Sustainable Wearables for Health Applications and Beyond via Uncertainty-Aware Energy Management

Dina Hussein, **Chibuike E. Ugwu***, **Ganapati Bhat**, **Janardhan Rao Doppa**
Washington State University
{dina.hussein, chibuike.ugwu, ganapati.bhat, jana.doppa}@wsu.edu

Abstract

Achieving good health and well-being through lower mortality rates of non-communicable diseases and early warning of health risks are key goals of United Nations (UN). Wearable internet of things (IoT) are one of the most promising technology to achieve these goals through their ubiquitous monitoring of key health indicators and in-situ data processing. However, small form-factor of wearable devices constrains the battery capacity, thus requiring frequent recharging or battery replacements, which lowers their adoption rate and benefits. Augmentation of battery energy by scavenging ambient sources, such as light, is a promising solution to improve operating lifetime of IoT devices. However, ambient energy sources are highly uncertain, making energy management (EM) challenging. To handle these challenges, this paper presents a novel uncertainty-aware EM approach. First, we develop a conformal prediction-based method for future energy harvest (EH) that provides small uncertainty regions with provable coverage guarantees (true output vector is within the region). The EH uncertainty regions are then leveraged in an EM algorithm that uses overhead-aware sampling to evaluate the quality of multiple decisions with varying EH before making a decision using a lightweight machine learning model. Experiments on two diverse real-world datasets with 10 users show that conformal prediction achieves more than 90% coverage with tight prediction intervals; and the EM algorithm produces decisions that are, on average, within 2 Joules of an optimal Oracle.

1 Introduction

The United Nations (UN) good health and well-being social development goal (SDG) aims to reduce premature mortality from non-communicable diseases (goal 3.4) and strengthen the capacity for early warning and health risk reduction (goal 3.D). Achieving these goals will require development of low-cost and pervasive technologies that can monitor key health parameters and provide real-time analytics to users [WHO, 2022; Latif *et al.*, 2018; Amu *et al.*, 2023; Hussein *et al.*, 2022a; Hussein *et al.*, 2025]. Wearable Internet of Things (IoT) devices offer great potential to achieve the SDGs by enabling continuous monitoring of symptoms in free-living environments and provide early warning on health risks [Hussein *et al.*, 2024a]. Wearable devices can also help in chronic disease management through monitoring of symptoms [Atzori *et al.*, 2010; Espay *et al.*, 2016; Daneault, 2018; Hussein *et al.*, 2025]. However, widespread adoption of wearable devices in health settings has been limited due to *small battery capacities and the need for frequent recharging* [Hussein *et al.*, 2022b; Mamun *et al.*, 2022].

Energy harvesting and management are promising technologies to improve the self-sustainability of IoT devices [Kansal *et al.*, 2007; Odema *et al.*, 2021; Tuncel *et al.*, 2020]. Indeed, usage of renewable energy sources aligns with the goal of ensuring access to affordable, reliable, sustainable and modern energy SDG 7.1. Reducing dependency on the grid for recharging of wearable devices will lead to wider adoption, thus improving health outcomes. It can also lower the cost of maintenance by minimizing battery replacements. However, ambient energy sources are *highly uncertain* with wide variations across locations, time, seasons, and user behavior [Kansal *et al.*, 2007]. Accounting for EH variations at runtime is critical to ensure device sustainability.

Attaining self-sustainability in IoT devices requires accurate forecasting of future energy availability and effective management of the device's energy resources. Predicting future energy supply is essential to capture the variability of energy sources within the system. Energy management (EM) algorithms must incorporate the uncertainty in future energy to support optimal decision-making, thereby promoting self-sustainability and improving application performance [Hussein *et al.*, 2024b]. To address these needs, this paper presents novel techniques for energy prediction with *formally guaranteed uncertainty bounds* and *uncertainty-aware* EM.

An effective energy harvest (EH) predictor must meet two essential requirements: it should deliver high prediction accuracy relative to the ground-truth, and it should generate uncertainty bounds for each prediction that ensure both high coverage (the true value falls within the interval) and narrow prediction intervals (tight uncertainty bounds). To fulfill these requirements by design, we introduce an upper-bound calibrated multi-target conformal prediction (UC-MTCP) approach aimed at producing compact uncertainty regions for

EH prediction vectors while maintaining coverage guarantees. The core concept of UC-MTCP is to extend the idea of single-target conformal prediction [Romano *et al.*, 2019] to multiple target variables, each corresponding to a future EH interval, and apply an upper bound correction step during calibration to construct tight uncertainty regions. We provide theoretical guarantees for coverage of UC-MTCP.

Forecasting future EH alone is not adequate to ensure the sustainable operation of IoT devices. This is due to the necessity of integrating both the predictions and their associated uncertainty bounds into the EM algorithm to determine the appropriate energy allocation or budget for each decision interval. To address this, we design a constrained optimization problem that derives energy allocation bounds, aiming to fulfill application requirements, maximize application performance, and preserve energy sustainability, while incorporating uncertainty in EH predictions.

The EM algorithm employs overhead-aware sampling to evaluate multiple EM decisions in each decision interval. The sampling process is designed to span the entire uncertainty region without surpassing the overhead budget allocated for EM. Finally, EM decisions must be combined to obtain a single energy budget that closely approximates the optimal decisions of an Oracle with access to ground-truth future EH values. To achieve this, we employ a lightweight machine learning model that takes candidate decisions from the sampling step and outputs a decision that aims to mimic the Oracle's performance for each time horizon.

We assess the effectiveness of the proposed approach using EH and activity data drawn from two diverse real-world datasets [Alemdar *et al.*, 2013; Sztyler *et al.*, 2016]. To generate ground truth EH data, activity information is combined with solar irradiation data under both outdoor and indoor conditions. The evaluation demonstrates that our proposed approach yields decisions that are on average within 2 Joules (J) of those made by an optimal Oracle. End-to-end implementation on an IoT device shows that our approaches consume about 88 mJ in the worst-case scenario with lower energy in typical scenarios. Compared to baselines, the proposed approaches also achieve at least 25% improvement in the utility.

Deployment Plan: Our immediate next step is to deploy the wearable IoT devices in real-world settings. The first step will validate the EH prediction and management approaches in outdoor settings that have lower degree of uncertainty. Once validated in outdoor conditions, the wearable devices will be deployed in health settings, such as movement disorder monitoring, to validate the proposed approach with user data. These validations will be the next step for deployment of the EM and health algorithms across the world.

