure sensors is smaller compared to accelerometers and the sensor used for the evaluation can be considered to be at the low end of the spectrum in terms of accuracy: In the relevant temperature and measurement range, the BME280 we evaluated BMAR with has an absolute and relative accuracy of  $\pm 100Pa$  and  $\pm 12Pa$ , respectively [47]. This compares to  $\pm 30Pa$  and  $\pm 6Pa$  for the recently released BMP581 (Bosch Sensortec) [48] as well as  $\pm 50Pa$  and  $\pm 3Pa$  for the inexpensive SPL07-003 (Goertek Microelectronics Inc.) [29]. This suggests that our results likely generalize to most available pressure sensors, including inexpensive commodity barometric pressure sensors.

We expect a better absolute accuracy to improve the rejection of non-simultaneous recordings and a better relative accuracy to improve the pre-alignment of the recordings. In the future, we plan to characterize the specific impact of sensor characteristics (accuracy, sampling rate, resolution, noise, and RTC synchronization) on the synchronization error.

### 7.3 Needed Overlap of Cross-sensor Signal Recordings

In our current design, BMAR aligns signals of a specified duration  $l_1$  to a recording of length  $l_2$ , which fully contains the former (Section 6.3). However, our method also works in scenarios where signals only *partially* overlap. In these, the effectively smaller overlap decreases the amount of data for calculating synchronization cost and may lead to a higher risk of misalignment. Considering a minimum overlap of  $l' \leq l_1$ , the expected error is in the range of our evaluation results for total recording lengths between  $l_1$  and  $l'$ .

### 7.4 Limitations of Accelerometer-based Refinements

The improvement in refining synchronization based on accelerometer signals depends on the signal characteristics. For our method, we cannot give worst-case guarantees; if recordings contain no events or asynchronous patterns, our refinement may worsen the result of the pre-alignment.



To prevent potentially negative effects of our refinement, BMAR limits the range within which pre-alignment sequences may be refined as part of our synchronization method. Since the loss of the accelerometer-based alignment does not generally decrease monotonically to a correct global minimum, the refinement may worsen the result of a pre-alignment that was successful but with an error larger than the range of the refinement. 10a shows how our refinement for short recordings during inactive sessions is close to a random offset within alignment tolerance and only decreases for longer recordings. Future work could improve our approach by leveraging the pre-alignment confidence and suitability of the accelerometer data to dynamically determine the range of the refinement.

### 7.5 Devices across Multiple Users and Locations

BMAR is currently designed to synchronize signals captured on devices worn by a single person. However, our method of aligning signal traces based on air pressure recordings could extend to synchronizing signal streams captured across multiple people and beyond the realm of wearable devices. We expect the method to produce accurate results if all devices are in the same location or move collectively (i.e., in a car, train, or elevator), but we anticipate a negative impact of movements that are spread out between devices, such as when groups of people who wear them do not transition between building floors at the same time. Depending on the setting, the use of different sensing modalities may be appropriate such as sound or temperature. We tested the latter early on in our project but found that for wearable devices, the impact of clothing and body temperature makes cross-device synchronization challenging. It may, however, be a viable option when synchronizing non-wearable devices.

### 7.6 Limitations of BMAR's Evaluation: Single Participant vs. Breadth of Activities for More Insight

We acknowledge that in our evaluation, multiple devices collected data on a single participant only. While such an evaluation may be uncommon, it is important to point out that BMAR is not specific to human factors or capabilities; rather, the performance of the method depends on the type of activities. Therefore, the breadth of assessed activities and environments is considerably more important in evaluating our method than the number of participants, each of which would perform activities only slightly differently.

In our evaluation, the sensor devices were affixed to five locations on the participant's body, without a strict enforcement of inter-device distance, which covered a range of device constellations and variability similar to varying body heights across participants. Our results showed no notable impact of the body location on the accuracy of trace alignment.

