Addressing Biosignal Data Sharing Security Issues with Robust Watermarking

Vishwa Goudar and Miodrag Potkonjak
Computer Science Department
University of California, Los Angeles

***Abstract*—One of the most important infrastructure requirements in the domain of remote health monitoring BASNs is the secure collection and dissemination of the user's medical data. Data security desiderata in this application domain are not limited to ensuring the confidentiality and integrity of medical data that has been logged to a data sink. Requirements also arise from the need to provide the data owner (BAN user / patient) and the data consumers (healthcare providers, insurance companies, medical research facilities) secure control over the data as it is shared between these various stakeholders. Here, we study a robust watermarking technique to embed security information into biosignal data such that the semantic fidelity of the data is unaffected, while simultaneously ensuring that the watermark is not easily erased or corrupted by malicious data consumers. In doing so, we address three use-cases: proof of ownership, wherein the data owner can prove that she/he is the originator of the data; data tracking, wherein the data owner can trace unauthorized sharing of her/his biosignal data; and content authentication, wherein the data owner can prove whether the biosignal data has been maliciously altered. Based on experimentally collected datasets from a gait-stability monitoring BASN, we show that the embedding of 800 bit watermarks can be achieved robustly and effectively, with near-imperceptible changes to the signal waveform and no loss in the the signal's diagnostic quality.**

Keywords*—Body Area Networks, Watermarking, Medical Data Sharing, Medical Data Security.

I. Introduction

The rising maturity of Body Area Sensor Networks (BASNs) and their underlying technologies is increasing the feasibility of their large-scale deployment over the long term in a variety of domains including health, military and sport. Self-contained, battery-powered embedded systems equipped with a diverse set of sensors, and capable of wireless communication, BASNs promise continuous and nearly-unobtrusive monitoring of the wearer's physiological and bio-mechanical signals and the inference of her/his activity, behaviors and health therefrom. However, several requirements are still being tackled towards successful large-scale and long-term adoption of BASNs. These include improvements in the design and configuration of existing systems so they may afford as yet unachieved capabilities and improved performance, spanning improved sensing, signal processing, communication, data storage and power-management [1][2][3][4][5][6], as well as, enhanced infrastructure support for the deployed BASNs and management of the data they collect.

One of the most important infrastructure requirements in the domain of remote health monitoring BASNs is the secure collection and dissemination of user's medical data [1]. Data security desiderata in this application domain are not limited to ensuring the confidentiality and integrity of medical data that has been logged to a data sink. Requirements also arise from the need to provide the data owner (BASN user / patient) and the data consumers (healthcare providers, insurance companies, medical research facilities) secure control over the data as it is shared between these various stakeholders.

978-1-4799-4657-0/14/$31.00 © 2014 IEEE

Several benefits of adopting BASNs for remote health monitoring exist beyond monitoring patient health while she/he is under medical observation. The detailed data that is continuously collected by the BASN is logged to the patient's health records and can see multiple uses thereafter. It can be used by the BASN user's healthcare providers to monitor long-term trends in the patient's health, with the temporally and diagnostically rich measurements offered by BASNs enabling shortened time-to-treatment and reduced frequency and duration of hospitalization [7]. It can also be used for insurance purposes - as supporting documentation in insurance claims, or as measures in assessing coverage and indemnity requirements. Data collected periodically with BASNs over long periods may also be selectively shared in lieu of compensation with medical researchers conducting clinical studies.

Although these use-cases for BASN collected patient health data promise to propel preventative healthcare, several security concerns emerge from the necessity to share large amounts of medical data. For one, the infrastructure may need to provide the ability to establish that a dataset was collected from a specific patient. With this ability, the data owner (or patient) can prove ownership of her/his data if the ownership ever comes into question. Similarly, the data consumer can utilize this feature to track ownership of the datasets he has access to. Second, having shared her/his data with different consumers, the data owner may want to track the movement of this data even after it has changed hands several times. With this facility, the data owner can trace un-authorized sharing of her/his data. Third, data owners may be required to verify that their data has not been altered before being shared between consumers, a process known as content authentication.

While cryptography-based systems have been proposed to address these security concerns, the security of these solutions itself lies in the strength of their respective key-management protocols. Further, a malicious consumer may come into possession of the decryption keys through legitimate channels only to then illegitimately claim ownership of the data and subsequently share it with other consumers, possibly after altering it. Instead, we offer digital watermarking as an alternative solution for these problems. Watermarking is defined as the practice of imperceptibly altering a Work to

embed a message about that Work [8]. As we shall see, embedding identifying information regarding the data owner and/or summary information regarding the physiological signal into the signal itself enables proof of ownership, data tracking and content authentication. Embedding this information in the data not only obviates the need for key management but also increases overall security and reduces the transmission and storage overhead necessary to achieve data security.

Furthermore, watermarking supports these capabilities while constraining the extent to which the resulting changes to the measured physiological signal are perceivable. This translates to minimal changes to the signal waveform. More importantly, watermarks may also be applied without compromising the diagnostic quality of the medical data. Here, it is important to note that diagnostic quality of the data isn't based on the extent to which the watermarked and original data measurements differ. Rather, it is based on the signal's semantics - the data consumer's ability to accurately estimate medical diagnostic metrics and the signal features that they depend upon, from the watermarked data. While pertinent functional tools and tests will inform the selection of those medical diagnostic metrics, in the context of mobile health monitoring these metrics are estimated from features of the physiological and/or biomechanical signals sampled by BASNs. Therefore, as long as the data can be watermarked without altering the signal features and diagnostic metrics that are derived from it, the diagnostic value of the data is left unaffected by the watermarking process.