In summary, this paper makes the following contributions:

- A novel upper-bound calibrated multi-target conformal prediction (UC-MTCP) method, enabling reliable uncertainty quantification and guaranteed coverage across multiple future energy harvesting (EH) time steps. To our knowledge, this is the first application of UC-MTCP with such upper-bound calibration for EH forecasting.
- A novel energy management (EM) strategy that incorporates uncertainty in EH predictions to make informed, optimized runtime decisions.
- Extensive experiments on two real-world activity datasets featuring both indoor and outdoor EH scenarios. Results demonstrate that our method outperforms existing baselines by at least 25% in application quality.

2 Problem Setup and Related Work

Problem Setup: We focus on a wearable IoT system composed of multiple sensors and harvesters, as illustrated in Figure 1. We assume that the device integrates a rechargeable battery for energy supply. Additionally, we consider that wearable devices in this study utilize EH from body motion and ambient light as a means to supplement the battery.

Next, we consider an EM framework where decisions are made over a fixed time horizon of one day, with each day partitioned into T equal-length time intervals [Huynh *et al.*, 2008]. The objective is to develop an EH prediction function f that, given a set of input features X, produces *multiple* EH predictions spanning intervals t through $t + H$, where H is the number of future time intervals. Beyond generating these predictions, the function f must also output reliable uncertainty bounds that satisfy user-specified coverage guarantees at runtime. The purpose of the EM policy π is to take EH predictions from f as input and determine an energy budget for interval t that enhances application performance while ensuring energy sustainability through energy-neutral operation (ENO) [Kansal *et al.*, 2007]. In doing so, the policy π must incorporate uncertainty by strategically sampling from within the uncertainty region when making decisions.

Related Work: Precise prediction of future EH plays a critical role in enabling optimal energy allocation. Prior research has explored both analytical and machine learning techniques to tackle this challenge [Yamin and Bhat, 2023; Cammarano *et al.*, 2012; Piorno *et al.*, 2009; Wan *et al.*, 2011; Tuncel *et al.*, 2020], with an emphasis on generating single-point predictions of future EH. Although these methods offer estimations, they lack uncertainty quantification (UQ), which is a crucial requirement for reliable EM decisions in wearable IoT systems, particularly within health applications.

Studies have also investigated the use of ensemble methods to generate reliable prediction intervals by leveraging both the mean (μ) and variance (σ) estimators [Roy and Larocque, 2020; Shrestha and Solomatine, 2006]. While these methods aim to form effective prediction intervals, they frequently result in overly conservative bounds and often fall short of ensuring the required coverage, which is an essential aspect for practical EM applications. Consequently, there is still a pressing need for approaches that can deliver compact uncertainty regions while offering provable coverage guarantees.

We propose leveraging the conformal prediction (CP) framework [Vovk *et al.*, 2005] to address the shortcomings of existing EH prediction approaches, particularly in the area of UQ. CP offers a distribution-free UQ framework for constructing prediction intervals that provide finite-sample, distribution-free coverage for any given predictive model [Vovk *et al.*, 2005; Romano *et al.*, 2019; Romano *et al.*, 2020; Tibshirani *et al.*, 2019; Regression, 2017; Ghosh *et al.*, 2023a; Ghosh *et al.*, 2023b; Shi *et al.*, 2024; Shahrokhi *et al.*, 2025].

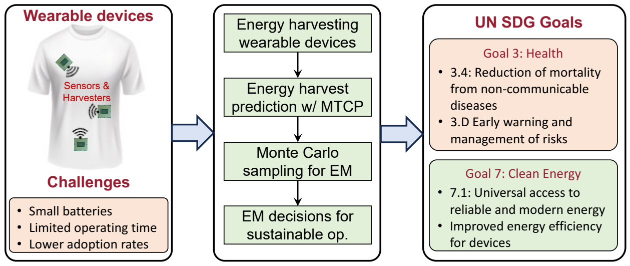

Figure 1: Overview of our uncertainty-aware EM approach with alignment of UN goals.

CP constructs prediction intervals by calculating a nonconformity score between predictions and ground truth, evaluating these scores on calibration data, and using a quantile threshold derived from a specified error rate (say 5%). Driven by the need for compact uncertainty regions in EH predictions, this work explores the under-studied application of CP in the context of multi-target regression tasks.

EM for wearable IoT devices has been extensively explored in prior work [Kansal *et al.*, 2007; Vigorito *et al.*, 2007; Basaklar *et al.*, 2022; Geisler *et al.*, 2017; Bhat *et al.*, 2017]. Existing methods for EM include linear programming, model predictive control, and reinforcement learning. The EM algorithms aim to maximize device performance while maintaining ENO. However, these methods overlook the uncertainty in future EH during decision-making, making them ill-equipped to adapt to abrupt changes in EH patterns (e.g., transitions from cloudy to sunny days). The proposed work addresses the limitations of prior approaches by reasoning about uncertainty in multiple future decision intervals, resulting in enhanced robustness compared to baseline strategies.

3 Uncertainty Quantification for EH

This section provides a brief background on conformal prediction (CP) followed by our proposed use-inspired CP approach for uncertainty quantification in EH with guarantees.

3.1 Background on Conformal Prediction

Conformal Prediction (CP) [Romano *et al.*, 2019; Vovk *et al.*, 2005; Romano *et al.*, 2020; Gibbs *et al.*, 2023; Tibshirani *et al.*, 2019] is a general framework for uncertainty quantification that provides rigorous coverage guarantees for any black-box predictive model. Given a pre-trained predictive model f, n new calibration examples with p input features and a univariate target $\{(X_i, Y_i)\}_{i=1}^n$ where each feature vector X_i belongs to the input space $\mathcal{X} \subseteq \mathbb{R}^p$, and each target variable Y_i belongs to the output space $\mathcal{Y} \subseteq \mathbb{R}$ (which were not used during training), the goal is to construct a prediction set for a new test input X_{n+1} such that the corresponding target variable Y_{n+1} is contained within the set with a user-specified confidence level α (say 5%). The CP framework requires a non-conformity scoring function $S : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ that quantifies the discrepancy between the predicted value $f(X_i)$ and the corresponding true output Y_i (e.g., absolute residual) $S(X_i, Y_i) = |Y_i - f(X_i)|$. Instead of producing a single-point prediction $f(X_{n+1})$, CP constructs a prediction interval/region $\mathcal{R}(X_{n+1}) \subseteq \mathbb{R}$ that confidently ensures that the true output Y_{n+1} falls within it with a probability of at least $1 - \alpha$ on average, as shown in Equation 1, providing a marginal measure of reliability for the model's predictions.

$$\mathbb{P}[Y_{n+1} \in \mathcal{R}(X_{n+1})] \geq 1 - \alpha. \tag{1}$$

For the tasks that involve predicting multiple continuous target variables simultaneously – also known as *multi-target regression* – the target variable is extended to $Y_i \in \mathbb{R}^d$, where d is the number of targets. Instead of constructing a single prediction interval, CP now constructs a d-dimensional prediction region $\mathcal{R}(X_i) \subseteq \mathbb{R}^d$, ensuring that the true response vector Y_i is contained within this region with high probability. This is the focus of this work as we need to quantify the uncertainty for EH prediction for horizon $H > 1$.

