

Article

# Wireless Body Area Network Control Policies for Energy-Efficient Health Monitoring

Yair Bar David <sup>1</sup>, Tal Geller <sup>1</sup>, Ilai Bistriz <sup>2,\*</sup>, Irad Ben-Gal <sup>1</sup>, Nicholas Bambos <sup>2,\*</sup> and Evgeni Khmelnitsky <sup>1</sup>

<sup>1</sup> Department of Industrial Engineering Tel Aviv University, Ramat-Aviv, Tel-Aviv, 69978, Israel; yairbd2@gmail.com (Y.B.D.); taljgeller@gmail.com (T.G.); bengal@tauex.tau.ac.il (I.B.G.); xmel@tau.ac.il (E.K.)

<sup>2</sup> Department of Electrical Engineering Stanford University, Stanford, CA 94305, USA

\* Correspondence: bistriz@stanford.edu (I.B.); bambos@stanford.edu (N.B.). Tel.: (+972-36406910, I.B.G.)

**Abstract:** Wireless body area networks (WBANs) have strong potential in the field of health monitoring. However, the energy consumption required for accurate monitoring limits the time between battery charges of the wearable sensors, which is a key performance factor (and can be critical in the case of implantable devices). In this paper, we study the inherent trade-off between the power consumption of the sensors and the probability of misclassifying a patient's health state. We formulate this trade-off as a dynamic problem, in which at each step we can choose to activate a subset of sensors that provide noisy measurements of the patient's health state. We assume that the (unknown) health state follows a Markov chain, so our problem is formulated as a partially observable Markov decision problem (POMDP). We show that all the past measurements can be summarized as a belief state on the true health state of the patient, which allows tackling the POMDP problem as an MDP on the belief state. We then empirically study the performance of a greedy one-step look-ahead policy compared to the optimal policy obtained by solving the dynamic program. For that purpose, we use an open-source Continuous Glucose Monitoring (CGM) data set of 232 patients over six months and extract the transition matrix and sensor accuracies from the data. We find that the greedy policy saves ~50% of the energy costs while reducing the misclassification costs by less than 2% compared to the most accurate policy possible that always activates all sensors. Our sensitivity analysis reveals that the greedy policy remains nearly optimal across different cost parameters and a varying number of sensors. The results also have practical importance, because while the optimal policy is too complicated, a greedy one-step look-ahead policy can be easily implemented in WBAN systems.

**Keywords:** wireless body area networks; controlled sensing; energy efficiency; partially observable Markov decision processes (POMDPs); remote health monitoring

## 1. Introduction

This paper develops and studies control policies for wireless body area networks (WBANs). The operation of WBANs involves activating a set of sensors that are located on or near a patient (e.g., wearable), while collecting data related to his physiological and mental activities. The sensing process is managed by a computing unit, which selects a subset of sensors to be activated at a given time, analyses the obtained measurements, and transmits the information and initial results to a central unit (server) for further analysis and decision making. WBANs operation allows, in particular, estimating the physical and health condition of a patient and monitoring the evolution of these health conditions over time and varying conditions [1], [2], [3], [7].

Recent research and practical developments in sensors and WBAN technology allow efficient implementations of WBANs in modern real-world applications [3], [4], [7]. One of the main factors that impacts the efficiency of such systems in the context of health monitoring is energy consumption as well as the associated trade-off between the amount of required energy to operate the system and the accuracy of health state estimation based

on the collected information [1], [2]. This paper proposes a model and algorithms that account for both factors: the consumption of energy by the system and the level of accuracy that it provides. The optimization techniques in WBANs regularly deal with optimizing the hardware and communication components, such as wireless communication protocols between the sensors and the controlling unit, to reduce energy consumption [8], [9], [17]. Another stream of research has aimed to increase energy efficiency by optimizing the controlling algorithms, which usually concern sensor selection based on the information gained through the system's activities [3], [10]; reducing the volume of the transmitted data in the system [15]; or optimizing the resource allocation in the system [16], [18], [19].

Partially observable Markov decision processes (POMDP) have been used to model the behavior and information transitions in WBANs [2], [5], [7], [11]. A POMDP is defined by a set of states, a set of actions, the conditional transitions between states, a cost function, and a set of observations [6].

Prior work on WBAN control used simplifying assumptions such as allowing only homogeneous or identical sensors and "perfect sensing information" [8], [9], [10]. We extend the POMDP model to allow multiple non-identical noisy sensors, which yield a vector of sensor outputs at each epoch. To convert our POMDP into an MDP, we show that all the information captured in all the past outputs can be summarized using a belief vector, which is a probability vector over the possible health states. This belief state is the state of our MDP.

Our goal is to design an efficient dynamic policy that the controller will use to select a subset of sensors at each epoch, which maintains highly accurate knowledge of the patient's health state while saving as much energy as possible.

By discretizing the continuous belief state, we can numerically approximate the optimal policy for our dynamic program with an arbitrary degree of accuracy. However, computing the optimal policy has a very high computational complexity that renders it impractical. Even if the optimal policy can be computed offline given the patient parameters, it requires storing very large policy tables on the controller, which should be kept as a simple and efficient device. The greedy one-step look-ahead policy is generally a suboptimal compromise in cases where the optimal policy is not practical. Its performance compared to the optimal policy depends on the problem parameters, among them the Markov transition matrix between the health states, the sensor accuracies, and the cost function parameters.

We study the empirical performance of the greedy policy using an open-source Continuous Glucose Monitoring (CGM) dataset, which includes the measurements of 232 patients from different age groups and backgrounds over ~6 months. We estimate the transition matrix and sensor accuracies from the data and then evaluate the greedy policy over these 232 patients. Our empirical study includes sensitivity analysis for the cost parameters and the number of sensors that are external to the data. We also study the performance of the greedy policy using synthetic simulations over a wider range of parameters.

The encouraging results show that while the greedy algorithm is suboptimal in general, it is a very appealing choice for CGM. The greedy policy saves ~50% of the energy costs while reducing the accuracy cost by less than 2% (i.e., reducing the numerical cost by a factor of 0.98, and not the accuracy itself). This suggests that in certain cases the simple and efficient greedy policy can be used for WBAN control in practice while losing very little compared to the optimal intractable policy.

## 2. Health Sensing Model

### 2.1. Patient Health States

We first discuss the health state transitions. We assume a finite and ordered set of patient health states, denoted by  $\mathcal{H} = \{h_1, h_2, \dots, h_j\}$ , such that for each pair  $h_j, h_{j'}$ , where  $j < j'$ ,  $h_j$  is considered less healthy than  $h_{j'}$ . Accordingly,  $h_j$  is considered the healthiest state and  $h_1$  is considered the least healthy state. In practice, the health states could be defined, for example, by the known states of a certain disease or the treatment stages the

patient undergoes. In some applications,  $h_1$  can be considered a terminal health state in which the monitoring of the individual is no longer relevant (e.g., hospitalization, patient mortality, etc.). We assume that the health state is a Markov chain. The transition probability between any two states during two consecutive epochs,  $h^t, h^{t+1} \in \mathcal{H}$ , is given by a transition matrix  $\mathbf{T}$ :

$$\mathbf{T} = \mathbf{T}_{ij} = Pr(h^{t+1} = h_j | h^t = h_i) \quad (1)$$

In practice, the transition probabilities can be estimated using historical clinical data. Therefore, we assume that  $\mathbf{T}$  is known to the controller and is fixed throughout the sensing (monitoring) period.