Most importantly, our evaluation has shown that the most challenging activity for our cross-sensor synchronization approach is no activity (i.e., resting and sleeping). While we believe that our results generalize and an evaluation with multiple participants would not change the findings, we acknowledge that for assessing this conclusively, our evaluation could be extended to multiple participants in the future.

### 7.7 Speculative Comparison with Online Synchronization

Online synchronization through a wireless network may be a viable alternative to BMAR to reach higher accuracy (Section 2.1). A main advantage of BMAR is its low power consumption: At the sampling rates used in our evaluation, a barometric pressure sensor operates on less than  $35 \mu W$  [48] and the low-resolution accelerometer uses less than  $70 \mu W$  [54]. Of the commonly used wireless protocols, currently, only Bluetooth low energy reaches comparable values: Makara et al. synchronized two nodes with a power consumption of 133  $\mu W$  with a typical accuracy of 10 ms.

However, producing reliable implementations in practice is challenging, limiting the choice of SoC, and no prior work has shown synchronization of more than two nodes with power consumption in this order of magnitude. In such a scenario, a more powerful device such as a smartphone may be needed as central device for the other

nodes to remain in the stated range of power consumption while BMARs power requirements are unaffected by the number of devices.

Generally, most existing synchronization algorithms are optimized for low error—not low power consumption, which has thus often not been evaluated in prior work. Because most wearable devices already routinely sample acceleration and/or air pressure from built-in sensors, our method imposes no additional power consumption during runtime. Additionally, our method may be used with devices that are not engineered to be compatible to be used in the same wireless network.

## 8 CONCLUSION

We have presented BMAR, a novel method to synchronize sensor data across multiple wearable devices without the need for communication or explicit user interaction during runtime. BMAR leverages barometric pressure sensors and accelerometers, which are inexpensive, low-power, and commonly found in wearable devices. Therefore, they represent a simple solution to synchronizing wearable devices without the need for inter-device communication, allowing data recording and device operation to focus on low power consumption in miniaturized devices.

BMAR's two-stage process comprises pre-alignment, which we use to avoid latching on to local but not a global optimum in alignment, and refinement. We evaluated BMAR using custom-built, low-power wearable sensing devices that recorded a series of signals across a variety of activities over the course of up to 10 hours. Our method achieved accurate synchronization with a median error of 33.4 ms, which is low enough for most practical purposes. Even during activities with little to no motion such as sleeping, our method synchronized signals with an error of 3 s, all while reliably rejecting signals that were not concurrently recorded across devices.

We believe that BMAR is a practical approach for supporting cross-device data collections that build on unobtrusive miniaturized devices, focus on duration, and ease of use.