In this paper, we detail a linear programming based technique to watermark bio-signal data collected by BASNs, for the purposes of secure sharing of the data, such that (i) the watermarks are robust, (ii) the watermarks produce imperceptible alterations to the data, and (iii) the watermarks do not affect the diagnostic quality of the data. It is neccesary that the watermarking process produce robust watermarks so that they are not easily corrupted or erased by malicious data consumers, and in this sense the robustness of a watermark is key to the level of security that the watermarking technique will provide. We study the performance of our technique in the context of foot plantar pressure datasets collected from multiple subjects by a wireless gait stability monitoring BASN known as HERMES [9]. In this application context we evaluate the technique in terms of embedding effectiveness, i.e. the probability of successful watermarking, capacity, i.e. the number of bits that can be successfully embedded, and fidelity, i.e. the extent to which the original and watermarked works differ perceptually. We show that while altering the signal by less than 5%, upto 800 bits can be embedded robustly with high effectiveness ($> 99\%$) per stride of the HERMES user.

II. Related Work

The confidentiality and integrity of BASN collected patient data during exchange and storage within the network, as well as over the course of transmission outside the network for post-processing and storage at the data sink, have been identified as a key security issue for BASN deployments and addressed by a number of studies [10] [11]. Further, secure sharing of patient health records has been studied in the context of data encryption-based solutions [12] [13]. In [12], the authors proposed a patient-centric framework that

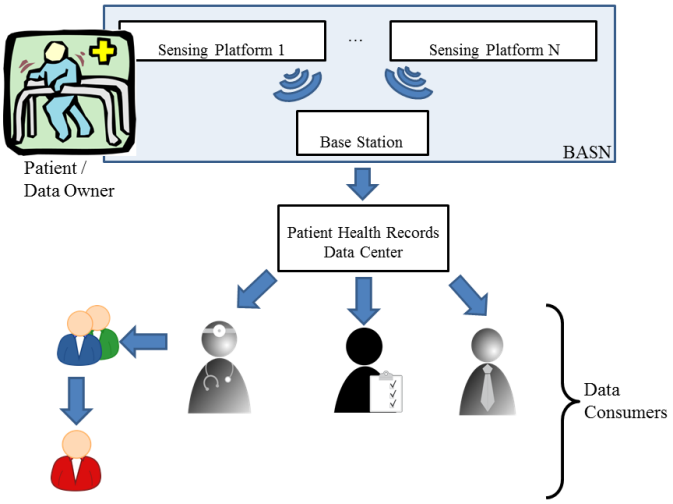

Fig. 1: Body Area Sensor Network Architecture.

provides data access control for patient health records stored on semi-trusted servers. The framework provides mechanisms based on Attribute Based Encrypyion (ABE) for scalable key management, flexible access for multiple classes of data consumers, and efficient user revocation. In [14], the authors propose a detailed data provenance framework to collect and share provenance metadata for patient health records, to help data consumers verify the accuracy and authenticity of the data and track its origins and changes made to it. However, the framework does not define and address a threat model.

Watermarking techniques have been thoroughly studied as a means to achieve proof of ownership and transaction tracking [8]. Primarily, these techniques have been applied to image, video and audio data, and include least significant bit alteration methods, quantization methods and signal decomposition based methods, where the general idea is to decompose the data in the spatial and/or spectral domains and alter the coefficients of the decomposed signal in a manner that maximizes the embedded information while curtailing the visually and/or audibly perceivable changes. These watermarking techniques have been extended to the domain of medical images, as well as physiological signal measurements including ECG and EEG signals. The authors of [15] compare and contrast three watermarking techniques with regards to their ability to verify EEG signal integrity after noise contamination resulting from communication. The authors of [16] propose an LSB watermarking scheme in support of proof-of-ownership for ECG signals. However, LSB watermarks provide poor robustness to malicious alterations. The authors of [17] describe a spread spectrum watermarking scheme that embeds robust and imperceptible watermarks into ECG signals. However, such a scheme addresses security considerations only during communication of the data rather that over the course of sharing it. To our knowledge, our work enabling robust watermarking of bio-signal data under perceptual and semantic constraints, is the first to address security issues arising from data sharing rather than communication.

Fig. 2: The HERMES smart shoe platform.

III. Preliminaries

A. System Architecture

Fig. 1 illustrates a typical BASN architecture in terms of its data flow. A BASN typically consists of several body worn sensors/sensing platforms measuring physiological and/or biomechanical signals continuously and wirelessly transmitting the measurements over a short range to a base station such as a smart phone. The data is then forwarded by the base station to a data sink that logs it to the patient's health records. A commmon use-case is one where the sensing platforms and the base-station collaboratively apply signal processing techniques to detect "anomalous" health events and alert the patient and/or the healthcare provider(s) [18]. Aside from such use-cases involving real-time data consumption, the data may also be consumed offline by its owner, her/his healthcare providers, insurance providers, family, research orgianizations, etc. Further, the data may be shared between data consumers as well, after authorization has been provided by the data owner.

We shall study the performance of our watermarking technique in the context of a gait stability monitoring BASN called HERMES [9]. HERMES is a battery powered smart-shoe composed of 99 passive-resistive pressure sensors located across the shoe insole, that are used to monitor foot plantar pressure as the subject ambulates (see Fig. 2). Typically, HEREMES continously samples each of the plantar pressure sensors at 50Hz as the patient ambulates, and forwards the data to a smart phone via a low-power Bluetooth radio. Features extracted from the measurements of HERMES's sensors can be used to compute several functional gait metrics. For example, the foot contact metric of the GARS-M gait assessment scale [19] evaluates the heel strike angle, which can be estimated by HERMES based on the difference in time between heel strike and forefoot strike. Similarly, the staggering metric from the GARS-M scale determines laterally directed loss of balance, and can be monitored based on the plantar pressure difference between lateral areas of the foot. Other metrics that can be gleaned include the inter foot-strike interval that is relevant to the step symmetry and continuity measures of the Tinetti gait assessment tool [19]. Finally, the spatial average of peak plantar pressure is of interest in monitoring diabetic patients for plantar ulcers. Signal processing for functional gait metric computations based on the data collected by HERMES either occurs at the base station (smart phone), or may be performed offline.