## 2.2. Sensors

Consider a WBAN that consists of  $N$  sensors operating over a time horizon of  $L$  epochs. Each sensor may be responsible for collecting different types of information (e.g., different vital signs). The information collected by the sensors is used to estimate the patient's health state, which is unknown to the controller at all times. At each decision epoch,  $t$ , the controller chooses to activate some subset of the  $N$  sensors from the action set  $\mathcal{S}$  of all possible activation subsets. Since some combinations of sensors cannot be realized due to physical or physiological limitations, the size of  $\mathcal{S}$  might be smaller than  $2^N$ . Naturally, reducing the size of the set  $\mathcal{S}$  reduces the complexity of the controller's decision-making. The sensor activation status at epoch  $t$  is denoted by  $\mathbf{s}^t = (s_1^t, s_2^t, \dots, s_N^t)$ , where  $s_n^t = 1$  refers to an activated sensor, and  $s_n^t = 0$  refers to a deactivated sensor,  $n = 1, \dots, N$ . The active sensors are selected to optimize a cost function (defined in Section 2.3) that considers both the energy consumption and misclassification costs. For ease of notation, the time index may be omitted when discussing the general time-independent properties.

Let  $l_n^t$  denote the output of sensor  $n$  at time  $t$ ,  $\mathbf{l}^t$  the output vector of all the sensors at  $t$ ,  $L(\mathbf{s}^t)$  the set of all the possible output vectors for a sensor activation vector  $\mathbf{s}^t$  and  $\{\mathbf{l}^{\bar{t}}\}_{\bar{t}=1}^t$  a sequence of sensor outputs during  $t$  epochs. We assume that given a patient's health state, the probability of a certain sensor output is known. For example, given the particular health state of a patient with diabetes, the probability that a blood sugar level sensor will return a certain value may be obtained from a known distribution.

WBAN sensors are typically power-constrained and low-cost. To account for this, we consider detectors that can only sense whether a threshold has been crossed. Hence, each sensor can only provide a binary output of either "1" or "0", i.e.,  $l_n^t \in \{0, 1, \emptyset\}$ , where  $\emptyset$  denotes a deactivated sensor. For example, assume there are three sensors,  $s_1$ ,  $s_2$  and  $s_3$ , and only  $s_1$  and  $s_2$  have been activated, implying  $\mathbf{s} = (1, 1, 0)$ . Then,  $L(\mathbf{s}) = \{(1, 1, \emptyset), (1, 0, \emptyset), (0, 1, \emptyset), (0, 0, \emptyset)\}$ .

We define  $\mathbf{A}^{N \times J}$  as the accuracy matrix. The sensor accuracy  $A_{nj}$  is modeled as the probability of observing an output of "1" from sensor  $n$  given that the individual is currently in health state  $h_j$ :

$$A_{nj} = Pr(l_n = 1 | h_j) \quad \forall n = 1, \dots, N; \forall j = 1, \dots, J \quad (2)$$

Accordingly, the complementary probability of a signal of "0" is  $1 - A_{nj}$ .

**Definition 1.** Define the output function of sensor  $n$  as:

$$f_n(l_n, j) = Pr(l_n | h = h_j) = \begin{cases} A_{nj} & l_n = 1 \\ 1 - A_{nj} & l_n = 0. \\ 1 & l_n = \emptyset \end{cases} \quad (3)$$

We assume independence between the sensor outputs given a patient's true health state. This is a natural assumption since it is only the health state that couples the sensors'

outputs and each is affected by its own local noise. Therefore, we may define the output function of all sensors as:

$$f(\mathbf{l}, j) = Pr(\mathbf{l} | h = h_j) = \prod_{n=1}^N Pr(l_n | h = h_j) = \prod_{n=1}^N f_n(l_n, j) \quad (4)$$

Note that since a deactivated sensor ( $l_n = \emptyset$ ) provides no information, its output function should return 1 in order to neutralize its impact on the probability of the output of the activated sensors.

### 2.3. Power and Misclassification Costs

Our goal is to develop algorithms that account for both the cost of activating a set of sensors, and the misclassification cost associated with possible errors in determining the patient's health state. The cost of activating a set of sensors,  $\mathbf{s}$ , is often straightforward and denoted by  $C(\mathbf{s})$ . The formulation of the misclassification cost is often more complicated and is related to false positive and false negative classification errors. In case where a patient's actual health state is better than the estimated one, the system incurs false positive error cost, which is proportional to the probability of such error. Similarly, false negative error cost is proportional to the probability that the actual health state of the patient is worse than the estimated one. We denote the false positive and false negative cost parameters by  $C_{FP}$  and  $C_{FN}$ , respectively. Since the actual health state is unknown and random, the system may incur both false positive and false negative costs at the same time. The following definition formulates the considered misclassification cost.

**Definition 2.** The misclassification cost per health state,  $h_j$ , denoted by  $\rho_j(\{\mathbf{l}^t\}_{t'=1}^t)$ , is defined as:

$$\rho_j(\{\mathbf{l}^t\}_{t'=1}^t) = C_{FP} \sum_{j'=1}^{j-1} Pr(h^{t+1} = h_{j'} | \{\mathbf{l}^t\}_{t'=1}^t) + C_{FN} \sum_{j'=j+1}^J Pr(h^{t+1} = h_{j'} | \{\mathbf{l}^t\}_{t'=1}^t) \quad (5)$$


and the total estimated misclassification cost is defined as:

$$\rho(\{\mathbf{l}^t\}_{t'=1}^t) = \sum_{j=1}^J Pr(h^{t+1} = h_j | \{\mathbf{l}^t\}_{t'=1}^t) \cdot \rho_j(\{\mathbf{l}^t\}_{t'=1}^t) \quad (6)$$

As indicated above, we multiply the probability that the patient is in a state worse than a specific health state  $h_j^t$  by the false-positive cost parameter  $C_{FP}$ , and we multiply the probability that the patient is in a state better than  $h_j^t$  by the false negative cost parameter  $C_{FN}$ . A more sophisticated cost structure can be assumed. In particular, health state's specific costs and/or dynamic costs can be incorporated in the model suggested in this paper. This will probably lead to the development of more accurate health monitoring algorithms.

### 2.4. Motivating Example

As we show in Section 6, by wisely choosing the subset of sensors to activate, much energy can be saved with only a small reduction in the average misclassification cost compared to activating all sensors. In this subsection, we illustrate this idea given a simple example. We examine a scenario where patient's health condition can be classified into one of three health states,  $\mathcal{H} = \{h_1, h_2, h_3\}$ . We assume that for each epoch  $t$ , the patient may either transit to the neighboring less health state or remain in the current state. For  $h_1$ , the patient may either remain in  $h_1$  or transit to the healthiest state  $h_3$ . The Markov chain is depicted in Figure 1, where the probabilities  $0 < \tau_1, \tau_2, \tau_3 < 1$ .



**Figure 1.** Schematic description of the Markov chain.

We also assume three completely accurate sensors, where each sensor can perfectly detect a particular health state. The sensor accuracy matrix is  $\mathbf{A}^{N \times J} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ .

Given the patient occupying  $h_1$ , the first sensor would output a signal “1”, while the other two sensors would output “0”. Similarly, given the patient occupying  $h_2$  or  $h_3$  the corresponding sensor would output “1”, while the other two would output “0”.

The optimal policy in this scenario only requires activating a single sensor throughout the sensing period (except for at  $t_0$ ) and still provides perfectly accurate sensing. At first, the controller activates all sensors to identify the initial health state of the patient. From there on, at each epoch, the controller only activates the sensor that corresponds to the current health state to identify whether the patient has transitioned to the next health state or remained in the current health state. Compared to the naïve solution of constantly activating all sensors, the optimal solution results in a significant reduction of energy consumption (a single sensor is activated) while preserving perfect knowledge of the patient’s health state. Even with imperfect but still fairly accurate sensors, activating a single sensor in this manner will result in a small misclassification cost while reducing the energy consumption dramatically.