## REFERENCES

- [1] Parastoo Alinia, Chris Cain, Ramin Fallahzadeh, Armin Shahrokni, Diane Cook, and Hassan Ghasemzadeh. 2017. How Accurate Is Your Activity Tracker? A Comparative Study of Step Counts in Low-Intensity Physical Activities. *JMIR Mhealth Uhealth* 5, 8 (Aug. 2017), e106.
- [2] Kioxia America. 2019. TH58CYG3S0HRAIJ. [https://static6.arrow.com/arrowpdfconversion/dd02224cfed190f11368fa73736fc9a9b13d6528/th58cyg3s0hrajij\\_datashet\\_en\\_20191001.pdf](https://static6.arrow.com/arrowpdfconversion/dd02224cfed190f11368fa73736fc9a9b13d6528/th58cyg3s0hrajij_datashet_en_20191001.pdf). Last accessed: 2022-12-09.
- [3] Byeong Wan An, Jung Hwal Shin, So-Yun Kim, Joohee Kim, Sangyoon Ji, Jihun Park, Youngjin Lee, Jiuk Jang, Young-Geun Park, Eunjin Cho, Subin Jo, and Jang-Ung Park. 2017. Smart Sensor Systems for Wearable Electronic Devices. *Polymers* 9, 8 (2017). <https://doi.org/10.3390/polym9080303>
- [4] Farzad Asgarian and Khalil Najafi. 2022. BlueSync: Time Synchronization in Bluetooth Low Energy With Energy-Efficient Calculations. *IEEE Internet of Things Journal* 9, 11 (2022), 8633–8645. <https://doi.org/10.1109/JIOT.2021.3116921>
- [5] David Bannach, Oliver Amft, and Paul Lukowicz. 2009. Automatic Event-Based Synchronization of Multimodal Data Streams from Wearable and Ambient Sensors. In *Smart Sensing and Context (Lecture Notes in Computer Science)*, Payam Barnaghi, Klaus Moessner, Mirko Presser, and Stefan Meissner (Eds.). Springer, Berlin, Heidelberg, 135–148. [https://doi.org/10.1007/978-3-642-04471-7\\_11](https://doi.org/10.1007/978-3-642-04471-7_11)
- [6] Rummana Bari, Md Mahbubur Rahman, Nazir Saleheen, Megan Battles Parsons, Eugene H Buder, and Santosh Kumar. 2020. Automated detection of stressful conversations using wearable physiological and inertial sensors. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 4, 4 (2020), 1–23.
- [7] André Bideaux, Bernd Zimmermann, Stefan Hey, and Wilhelm Stork. 2015. Synchronization in wireless biomedical-sensor networks with Bluetooth Low Energy. *Current Directions in Biomedical Engineering* 1, 1 (Sept. 2015), 73–76. <https://doi.org/10.1515/cdbme-2015-0019> Publisher: De Gruyter Section: Current Directions in Biomedical Engineering.
- [8] Rajendra Singh Bisht, Sourabh Jain, and Naveen Tewari. 2021. Study of Wearable IoT devices in 2021: Analysis & Future Prospects. In *2021 2nd International Conference on Intelligent Engineering and Management (ICIEM)*. 577–581. <https://doi.org/10.1109/ICIEM51511.2021.9445334>
- [9] Richard Bloss. 2015. Wearable sensors bring new benefits to continuous medical monitoring, real time physical activity assessment, baby monitoring and industrial applications. *Sensor Review* 35, 2 (Jan. 2015), 141–145.
- [10] Alan Kevin Bourke, Espen Alexander F Ihlen, Ronny Bergquist, Per Bendik Wik, Beatrix Vereijken, and Jorunn L Helbostad. 2017. A Physical Activity Reference Data-Set Recorded from Older Adults Using Body-Worn Inertial Sensors and Video Technology-The ADAPT