3.2 Upper-bound Calibrated Multi-Target CP

We propose a lightweight upper-bound calibrated multi-target conformal prediction (UC-MTCP) approach designed to construct uncertainty regions for EH with guaranteed theoretical and empirical coverage. Our approach has minimal computational overhead which is a critical need for resource-constrained devices.

Given a user-specified error rate, α, and a prediction horizon H, UC-MTCP employs the Conformalized Quantile Regression (CQR) [Romano *et al.*, 2019] method to construct base prediction intervals for each target horizon $h \in \{1, \ldots, H\}$. The adjusted error rate for each target is defined as $\beta = \alpha/H$. The prediction interval for each target is constructed as:

$$\mathcal{R}_h(X) = [L, U] = [\hat{q}^h_{\beta/2}(X), \hat{q}^h_{1-\beta/2}(X)],$$

where $\hat{q}^h_{\beta/2}(X)$ and $\hat{q}^h_{1-\beta/2}(X)$ represent the lower and upper quantile estimates for the h-th target variable, respectively.

Beyond achieving target coverage, it is important to have an upper bound that is as low as possible to avoid excessive energy allocation. This is particularly important when the actual energy harvest is significantly lower than the predicted upper bound. By adjusting the upper bound, we can mitigate the risk of energy failure caused by an overestimated allocation relative to the true energy harvest.

Motivated by this practical need, we incorporate an *excess-aware upper bound calibration step*, which regulates the base upper bound by exploiting overly conservative single-target horizon predictions where the empirical coverage significantly exceeds the target coverage level. To achieve this, we evaluate the base model $\mathcal{R}_h$ on a new upper-bound calibration dataset $\mathcal{D}_{cal2} = \{(X_i, Y_i)\}$, (which was not used for training and base CP calibration). For each target h and error rate β, the process begins by computing standard CP intervals. This involves constructing $[L(X_i), U(X_i)]$ for all $X_i \in \mathcal{D}_{cal2}$ using $\mathcal{R}_h$ and identifying valid cases set $\mathcal{D}_{valid}$

$$\mathcal{D}_{valid} = \{(X_j, Y_j) : L(X_j) \leq Y_j \leq U(X_j) \, \forall \, j \in \mathcal{D}_{cal2}\}$$

For these valid cases, we compute the absolute difference between the ground truth EH and upper bound as

$$E_j = |Y_j - U(X_j)| \textbf{ for all } j \in \mathcal{D}_{valid}$$

Next, the empirical coverage for the upper-bound calibration dataset $\mathcal{D}_{cal2}$ is determined as

$$\hat{C}_{UB} = \frac{\sum_{i=1}^{|\mathcal{D}_{cal2}|} \mathbf{1}\left[L(X_i) \leq Y_i \leq U(X_i)\right]}{|\mathcal{D}_{cal2}|}$$

If the empirical coverage exceeds the desired level (i.e. $\hat{C}_{UB} > (1-\beta)$), the excess is computed as

$$\texttt{excess} = \max(0, \hat{C}_{UB} - (1-\beta))$$

An adjustment factor (AF) is then derived as the "excess"-th percentile of the computed errors

$$AF = \text{percentile}(\{E_j\}_{j=1}^{|\mathcal{D}_{valid}|}, \texttt{excess} \times 100)$$

Finally, for a new test input X_{test} with a base prediction interval $[L(X_{\text{test}}), U(X_{\text{test}})]$, the adjusted prediction interval $\mathcal{R}_h = [L(X_{\text{test}}), U'(X_{\text{test}})] : U'(X_{\text{test}}) = U(X_{\text{test}}) - AF$.

Given the domain knowledge of the non-negativity constraint of EH (energy is not negative), we further reduce the uncertainty regions. Specifically, the prediction interval for each target variable is adjusted to fall in the positive region:

$$\mathcal{R}_h^{+}(X_{\text{test}}) = \left[\max\{0, L(X_{\text{test}})\}, \max\{0, U'(X_{\text{test}})\}\right],$$

This constraint leads to tighter uncertainty regions while maintaining the required coverage levels. The final constrained H-dimensional uncertainty region becomes: $\mathcal{R}^{+}(X_{\text{test}}) = \prod_{h=1}^{H} \mathcal{R}_h^{+}(X_{\text{test}})$

Coverage Guarantee of UC-MTCP: The following theorem provides the marginal coverage guarantee for our UC-MTCP approach: the true EH values lie within the predicted uncertainty region with a high probability of at least $(1-\alpha)$. Proof is presented in the Appendix.

THEOREM 1. *Given that all targets $h \in \{1, \cdots, H\}$ preserve the base single target CP coverage at error rate $\beta = \alpha/H$, then for a new test input X_{test}, the uncertainty regions from UC-MTCP $\mathcal{R}^{+}(X_{test})$ covers the true target $Y_{test} \in \mathbb{R}^H$ with a probability of at least $1 - \beta H$, where $\beta H = \alpha$*

$$\mathbb{P}[Y_{\text{test}} \in \mathcal{R}^{+}(X_{\text{test}})] \geq 1 - \alpha \tag{2}$$

4 Uncertainty-Aware Energy Management

4.1 Problem Formulation

The EM problem uses the following formulation introduced in [Kansal *et al.*, 2007] to achieve energy-neutral operation:

$$\max.\ Q(\mathbf{E_A}) = \sum_{t=0}^{T-1} \beta^t \ln\left(\frac{E_A^t}{M_E}\right) \quad \text{s. t.} \tag{3}$$

$$E_B^{t+1} = E_B^t + \eta\mathcal{E}_H^t - E_A^t - E_o,\ 0 \leq t \leq T-1 \tag{4}$$

$$E_B^{t+1} \geq E_{min} \quad 0 \leq t \leq T-1 \tag{5}$$

$$E_B^T \geq E_{target} \tag{6}$$

Objective: The goal of the EM problem is to determine the energy allocation (budget) E_A^t for each interval t such that the total quality of service Q is maximized. The quality Q is captured by the sum of logarithms of energy allocation E_A^t scaled by a constant factor M_E at each interval. The scaling factor ensures that the quality is positive only when the device has sufficient energy allocated to achieve a minimum level of performance through sampling and data processing.