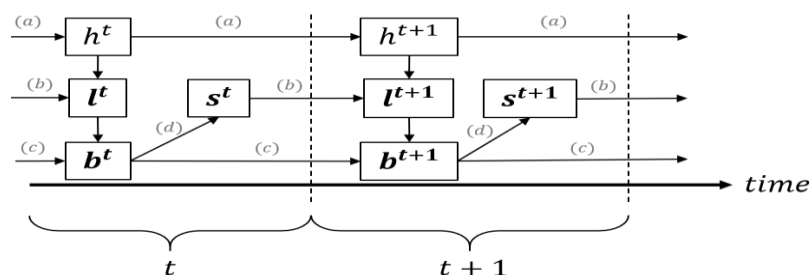
The optimal policy in the above scenario is intuitive and easy to guess. Once the sensor accuracy matrix is not deterministic or the transition matrix has a different structure, computing the optimal solution amounts to solving the POMDP which in general is highly complicated. As done in the literature, the first step to solve the POMDP is to summarize the information captured in the past measurements using a compact belief state on the health state of the patient. This converts the POMDP into a belief state MDP for which we can, in principle, compute the optimal policy by dynamic programming.

### 3. Belief States

The patient’s actual health state is unknown to the controller, and our goal is to monitor and estimate this state. Therefore, the controller uses the sensors’ outputs to produce a probability distribution over the set of health states. This distribution is defined as a belief state.

**Definition 3.** The belief state  $\mathbf{b}^t = (b_1^t, \dots, b_j^t)$  is defined s.t.  $b_j^t = Pr(h^t = h_j | \{\mathbf{l}^{t'}\}_{t'=1}^t)$  for every  $j$ .

The belief state is the probability distribution of the health states of the patient given the past measurements. Next we prove that tracking belief states instead of all past measurements comes with no loss of optimality (i.e., it is a sufficient statistic) and that it can be recursively updated. During the monitoring period, the belief states evolve based on the outputs,  $\mathbf{l}^t$ , received from the activated sensors,  $\mathbf{s}^{t-1}$ . Below is a schematic description of the order of events in two specific subsequent epochs:



**Figure 2.** Order of events throughout sensing epochs.

As shown in Figure 2, health states  $h^t$  evolve according to the Markov chain transition matrix as an independent process (a). The subset of activated sensors  $\mathbf{s}^t$  at each epoch provide outputs  $\mathbf{l}^{t+1}$  during the following epoch  $t + 1$  (b). The activated sensor outputs also depend on the health state. Using the activated sensor outputs  $\mathbf{l}^{t+1}$  and the belief state  $\mathbf{b}^t$ , we update the belief state  $\mathbf{b}^{t+1}$  (c). Finally, given the updated belief state, the controller selects the subset of activated sensors  $\mathbf{s}^{t+1}$  at the next epoch (d).

### 3.1. Markovian Property of Belief States

In this subsection, we show that the past sensor measurements can be summarized as a belief state vector, which is a Markov chain. This converts the POMDP into a belief state MDP. Note that in contrast to a generic POMDP, our cost function is by itself an average taken over the belief states.

**Lemma 1.** The belief state  $\mathbf{b}^t$  is Markovian. The transition probabilities are given by

$$Pr(b_j^{t+1} | \{\mathbf{l}^{t'}\}_{t'=1}^t) = Pr(b_j^{t+1} | \mathbf{b}^t) = \sum_{\mathbf{l} | \mathbf{b}^{t+1} = \tau(\mathbf{b}^t, \mathbf{l})} Pr(\mathbf{l}^{t+1} = \mathbf{l}) \quad (7)$$

where

$$b_j^{t+1} = \tau(\mathbf{b}^t, \mathbf{l}) \triangleq \frac{\sum_i T_{ij} b_i^t f(\mathbf{l}^{t+1}, j)}{\sum_k \sum_m f(\mathbf{l}^{t+1}, k) T_{mk} b_m^t} \quad (8)$$

**Proof:** Observe that

$$\begin{aligned} Pr(h^{t+1} = j | \{\mathbf{l}^{t'}\}_{t'=1}^t) &= \sum_i Pr(h^{t+1} = j | h^t = i, \{\mathbf{l}^{t'}\}_{t'=1}^t) \cdot Pr(h^t = i | \{\mathbf{l}^{t'}\}_{t'=1}^t) \stackrel{(a)}{=} \\ &\sum_i Pr(h^{t+1} = j | h^t = i) \cdot Pr(h^t = i | \{\mathbf{l}^{t'}\}_{t'=1}^t) \triangleq \sum_i T_{ij} b_i^t \end{aligned} \quad (9)$$

where (a) uses the assumption that  $h$  is a Markov chain. In a similar manner, given the belief state  $\mathbf{b}^t$ , the next set of sensor outputs,  $\mathbf{l}^{t+1}$ , is independent of the past sensor outputs:

$$\begin{aligned} Pr(\mathbf{l}^{t+1} | \{\mathbf{l}^{t'}\}_{t'=1}^t) &= \sum_j Pr(\mathbf{l}^{t+1} | \mathbf{l}_{1:t}, h^{t+1} = j) \cdot Pr(h^{t+1} = j | \{\mathbf{l}^{t'}\}_{t'=1}^t) \\ &= \sum_j Pr(\mathbf{l}^{t+1} | h^{t+1} = j) \cdot Pr(h^{t+1} = j | \{\mathbf{l}^{t'}\}_{t'=1}^t) \\ &= \sum_j f(\mathbf{l}^{t+1}, j) \cdot Pr(h^{t+1} = j | \{\mathbf{l}^{t'}\}_{t'=1}^t) = \sum_j \sum_i f(\mathbf{l}^{t+1}, j) T_{ij} b_i^t \end{aligned} \quad (10)$$

This shows that the belief state summarizes all relevant information given past sensor outputs. Furthermore, the belief state can be recursively updated using the following equation:

$$\begin{aligned} b_j^{t+1} &= Pr(h^{t+1} = j | \{\mathbf{l}^{t'}\}_{t'=1}^{t+1}) = \frac{Pr(h^{t+1} = j, \mathbf{l}^{t+1} | \{\mathbf{l}^{t'}\}_{t'=1}^t)}{Pr(\mathbf{l}^{t+1} | \{\mathbf{l}^{t'}\}_{t'=1}^t)} \\ &= \frac{Pr(h^{t+1} = j | \{\mathbf{l}^{t'}\}_{t'=1}^t) \cdot Pr(\mathbf{l}^{t+1} | h^{t+1} = j, \{\mathbf{l}^{t'}\}_{t'=1}^t)}{Pr(\mathbf{l}^{t+1} | \{\mathbf{l}^{t'}\}_{t'=1}^t)} \stackrel{(a)}{=} \\ &\frac{\sum_i T_{ij} b_i^t f(\mathbf{l}^{t+1}, j)}{Pr(\mathbf{l}^{t+1} | \{\mathbf{l}^{t'}\}_{t'=1}^t)} \stackrel{(b)}{=} \frac{\sum_i T_{ij} b_i^t f(\mathbf{l}^{t+1}, j)}{\sum_k \sum_m f(\mathbf{l}^{t+1}, k) T_{mk} b_m^t} \end{aligned} \quad (11)$$

where (a) follows from (9) and, since  $b_j^t = Pr(h^t = j | \{\mathbf{l}^{t'}\}_{t'=1}^t)$  by definition, (b) follows from (10).

Lemma 1 can be straightforwardly applied to simplify the calculation of the misclassification costs:

**Corollary 1.** The misclassification costs can be written as

$$\rho_j(\mathbf{b}^{t+1}) = \rho_j(\mathbf{b}^t, \mathbf{l}) = C_{FP} \sum_{j'=1}^{j-1} b_{j'}^{t+1}(\mathbf{b}^t, \mathbf{l}) + C_{FN} \sum_{j'=j+1}^J b_{j'}^{t+1}(\mathbf{b}^t, \mathbf{l}) \quad (12)$$

so the controller can calculate the total estimated misclassification cost as follows:

$$\rho(\mathbf{b}^t, \mathbf{l}) = \sum_{j=1}^J b_j^{t+1}(\mathbf{b}^t, \mathbf{l}) \cdot \rho_j(\mathbf{b}^t, \mathbf{l}) \quad (13)$$

#### 4. Sensor Activation Control

The Markovity of the belief states allows us to pose our problem as a POMDP with the belief state as the system state. As opposed to tracking all past measurements, the belief state maintains its dimension over time. This enables us to formalize the dynamic programming equations, resulting in a tractable approximation of the optimal policy. In this section, we discuss two control policies that are based on the belief state MDP. The first is an optimal policy (Section 4.1) and the second is a one-step look-ahead greedy policy (Section 4.2). We prove that with accurate sensors that are all activated or deactivated together, the optimal policy can be computed accurately and efficiently (Section 4.3). We denote the minimum total expected cost of the path that starts at  $\mathbf{b}^t$  by the value function,  $V(\mathbf{b}^t)$ , that is, the “cost to go” [13].