B. Threat Model

In the three security use-cases that we address, namely proof of ownership, data tracking and content authentication,

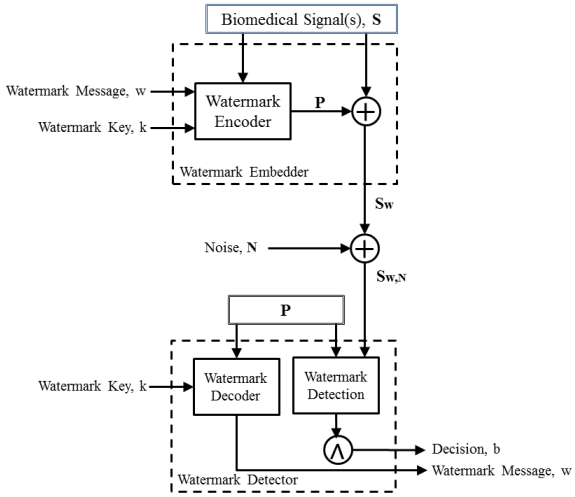

Fig. 3: Watermarking System.

our goal is to thwart the effort of a malicious data consumer attempting to corrupt or erase the embedded watermark and possibly replace it with his own. By successfully achieving this in the proof of ownership and data tracking use-cases, the malicious data consumer gains the ability to illegally share data that is not owned by him and cover his tracks while doing so. In the case of content authentication, the malicious consumer can alter the data without leaving any evidence of the transformation. For example, this can be used by a disingenuous healthcare provider looking to avoid liability lawsuits, or by a duplicitous insurance provider looking to deny legitimate claims. Therefore it is important to validate the robustness claim of our watermarking procedure.

Towards this end, we consider two types of threats. First, we consider a malicious data consumer that adds gaussian noise to the watermarked signal in an attempt to corrupt the watermark and decrease the likelihood of successfully detecting it. This will hinder the ability of a data owner to prove ownership or track shared data. Seconds, we consider a malicious consumer who attempts to reverse engineer our technique with the aid of the watermarked signals, to gain access to the watermark encoding. Thereafter, he may remove the watermark and arrive at the original signal, which he is then free to share in an untraceable manner.

IV. Robust Watermarking of Biomedical Signals

Fig. 3 outlines the blind watermarking system we will use to address the proof of ownership, data tracking and content authentication requirements for biomedical signal data collected by BASNs. The system consists of a watermark embedder used by the data owner to imperceptibly and robustly embed a watermark, and a watermark detector that is then used to test for the existence of the watermark in biomedical signal datasets, as required. The embedder uses an encoding algorithm to generate an encoding vector, P, for the watermark message, w, that is most suitable for the biomedical signal vector S given a secret watermark key k. Suitability of the encoding is defined in terms of the robustness and imperceptibility afforded by it once it is embedded into the original

signal. The embedder then embeds the encoding vector into the original signal vector, S, to generate an altered biomedical signal, S_w, which now includes the watermark.

As per the first threat in our threat mode, over the course of sharing the watermarked signals, a malicious consumer may further alter the signal into $S_{w,N}$ by adding noise to it in an attempt to corrupt the watermark. Once this version of the signal needs to be validated by checking for the existence of the data owner's watermark, the watermark detector accepts the publicly available watermarked signal ($S_{w,N}$) and the watermark encoding vector (P) from the data owner. Following this, it (i) confirms that the encoding vector exists in the signal under test ($S_{w,N}$) via decision bit b, and (ii) decodes the encoding vector (P) with the help of the watermark key (k) into the data owner's watermark message (w). Here, the decision bit (b) indicates whether the encoded vector exists in the signal under test or not. Note that encoding vector P and the original signal S are kept secret from the data consumers. Also note that the detection system here is a blind detector. That is, the original signal S is never disclosed during the decoding process.

To enable proof of ownership, a watermark identifying the data owner is embedded into the medical signal. Proving ownership is as simple as proving that the data owner's identifier is embedded in the signal when it is tested. As long as the watermarking procedure produces robust watermarks, the watermark detector should indicate the watermark's presence. Data tracking is achieved with the data owner embedding distinct identifiers in each shared copy of her/his biomedical signal data. For example, data shared with one's healthcare provider is embedded with a different watermark message than the one shared with the insurance provider or with one's friends. If the data ends up in the possession of an unauthorized party, the data owner can iterate through the set of watermark messages she/he has used and identify the legitimate but malicious data consumer that released the patient's data without her/his authorization. Finally, content authentication is achieved by embedding the data with a watermark comprised of a cryptographic hash of the original signal. On successful detection of the watermark by the watermark detector, authentication of the data is achieved by removing the watermark from the watermarked signal and comparing it to a cryptographic hash of the signal that remains. If this is, in fact, the original signal, the watermark and hash will match. If the signal has been altered, the watermark and the hash will differ.