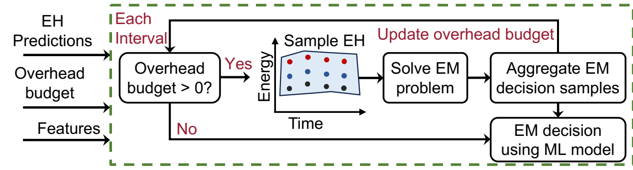

Figure 2: Overhead-aware sampling for energy management.

Constraints: The EM problem is constrained by the battery dynamics through EH and minimum battery levels. Specifically, the first constraint in Equation 4 captures the device battery dynamics whereby the battery level at the beginning of any interval $t+1$ is given by the battery level in t, stochastic EH $\mathcal{E}_H^t$ with efficiency η, energy decision E_A^t in interval t, and EM overhead E_o. The next constraint specifies that the device battery level must always be greater than or equal to E_{min} to account for any emergency situations, such as falls. Finally, Equation 6 ensures ENO operation for the device by constraining the battery level at the end of a horizon(E_B^T) to be greater than a target value of E_{target}.

Solution Approach: Concave logarithmic objective and linear constraints result in a convex optimization problem to achieve ENO in wearable devices [Boyd and Vandenberghe, 2004]. Common approaches to solve the problem include interior-point methods [Boyd and Vandenberghe, 2004] or with iterative solvers. It is useful to leverage iterative solvers in wearable devices since they can be fine-tuned to balance accuracy and decision-making overhead. Consequently, we propose to utilize the iterative gradient projection (IGP) algorithm [Karakoç *et al.*, 2020] to solve the EM problem.

Even with an iterative algorithm, it is not practical to obtain optimal EM decisions at runtime due to lack of knowledge and uncertainty in future EH. However, an upper bound on the quality of EM decisions can be obtained by solving the EM problem *offline* with actual EH values. The optimal solutions are used as an Oracle policy to evaluate the effectiveness of real-world EM policies that use estimates of future EH.

4.2 Overhead-Aware Sampling of EH Uncertainty

The EM algorithm on the wearable device must carefully handle the uncertainty in EH so that the EM decisions maximize the application QoS while maintaining ENO. Algorithm 1 and Figure 2 show the overhead-aware sampling to handle the uncertainty in future EH. The algorithm takes EH predictions and uncertainty regions from UC-MTCP, EM budget, and current battery level as inputs at the beginning of each interval t. The uncertainty region must be sampled uniformly to ensure that the EM decisions are not skewed by a small region of the EH predictions. The EM algorithm must also balance the granularity of sampling with the overhead of performing EM evaluation. Therefore, the EM algorithm starts with an energy budget for decision-making and samples the uncertainty region randomly until the budget is used.

Each H-dimensional sample from the uncertainty region provides prediction of future EH for intervals t to $t+H$.

Algorithm 1: Overhead-Aware Sampling for EM Decision-Making

```
Input: UC-MTCP EH uncertainty region for intervals t to t + H, Overhead budget E_o, Features X and Battery E_B^t
M_{E_A} ← Initialize decision matrix
while E_o ≥ 0 do
    [Ê_H^t, ..., Ê_H^{t+H}] ← Sample from uncertainty region
    [Ẽ_A^t, ..., Ẽ_T^{T-1}] ← Solve EM problem with EH samples
    M_{E_A} ← [M_{E_A}; [Ẽ_A^t, ..., Ẽ_T^{T-1}]]
    E_o ← E_o − E_IGP
end
Use M_{E_A} and X as input for ML model to obtain EM decision
return EM decision
```

These predictions are used in the EM problem to obtain potential energy allocations for intervals t to $T-1$. The decisions are then stored in a matrix so that all potential decisions can be used to make the final EM decision. The EH sampling and EM evaluation are executed until the energy budget for EH sampling is exhausted. Finally, the matrix with candidate EM decisions from the EH sampling is provided as input to the proposed ML model for making the final EM decision.

4.3 Lightweight ML Model for Oracle Tracking

The EM algorithm running on wearable IoT devices must make an overall energy allocation decision using the individual candidate decisions as a result of the EH sampling. A naive approach is to randomly choose a decision from the overhead-aware sampling. However, random decisions do not adequately account for the uncertainty in EH or error with respect to the Oracle. Therefore, the proposed approach trains a lightweight ML model that aims to follow Oracle decisions.

Input Features: As described above, the overhead-aware sampling of EH provides potential EM decisions for intervals t to $T-1$. In particular, different levels of EH uncertainty are considered in each row of the decision matrix, thus they can be used as a key feature for decision-making. We utilize the average of EM decisions as features since the sampling may differ in each interval due to changes in overhead budget. The ML model also utilizes the features used to UC-MTCP in EM decision-making since they provide insight into past EH values. Overall, the ML model utilizes EH prediction features X and $\hat{E}_A^{t+1} - \hat{E}_A^{t+H}$, where $\hat{E}_A^{t+1}$ represents the average of potential EM decisions as features in EM decision-making.

Training Target: The goal of the ML model is to obtain decisions that are close to the Oracle without knowledge of future EH. Consequently, we utilize the Oracle decisions as the training target for the ML model. The Oracle is obtained offline using actual EH values in the EM problem formulation.

Model Structure: We note that EM decisions can be obtained with any supervised learning algorithm. Fully connected neural networks are used in this work due to their low overhead and strong performance in diverse scenarios.

5 Experimental Results

5.1 Experimental Setup

Wearable IoT Device Model: Our device model centers on a wearable IoT device that includes with five energy harvesters and various sensors. It incorporates piezoelectric sensors positioned at four locations on the knees and elbows to harvest motion energy. Furthermore, the device has an SP3-37 [FlexSolarCells, 2013] flexible photovoltaic (PV) cell for light EH.

Datasets: We leverage activity data from two publicly available datasets: 1) the ARAS dataset [Alemdar *et al.*, 2013] and 2) the Mannheim dataset [Sztyler *et al.*, 2016] for evaluation in both indoor and outdoor environments. Activity data for four users from two houses over 30 days are included in the ARAS dataset, while data from six users for two weeks are included in the Mannheim dataset.