##### 4.1. Optimal Policy

The Markov property of the belief states proven in Lemma 1 shows that our problem can be formulated as a POMDP over a finite time horizon of  $L$  epochs. Hence, the optimal solution is given by the following dynamic programming equation:

$$\begin{aligned} V^o(\mathbf{b}^t) &= \min_{\mathbf{s}^t} \left\{ (1 - \omega)C(\mathbf{s}^t) + \int_b \left( Pr(\mathbf{b}^{t+1} = b | \mathbf{b}^t) V^o(\mathbf{b}^{t+1}) + \omega \rho_j(b) \right) db \right\} \\ &\stackrel{(a)}{=} \min_{\mathbf{s}^t} \left\{ (1 - \omega)C(\mathbf{s}^t) + \sum_{\mathbf{l} \in L(\mathbf{s}^t)} \left( Pr(\mathbf{l}^{t+1} = \mathbf{l} | \mathbf{b}^t) \left( V^o(\mathbf{b}^{t+1}(\mathbf{b}^t, \mathbf{l})) + \omega \rho_j(\mathbf{b}^t) \right) \right) \right\} \end{aligned} \quad (14)$$

where (a) follows from (10), which also shows that

$$Pr(\mathbf{b}^{t+1}(\mathbf{b}^t, \mathbf{l}) = b | \mathbf{b}^t) = Pr(\mathbf{l}^{t+1} = \mathbf{l} | \mathbf{b}^t) \quad (15)$$

Here  $\omega$  is the weight the system designer assigns on misclassification costs vs power costs. Lower values of  $\omega$  result in emphasizing power cost reduction while higher values result in minimizing misclassification costs. The default value of  $\omega$  is 0.5.

From (14), we can extract an optimal policy by selecting the  $\mathbf{s}^t$  that minimizes the value function for the belief state, i.e.,  $\underset{\mathbf{s}^t}{\operatorname{argmin}} V^o(\mathbf{b}^t)$ . The equation (14) can be solved using the value iteration method [13].

Our problem can be formulated with an infinite horizon by adding a discount factor to (14). Alternatively, we can assume a terminal health state, which acts as an absorbing state in the Markov chain. In practice, this state occurs when the WBAN health sensing is no longer relevant: for example, when the patient arrives at an emergency room and requires medical intervention, hospitalization, mortality, etc.

##### 4.2. One-Step Look-Ahead Greedy Policy

The greedy algorithm is a one-step look-ahead policy that at each epoch finds which sensors to activate in order to minimize the immediate cost incurred. The value function of this greedy policy is:

$$V^G(\mathbf{b}^t) = \min_{\mathbf{s}^t} \{ (1 - \omega) \cdot C(\mathbf{s}^t) + \omega \cdot \sum_{\mathbf{l} \in L(\mathbf{s}^t)} Pr(\mathbf{l}^{t+1} = \mathbf{l} | \mathbf{b}^t) \cdot \rho(\mathbf{b}^t, \mathbf{l}) \} \quad (16)$$

#### 4.3. Accurate Sensor Use Case

In this subsection, we consider the case of accurate sensors, which can either all be activated or deactivated at a certain epoch. Although in general, sensors are not accurate, it is reasonable to assume that sensors manufactured for tracking health states will be fairly accurate. The action space is  $\mathcal{S} = \{0,1\}$ , such that all sensors are activated if  $a = 1$  and idle if  $a = 0$ . The sensor accuracies are  $A_{nj} \in \{0,1\}$  for all  $n, j$ . In addition, we assume that the sensor accuracy matrix is invertible, which can be achieved by design. The result of these assumptions is that each possible sensor output vector maps to a single health state. We note that when a single mapping exists between each possible sensor output vector and each health state, the problem can be solved regardless of  $A$ 's structure. The following lemma shows that in this case, the number of effective belief states is finite and linear in both the time horizon  $L$  and the number of health states  $J$ . Hence, the dynamic programming problem can be accurately solved.

**Lemma 2.** Let  $\mathcal{S} = \{0,1\}$ , and let the sensors be fully accurate, where  $A_{nj} \in \{0,1\}$  for all  $n, j$ . Assume that  $A$  is invertible. Then, starting from the belief state  $\mathbf{b}^0$ , only  $L \cdot (J + 1)$  belief states are reachable.

*Proof of Lemma 2.* Let  $h^t$  be the health state at time  $t$ , and  $\mathbf{h}^t$  be the standard vector that has one at the  $h^t$ -th component. If the matrix  $A$  is invertible, then we can reconstruct the health state as  $h^t = A^{-1}\mathbf{l}$ . Hence  $f(\mathbf{l}, k) = \mathbf{1}_{\{k=A^{-1}\mathbf{l}\}}$ . For  $j \neq h^{t+1}$ , we have  $b_j^{t+1} = 0$ , and for  $j = h^{t+1}$

$$b_j^{t+1} = \frac{\sum_i T_{ij} b_i^t f(\mathbf{l}^{t+1}, j)}{\sum_k \sum_m f(\mathbf{l}^{t+1}, k) T_{mk} b_m^t} = \frac{\sum_i T_{ij} b_i^t \mathbf{1}_{\{j=A^{-1}\mathbf{l}\}}}{\sum_k \sum_m \mathbf{1}_{\{k=A^{-1}\mathbf{l}\}} T_{mk} b_m^t} = \frac{\sum_i T_{ij} b_i^t}{\sum_m T_{mk} b_m^t} = 1$$

If the sensors are not activated (i.e.,  $a = 0$ ),  $f(\mathbf{l}, k) = 1$  for all  $\mathbf{l}$ , hence,

$$b_j^{t+1} = \frac{\sum_i T_{ij} b_i^t}{\sum_k \sum_m T_{mk} b_m^t} \stackrel{(a)}{=} \sum_i T_{ij} b_i^t \Rightarrow \mathbf{b}^{t+1} = \mathbf{T} \cdot \mathbf{b}^t$$

where (a) follows since the denominator sums over the entire probability space. To summarize:

$$\mathbf{b}^{t+1} = \begin{cases} \mathbf{T} \cdot \mathbf{b}^t & a = 0 \\ \mathbf{h}^{t+1} & a = 1 \end{cases}$$

which means that if all sensors are activated, we can perfectly estimate the health state  $h^t$ . Hence, only belief states  $\mathcal{B}_0 = \{\mathbf{b} | \mathbf{b} = \mathbf{T}^m \cdot \mathbf{b}^0 \text{ or } \mathbf{b} = \mathbf{T}^m \cdot \mathbf{h} \text{ for } m \in \mathbb{N}, \mathbf{h} \in \mathcal{J}\}$  are reachable from  $\mathbf{b}^0$ , where  $\mathcal{J}$  is the set of  $J$  pure health states. The set  $\mathcal{B}_0$  can be computed in advance before solving the dynamic programming equation.

The implication of the above result is that the dynamic programming problem can be accurately solved with complexity  $O(L^2J)$ . To implement this computation, we also exploit the fact that each belief state  $\mathbf{b}^t$  can transition into two possible belief states:  $\mathbf{h}^{t+1}$  and  $\mathbf{T} \cdot \mathbf{b}^t$ , where  $\mathbf{T}$  is the transition matrix.