Before we delve into details regarding the construction of such a watermarking procedure, we highlight an important property of most biomedical signals which we leverage in our watermark encoding algorithm. Bio-medical signals are often segmentable in a straight-forward manner since they arise from human physiology. For example, ECG signals are based on heart-beats and have a distinct waveform (e.g. QRS complex in ECG) that is well-studied and whose origin is well-understood. The same is true of EEG waveforms. In the case of bio-mechanical measurements, human movement and gait usually consists of periodic and repeated sequences of movements which can be used to segment spatio-temporal bio-signals comprised of accelerometer, gyroscope and pressure/force measurements at the joints and limbs. For example, the foot plantar pressure measurements of HERMES can be attributed to the well-studied stance and swing gait phases. Stance is the phase of human gait when the foot is in contact with the ground, while swing is the phase when the foot is in the air. Therefore, the pressure measurements at all sensors can be segmented based on these phases. Consequently, we denote sensor s_i's measurement at epoch j of the segmented signal as sample s_{ij}. Encoding a watermark message (w) into the samples can be equated to perturbing each sample s_{ij} by some amount p_{ij}:

$$s_{w_{ij}} = s_{ij} + p_{ij} \tag{1}$$

Now, there are three requirements for the watermarking procedure. **First**, the watermark should be robustly embedded in the original signal. Recall that this means the watermark is detectable even if the watermarked signal is altered maliciously by the addition of gaussian noise. To satisfy this requirement, robust detection of the watermark will be achieved with a linear-correlation based detection procedure. In other words, detection of the watermark depends on how well correlated the watermarked signal and the watermark are. The detection metric z is calculated as:

$$z = S_t \cdot P \tag{2}$$

where S_t is the signal under test for the watermark encoded as P. The linear-correlation based technique evaluates the existence of the watermark in the tested signal based on the detection metric, z. Specifically, if the metric value is greater than some threshold τ, positive detection occurs. Otherwise negative detection occurs. However, there is always a chance that the watermark detector will produce incorrect results. Either the detector will indicate the presence of the watermark when it does not exist, which is called a false positive. False positives can occur when the original signal happens to be well correlated with the encoded watermark by chance, driving the detection metric beyond the threshold. Alternatively, the detector may fail to detect the watermark even though the signal has includes the watermark, which is known as a false negative. This occurs when the original signal exhibits a strong negative correlation with the encoded watermark by chance, and reduces the detection metric below the threshold, thereby eluding detection. Therefore, the false positive and false negative rates play and imporant role in achieving an effective and robust watermarking scheme, and must be taken into consideration by the encoding algorithm.

A. Thwarting the Watermark Corruption Threat

If $S_t = S_w$, then equations (3a) - (3c) outline why the detection metric has a high value. Since P is highly correlated with itself, the value of $(P \cdot P)$ is high. In contrast, an arbitrary original signal S is unlikely to correlate well with P especially in high-dimensional space. Therefore, $(S \cdot P)$ is very likely to have a small value. Further, if S_t does not embed the watermark message w that has been encoded to P, then z is highly likely to evaluate to a low value since S_t and P will exhibit low correlation. Therefore, z is most likely to evaluate to a high value only when the tested signal S_t embeds the encoded watermark P.

$$z = S_w \cdot P \tag{3a}$$
$$= (S + P) \cdot P \tag{3b}$$
$$= (S \cdot P) + (P \cdot P) \tag{3c}$$

This procedure also yields robust watermarks for the reasons illustrated in equations (4a) - (4c). Here, the tested signal S_t is a maliciously altered version of the watermarked signal with the goal of erasing the watermark. Assuming the malicious alteration is in the form of additive gaussian noise, the watermark should still be detectable. The reason is similar to our discussion above - additive gaussian noise is unlikely to correlate well with the encoded watermark P. Therefore, although the addition of such noise might change the detection metric z, as long as $(S_w \cdot P)$ is above threshold, this change is unlikely to cause a false negative.

$$z = S_{w,N} \cdot P \tag{4a}$$
$$= (S + P + N) \cdot P \tag{4b}$$
$$= (S \cdot P) + (P \cdot P) + (N \cdot P) \tag{4c}$$

B. Assessing the Watermark Reverse Engineering Threat

In the reverse engineering threat, we assume that the attacker may or may not have knowledge regarding the watermark message (w), but he does not know the encoding vector P. He also does not have access to the original bio-signals (S) or to the watermark key (k). Of course, he is assumed to have access to the watermarked signal (S_w). Therefore, in reverse engineering the watermarking encoding, the attacker must among other things find the encoding vector (P) subject to the constraints in equations (5a) - (5b). Given that P is an unknown, the constraint in equation (5a) is non-linear. Furthermore, it is non-convex, making the problem difficult to solve in polynomial time. Consequently, the linear-correlation based detector makes it very difficult for an attacker to successfully reverse engineer the encoding vector P. Further, as shall see, a lack of access to the watermark key (k) will pose another significant hurdle to the attacker as he attempts to reverse engineer the encoded vector.

$$(S_w - P) \cdot P \leq \tau \tag{5a}$$
$$S_w \cdot P \geq \tau \tag{5b}$$

The **second** requirement for a watermarking procedure is that the watermark must be successfully embedded in the original signal. If it cannot be embedded, the procedure must indicate this. Further, the larger the watermark that can be successfully embedded, the stronger the procedure's ability to address the use cases. For example, a larger watermark in the content authentication use-case can support a larger cryptographic hash of the original signal making for a much stronger guarantee regarding malicious alterations to the biosignal data. As we shall see, tradeoffs exist between the size of the watermark that can be successfully embedded and the embedding effectiveness. In other words, the larger the watermark, the more difficult it is to successfully embed.