The user activity data for both datasets are limited in duration to perform a comprehensive evaluation of the proposed approach in different conditions. Therefore, both datasets are extended to cover a six-year period (2015–2020) by shuffling and augmenting the original activity data. Typical EH levels from motion and ambient light irradiance are combined with the activity data to obtain ground truth EH values for each interval [Tuncel *et al.*, 2020; Andreas and Stoffel, 1981].

Data Splitting: The energy dataset is divided into three sets: training (2015), validation (2016), and testing (2017–2020).

Proposed CP Parameters: The UC-MTCP configuration includes the base CQR prediction model architecture, prediction horizon H, and error rate $\alpha = 0.1$ corresponding to a 0.9 expected marginal coverage level. The neural network model is composed of three layers, each containing 64 units. For each prediction horizon $H \in \{2, 3, 4, 5\}$, and each target $h \in \{1, ..., H\}$, we set the error rate to $\beta = 0.1/H$.

Proposed EM Parameters: The EM parameters consist of the decision intervals, battery constraints, and ML model structure. The EM horizon covers a full 24-hour period, segmented into 24 decision intervals corresponding to each hour of the day. The battery's energy goal is defined as 100 J, with a minimum energy threshold of 10 J. The framework employs a discount factor γ of 0.99. The QoS ME is set to eight. The ML model is implemented as a fully connected neural network comprising three hidden layers, each with 32 neurons.

Evaluation Metrics: The UC-MTCP framework is assessed based on the coverage and the size or area of the resulting uncertainty region. Subsequently, we analyze the effectiveness of individual EM decisions by benchmarking them against the Oracle. The evaluation metric for both the proposed EM methods is the mean absolute error (MAE) relative to the Oracle. Lastly, the cumulative quality of service Q over the horizon is used as a measure of overall performance.

5.2 Baseline Methods for Comparison

Baseline Point Prediction Method: EM algorithms typically use point predictions of EH to guide decisions through iterative algorithms. Although various point prediction techniques can be applied, this approach adopts a simple and computationally efficient strategy by using the average of actual EH values from the past three days as the point forecast.

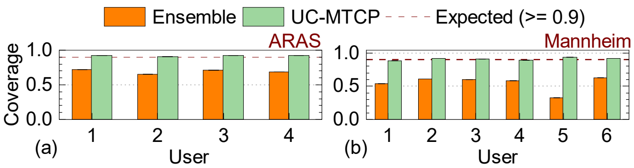

Figure 3: Comparison of coverage for Ensemble and UC-MTCP uncertainty regions for ARAS and Mannheim datasets across all users in Horizon 4. Similar results are observed for all other prediction horizons for both datasets.

Baseline uncertainty region with Ensemble Method: The ensemble approach employs a set of random forests to generate EH predictions. These individual forecasts are aggregated to construct prediction intervals based on the mean (μ) and variance (σ) of predictions. Specifically, the interval bounds are calculated as: $(\mu \pm \sigma)$. The overall uncertainty region across all prediction horizons is then determined by taking the H-dimensional intersection of the prediction intervals.

Random Choice Energy Decision (UC-MTCP-Random): A simple baseline approach involves selecting the final EM decision at random from the range of possible allocations generated through EH sampling. Accordingly, UC-MTCP-Random acts as a reference method by randomly picking an energy allocation within the sampled range. This baseline is used to emphasize the advantages of integrating ML-based decision-making in the proposed framework.

5.3 Evaluation of the UC-MTCP Approach

An essential component of forecasting future EH is the ability to effectively handle prediction errors and quantification of the uncertainty. Ensuring that the achieved coverage aligns with or surpasses a user-defined threshold is crucial for dependable and precise energy allocation decisions. To this end, we assess the coverage performance across various prediction horizons H. Figure 3 presents the coverage results for both the proposed UC-MTCP framework and the ensemble-based method. For a consistent evaluation, both approaches enforce the constraint that uncertainty regions remain non-negative, reflecting the inherently non-negative nature of EH values. As seen in Figure 3, the ensemble method falls short of the target 0.9 coverage level, whereas UC-MTCP consistently meets this threshold across all users and datasets, demonstrating the robustness and reliability of the proposed solution.

5.4 Evaluation of the Proposed EM Approach

This section analyzes benefits provided by accounting for uncertainty using UC-MTCP in comparison to baseline approaches. We first compare the proposed EM against point prediction and random choice methods. Then, we perform additional comparisons against point prediction since most prior EM methods use point predictions for EM.

Comparison of Error with Different EM Approaches: Wearable device must provide reliable operation to the user across multiple years as symptoms of a health condition progress. Moreover, EH patterns fluctuate over the year due to seasonal variations. Hence, it is important to demonstrate that our method remains effective regardless of the time of year or specific year of deployment. To evaluate this, we examine the accuracy of EM decisions across different years in Figures 4 and 5. These figures depict the average MAE along with the standard deviation for each year of evaluation for ARAS and Mannheim datasets, respectively. The results reveal that EM decisions based on UC-MTCP consistently outperform all baseline methods across every year analyzed, highlighting its ability to adapt to evolving EH patterns. This contrasts sharply with other methods that do not integrate UC-MTCP into the EM process. Overall, the findings emphasize the robustness of UC-MTCP in delivering accurate forecasts and demonstrate the value of incorporating uncertainty-aware predictions into EM decision-making.

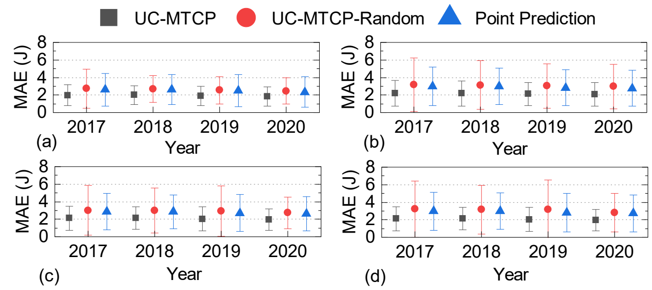

Figure 4: Yearly Mean Absolute Error (Mean and standard deviation) of UC-MTCP, UC-MTCP-Random, and Point Prediction with different users in ARAS dataset.

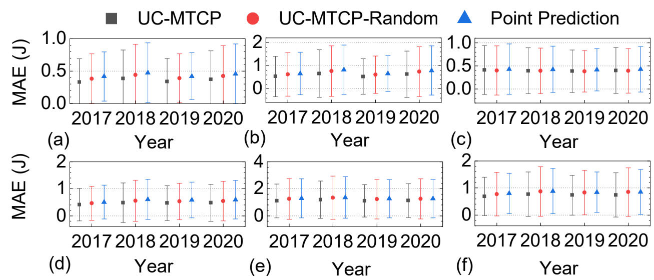

Figure 5: Yearly Mean Absolute Error (Mean and standard deviation) of UC-MTCP, UC-MTCP-Random, and Point Prediction with different users in Mannheim dataset.