In the general case of inaccurate sensors, the update function of  $\mathbf{b}^{t+1}$  is not quadratic or even convex. Considering the lack of a convenient structure for the problem, one needs to compute  $V(\mathbf{b}^{t+1})$  for every belief  $\mathbf{b}^{t+1}$  in the  $J$ -dimensional simplex. In practice, this means that only an approximate solution can be obtained by discretizing the belief space. Therefore, the proof of Lemma 2 implies that the general scenario with inaccurate sensors or a complicated action space  $\mathcal{S}$  can only be solved numerically.

#### 4.4. Belief State Discretization

The dynamic programming formulation in (14) explores all the possible future states over a finite period, either using a discount factor or a terminal health state. In practice, a continuous belief state results in a very large state space, which results in exponentially growing computation. In this section we present a numerical procedure, which discretizes the belief state space and, accordingly, approximates the value function,  $V^o$ . The procedure first defines a finite set of valid belief state vectors. The vectors can be related to different medical treatments or protocols, as mapped by medical professionals. Then, given an actual belief state vector, the proposed procedure finds the vector from the finite set which is closest to the actual one. Finally, the procedure calculates the value function for each valid belief state vector. The steps of the procedure are the following:

1. Set a level of discretization and define the set  $B$ , which contains all valid belief state vectors.
2. Given an actual vector,  $\mathbf{b}$ , calculate its distance  $\sum_j |\tilde{b}_j - b_j|$  to each valid vector,  $\tilde{\mathbf{b}} \in B$
3. Return the vector,  $\tilde{\mathbf{b}}^*$  which is the closest to  $\mathbf{b}$ .

For example, if the level of discretization is selected to be 0.2, and  $\mathbf{b} = [0.17, 0.35, 0.40, 0.08]$ , then the above procedure returns  $\tilde{\mathbf{b}}^* = [0.20, 0.40, 0.40, 0.00]$  as the closest valid vector. After discretizing the state space, we use (14) and calculate  $V^o(\tilde{\mathbf{b}})$  for each valid state  $\tilde{\mathbf{b}} \in B$ .

The fine discretization required to guarantee a good approximation of the optimal policy leads to high dimensional belief space  $B$  (as we demonstrate numerically in Subsection 6.2). As a result, implementing the optimal policy is complicated since computational load grows exponentially with the dimension of  $B$ . Even solving the dynamic programming offline does not avoid the high complexity since it requires instead storing large policy tables that map the discretized states to actions. This complexity can be somewhat mitigated by using more advanced discretization methods.

A more sophisticated method for improving the discretization is to create a set  $B$ , where each  $\mathbf{b} \in B$  is selected so that the discretized belief state space reflects the actual belief states that patients tend to follow. This approach can boost performance without needing to increase the discretized belief state space  $B$  by employing a non-uniform discretization. We leave this method for further research.

To avoid discretization, function approximation can be employed, which enables solving the POMDP over a continuous state space [14], although still with high complexity. Since our focus is to study the performance of the greedy policy, we use the simple discretization method described above to provide a “very close to optimal” benchmark.

## 5. Empirical Analysis on Glucose Data

### 5.1. Use-Case Description


We now demonstrate the implementation of the WBAN controlling solution in a real-world use-case. We use an open-source data set of continuous glucose measurements collected from a group of diabetes type-1 patients [20]. The measurements were recorded using a wireless glucose level sensor placed on the patients for extended periods. Using our proposed greedy control policy, we aim to reduce energy consumption while maintaining high monitoring accuracies, thus extending the sensors’ lifespans and allowing longer sensing periods. The model parameters (sensor accuracies, transition probabilities,

health state trajectories) were extracted from the data, allowing us to evaluate the effectiveness of the greedy policy using real-life glucose measurements, that represent a potential application.

### 5.2. The Data

The raw data contains glucose measurements for a group of 232 heterogenic patients from all ages with type 1 diabetes and glycated hemoglobin (HbA1c) 7.0% – 10.0. For most patients, the measurements were collected over 6 months using three different wireless real-time continuous glucose monitoring (RT-CGM) systems.

Figure 3 displays the glucose level measurements over time for 3 randomly selected patients over a period of 12 hours. For the majority of patients, the time gap between subsequent measurements is 5 minutes. However, about 16.9% of the records have abnormal time gaps, which could be caused due to malfunctions in the measurements collection, a patient which has removed the sensor, etc. In order to overcome the difficulties when dealing with unequal gaps, we have checked what is the most common time gap per patient, and set it as the patient's time frame for measurements collection.



**Figure 3.** Glucose levels of randomly selected patients over time.

### 5.3. Data Modelling

#### Health States

In order to apply the POMDP based model, we first define the set of health states each patient can occupy. To this end, we discretize the range of glucose measurements by dividing the entire range of glucose values into 3 smaller ranges, each defined as a health state. The defined ranges are based on medical standards [21]. We note that the defined range structure can be designed according to varying medical needs. Specifically,  $h_3 \leq 100$ ,  $100 < h_2 \leq 180$  and  $180 < h_1$ .

#### Sensors and Actions

The glucose data was measured using a single sensor, which was activated every 5 minutes (defined as a single epoch). We define a decision time-frame as  $k$  consecutive epochs, for which the controller must select a sample plan, i.e. during which epochs is the sensor activated. We assume that during the decision time-frame, the sample plan remains fixed. In the POMDP based model, the set of actions the controller may take contains a set of acceptable combinations of activated sensors. Since the data was obtained using a single sensor per each patient, we define the set of actions as the number of epochs during the decision time frame in which the sensor will be activated. This definition substitutes the assumption that the system consists of multiple sensors and allows us to define multiple actions. We note that the actions defined above could be described as selecting the sensor's frequency throughout the decision epoch or activating sensors with different sampling frequencies. We also note that the set of actions will not necessarily consist of all possible epoch combinations throughout the decision time frame  $k$ . For example, given  $k = 6$ , the decision time-frame will consist of 6 epochs (30 minutes). We may define the set of actions as  $S = \{(4), (2, 6), (1, 3, 5)\}$ , where each tuple contains the index of epochs for which the sensor will carry out a measurement. The first tuple can be expressed as setting

the sensor to a low frequency setting, whereas the last may be expressed as a high frequency setting. Intuitively, the higher the sensing frequency, more energy is consumed in creating a final measurement for the  $k$  epochs. The collected measurements are then aggregated to a single observation using an aggregation function, and discretized according to the states defined in the previous sections to provide sensor outputs  $l \in \{1, 2, 3\}$ . The aggregation function may be defined, for example, as an average or maximum over the collected records.

This serves as the output obtained by the sensor.

In principle, more advanced aggregation functions could be designed to average all observation throughout the decision epoch, unless there exists an observation which is relatively far from the other observations throughout the decision epoch. This allows to design the system so that extreme observations are handled differently than decision epochs where the observation values are relatively close to each other, promising that extreme glucose measurements are not missed due to the averaging of observations in the decision epoch. For example, this could be done by considering the number of STDs each observation is from the average of all observations in the decision epoch (a process known as SPC). If, for a certain observation, the absolute distance in STDs from the average crosses a predefined threshold, the observation would be considered extreme, thus solely defining the patient's health state in that decision epoch (unlike the average of all observations). The higher the number of STDs, the more tolerant the controller will be to high value observations.

We note that  $k$  can be selected by the system designer according to varying system requirements. Smaller values of  $k$  may provide better monitoring performance, due to the fact that controller decision will be made more frequently. On the other hand, small values would require higher computational resources.

#### Transition Matrix

Using a portion of the patient's measurements, the transition matrix can be estimated based on actual health state transitions. The conditional probability  $T_{ij}$ , is calculated as the ratio between the number of transitions between  $h_i$  and  $h_j$  and the number of transitions from  $h_i$  to any of the other health states, over the selected portion. The calculated transition matrix may be used for future monitoring.

#### Sensor Accuracies

The sensor accuracy matrix is also calculated using a portion of the patient's data. For each sensor  $n$  we calculate the probabilities to obtain a certain output, given the patient is in a certain health state. Since each sensor may provide an output  $l \in \{1, 2, 3\}$ , the sensor accuracy matrix per sensor  $n$  will be of the form  $A_n^{3 \times 3}$ . For example, given the patient is currently occupying  $h_1$ , the probability to obtain an output value  $l = 2$  from sensor  $n = 3$  is  $(A_3)_{1,2}$ .