Third, the embedding should only result in imperceptible alterations of the original signal. As we have discussed in section I, imperceptible embedding in the context of biomedical signals translates to imperceptible changes to the signal waveform, which we shall refer to as signal waveform imperceptibility. It also translates to no loss in the diagnostic value of the signal data which we call semantic imperceptibility. The constraints on semantic imperceptibility depend on the diagnostic metrics and their underlying signal features which will be derived from the data. The goal of semantic imperceptibility is to ensure that the diagnostic metrics that will be derived from the watermarked signals are exactly the same as those derived from the original signals, thereby yielding no compromise to diagnostic fidelity by the watermarking process.

V. Watermark Encoding Algorithm

We now describe a linear-programming approach for watermark encoding that addresses these three requirements for a biomedical signal watermarking scheme.

A. Constraint: Embedding the Watermark

To address the successful embedding requirement, it is necessary that the watermark message w be correctly encoded by the vector P. This is achieved by the constraint in equations (6a) - (6b). The watermark message is encoded as a set of perturbations over all sensors and for all epochs of the segmented signal. The constraints in equations (6a) - (6b) requires that the sign of perturbation at position m specified by the watermark key (k) be constrained by the bit m of the binary encoding for the watermark message (w). In other words, k acts as a seed to a random number generator (pos) that generates (sensor, epoch) pairs indicating the position of perturbation for each bit in the watermark message. The perturbation at this (sensor, epoch) position must be positive if the corresponding bit is 1 and negative otherwise. Note that there may be several positions not generated by pos where the perturbations are free to attain positive or negative values to accomodate other constraints in the Linear Program (LP). Also note that an attacker will not be privy to the value of k, as it is kept secret by the data owner. Therefore, as we have stated, the lack of this knowledge poses an additional hurdle to an attacker attempting to reverse engineer the data owner's encoding vector (P).

$$p_{pos_{k,w,m}} > 0 \quad if \quad w_m = 1 \quad \forall m \in [1, log(w)] \tag{6a}$$
$$p_{pos_{k,w,m}} < 0 \quad if \quad w_m = 0 \quad \forall m \in [1, log(w)] \tag{6b}$$

B. Constraint: Signal Waveform Imperceptibility

We address imperceptibility of the changes to the original signal waveform in two ways. First, each perturbation p_{ij} should be within a small percentage of the original sample s_{ij} (equation (7)). In other words, the perturbations are bounded by $[100 - \epsilon \times 100, 100 + \epsilon \times 100]\%$ of the original samples.

$$-\epsilon s_{ij} \leq p_{ij} \leq \epsilon s_{ij} \quad \forall i, j \tag{7}$$

The second signal waveform imperceptibility constraint applies to signal smoothness after perterbuation. Here, a bound

is applied to the first derivative of the altered signal (equation (8a)). Equations (8b) - (8c) translate this constraint to LP form. Together, these two sets of constraints limit the perceptibility of changes to the signal waveform.

$$|p_{ij-1} - p_{ij}| \leq \Delta \qquad \forall i, j > 1 \tag{8a}$$
$$p_{ij-1} - p_{ij} \leq \Delta \qquad \forall i, j > 1 \tag{8b}$$
$$p_{ij} - p_{ij-1} \leq \Delta \qquad \forall i, j > 1 \tag{8c}$$

C. Constraint: Semantic Imperceptibility

We address semantic imperceptibility in the context of the maximum amplitude and foot contact signal features of plantar pressure signals measured by HERMES. The maximum amplitude feature is computed as the maximum pressure observed at each sensor over the stance phase. The peak pressure diagnostic metric is then derived as the mean maximum amplitude taken over all sensors. Similarly, the staggering diagnostic metric is derived as the average difference between the maximum amplitude features of lateral areas of the foot sole. To prevent loss of semantic accuracy for these metrics, the maximum amplitude signal feature is computed from the watermarked signal (equation (9a)), and the peak pressure and staggering diagnostic metrics derived from these signal features are constrained to the corresponding metric values derived from the original signals (equations (9b) - (9c)). Here, A and B correspond to the maximum amplitude and staggering diagnostic metrics based on the original signals, and a_i and b_i are the corresponding coefficients for each sensor i.

$$\max_j(s_{ij} + p_{ij}) = M_i \qquad \forall i \tag{9a}$$
$$\sum_i (a_i M_i) = A \tag{9b}$$
$$\sum_i (b_i M_i) = B \tag{9c}$$

Note that equation (9a) is not in LP form and must be transformed before input to an LP solver (equation (10)). However, equation (10) does not provide the true maximum (M_i) of the perturbed samples. To arrive at the true maximum, m_i must be minimized, and this will be addressed in the following subsection.

$$s_{ij} + p_{ij} \leq m_i \qquad \forall i, j \tag{10}$$

The foot contact signal feature is defined as the sampling epoch, of the segmented signal, at which foot contact is sensed at a sensor. Note that foot contact at different sensors will be different. For example, contact at the heel occurs much before contact at the toes. Contact is sensed at a sensor when its pressure measurements exceed some threshold tr. To guarantee that the foot contact based diagnostic metrics, namely heel strike angle and inter-footstrike interval, are left unaltered after the watermark has been embedded, we ensure that at each sensor the perturbed signal at all epochs prior to the foot contact epoch is less than the threshold tr (equation (11a)), while the perturbed signal at the actual foot contact epoch is greater than or equal to tr (equation (11b)). Here, the foot contact epoch derived from the original signals is denoted t_{fc}. In this way, the foot contact signal feature derived from the watermarked signal will be identical to that derived from the original signal. Consequently, all diagnostic metrics that depend on this feature will also be identical before and after watermarking.

$$s_{ij} + p_{ij} < tr \qquad \forall i, j < t_{fc} \tag{11a}$$
$$s_{it_{fc}} + p_{it_{fc}} \geq tr \qquad \forall i \tag{11b}$$