- Study Data-Set. *Sensors (Basel)* 17, 3 (March 2017).
- [11] Elisa Bruno, Sebastian Böttcher, Pedro F. Viana, Marta Amengual-Gual, Boney Joseph, Nino Epitashvili, Matthias Dümpelmann, Martin Glasstetter, Andrea Biondi, Kristof Van Laerhoven, Tobias Loddenkemper, Mark P. Richardson, Andreas Schulze-Bonhage, and Benjamin H. Brinkmann. 2021. Wearable devices for seizure detection: Practical experiences and recommendations from the Wearables for Epilepsy And Research (WEAR) International Study Group. *Epilepsia* 62, 10 (2021), 2307–2321. <https://doi.org/10.1111/epi.17044> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/epi.17044>
- [12] Dario Camuffo, Chiara Bertolin, Phil D. Jones, Richard Cornes, and Emmanuel Garnier. 2010. The earliest daily barometric pressure readings in Italy: Pisa AD 1657-1658 and Modena AD 1694, and the weather over Europe. *The Holocene* 20, 3 (May 2010), 337–349. <https://doi.org/10.1177/0959683609351900> Publisher: SAGE Publications Ltd.
- [13] Meredith A Case, Holland A Burwick, Kevin G Volpp, and Mitesh S Patel. 2015. Accuracy of Smartphone Applications and Wearable Devices for Tracking Physical Activity Data. *JAMA* 313, 6 (Feb. 2015), 625–626.
- [14] A. Choi and H. Shin. 2017. Photoplethysmography sampling frequency: pilot assessment of how low can we go to analyze pulse rate variability with reliability? *Physiological Measurement* 38, 3 (Feb. 2017), 586. <https://doi.org/10.1088/1361-6579/aa5efa> Publisher: IOP Publishing.
- [15] Jordana Dahmen, Alyssa La Fleur, Gina Sprint, Diane Cook, and Douglas L. Weeks. 2017. Using wrist-worn sensors to measure and compare physical activity changes for patients undergoing rehabilitation. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. 667–672. <https://doi.org/10.1109/PERCOMW.2017.7917643>
- [16] Peter H Dana and Bruce M Penrod. 1990. The role of GPS in precise time and frequency dissemination. *GPS World* 1, 4 (1990), 38–43.
- [17] Walteneus Dargie. 2009. Analysis of time and frequency domain features of accelerometer measurements. In *2009 Proceedings of 18th International Conference on Computer Communications and Networks*. IEEE, 1–6.
- [18] Analog Devices. 2018. ADXL355, Low Noise, Low Drift, Low Power, 3-Axis MEMS Accelerometers. [https://www.analog.com/media/en/technical-documentation/data-sheets/adxl354\\_355.pdf](https://www.analog.com/media/en/technical-documentation/data-sheets/adxl354_355.pdf). Last accessed: 2022-12-09.
- [19] Renesas Electronics. 2022. DA1469x. <https://www.renesas.com/us/en/document/dst/da1469x-datasheet?r=1606281>. Last accessed: 2022-12-09.
- [20] Jeremy Elson, Lewis Girod, and Deborah Estrin. 2002. Fine-Grained Network Time Synchronization Using Reference Broadcasts. In *5th Symposium on Operating Systems Design and Implementation (OSDI 02)*. USENIX Association. <https://www.usenix.org/conference/osdi-02/fine-grained-network-time-synchronization-using-reference-broadcasts>
- [21] Saurabh Ganerwal, Ram Kumar, and Mani B. Srivastava. 2003. Timing-sync protocol for sensor networks. In *Proceedings of the 1st international conference on Embedded networked sensor systems (SenSys '03)*. Association for Computing Machinery, New York, NY, USA, 138–149. <https://doi.org/10.1145/958491.958508>
- [22] Yan Gao, Yang Long, Yu Guan, Anna Basu, Jessica Baggaley, and Thomas Ploetz. 2019. Towards Reliable, Automated General Movement Assessment for Perinatal Stroke Screening in Infants Using Wearable Accelerometers. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 1, Article 12 (mar 2019), 22 pages. <https://doi.org/10.1145/3314399>
- [23] Shkurta Gashi, Elena Di Lascio, and Silvia Santini. 2019. Using unobtrusive wearable sensors to measure the physiological synchrony between presenters and audience members. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 1 (2019), 1–19.
- [24] Unmesh Ghoshdastider, Reinhard Viga, and Michael Kraft. 2014. Wireless time synchronization of a collaborative brain-computer-interface using bluetooth low energy. In *SENSORS, 2014 IEEE*. 2250–2254. <https://doi.org/10.1109/ICSENS.2014.6985489>
- [25] Unmesh Ghoshdastider, Reinhard Viga, and Michael Kraft. 2015. Experimental evaluation of a pairwise broadcast synchronization in a low-power Cyber-physical system. In *2015 IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNet)*. 50–52. <https://doi.org/10.1109/WISNET.2015.7127399>
- [26] Thomas Gilbert, Sally Day, Antonia F De C Hamilton, and Jamie Ward. 2022. A Simple Method for Synchronising Multiple IMUs using the Magnetometer. In *Proceedings of the 2022 ACM International Symposium on Wearable Computers (ISWC '22)*. Association for Computing Machinery, New York, NY, USA, 100–102. <https://doi.org/10.1145/3544794.3558466>
- [27] Martin Gjoreski, Hristijan Gjoreski, Mitja Luštrek, and Matjaž Gams. 2016. How Accurately Can Your Wrist Device Recognize Daily Activities and Detect Falls? *Sensors (Basel)* 16, 6 (June 2016).
- [28] Waqas Ikram, Ivan Stoianov, and Nina F Thornhill. 2010. Towards a radio-controlled time synchronized wireless sensor network: A work in-progress paper. In *2010 IEEE 15th Conference on Emerging Technologies & Factory Automation (ETFA 2010)*. IEEE, 1–4.
- [29] Goertek Microelectronics Inc. 2022. SPL07-003, Digital pressure sensor. <https://media.digikey.com/pdf/Data%20Sheets/Goertek%20Microelectronics%20PDFs/SPL07-003.pdf>. Last accessed: 2023-02-14.
- [30] ECS Inc International. 2021. ECX-16, 32.768 KHz SMD Tuning Fork Crystal. <https://ecstxtal.com/store/pdf/ECX-16.pdf>. Last accessed: 2022-12-09.
- [31] Aida Kamišalić, Iztok Fister, Muhamed Turkanović, and Sašo Karakatić. 2018. Sensors and Functionalities of Non-Invasive Wrist-Wearable Devices: A Review. *Sensors* 18, 6 (2018). <https://doi.org/10.3390/s18061714>