Energy Management with Point Predictions: Most, if not all, EM approaches utilize point predictions to make decisions. We evaluate the performance of the proposed method against the point prediction approach in Figure 6. This figure presents a comparison of the average energy allocation and the absolute value of the average cumulative QoS for both methods, using data from user one in the ARAS and Mannheim datasets. Each bar corresponds to a specific EH range, enabling analysis across scenarios with varying energy availability – from low to high EH conditions. The results reveal significant differences in energy efficiency and QoS between the two approaches. Notably, the proposed method achieves more consistent QoS across all percentiles compared to point prediction. In contrast, the point prediction approach often results in lower utility, even when consuming more energy. This inefficiency arises from its inability to handle uncertainty, leading to over-allocation and, ultimately, battery depletion. These findings underscore the critical role of incorporating uncertainty in EH for EM decisions.

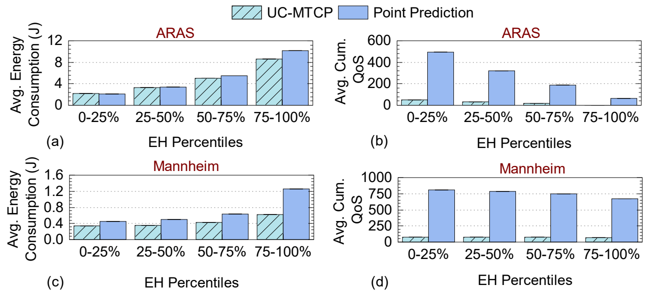

Figure 6: Average energy consumption and absolute of average cumulative utility of UC-MTCP and point prediction with user one in ARAS and Mannheim datasets. Lower cumulative values lead to better quality since the logarithms are negative.

5.5 Implementation Overhead

We utilize the TI-CC2652R microcontroller [Texas Instruments Inc., 2018] to evaluate the energy consumption and execution time of the proposed algorithms: UC-MTCP, the IGP algorithm, and the ML model. Both UC-MTCP and ML models are implemented using multi-layer perceptrons. Experimental results indicate that UC-MTCP requires about 58 ms for execution, 0.6 mJ of energy per operation. The IGP algorithm imposes a higher computational burden due to its gradient-based approach. Each IGP evaluation takes around 300 ms and consumes 8.6 mJ of energy. The ML model takes 188 ms for a single inference while consuming 2.0 mJ energy.

These overheads are minor relative to the operational timelines and energy budgets typical of IoT systems. All algorithmic overheads are accounted for in the optimization formulation given in Equation 4. In summary, the overhead of these algorithms remains within feasible limits for deployment on resource-constrained IoT devices, supporting their practical applicability in real-world scenarios.

6 Social Impact and Deployment Plan

The proposed research will lead to significant social impact in healthcare and energy sustainability across the world, especially in underdeveloped parts. The social impact will be achieved through the key metrics defined in the following.

Societal Health Benefits: Prior studies have shown that chronic diseases such as Parkinson's disease or type 2 diabetes could be better managed in remote populations if patients had access to devices that provide continuous monitoring of health parameters [Mattison *et al.*, 2022]. However, lack of sustainable energy and frequent battery recharging lead to unreliable data [Phillips *et al.*, 2018; Espay *et al.*, 2016; Maetzler *et al.*, 2016]. The proposed approach precisely addresses this by enabling self-sustainable operation.

We also analyze the number of potential patients impacted by chronic diseases in underdeveloped parts of the world, as highlighted by UN SDG goals [Amu *et al.*, 2023]. The study in [WHO, 2022] estimates more than 47 million people in Africa will suffer from chronic diseases like diabetes by 2045. Recent studies have said that adverse outcomes in patients can be better managed with wearable technology and data analytics [Maha *et al.*, 2024; Latif *et al.*, 2018; WHO, 2022]. The proposed approaches will make a direct impact on this population by enabling continuous monitoring of symptoms with health algorithms on wearable devices.

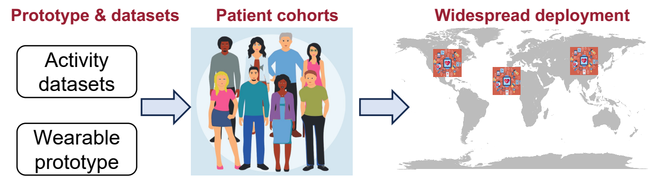

Figure 7: Deployment plan for the proposed energy management approach for health monitoring with wearable devices.

Energy Benefits: Widespread adoption of wearable devices will increase demands of energy from the grid. While the energy needs of each device are small, cumulative needs of millions of devices can lead to increased demands from the grid. To this end, the EH from body motion and ambient light will greatly reduce energy needs and improve sustainability.

Assuming that each device has a battery of 100 J and we deploy 50 million devices in the field, energy savings of 50 J per device per day through ambient sources will lead to cumulative energy savings of 21,000 units of energy per month. This energy can be used to potentially power 1400 homes in Eastern Africa for a month while reducing costs and carbon emissions significantly [Dagnachew *et al.*, 2023].

6.1 Deployment Plan

The current study focuses on employing datasets for EH and validating a proof-of-concept on a wearable device prototype, as shown in Figure 7. Utilizing prior datasets helps us in evaluating the proposed approaches before deploying in real-world scenarios. Based on positive results in this paper, we plan to work with domain experts in healthcare and the WSU School of Medicine to identify a cohort of patients to deploy the wearable technology for health monitoring. After a successful pilot study, we will work with wearable device manufacturers to create products in commercial form-factors so that they can be deployed to intended populations across the world, especially in areas with greater societal needs. This three step deployment plan from datasets to pilot studies to commercial products will ensure that the devices provide reliable, sustainable, and accurate metrics to patients.

7 Conclusion

Wearable and IoT devices offer significant promise in areas such as health monitoring and digital agriculture. Ambient EH presents a promising approach to prolong battery life in these systems. Effective EM, in turn, depends on accurate forecasting of future EH as well the uncertainty tied to those predictions. To address this, the paper introduces a novel EM framework that utilizes conformal prediction (CP) to generate uncertainty-aware EH forecasts across multiple future intervals. To enhance decision-making, a lightweight ML model is employed to select the optimal EM decision from a set of CP-derived uncertainty regions. The proposed approach outperforms baseline methods by ensuring reliable EM allocations, even under highly variable EH scenarios.