D. Objective Function: Maximizing Robustness

Finally, the robustness requirement is addressed by the objective function. Recall from our discussion in section IV that minimizing false positives and false negatives translates to ensuring that the local-correlation based detection metric (z) is greater than τ for a watermarked signal and less that τ for an un-watermarked signal. To minimize the chance that a noise vector (n) introduced into the watermarked signal decreases its z value below τ, we must maximize the distance or difference between $(S \cdot P)$ and $(S_w \cdot P)$. Picking a threshold value τ that is close to but larger than $(S \cdot P)$ will then make it difficult for a random noise vector to yield a false negative. In other words, we must maximize $(P \cdot P)$ (equations (4a) - (4c)). In LP form, this is similar to saying that we must maximize $\sum_{i,j} |p_{i,j}|$.

Although thus far we have considered generating the encoding vector P individually for each signal segment, τ imposes a unified constraint across segments. In other words, $(S \cdot P)$ and $(S_w \cdot P)$ must be separated by τ regardless of the segment under consideration. To retain the ability to solve the LP individually for each segment, we additionally force $(S \cdot P)$ to be as close to some constant X as possible (equations (12a) - (12b)). In doing so, $(S \cdot P)$ is tied as close to X as possible as the value of $\sum_{i,j} |p_{i,j}|$ is maximized, thus driving $(S_w \cdot P)$ as far away from X as possible. The best value for X can be arrived at by a line search within bounds generated based on the values of S and ϵ. Note that the objective function in equation (12b) not only aims to limit false positive and false negative rates, but also addresses the maximum amplitude related constraint from the previous subsection.

$$\sum_{i,j} s_{ij} p_{ij} \geq X \tag{12a}$$
$$minimize: \quad \sum_{i,j} s_{ij} p_{ij} - \sum_{i,j} |p_{i,j}| + \sum_i m_i \tag{12b}$$

VI. Validation

We validate the ability of our linear programming based approach to robustly watermark biomedical signals by applying the technique to three plantar pressure datasets collected with HERMES. The datasets correspond to plantar pressure measurements of male and female subjects with different weights and gait profiles, collected over a few hundred footsteps as the subjects walked. The LP is applied independently to the plantar pressure signals at each footstep of each subject. Note that the number of available watermarkable positions is limited by the number of sensors and epochs per foot steps. Generally, a

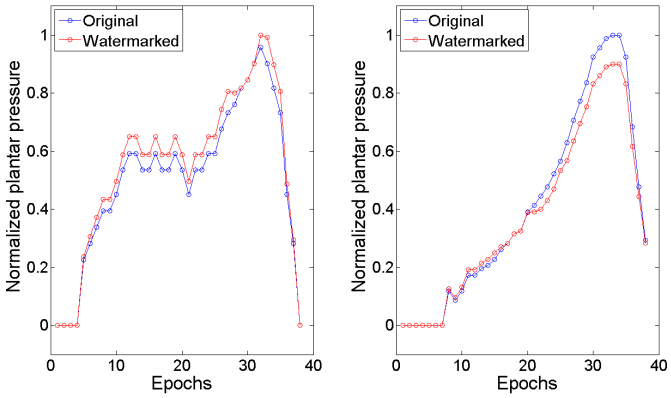

Fig. 4: Original vs. watermarked plantar pressure measurements for 2 sensors over the stance phase of a single footstep from dataset 1.

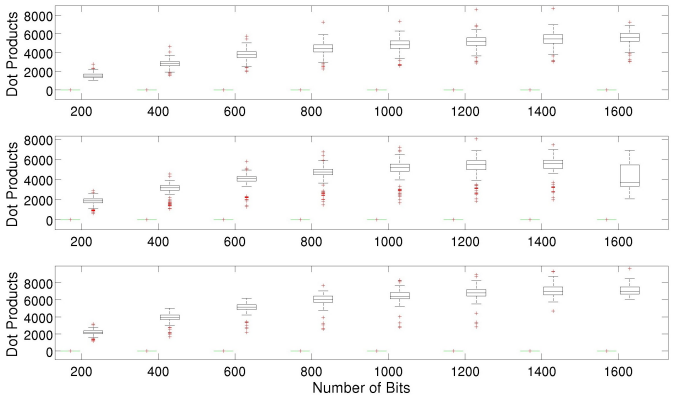

Fig. 5: Boxplots of detection metric (z) before (green) and after (black) embedding encoded watermark, as the size of the watermark message is varied. Each panel represents one dataset, ordered from top to bottom.

vast majority of steps last fewer than 50 epochs. Further, most sensors experience plantar pressure for a minority of these epochs, experiencing 0 pressure at most epochs during which no perturbation is possible (see equation (7)). Consequently, each of the datasets exposes between 1500 and 2400 embeddable positions for each signal segment or foot step.

Fig. 4 shows the watermarked signals from 2 sensors over a single footstep of one of the plantar pressure datasets, where w is 200 bits long, ϵ is 10%, Δ is $\delta \times max(s_{ij-1}, s_{ij})$, and δ is 0.05. Note that we will continue with this definition of Δ for the rest of our validations. Several observation with regards to the requirements for the watermarking procedure are apparent in Fig. 4. The signal waveforms are very similar before and after watermarking for both sensors. Robust watermarking was achieved even at ϵ values of 1%. However, a larger value was used in Fig. 4, so the original and watermarked signals could be distinguished from one another. Further, Fig. 4 also shows that the semantic imperceptibility requirement for the foot contact signal feature dictates minimal changes to the waveforms at the initial epochs prior to foot contact, which we observe in both plots. In contrast, the LP solver achieves imperceptibility of the maximimum amplitude signal feature in a different way. The maximum amplitude of the watermarked signal is less than that of the original signal at some sensors (Fig. 4, right plot), but more than that of the original signal at other sensors (Fig. 4, left plot). However, these difference are balanced out to achieve the overall effect of identical diagnostic metrics before and after watermarking.