In order to calculate the probability to receive a specific output  $l$  given the patient is in health state  $j$ , we calculate the ratio between the number of epochs where the patient was in health state  $j$  and the sensor returned an output  $l$ , divided by the number of epochs the patient was in health state  $j$ .

#### 5.4. Model Parameter Extraction

Initially, we have divided the data for each patient to train and test sets to 75% and 25% of the records, respectively. Using the training set, we first extract the transition matrix and sensor accuracies per patient. We then use all parameters needed for the POMDP based model (set of health states  $H$ , transition matrix  $T$ , sensor accuracies matrix  $A$ , set of possible actions  $S$ , cost parameters  $C_{FN}$ ,  $C_{FP}$ ) to apply a policy over the test set data. Finally, we compare the overall performance of the policy to a baseline. The baseline is defined as an inclusive policy - the sensor is activated at all epochs. The collected measurements are then aggregated using the aggregation function, thus creating a baseline health state prediction, for each  $k$  epochs. To compare between policies, we will examine

the Total Cost, the power cost and the health state classification accuracy (calculated using the most likely health state, according to the belief state vector).

We set the following parameters:  $k = 6$ ,  $C_{FN} = 100$ ,  $C_{FP} = 50$  and  $S = \{(4), (2,6), (1,3,5)\}$ , where the last one is the set of possible actions (the epoch indices in which the sensor is activated).

The activation costs are defined to be the number of epochs where the original sensor is activated during the decision time frame of  $k$  epochs. For example, if the sensor is activated for 3 epochs out of a 6-epoch decision time frame, the activation cost will be 3.


### 5.5. Policy Comparison

Table 2 compares the overall costs of the greedy policy and the naïve (inclusive) solution, averaged over all patients:

**Table 1.** Activation costs and accuracy results.

|                   | Activation Costs | Accuracy      |
|-------------------|------------------|---------------|
| Naïve             | 100%             | 85.1%         |
| Greedy            | 48.33%           | 83.5%         |
| <b>Difference</b> | <b>-51.67%</b>   | <b>-1.88%</b> |

We observe that the activation costs were reduced by 51.67%, whereas the accuracy was reduced by a small factor of 1.88%, resulting in an accuracy score of 83.5%.




**Figure 4.** : Activation and misclassification costs obtained by the different proposed policies.

Figure 4 compares the performance of value iteration, greedy, and inclusive. Clearly, the inclusive policy inflicts higher activation and misclassification costs, compared to the greedy and value iteration solutions. In addition, the value iteration solution obtained results similar to the greedy solution. This is probably caused due to the discretization of the belief states and the fact that the sensor accuracies are high enough for the greedy solution to obtain a near optimal solution. For the scenario presented here and the defined system parameters, the greedy solution provides good performance.

### 5.6. Sensitivity Analysis

Our results from Subsection 5.5 suggest that the greedy policy is nearly optimal for the scenario we extracted from the real Glucose data. To establish this statement, we now conduct a sensitivity analysis of two main system parameters. As our benchmark, we use the discretized dynamic programming solution with 15 discretization levels, since our results in Subsection 6.2 show that the performance barely improves with more than 15 discretization levels. We start by varying the parameter  $\omega$  between 0 and 1, thus changing the weight of the two cost components in equation (16). The graph below demonstrates

the cost ratio between the greedy and value iteration methods. Each point considers the total costs obtained over all patients, and the error bars show the standard deviation of this ratio over the patients.




**Figure 5.** Ratio between total costs obtained by greedy and value iteration methods obtained over all.

patients for varying  $\omega$  values, for the following set parameters:  $k = 6$ ,  $C_{FN} = 100$ ,  $C_{FP} = 50$


One can observe that the performance (total cost) of the greedy policy is never far behind the optimal policy. For the extreme case of  $\omega = 1$ , in which the controller will aim at minimizing misclassification costs only, the greedy method obtains costs of approximately 13.5% higher, compared to the value iteration method. However, for medical applications, the regime of interest is when the misclassification cost is more important than the energy cost. This follows since very small misclassification probabilities are required in practice. When  $\omega \leq 0.5$ , the cost of the greedy policy is never more than 2% more than that of the optimal policy. Hence, the greedy policy is nearly optimal for the practical regime of interest in WBAN health monitoring systems in this case.

We next demonstrate the trade-off between the misclassification cost and the activation cost in the Glucose use-case, by varying the parameter  $\omega$  between 0 and 1, thus changing the weight of the two cost components in equation (16). We note that although the policy is obtained given the weighting parameter, the calculated costs reflect the true costs obtained when applying the solution to the test set. Lower values of  $\omega$  result in emphasizing power cost reduction while higher values result in minimizing misclassification costs. The obtained data points have a high fit of  $R^2 = 0.997$  to  $f(x) = \frac{145.09}{276.96x+27.65} - 0.07$ . This high fit in Figure 6 suggests that  $f(x)$  can be used by the system designer to adjust  $\omega$  for the Glucose monitoring case.



**Figure 6.** Activation costs vs. misclassification costs per  $\omega$ , averaged over all patients.

Another main system parameter is the length of the decision epoch  $k$ . Higher values of  $k$  require lower computational resources but will introduce more noise into the system, caused by the aggregation of more raw observations. Figure 7 shows that the greedy policy is nearly optimal for a wide ranges of values of  $k$ .



**Figure 7.** Ratio between total costs obtained by greedy and value iteration methods obtained over all patients for  $k$  values, for the following set parameters:  $\omega = 0.5$ ,  $C_{FN} = 100$ ,  $C_{FP} = 50$

## 6. Empirical analysis and sensitivity analysis on Synthetic DATA


In this section, we provide simple numerical examples that will allow us to study the behavior of the greedy policy compared to the optimal policy in more detail, and to provide more intuition and interpretability for the proposed greedy policy. For the glucose data, the transition matrix and sensor accuracies were estimated for the data so they do not represent a degree of freedom. Here all parameters are arbitrary, allowing for a more general sensitivity analysis.

### 6.1. WBAN Dynamics for the Greedy Policy

We define the model parameters as follows:  $J = 4$  (the number of health states),  $N = 5$  (the number of sensors). Additional parameters are as follows:


- $\mathbf{T} = \begin{bmatrix} .90 & .04 & .04 & .02 \\ .02 & .90 & .04 & .04 \\ .01 & .01 & .90 & .08 \\ .01 & .01 & .03 & .95 \end{bmatrix}$  is the transition matrix, from state  $i$  (row) to state  $j$  (column)
- $\mathbf{A} = \begin{bmatrix} .99 & .01 & .50 & .50 \\ .50 & .90 & .10 & .50 \\ .50 & .50 & .90 & .10 \\ .90 & .10 & .50 & .50 \\ .60 & .40 & .60 & .50 \end{bmatrix}$  is the sensor accuracy matrix.
- $C_s = [5 \ 1 \ 1 \ 1 \ 0.2]$  denotes the sensors' activation costs,
- $C_m = [50 \ 100]$  denotes the misclassification costs ( $FP, FN$ )

The transition matrix  $\mathbf{T}$  was chosen based on the reasonable assumption that the probability of remaining in a certain health state is larger than the probability of transitioning between health states. In practice, the transition probabilities can be estimated using historical clinical data like we did for the Glucose data in Section 5. The cost parameters were selected to reflect a clear trade-off between the misclassification costs and the sensors' activation costs. In practice, the sensor activation costs can be based on the cost of recharging the sensors, and the misclassification costs can be estimated based on the cost of the medical care needed due to health state misclassification.



**Figure 8.** The unknown health state and the most likely health state (using the greedy policy) over 20 epochs.

We simulate the greedy algorithm described in Section 4.2 to derive a greedy policy. The most likely state shown in Figure 8 is generated by  $h^{ML} = \text{argmax}_j(\mathbf{b}_j^t)$ . One can observe that the greedy policy generally provides accurate predictions of the actual health states in the proposed example.