- [32] Sunwook Kim and Maury A. Nussbaum. 2013. Performance evaluation of a wearable inertial motion capture system for capturing physical exposures during manual material handling tasks. *Ergonomics* 56, 2 (2013), 314–326. <https://doi.org/10.1080/00140139.2012.742932> PMID: 23231730. arXiv:<https://doi.org/10.1080/00140139.2012.742932>
- [33] Heike Leutheuser, Dominik Schuldhaus, and Bjoern M Eskofier. 2013. Hierarchical, multi-sensor based classification of daily life activities: comparison with state-of-the-art algorithms using a benchmark dataset. *PLoS one* 8, 10 (2013), e75196.
- [34] Zan Li, Torsten Braun, and Desislava C. Dimitrova. 2015. Methodology for GPS Synchronization Evaluation with High Accuracy. In *2015 IEEE 81st Vehicular Technology Conference (VTC Spring)*. 1–6. <https://doi.org/10.1109/VTCSpring.2015.7145929> ISSN: 1550-2252.
- [35] Xi Liu, Lei Liu, Steven J. Simske, and Jerry Liu. 2016. Human Daily Activity Recognition for Healthcare Using Wearable and Visual Sensing Data. In *2016 IEEE International Conference on Healthcare Informatics (ICHI)*. 24–31. <https://doi.org/10.1109/ICHI.2016.100>
- [36] Eduardo José da S. Luz, David Menotti, and William Robson Schwartz. 2014. Evaluating the use of ECG signal in low frequencies as a biometry. *Expert Systems with Applications* 41, 5 (April 2014), 2309–2315. <https://doi.org/10.1016/j.eswa.2013.09.028>
- [37] Dmytro Makara, Vladyslav Tsybul'nyk, and Taras Kurnyts'kyi. 2019. Power Efficient Clock Synchronization in Bluetooth-Based Mesh Networks. In *Ambient Intelligence (Lecture Notes in Computer Science)*, Ioannis Chatzigiannakis, Boris De Ruyter, and Irene Mavrommati (Eds.). Springer International Publishing, Cham, 14–26. [https://doi.org/10.1007/978-3-030-34255-5\\_2](https://doi.org/10.1007/978-3-030-34255-5_2)
- [38] Karandeep Malhi, Subhas Chandra Mukhopadhyay, Julia Schnepfer, Mathias Haefke, and Hartmut Ewald. 2012. A Zigbee-Based Wearable Physiological Parameters Monitoring System. *IEEE Sensors Journal* 12, 3 (2012), 423–430. <https://doi.org/10.1109/JSEN.2010.2091719>
- [39] Miklós Maróti, Branislav Kusy, Gyula Simon, and Ákos Lédeczi. 2004. The flooding time synchronization protocol. In *Proceedings of the 2nd international conference on Embedded networked sensor systems (SenSys '04)*. Association for Computing Machinery, New York, NY, USA, 39–49. <https://doi.org/10.1145/1031495.1031501>
- [40] David Mills. 1992. *Network time protocol (version 3) specification, implementation and analysis*. Technical Report.
- [41] Stylianos Paraschiakos, Ricardo Cachucho, Matthijs Moed, Diana van Heemst, Simon Mooijaart, Eline P Slagboom, Arno Knobbe, and Marian Beekman. 2020. Activity recognition using wearable sensors for tracking the elderly. *User Modeling and User-Adapted Interaction* 30, 3 (July 2020), 567–605.
- [42] Sungmee Park and Sundaresan Jayaraman. 2017. The wearables revolution and Big Data: the textile lineage. *The Journal of The Textile Institute* 108, 4 (April 2017), 605–614. <https://doi.org/10.1080/00405000.2016.1176632> Publisher: Taylor & Francis \_eprint: <https://doi.org/10.1080/00405000.2016.1176632>
- [43] Thomas Plotz, Chen Chen, Nils Y. Hammerla, and Gregory D. Abowd. 2012. Automatic Synchronization of Wearable Sensors and Video-Cameras for Ground Truth Annotation – A Practical Approach. In *2012 16th International Symposium on Wearable Computers*. 100–103. <https://doi.org/10.1109/ISWC.2012.15>
- [44] Carl C Rheinländer and Norbert Wehn. 2016. Precise synchronization time stamp generation for Bluetooth low energy. In *2016 IEEE SENSORS*. IEEE, 1–3.
- [45] Thomas Rietveld, Riemer J. K. Vegter, Rienk M. A. van der Slikke, Aldo E. Hoekstra, Lucas H. V. van der Woude, and Sonja de Groot. 2019. Wheelchair mobility performance of elite wheelchair tennis players during four field tests: Inter-trial reliability and construct validity. *PLoS ONE* 14, 6 (06 2019), 1–16. <https://doi.org/10.1371/journal.pone.0217514>
- [46] Óscar Seijo, Jesús Alberto López-Fernández, Hans-Peter Bernhard, and Iñaki Val. 2020. Enhanced Timestamping Method for Subnanosecond Time Synchronization in IEEE 802.11 Over WLAN Standard Conditions. *IEEE Transactions on Industrial Informatics* 16, 9 (2020), 5792–5805. <https://doi.org/10.1109/TII.2019.2959200>
- [47] Bosch Sensortec. 2022. BME280, Combined humidity and pressure sensor. <https://www.bosch-sensortec.com/media/boschsensortec/downloads/datasheets/bst-bme280-ds002.pdf>. Last accessed: 2022-12-09.
- [48] Bosch Sensortec. 2022. BMP581, Barometric Pressure Sensor. <https://www.bosch-sensortec.com/media/boschsensortec/downloads/datasheets/bst-bmp581-ds004.pdf>. Last accessed: 2022-12-09.
- [49] Shaban Shabani, Alan K. Bourke, Amir Muaremi, Jens Praestgaard, Kate O’Keeffe, Rob Argent, Martin Brom, Celeste Scotti, Brian Caulfield, and Lorcan C. Walsh. 2022. An Automatic Foot and Shank IMU Synchronization Algorithm: Proof-of-concept. In *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. 4210–4213. <https://doi.org/10.1109/EMBC48229.2022.9871162> ISSN: 2694-0604.
- [50] M.L. Sichitiu and C. Veerarittiphan. 2003. Simple, accurate time synchronization for wireless sensor networks. In *2003 IEEE Wireless Communications and Networking, 2003. WCNC 2003., Vol. 2*. 1266–1273 vol.2. <https://doi.org/10.1109/WCNC.2003.1200555> ISSN: 1525-3511.
- [51] Kasun Somaratne, F John Dian, and Amirhossein Yousefi. 2018. Accuracy analysis of time synchronization using current consumption pattern of BLE devices. In *2018 IEEE 8th Annual Computing and Communication Workshop and Conference (CCWC)*. IEEE, 841–844.
- [52] Andreas Spilz and Michael Munz. 2021. Novel Approach To Synchronisation Of Wearable IMUs Based On Magnetometers. <https://doi.org/10.48550/arXiv.2107.03147> arXiv:2107.03147 [cs, eess].
- [53] Sabarish Sridhar, Prasant Misra, and Jay Warrior. 2015. CheepSync: A Time Synchronization Service for Resource Constrained Bluetooth Low Energy Advertisers. *CoRR* abs/1501.06479 (2015). arXiv:1501.06479 <http://arxiv.org/abs/1501.06479>