We studied the robustness of our watermarking procedure by varying the embedding capacity while watermarking all plantar pressure signals for all footsteps in all datasets. For a given embedding capacity, the embedding procedure was repeated 10 times for each foot step of each dataset, as the watermark message w, and therefore its binary encoding, was randomly varied. We set ϵ to 0.05, δ to 0.025 and X to 0. Fig. 5 shows boxplots of the detection metric before watermarking (i.e. $z = S \cdot P$, in green) and after watermarking (i.e. $z = S_w \cdot P$, in black). As expected, the objective function ties $S \cdot P$ to 0 while $P.P$ increases with the embedding capacity. The general increasing trend for the separation of the detection metrics before and after embedding are a result of more embedding positions afforded with an increasing capacity. On the other hand, this increase tapers off at the upper end of the tested range, as the smoothness constraint (equation (8a)) limits the extent to which the perturbations can be close to ϵ.

Fig. 6 plots the embedding efficiency as the embedding capacity was varied. Note that at 1600 bits, we are at two-thirds or more of the maximum capacity available (1500-2400 bits). At this capacity, between 18% to 90% of the segments allow successful embedding. In general, we observe that after about 1000 bits, the embedding efficiency begins to fall. In other words, between 40% and 50% of maximum capacity can be comfortably supported before embedding efficiency begins to degrade.

While the watermark key in the previous set of experiments was fixed, we next tested for robustness with respect to the embedding positions. To accomplish this, six random keys were generated and tested for embedding capacities of 300 and 1300, while values of ϵ, δ and X were retained. Again, each segment was tested 10 times with different watermark messages (w). Fig. 7 shows the robustness and embedding efficiency results for these sets of experiments. We observed

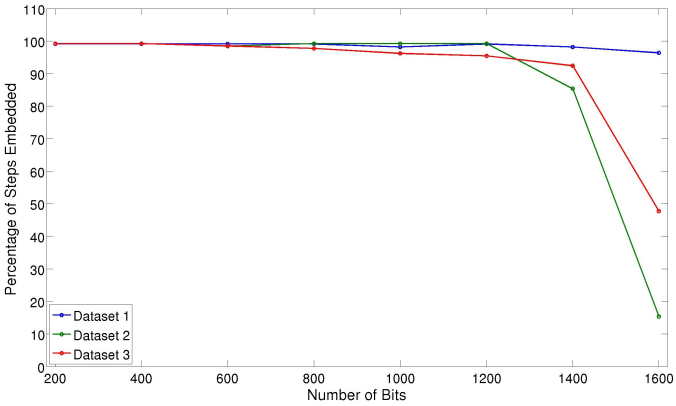

Fig. 6: Embedding Effectiveness for each dataset as the size of the watermark message is varied.

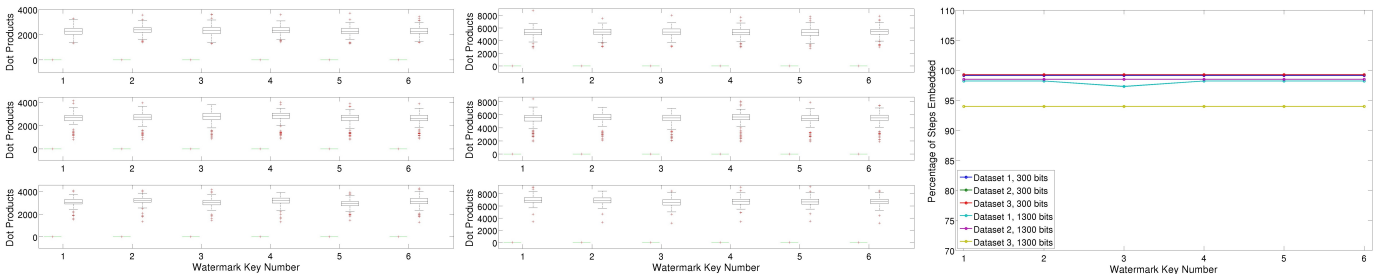

(a) (b) (c)

Fig. 7: Robustness and Embedding Efficiency results when the watermark key k is varied. The plot on the left (middle) shows boxplots for an embedding capacity of 300 (1300) bits. The plot on the right compares embedding efficiency for both these capacities across all datasets.

that for both embedding capacities, the difference between the pre- and post-embedding detection metrics is similar across all watermark keys tested. We also observed that the embedding position does not significantly impact embedding efficiency.

VII. Discussion

An important performance characteristic of a watermarking approach is its embedding capacity. For the use-cases studied, we have shown that the embedding capacity depends, to a large extent, on the sample size and waveform of the signal that is being watermarked. If each bio-signal segment afforded fewer embeddable positions, it is obvious that fewer bits can be embedded. The strength of our approach however is that for the application context tested (foot plantar pressure measurements), up to half of those embeddable positions could be robustly embedded with high efficiency. Another important factor is the number of constraints. As the number of signal features that are required for diagnostic purposes increase, the number of constraints will increase. As a result, the solution space will be more restricted producing poorer robustness, efficiency, or both.

While it turned out that the robustness performance of our technique in the context of plantar pressure measurements was consistently high, it is conceivable that this will not always be the case. Depending on the signal being watermarked and the impact of other constraints applied, it may be difficult to find encoding vectors that poorly correlate with the original signal. Here, a statistical approach must be taken to arrive at an appropriate τ based on the ROC curve, such that the false positive and false negative rates are minimized. This is where the line search on X is likely to come in handy, although we didn't need to perform this further optimization for our validation study.

VIII. Conclusion

In this work, we have addressed security concerns arising from the sharing of BASN collected biosignal measurements with and between medical data consumers. Specifically, three use-cases are addressed: proof of ownership, data tracking and content authentication. We propose the use of linear-correlation based detection method that we show to be able to successfully thwart watermark corruption and erasure threats. Further, we propose an LP formulation to encode watermarks in a robust manner while adhering to imperceptibility constraints that limit significant alterations to the signal waveform and prevent loss of diagnostic fidelity. We study the performance of this technique in the context of foot plantar pressure datasets collected by a gait stability monitoring BASN known as HERMES. Results of our validation studies indicate that upto 800 bits can be succefully embedded in a robust and effective manner into signal segments sampled by 99 sensors for upto 1 second at 50Hz.

References

[1] M. Hanson, H. Powell, A. Barth, K. Ringgenberg, B. Calhoun, J. Aylor, and J. Lach, "Body area sensor networks: Challenges and opportunities," *Computer*, vol. 42, no. 1, pp. 58 –65, jan. 2009.

[2] V. Goudar and M. Potkonjak, "Fault-tolerant and low-power sampling schedules for localized basns," *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on*, vol. 3, no. 1, pp. 86–95, March 2013.

[3] V. Goudar, Z. Ren, P. Brochu, M. Potkonjak, and Q. Pei, "Optimizing the output of a human-powered energy harvesting system with miniaturization and integrated control," *Sensors Journal, IEEE*, vol. PP, no. 99, pp. 1–1, 2013.

[4] V. Goudar, Z. Ren, P. Brochu, Q. Pei, and M. Potkonjak, "Optimizing the configuration and control of a novel human-powered energy harvesting system," in *Power and Timing Modeling, Optimization and Simulation (PATMOS), 2013 23rd International Workshop on*, Sept 2013, pp. 75–82.

[5] V. Goudar and M. Potkonjak, "Power constrained sensor sample selection for improved form factor and lifetime in localized bans," in *Proceedings of the Conference on Wireless Health*, 2012, pp. 5:1–5:8.

[6] V. Goudar, Z. Ren, P. Brochu, M. Potkonjak, and Q. Pei, "Driving low-power wearable systems with an adaptively-controlled foot-strike scavenging platform," in *17th International Symposium on Wearable Computers (ISWC)*, 2013.

[7] P. Pawar, V. Jones, B.-J. F. Van Beijnum, and H. Hermens, "A framework for the comparison of mobile patient monitoring systems," *J. of Biomedical Informatics*, vol. 45, no. 3, pp. 544–556, Jun. 2012.

[8] I. Cox, M. Miller, J. Bloom, J. Fridrich, and T. Kalker, *Digital Watermarking and Steganography*, 2nd ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2008.

[9] H. Noshadi, S. Ahmadian, H. Hagopian, J. Woodbridge, F. Dabiri, N. Amini, M. Sarrafzadeh, and N. Terrafranca, "HERMES - mobile balance and instability assessment system," in *Proc. BIOSIGNALS*, 2010, pp. 264–270.

[10] M. Li, W. Lou, and K. Ren, "Data security and privacy in wireless body area networks," *Wireless Communications, IEEE*, vol. 17, no. 1, pp. 51–58, 2010.

[11] M. Li, S. Yu, J. D. Guttman, W. Lou, and K. Ren, "Secure ad hoc trust initialization and key management in wireless body area networks," *ACM Trans. Sen. Netw.*, vol. 9, no. 2, pp. 18:1–18:35, Apr. 2013.

[12] M. Li, S. Yu, Y. Zheng, K. Ren, and W. Lou, "Scalable and secure sharing of personal health records in cloud computing using attribute-based encryption," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 24, no. 1, pp. 131–143, 2013.

[13] J. H. Lim, A. Zhan, E. Goldschmidt, J. Ko, M. Chang, and A. Terzis, "Healthos: A platform for pervasive health applications," in *Proceedings of the Second ACM Workshop on Mobile Systems, Applications, and Services for HealthCare*, 2012, pp. 4:1–4:6.

[14] A. Prasad, R. Peterson, S. Mare, J. Sorber, K. Paul, and D. Kotz, "Provenance framework for mhealth," in *Communication Systems and Networks (COMSNETS), 2013 Fifth International Conference on*, 2013, pp. 1–6.

[15] X. Kong and R. Feng, "Watermarking medical signals for telemedicine," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 5, no. 3, pp. 195–201, 2001.

[16] A. Ibaida, I. Khalil, and R. van Schyndel, "A low complexity high capacity ecg signal watermark for wearable sensor-net health monitoring system," in *Computing in Cardiology, 2011*, 2011, pp. 393–396.

[17] S. Kaur, O. Farooq, R. Singhal, and B. Ahuja, "Digital watermarking of ecg data for secure wireless commuication," in *Recent Trends in Information, Telecommunication and Computing (ITC), 2010 International Conference on*, 2010, pp. 140–144.

[18] V. Goudar and M. Potkonjak, "A semantically-adaptive strategy for energy-efficiency in wireless medical monitoring devices," in *Sensors, 2013 IEEE*, Nov 2013, pp. 1–4.

[19] A. Wendy and M. Moffa-Trotter, "Functional tools for assessing balance and gait impairments," *Topics in Geriatric Rehabilitation*, vol. 15, no. 1, pp. 66–83, 1999.