Acknowledgments

This work was supported in part by NSF CAREER award CNS-2238257 and in part by the AgAID AI Institute for Agriculture Decision Support, supported by the United States Department of Agriculture (USDA) - National Institute of Food and Agriculture (NIFA) award #2021-67021-35344.

Contribution Statement

Dina Hussein and Chibuike E. Ugwu contributed equally to this work.

References

[Alemdar *et al.*, 2013] Hande Alemdar, Halil Ertan, Ozlem Durmaz Incel, and Cem Ersoy. ARAS Human Activity Datasets in Multiple Homes with Multiple Residents. In *Int. Conf. Perv. Comput. Tech. Health. and Work.*, pages 232–235, 2013.

[Amu *et al.*, 2023] Hubert Amu, Theodora Yayra Brinsley, Frank Oppong Kwafo, Selasi Amu, and Luchuo Engelbert Bain. Improving Investment in Chronic Disease Care in Sub-Saharan Africa is Crucial for the Achievement of SDG 3.4: Application of the Chronic Care Model. *Archives of Public Health*, 81(1):169, 2023.

[Andreas and Stoffel, 1981] A. Andreas and T. Stoffel. NREL Solar Radiation Research Laboratory (SRRL): Baseline Measurement System (BMS); Golden, Colorado (Data); NREL Report No. DA-5500-56488, 1981. Accessed 3/20/2023.

[Atzori *et al.*, 2010] Luigi Atzori, Antonio Iera, and Giacomo Morabito. The Internet of Things: A Survey. *Comput. Netw.*, 54(15):2787–2805, 2010.

[Basaklar *et al.*, 2022] Toygun Basaklar, Yigit Tuncel, and Umit Y Ogras. tinyMAN: Lightweight Energy Manager using Reinforcement Learning for Energy Harvesting Wearable IoT Devices. In *tinyML Res. Symp.*, 2022.

[Bhat *et al.*, 2017] Ganapati Bhat, Jaehyun Park, and Umit Y Ogras. Near-Optimal Energy Allocation for Self-Powered Wearable Systems. In *ICCAD*, pages 368–375, 2017.

[Boyd and Vandenberghe, 2004] Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[Cammarano *et al.*, 2012] Alessandro Cammarano, Chiara Petrioli, and Dora Spenza. Pro-Energy: A Novel Energy Prediction Model for Solar and Wind Energy-Harvesting Wireless Sensor Networks. In *ACM MobiHoc*, 2012.

[Dagnachew *et al.*, 2023] Anteneh G Dagnachew, Su-Min Choi, and Giacomo Falchetta. Energy Planning in Sub-Saharan African Countries Needs to Explicitly Consider Productive Uses of Electricity. *Scientific Reports*, 13(1):13007, 2023.

[Daneault, 2018] Jean-Francois Daneault. Could Wearable and Mobile Technology Improve the Management of Essential Tremor? *Frontiers in neurology*, 9:257:1–257:8, 2018.

[Espay *et al.*, 2016] Alberto J Espay, Paolo Bonato, Fatta B Nahab, Walter Maetzler, John M Dean, Jochen Klucken, Bjoern M Eskofier, Aristide Merola, Fay Horak, Anthony E Lang, et al. Technology in Parkinson's Disease: Challenges and Opportunities. *Movt. Disorders*, 31(9):1272–1282, 2016.

[FlexSolarCells, 2013] FlexSolarCells. SP3-37 Datasheet, 2013. https://bit.ly/3dcJ0lK, Accessed 3/28/2021.

[Geisler *et al.*, 2017] Matthias Geisler, Sebastien Boisseau, Matthias Perez, Pierre Gasnier, Jerome Willemin, Imene Ait-Ali, and Simon Perraud. Human-Motion Energy Harvester for Autonomous Body Area Sensors. *Smart Materials and Structures*, 557(1):012024, 2017.

[Ghosh *et al.*, 2023a] Subhankar Ghosh, Taha Belkhouja, Yan Yan, and Janardhan Rao Doppa. Improving Uncertainty Quantification of Deep Classifiers via Neighborhood Conformal Prediction: Novel Algorithm and Theoretical Analysis. In *AAAI*, 2023.

[Ghosh *et al.*, 2023b] Subhankar Ghosh, Yuanjie Shi, Taha Belkhouja, Yan Yan, Jana Doppa, and Brian Jones. Probabilistically Robust Conformal Prediction. In *UAI*, 2023.

[Gibbs *et al.*, 2023] Isaac Gibbs, John J Cherian, and Emmanuel J Candès. Conformal Prediction with Conditional Guarantees. *ArXiv Preprint ArXiv:2305.12616*, 2023.

[Hussein *et al.*, 2022a] Dina Hussein, Taha Belkhouja, Ganapati Bhat, and Jana Doppa. Reliable Machine Learning for Wearable Activity Monitoring: Novel Algorithms and Theoretical Guarantees. In *ICCAD*, 2022.

[Hussein *et al.*, 2022b] Dina Hussein, Ganapati Bhat, and Jana Doppa. Adaptive Energy Management for Self-Sustainable Wearables in Mobile Health. In *AAAI*, 2022.

[Hussein *et al.*, 2024a] Dina Hussein, Taha Belkhouja, Ganapati Bhat, and Jana Doppa. Energy-Efficient Missing Data Imputation in Wearable Health Applications: A Classifier-aware Statistical Approach. In *IJCAI*, 2024.

[Hussein *et al.*, 2024b] Dina Hussein, Nuzhat Yamin, and Ganapati Bhat. EMI: Energy Management Meets Imputation in Wearable IoT Devices. *IEEE TCAD*, 43(11):3792–3803, 2024.

[Hussein *et al.*, 2025] Dina Hussein, Taha Belkhouja, Ganapati Bhat, and Jana Doppa. Sensor-Aware Data Imputation for Time-Series Machine Learning on Low-Power Wearable Devices. *ACM TODAES*, 30(1):1–27, 2025.

[Huynh *et al.*, 2008] Tâm Huynh, Mario Fritz, and Bernt Schiele. Discovery of Activity Patterns using Topic Models. In *Proc. of the 10th Int. Conf. on Ubiquitous Comput.*, pages 10–19, 2008.

[Kansal *et al.*, 2007] Aman Kansal, Jason Hsu, Sadaf Zahedi, and Mani B Srivastava. Power Management in Energy Harvesting Sensor Networks. *ACM Trans. Embedd. Comput. Syst.*, 6(4):32, 2007.