**Figure 9.** The number of sensors activated throughout the simulation using the greedy policy.

Figure 9 shows the number of sensors activated throughout the simulation. One can observe that the largest number of sensors is usually activated during health state transitions, i.e., in periods where the patient's health is relatively unstable. In addition, Table 1 summarizes the dynamics of the activated sensors regarding their different accuracies and costs. For example, even though sensor 5 incurs low costs, it is almost never activated due to the low accuracy it provides. Sensors 1 and 4 are informative regarding  $h_1$  and  $h_2$ . Although sensor 1 is very accurate compared to sensor 4, it is activated during a smaller number of epochs, which follows due to the lower costs incurred by sensor 4.

**Table 2.** Percent of epochs during which each sensor was activated throughout the simulation.

| Sensor Usage |     |     |     |    |
|--------------|-----|-----|-----|----|
| 1            | 2   | 3   | 4   | 5  |
| 28%          | 64% | 61% | 69% | 3% |

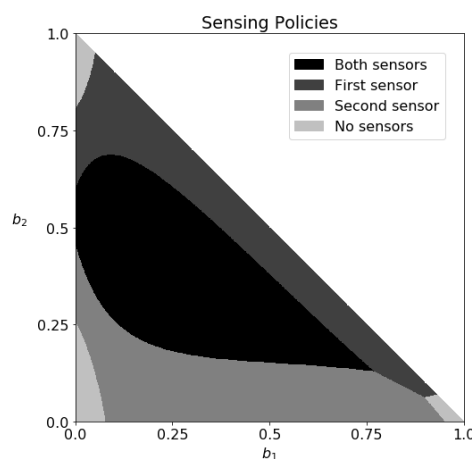
### 6.2. Comparison of the Greedy Policy to the Optimal Policy

We now compare the proposed greedy policy with the optimal dynamic programming approaches, presented in Sections 4.2 and 4.1, respectively.

For this section, we define a more compact problem, which will allow deeper analysis. The parameters are defined as follows:  $J = 3$ ,  $N = 2$ . Additional parameters are as follows:

- $T = \begin{bmatrix} .90 & .06 & .04 \\ .04 & .90 & .06 \\ .01 & .04 & .95 \end{bmatrix}$  and  $A = \begin{bmatrix} .90 & .10 & .50 \\ .50 & .90 & .10 \end{bmatrix}$
- $C_s = [2 \ 2]$  and  $C_m = [17 \ 34]$

The selected policy gives the number of sensors activated given a certain belief state vector.




**Figure 10.** Greedy sensing policy over the belief state space  $b$ . Since  $b_3 = 1 - (b_1 + b_2)$ , the belief space forms a simplex and may be described by only the belief over the first two states.

Figure 10 maps the belief states into the subset of activated sensors. We observe that the controller tends to activate both sensors in the mid-ranges of the belief states, i.e., in areas where the information concerning the patient's active health state is not definitive. As the belief states move towards the edges of the graph, fewer sensors are needed because there is higher certainty concerning the patient's active health state. In the far corners of the graph, no sensors are activated. We note that the distinction between the use of the first and second sensors, when activating only one sensor, is related to each sensor's ability to differentiate between a different pair of health states.

We now examine the different behaviors displayed by the proposed greedy policy and the dynamic programming policy. For this purpose, we use the average total costs obtained by both solutions over 30 random transition matrices as the main performance indicator. We note that all transition matrices share the property that the probabilities of remaining in certain health states are generally higher than those of transitioning to other health states. As discussed in Section 4.4, the optimal dynamic programming solution requires belief state discretization. Higher discretization resolutions will allow more accurate health state monitoring but will incur an exponentially higher computational load.

Figure 11 demonstrates the effect of varying the discretization resolution on the performance of the dynamic programming solution. The  $x$ -axis shows the "binning" resolution defined for the discretized belief state space. Higher discretization levels lead to higher accuracy when approximating actual belief states. We can see that for lower values of discretization levels, the greedy policy outperforms the dynamic programming formulation, while for higher levels, the value iteration solution outperforms the greedy solution and converges to an average total cost that is approximately 3% lower than that of the greedy solution. Overall, the greedy solution shows a maximum gap of less than 4% from the dynamic programming solution for high discretization levels.



**Figure 11.** Performance of the value iteration and greedy policies as a function of the discretization level, averaged over 100 i.i.d. Monte Carlo simulations.

We observe also that there is a downward trend up to a discretization level of 15, where the trend plateaus. From this graph, we can see that a discretization level of 15 (meaning each element in the belief state vector could be approximated by a factor of  $\frac{1}{15}$ ) provides sufficient performance while maintaining a lower computational load compared to larger discretization levels. Hence, in the rest of the simulations, we will use 15 as the discretization level.


### 6.3. Sensor Accuracy Sensitivity Analysis



We now analyze the effect of the sensor accuracies on the policies' performances. Generally, using sensors with poor accuracies introduces more noise into the systems that may hinder the performance. We define a group of accuracy matrices using a parameter  $\epsilon$  as follows:

$$A = \begin{bmatrix} 1 - \epsilon & \epsilon & \epsilon \\ \epsilon & 1 - \epsilon & \epsilon \end{bmatrix}, 0 \leq \epsilon \leq 0.5$$

According to the above definition, each of the two sensors differentiates between a pair of health states. Larger values of  $\epsilon$  result in weaker differentiation between the pairs of health states.



**Figure 12.** Total costs of the greedy policy and the value iteration policy over different sensor accuracy matrices  $A$  defined by  $0 \leq \epsilon \leq 0.3$  (step size 0.05), averaged over 100 i.i.d. Monte Carlo simulations.

As expected, higher values of  $\epsilon$  result in higher costs for all policies, whereas near-zero values provide low costs down to the point where the sensors are perfect, which results in identical low costs for all the policies (see Figure 12). For  $\epsilon$  values higher than 0.3, the sensors yield very low-accuracy outputs, causing the controller to deactivate the sensors throughout the sensing period. For the majority of examined  $\epsilon$  values, the proposed greedy policy was nearly optimal. Only for very inaccurate sensors with  $\epsilon \geq 0.25$  the total cost of the greedy policy is higher by more than 5% than that of the (approximately) optimal value iteration. This is encouraging since medical sensors are typically required to be highly accurate in practice, due to the possible implications.


#### 6.4. Transition Matrix Sensitivity Analysis

Finally, we demonstrate the effect of stability in the patient's health state trajectories on the policies' performances. Less stable health state trajectories, in which the patient tends to change health states often, will be reflected by a more equally distributed transition matrix. For that purpose, we define a group of transition matrices using a parameter  $\tilde{T}$ :

$$T = \begin{bmatrix} \tilde{T} & \frac{1 - \tilde{T}}{2} & \frac{1 - \tilde{T}}{2} \\ \frac{1 - \tilde{T}}{2} & \tilde{T} & \frac{1 - \tilde{T}}{2} \\ \frac{1 - \tilde{T}}{2} & \frac{1 - \tilde{T}}{2} & \tilde{T} \end{bmatrix}, 0 \leq \tilde{T} \leq 1$$

Note that  $\tilde{T}$  determines the stability of the health state trajectories. For lower values of  $\tilde{T}$ , the obtained transition matrix is more uniform. For the value  $\tilde{T} = 0.3$ , the policies yield the highest possible costs, since then  $T$  provides no information concerning the patient's potential health state trajectory. For higher values of  $\tilde{T}$ , the transition matrix  $T$

causes higher-stability health state trajectories. For the majority of  $\tilde{T}$  values examined, the greedy policy is nearly optimal, showing the largest gap when the total costs of both policies decline for higher values of  $\tilde{T}$ . For the most extreme value of  $\tilde{T} = 1.0$  (not shown in the left graph in Figure 13), for which the health state is static, the total cost of the greedy policy is 450% compared to the value iteration policy.