- [54] STMicroelectronics. 2011. LIS2DH, MEMS digital output motion sensor: ultra low-power high performance 3-axis “femto” accelerometer. <https://www.st.com/resource/en/datasheet/lis2dh.pdf>. Last accessed: 2022-12-09.
- [55] Tessa Strain, Katrien Wijndaele, Paddy C Dempsey, Stephen J Sharp, Matthew Pearce, Justin Jeon, Tim Lindsay, Nick Wareham, and Søren Brage. 2020. Wearable-device-measured physical activity and future health risk. *Nature Medicine* 26, 9 (Sept. 2020), 1385–1391.
- [56] Bharath Sundararaman, Ugo Buy, and Ajay D. Kshemkalyani. 2005. Clock synchronization for wireless sensor networks: a survey. *Ad Hoc Networks* 3, 3 (May 2005), 281–323. <https://doi.org/10.1016/j.adhoc.2005.01.002>
- [57] Francisco Tirado-Andrés and Alvaro Araujo. 2019. Performance of clock sources and their influence on time synchronization in wireless sensor networks. *International Journal of Distributed Sensor Networks* 15 (Sept. 2019), 155014771987937. <https://doi.org/10.1177/1550147719879372>
- [58] Roberto Vio and Willem Wamsteker. 2001. Limits of the Cross-Correlation Function in the Analysis of Short Time Series. *Publications of the Astronomical Society of the Pacific* 113, 779 (Jan. 2001), 86. <https://doi.org/10.1086/317967> Publisher: The University of Chicago Press.
- [59] Harry J Witchel, Cécilia Oberndorfer, Robert Needham, Aoife Healy, Carina E I Westling, Joseph H Guppy, Jake Bush, Jens Barth, Chantal Herberz, Daniel Roggen, Björn M Eskofier, Waqar Rashid, Nachiappan Chockalingam, and Jochen Klucken. 2018. Thigh-Derived Inertial Sensor Metrics to Assess the Sit-to-Stand and Stand-to-Sit Transitions in the Timed Up and Go (TUG) Task for Quantifying Mobility Impairment in Multiple Sclerosis. *Front Neurol* 9 (Sept. 2018), 684.
- [60] Yun C. Zhang, Shibo Zhang, Miao Liu, Elyse Daly, Samuel Battalio, Santosh Kumar, Bonnie Spring, James M. Rehg, and Nabil Alshurafa. 2020. SyncWISE: Window Induced Shift Estimation for Synchronization of Video and Accelerometry from Wearable Sensors. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 3, Article 107 (sep 2020), 26 pages. <https://doi.org/10.1145/3411824>
- [61] Yun C Zhang, Shibo Zhang, Miao Liu, Elyse Daly, Samuel Battalio, Santosh Kumar, Bonnie Spring, James M Rehg, and Nabil Alshurafa. 2020. Syncwise: Window induced shift estimation for synchronization of video and accelerometry from wearable sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–26.
- [62] Ya-Li Zheng, Xiao-Rong Ding, Carmen Chung Yan Poon, Benny Ping Lai Lo, Heye Zhang, Xiao-Lin Zhou, Guang-Zhong Yang, Ni Zhao, and Yuan-Ting Zhang. 2014. Unobtrusive sensing and wearable devices for health informatics. *IEEE Trans Biomed Eng* 61, 5 (May 2014), 1538–1554.