[Karakoç *et al.*, 2020] Nurullah Karakoç, Anna Scaglione, Angelia Nedić, and Martin Reisslein. Multi-Layer Decomposition of Network Utility Maximization Problems. *IEEE/ACM Trans. on Netw.*, 28(5):2077–2091, 2020.

[Latif *et al.*, 2018] Siddique Latif, Muhammad Yasir Khan, Adnan Qayyum, et al. Mobile Technologies for Managing Non-Communicable Diseases in Developing Countries. In *Mob. Appl. Solut. Soc. Inclus.*, pages 261–287. 2018.

[Maetzler *et al.*, 2016] Walter Maetzler, Jochen Klucken, and Malcolm Horne. A Clinical View on the Development of Technology-Based Tools in Managing Parkinson's Disease. *Movt. Disorders*, 31(9):1263–1271, 2016.

[Maha *et al.*, 2024] Chukwudi Cosmos Maha, Tolulope Olagoke Kolawole, and Samira Abdul. Harnessing Data Analytics: A New Frontier in Predicting and Preventing Non-Communicable Diseases in the US and Africa. *Comput. Sci. IT Res. J.*, 5(6):1247–1264, 2024.

[Mamun *et al.*, 2022] Abdullah Mamun, Seyed Iman Mirzadeh, and Hassan Ghasemzadeh. Designing Deep Neural Networks Robust to Sensor Failure in Mobile Health Environments. In *Proc. 44th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC), 2022*, pages 2442–2446. IEEE, 2022.

[Mattison *et al.*, 2022] Graeme Mattison, Oliver Canfell, Doug Forrester, Chelsea Dobbins, Daniel Smith, Juha Töyräs, and Clair Sullivan. The Influence of Wearables on Health Care Outcomes in Chronic Disease: Systematic Review. *J. Med. Internet Res.*, 24(7):e36690, 2022.

[Odema *et al.*, 2021] Mohanad Odema, Nafiul Rashid, and Mohammad Abdullah Al Faruque. Energy-Aware Design Methodology for Myocardial Infarction Detection on Low-Power Wearable Devices. In *ASPDAC*, pages 621–626, 2021.

[Phillips *et al.*, 2018] Siobhan M Phillips, Lisa Cadmus-Bertram, Dori Rosenberg, Matthew P Buman, and Brigid M Lynch. Wearable Technology and Physical Activity in Chronic Disease: Opportunities and Challenges. *Am. J. Prev. Med.*, 54(1):144–150, 2018.

[Piorno *et al.*, 2009] Joaquin Recas Piorno, Carlo Bergonzini, David Atienza, and Tajana Simunic Rosing. Prediction and Management in Energy Harvested Wireless Sensor Nodes. In *Int. Conf. Wireless Comm., Vehicular Tech., Info. Theory and Aerospace & Electron. Syst. Tech.*, pages 6–10, 2009.

[Regression, 2017] Quantile Regression. *Handbook of Quantile Regression*. CRC Press: Boca Raton, FL, USA, 2017.

[Romano *et al.*, 2019] Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized Quantile Regression. *Adv. Neural Inf. Process. Syst.*, 32, 2019.

[Romano *et al.*, 2020] Yaniv Romano, Matteo Sesia, and Emmanuel Candes. Classification with Valid and Adaptive Coverage. *NeurIPS*, 2020.

[Roy and Larocque, 2020] Marie-Hélène Roy and Denis Larocque. Prediction Intervals with Random Forests. *Stat. Methods Med. Res.*, 29(1):205–229, 2020.

[Shahrokhi *et al.*, 2025] Hooman Shahrokhi, Devjeet Raj Roy, Yan Yan, Venera Arnaoudova, and Janaradhan Rao Doppa. Conformal Prediction Sets for Deep Generative Models via Reduction to Conformal Regression, 2025.

[Shi *et al.*, 2024] Yuanjie Shi, Subhankar Ghosh, Taha Belkhouja, Jana Doppa, and Yan Yan. Conformal Prediction for Class-wise Coverage via Augmented Label Rank Calibration. In *NeurIPS)*, 2024.

[Shrestha and Solomatine, 2006] Durga L Shrestha and Dimitri P Solomatine. Machine Learning Approaches for Estimation of Prediction Interval for the Model Output. *Neural Networks*, 19(2):225–235, 2006.

[Sztyler *et al.*, 2016] Timo Sztyler, Josep Carmona, Johanna Völker, and Heiner Stuckenschmidt. Self-tracking Reloaded: Applying Process Mining to Personalized Health Care from Labeled Sensor Data. In *Trans. Petri Nets and other Models Concur.*, pages 160–180. Springer, 2016.

[Texas Instruments Inc., 2018] Texas Instruments Inc. CC2652R Microcontroller. [Online] https://www.ti.com/product/CC2652R, Accessed August 12, 2022, 2018.

[Tibshirani *et al.*, 2019] Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal Prediction under Covariate Shift. *Adv. Neural Inf. Process. Syst.*, 32, 2019.

[Tuncel *et al.*, 2020] Yigit Tuncel, Shiva Bandyopadhyay, Shambhavi V Kulshrestha, Audrey Mendez, and Umit Y Ogras. Towards Wearable Piezoelectric Energy Harvesting: Modeling and Experimental Validation. In *Proc. ISLPED*, pages 55–60, 2020.

[Vigorito *et al.*, 2007] Christopher M Vigorito, Deepak Ganesan, and Andrew G Barto. Adaptive Control of Duty Cycling in Energy-Harvesting Wireless Sensor Networks. In *Proc. Conf. on Sensor, Mesh and Ad Hoc Comm. and Networks*, pages 21–30, 2007.

[Vovk *et al.*, 2005] Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in A Random World*, volume 29. Springer, 2005.

[Wan *et al.*, 2011] ZG Wan, YK Tan, and C Yuen. Review on Energy Harvesting and Energy Management for Sustainable Wireless Sensor Networks. In *2011 IEEE 13th Int. Conf. on Commun. Technol.*, pages 362–367. IEEE, 2011.

[WHO, 2022] WHO. Deaths from Noncommunicable Diseases on the Rise in Africa. *Africa: WHO Afr*, 2022.

[Yamin and Bhat, 2023] Nuzhat Yamin and Ganapati Bhat. Uncertainty-aware Energy Harvest Prediction and Management for IoT Devices. *ACM Trans. Des. Autom. Electron. Syst.*, 28(5):1–33, 2023.