**Figure 13.** Total costs of the greedy policy and the value iteration policy over different transition matrices  $T$ , defined by  $0.3 \leq \tilde{T} \leq 1$  (step size 0.05), averaged over 100 i.i.d. Monte Carlo simulations.

## 7. Conclusions

We studied the energy consumption versus accuracy trade-off in WBAN health monitoring systems. In our setting, the controller decides which subset of sensors to activate at every epoch. The objective is to monitor the unknown patient's health state while minimizing the power and misclassification costs over time.

We studied the performance of the simple one-step look-ahead approach for sensor activation. This simple policy is much easier to implement and offers better interpretability than the optimal one based on dynamic programming. Furthermore, since the optimal policy works with a continuous belief state space, it must be discretized which degrades the performance and requires storing large policy tables.

Our extensive empirical study on Glucose levels of 232 patients over 6 months reveals that the proposed greedy policy saves 50% of the energy cost while losing only 1.8% of the misclassification cost compared to the most accurate policy possible that activates all sensors at all times. Our sensitivity analysis confirms that the greedy policy is nearly optimal as long as the weight of the misclassification cost is at least twice that of the energy cost. Since in medical application the misclassification probabilities must be very low, this weight regime includes the practical regime we would aim to use in practice.

Our encouraging findings suggest that the greedy policy can be an appealing policy to use in practice. Our sensitivity analysis suggests that the reason the greedy policy is so effective is that, as we estimate from the data, the health state transitions are not extremely slow and the sensor accuracies are not extremely bad.

Since this is typically the case for various medical applications, studying the performance of the greedy policy on other medical datasets will further establish the greedy policy as a practical health monitoring control for WBAN systems.

Furthermore, additional cost structures may be considered in improving the applicability and real-world value of the solutions. For example, costs that dynamically change over time or different misclassification costs for more severe health states could be used. Additional extensions include relaxing model assumptions, such as binary sensors and known POMDP parameters (transition matrix, patient health states, etc.). For example,

assuming the transition matrix is unknown to the controller, sensor outputs from a group of patients may be used to initially estimate a generalized transition matrix.

**Author Contributions:** Conceptualization, N.B., I.B.G and E.K.; methodology, all authors.; software, Y.B.D. and T.G.; validation, Y.B.D., T.G. and I.B.; formal analysis, all authors; investigation, Y.B. D., T.G. and I.B.; resources, N.B. and I.B.G.; data curation, Y.B. D., T.G.; writing—original draft preparation, Y.B. D., T.G.; writing—review and editing, I.B., N.B., I.B.G. and E.K.; visualization, Y.B. D., T.G.; supervision, N.B., I.B.G. and E.K.; project administration, N.B., I.B.G. and E.K.; funding acquisition, N.B. and I.B.G. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by the Koret Foundation grant for Smart Cities and Digital Living 2030.

**Acknowledgments:** The authors would like to acknowledge Dr. Dan Feldman and the Koret Foundation for their continuous support.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. Y. Bar David, T. Geller, A. Ward, D. Miller, E. Khmelnsky, N. Bambos and I. Ben-Gal, "Optimal Health Monitoring via Wireless Body Area Networks," Proceedings of the 57th IEEE Conference on Decision and Control (CDC), pp. 6800-6805, Miami Beach, FL, 2018.
2. D. Miller, Z. Zhou, N. Bambos and I. Ben-Gal, "Optimal Sensing for Patient Health Monitoring," 2018 IEEE International Conference on Communications (ICC), Kansas City, MO, 2018, pp. 1-7.
3. M. Chen, S. Gonzalez, A. Vasilakos, H. Cao, and V. C. M. Leung, "Body Area Networks: A Survey," Mobile Networks and Applications, vol. 16, no. 2, pp. 171-193, 2010.
4. M. Ghamari, B. Janko, R. Sherratt, W. Harwin, R. Piechockic, and C. Soltanpur, "A Survey on Wireless Body Area Networks for eHealthcare Systems in Residential Environments," Sensors, vol. 16, no. 6, pp. 831, 2016.
5. A. Panangadan, S.M. Ali and A. Talukder, "Markov decision processes for control of a sensor network-based health monitoring system," Proceedings of the National Conference on Artificial Intelligence, Vol. 20, No. 3, pp. 1529. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.
6. V. Krishnamurthy, Partially Observed Markov Decision Processes. Cambridge University Press, 2016.
7. Y. O. Mohammed and U.A Baroudi, "Partially observable Markov decision processes (POMDPs) and wireless body area networks (WBAN)", KSII Transactions on Internet and Information Systems, Vol. 7, No. 5, pp. 1036-1057, 2013.
8. E. Jovanov, "A survey of power efficient technologies for Wireless Body Area Networks", Proceedings of the 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Vancouver, BC, pp. 3628-3628, 2008.
9. J. Xing and Y. Zhu, "A Survey on Body Area Network", Proceedings of the 5th International Conference on Wireless Communications, Networking and Mobile Computing, Beijing, pp. 1-4, 2009.
10. D. Tian and N. D. Georganas. "A coverage-preserving node scheduling scheme for large wireless sensor networks", Proceedings of the 1st ACM International Workshop on wireless sensor networks and applications, pp. 32-41, Atlanta, Georgia, 2002.
11. X. Fei, A. Boukerche and R. Yu, "A POMDP Based K-Coverage Dynamic Scheduling Protocol for Wireless Sensor Networks", IEEE Global Telecommunications Conference GLOBECOM, Miami, FL, pp. 1-5, 2010.
12. D. E. Spackman, S. Kadiyala, P. J. Neumann, D. L. Veenstra, and S. D. Sullivan, "Measuring Alzheimer Disease Progression with Transition Probabilities: Estimates from NACC-UDS", Current Alzheimer Research, vol. 9, no. 9, pp. 1050-1058, Jan. 2012.
13. D. P. Bertsekas, Dynamic programming and optimal control, vols. 1, 2. Belmont: Athena Scientific, 1995.
14. A. Koppel, G. Warnell, E. Stump, P. Stone, and A. Ribeiro, "Policy Evaluation in Continuous MDPs with Efficient Kernelized Gradient Temporal Difference", IEEE Transactions on Automatic Control, 2017.
15. S. Sasirekha and S. Swamynathan, "Cluster-chain mobile agent routing algorithm for efficient data aggregation in wireless sensor network", Journal of Communications and Networks, vol. 19, no. 4, pp. 392-401, 2017.
16. H. Kim, G. P. Villardi and J. Ma, "Energy efficient dynamic resource allocation for wireless systems using receiver puncturing technique", Journal of Communications and Networks, vol. 20, no. 4, pp. 366-373, 2018.
17. L. Miao, J. Mao and C. G. Cassandras, "Optimal Energy-Efficient Downlink Transmission Scheduling for Real-Time Wireless Networks", IEEE Transactions on Control of Network Systems, vol. 4, no. 4, pp. 692-706, 2017.
18. C. G. Cassandras, T. Wang and S. Pourazarm, "Optimal Routing and Energy Allocation for Lifetime Maximization of Wireless Sensor Networks With Nonideal Batteries", IEEE Transactions on Control of Network Systems, vol. 1, no. 1, pp. 86-98, 2014.
19. X. Ren, K. H. Johansson, D. Shi and L. Shi, "Quickest Change Detection in Adaptive Censoring Sensor Networks", IEEE Transactions on Control of Network Systems, vol. 5, no. 1, pp. 239-250, 2018.

20. JDRF CGM study group. Jdrf randomized clinical trial to assess the efficacy of realtime continuous glucose monitoring in the management of type 1 diabetes: research design and methods. *Diabetes technology & therapeutics*, vol. 10, pp. 310–321, 2008.
21. Standards of medical care in diabetes-2006. *Diabetes Care*, 29 (suppl 1):s4–s42, 2006. ISSN 0149-5992. <https://care.diabetesjournals.org/content/29/suppl1=s4>: