INDIAN STATISTICAL INSTITUTE, KOLKATA

Implementing a Health Recommendation System from Wearable Data

A dissertation submitted in fulfillment of the requirements for the award of

Master of Technology

in

Cryptology and Security

By: Priti Pramanick
CrS2315

Primary Supervisor:
Debasis Sahoo

Secondary Supervisor:
Dr. Debrup Chakraborty
Indian Statistical Institute, Kolkata

Certificate

This is to certify that the dissertation entitled **"Implementing a health recommendation system from Wearable data"** submitted by **Priti Pramanick** to the **Indian Statistical Institute, Kolkata**, in fulfillment for the award of the degree of **Master of Technology in Cryptology and Security**, is a bonafide record of work carried out by her under my supervision and guidance.

The dissertation has fulfilled all the requirements as per the regulations of this institute and, in my opinion, has reached the standard needed for submission.

Debrup Chakraborty.

Dr. Debrup Chakraborty
Indian Statistical Institute, Kolkata

Acknowledgment

I would like to express my sincere gratitude to my supervisors, **Debasis Sahoo** and **Dr. Debrup Chakraborty**, for their invaluable guidance, support, and encouragement throughout the course of this work. Their insights and suggestions have been instrumental in shaping the direction and depth of this dissertation.

I am also deeply thankful to **Ritesh Bawri** for inspiring the core idea of using data to personalize health recommendations and for enabling access to real-world datasets that formed the foundation of this project.

I appreciate the support of the faculty and staff of the **Indian Statistical Institute, Kolkata**, for providing a stimulating academic environment and the necessary resources to carry out this research.

Lastly, I extend my heartfelt thanks to my family and friends for their continuous encouragement and moral support throughout my journey.

Priti Pramanick

Priti Pramanick
Indian Statistical Institute
Kolkata: 700108, India

Contents

Chapter 1

Introduction

The integration of wearable technologies into everyday life has significantly expanded the scope of personal health monitoring. Devices such as fitness trackers, smartwatches, and mobile health apps generate continuous streams of data on parameters like sleep quality, heart rate, physical activity, and nutrition. While these data hold great potential for improving health awareness, they are often underutilized due to the absence of accessible systems that can analyze and interpret them meaningfully for individuals.

Traditional healthcare is largely episodic, relying on periodic clinical visits that may overlook day-to-day health fluctuations. In contrast, wearable devices provide a more dynamic picture of an individual's well-being. Recent research highlights a growing interest in using such real-world, real-time data to support health-related insights, particularly through the application of machine learning and data engineering methods. However, most current solutions remain either research-focused, disease-specific, or not designed for everyday users.

This thesis presents the design and implementation of a health scoring and feedback system that brings advanced data processing techniques into a user-friendly framework. The system collects and integrates weekly data from multiple health domains—sleep, cardio, fitness, nutrition, and metabolism. It uses Python and MongoDB to process this data, calculate meaningful health scores, and generate personalized lifestyle recommendations using natural language APIs. The final output is delivered as a structured PDF report, offering users both a summary and actionable insights.

The system aims not just to display statistics, but to interpret personal health data in a way that supports better decision-making. This work contributes to the broader vision of preventive and participatory healthcare, where individuals are empowered with insights tailored to their unique profiles. By demonstrating a practical, scalable solution built on wearable data, this thesis bridges the gap between raw health metrics and personalized wellness guidance.

1.1 Importance of Wearable Data on Digital Health Systems

Wearable devices have fundamentally altered how people engage with their health, making self-monitoring both accessible and continuous. Gadgets such as smartwatches, fitness trackers, and sleep monitors are now widely used to capture diverse physiological and behavioral metrics—ranging from heart rate and step count to oxygen levels, calories burned, and sleep patterns. As their adoption grows, so too does the volume of user-generated health data, opening the door to a shift from reactive treatment to preventive care.

A key advantage of wearable technologies is their capacity for constant, passive, and non-invasive health monitoring. Unlike conventional healthcare models that depend on sporadic clinical check-ups, wearables offer a steady stream of real-time data that mirrors the individual's everyday habits and routines. This continuous monitoring can reveal subtle health shifts—such as disrupted sleep patterns or reduced physical activity—which may signal stress, illness, or other health concerns before they become critical.

When integrated into digital health ecosystems, wearable data becomes a crucial enabler of personalized care. By aggregating information over time, these devices help form a comprehensive view of an individual's health trajectory. Such insights support the development of adaptive scoring systems, behavior-based recommendation engines, and digital coaching solutions. This approach enhances not only chronic disease tracking but also general wellness and fitness planning, making health interventions more tailored and responsive.

Wearables also contribute to making healthcare more accessible. They allow for remote tracking of health indicators, a feature especially valuable for those living far from medical centers or with mobility challenges. By enabling individuals to take an active role in monitoring their own health, these tools encourage autonomy, awareness, and early action. In parallel, they supply healthcare professionals with supplemental data that can enrich medical consultations and treatment decisions.

Yet, raw data alone is insufficient. Without effective interpretation, users may find these metrics confusing or unhelpful. Hence, intelligent health systems must not only collect data but also analyze, contextualize, and present it in a way that leads to meaningful insights. Transforming unprocessed sensor outputs into personalized, actionable feedback is what unlocks the real value of wearable health technology.

In summary, wearable technology has emerged as a cornerstone of modern digital health systems. Its ability to provide granular, continuous, and personalized health data makes it indispensable for developing tools that aim to promote preventive care and enhance user engagement. The integration of this data into intelligent systems paves the way for scalable, accessible, and meaningful health interventions.

1.2 Challenges in Wearable Data Processing and Analysis

Despite the promise of wearable technology in health monitoring, utilizing the data it generates poses several technical and practical challenges. These arise from the nature of the data—often noisy, inconsistent, and device-dependent—as well as from the computational and ethical constraints surrounding real-time health analytics. For digital health systems to be both reliable and scalable, these issues must be addressed effectively.

One of the foremost challenges involves the variability in data accuracy and reliability across consumer-grade devices. Unlike clinical-grade monitors, many commercial wearables prioritize affordability and user comfort, which can compromise measurement precision. Sensor misalignment, improper wearing, battery constraints, and hardware limitations can lead to incomplete, duplicated, or misleading data. For example, sleep monitoring devices may fail to register brief awakenings or misclassify rest periods, resulting in skewed assessments.

Another major concern is data heterogeneity. Wearables from different manufacturers often use proprietary data formats, sampling rates, and feature definitions. Even when platforms provide APIs, the data exposed may differ in resolution, semantics, or update frequency. Metrics such as stress levels or VO max are often estimated using closed-source algorithms, limiting interpretability and hindering scientific validation. As a result, merging or comparing data across devices becomes difficult without a robust preprocessing strategy.

Synchronizing multiple data streams is also complex. Health metrics like step count, heart rate, calorie burn, and sleep data are often recorded at different frequencies and may exhibit time lags or misalignments. Aligning this temporally diverse data is essential for accurate multi-dimensional analysis—for instance, examining how late meals affect sleep quality or how exercise impacts resting heart rate trends.

From a system architecture perspective, the continuous and voluminous nature of wearable data demands scalable infrastructure. Traditional relational databases may fall short when managing large-scale time-series data, making NoSQL alternatives like MongoDB more suitable. However, these too require careful schema design, indexing strategies, and optimization for querying nested or unstructured data. Preprocessing stages—such as data cleaning, normalization, interpolation, and feature extraction—consume significant computational resources and are critical for downstream analysis.

Privacy and data protection remain key challenges. Wearables collect highly sensitive personal health information, raising concerns over user consent, data ownership, and potential misuse. Systems handling this data must adhere to strict privacy regulations such as GDPR or HIPAA and implement safeguards including encryption, anonymization, and secure access controls to maintain user trust.

Equally important is the challenge of making the data interpretable. Raw figures like step counts or heart rate values offer limited value unless contextualized. Translating sensor outputs into personalized, understandable feedback requires advanced methods like machine learning, natural language generation, and data visualization—each demanding domain expertise and careful validation.

In summary, while wearable devices offer a rich stream of health-related data, effectively converting this into actionable insight requires addressing multiple layers of complexity. These include ensuring data quality, handling heterogeneity, achieving temporal synchronization, building scalable pipelines, safeguarding privacy, and delivering meaningful feedback. Addressing these obstacles is essential to unlock the full potential of wearable data in personalized healthcare.

1.3 Scope and Organisation of the Thesis

The scope of this thesis lies in developing a digital health reporting system that leverages multi-dimensional data from wearable devices to generate personalized health advice. This involves the design of a modular and adaptable framework capable of processing heterogeneous health metrics—such as sleep patterns, cardiovascular activity, physical fitness, and metabolic parameters—to compute health scores and provide actionable recommendations. The system utilizes NoSQL databases like MongoDB for data storage and Python-based pipelines for data processing and integration. Although the current implementation focuses on one-week data windows and generates static PDF reports, the design is flexible enough to be extended for longer-term monitoring and dynamic interfaces in future iterations.

This thesis is organized into five chapters, each building upon the previous to guide the reader from foundational context to technical implementation and research insights.

- **Chapter 1: Introduction** – Introduces the research motivation, importance of wearable data in digital health systems, and the key challenges in processing such data. It also outlines the scope and structural organization of the thesis.
- **Chapter 2: Survey of Machine Learning Algorithms for Wearable Data** – Provides an overview of wearable data analysis techniques and surveys different machine learning approaches applied to this domain. It categorizes existing methods into three types based on their modeling strategies and use cases, offering a structured comparison of the current state-of-the-art.
- **Chapter 3: Wearable Data Collection and Preprocessing** – Describes the data acquisition strategy using MongoDB, including how raw data from fitness devices is stored, queried, cleaned, and formatted for downstream analysis. This chapter outlines the data schema and essential preprocessing steps critical for reliable health metric extraction.
- **Chapter 4: A New ML Framework for Personalised Health Advice From Wearable Data** – Introduces the core technical contribution of the thesis: a machine learning-based framework for computing health scores and generating personalized recommendations. The chapter details system design, scoring methodologies, and architectural flow. It also presents experimental results and discusses the implications of the framework in real-world health tracking scenarios.
- **Chapter 5: Conclusion and Future Direction** – Summarizes the major contributions of the work, reflects on its broader impact, and identifies promising directions for future research. These include expanding the framework to additional health domains, enhancing personalization via dynamic user modeling, and integrating advanced AI-driven recommendation strategies.

This structured progression ensures a comprehensive understanding of the system's development, from foundational analysis and data preparation to methodological innovation and practical outcomes.

Chapter 2

Survey of Machine Learning Algorithms for Wearable Data

The proliferation of wearable devices has resulted in an unprecedented volume of health-related data being generated in real time. These devices routinely track various physiological and behavioral signals such as heart rate, sleep duration, step count, oxygen saturation, and activity intensity. While this data offers rich potential for advancing personal health analytics, its effective utilization requires robust machine learning techniques capable of managing noisy, non-uniform, and temporally dense information [1, 2].

This chapter explores how machine learning (ML) has been leveraged to interpret wearable data and deliver meaningful, personalized insights. Recent research has shown that ML plays a central role in transforming continuous sensor readings into actionable outputs for health assessment, activity recognition, risk detection, and user-specific feedback [3]. Algorithms must handle data irregularities, individual variability, and high-frequency sampling—often without complete or consistently labeled datasets.

We begin by outlining the core machine learning paradigms applied to wearable data. Supervised learning approaches such as decision trees, support vector machines, random forests, and deep neural networks have been widely used for classification tasks, such as detecting physical activities or identifying signs of fatigue or illness [1]. Regression models are employed to estimate continuous outcomes, such as stress levels or biological age, based on patterns in physiological signals.

Unsupervised techniques, including clustering and dimensionality reduction, are useful in discovering latent structures in unlabeled data—grouping users by sleep behavior or uncovering anomalies in daily routines [2]. Hybrid approaches, including semi-supervised learning and ensemble methods, have also been adopted to address the scarcity of labeled data and improve model robustness.

Specific challenges arise when applying ML to wearable health data. Data imbalance, where certain events (e.g., falls, seizures) are rare, can bias models toward more frequent behaviors. Wearable datasets also vary significantly across users and devices, requiring models that generalize well while still allowing for personalization. Additionally, temporal alignment of multimodal data streams is often necessary to capture cause-effect relationships between metrics, such as how evening calorie intake influences sleep quality [3].

Scalability and interpretability are also crucial. As health systems move toward user-facing applications, models must not only be accurate but also transparent and computationally efficient. Interpretability becomes especially important when ML outputs feed into recommendation engines or health scoring systems, as users must be able to understand and trust the feedback they receive [1].

This chapter presents a survey of these ML methodologies, discussing their applicability to wearable health data and the limitations that remain in real-world deployment. Understanding these techniques provides the groundwork for building intelligent systems that compute personalized health scores and generate tailored advice—goals central to the thesis at hand.

2.1 Wearable Data Analysis

Wearable data analysis involves transforming continuous sensor outputs from devices such as fitness bands, smartwatches, sleep trackers, and biosensors into interpretable health insights. Wearables are now widely adopted for continuous health monitoring across clinical and consumer contexts [4]. These devices routinely collect time-series information that captures key physiological and behavioral indicators—including heart rate, sleep duration, activity levels, calorie expenditure, skin temperature, and more. Analyzing this data enables the creation of digital health systems capable of monitoring well-being, detecting deviations from normal patterns, and providing user-specific feedback.

The nature of wearable data introduces both opportunities and analytical complexity. It is typically high in volume, recorded frequently throughout the day, and accumulated over weeks or months. At the same time, it is susceptible to noise caused by motion artifacts, environmental interference, or improper wear. Moreover, physiological responses vary considerably among individuals, meaning that the same data point can signify different conditions based on a person's health background, age, or activity context.

A structured analysis pipeline begins with preprocessing—cleaning raw data to handle noise and missing values, aligning timestamps across multiple sensor streams, and resampling to standardize sampling intervals. Feature extraction follows, converting raw sequences into meaningful indicators such as sleep efficiency, variability in activity intensity, or average resting heart rate. These features can be statistical, temporal, or domain-specific, depending on the target application.

Contextual awareness is critical in wearable data interpretation. Physiological metrics must be analyzed in relation to the surrounding activity—for example, a heart rate of 90 bpm may be expected during light exercise but could signal restlessness during sleep. Classification models can distinguish between behavioral states like rest, activity, or sleep by combining accelerometer and heart rate signals, enabling finer segmentation of health patterns and trends.

Personalization plays a vital role in ensuring that feedback is meaningful. Unlike fixed clinical thresholds, personalized baselines help identify significant deviations in a user's typical patterns. For instance, a 20% reduction in a person's usual step count or sleep duration may be more informative than a general benchmark. Adaptive scoring systems and anomaly detection models are therefore increasingly tuned to individual profiles to reflect relative change rather than absolute measures.

Security and ethical considerations also shape wearable data analysis. As this data reflects sensitive aspects of users' daily lives, strong privacy safeguards—such as secure storage, data anonymization, and transparent consent processes—are essential. Furthermore, interpretation must be communicated in a way that supports informed user decisions, avoiding confusion or misinterpretation of health signals.

The integration of wearable data into broader health systems—through multi-sensor fusion and longitudinal modeling—enables a more comprehensive understanding of health. Linking sleep patterns, physical activity, and physiological trends enhances the ability to predict changes and tailor recommendations. In this context, machine learning offers the tools needed to identify patterns, detect anomalies, and generate proactive insights at scale.

In conclusion, wearable data analysis encompasses a series of technical and contextual steps aimed at deriving personalized, actionable health insights from real-time data streams. While challenges such as variability, noise, and heterogeneity remain, advances in data science and computing are continuously improving the quality and relevance of insights delivered through wearable-driven health platforms.

2.2 Machine Learning Approaches on Wearable Data

Machine learning has become central to transforming raw wearable data into meaningful insights. Due to the high dimensionality, variability, and temporal nature of these datasets, researchers have adopted a wide range of models—ranging from deep neural networks to probabilistic graphical models and, more recently, transformer-based frameworks [1, 2]. This section categorizes key methodologies into three broad classes, reflecting both classical and emerging directions in wearable analytics.

2.2.1 Deep Learning Methods

Deep learning models are widely used for wearable data due to their ability to learn complex, hierarchical representations from raw sensor inputs [5, 2]. These architectures often outperform traditional methods in tasks involving temporal or multivariate sensor streams.

- **CNN-LSTM hybrids** are particularly effective in activity and sleep classification. These models combine spatial feature extraction with temporal sequence modeling, achieving high accuracy in sleep staging tasks, with reported AUCs reaching 0.94 to 0.95 in benchmark studies [1].

- **Autoencoders and Deep Belief Networks (DBNs)** have shown utility in anomaly detection, learning typical behavioral patterns and flagging deviations that may indicate health issues such as arrhythmias or stress spikes [2].

- **Bayesian convolutional neural networks (Bayesian CNNs)** offer a probabilistic enhancement to traditional CNNs by modeling uncertainty—an important factor when dealing with noisy sensor environments or incomplete data.

- **Recurrent models**, including LSTMs and GRUs, are commonly used for time-series prediction in heart rate variability, stress monitoring, and sleep dynamics. Their memory capabilities allow them to capture long-term dependencies in physiological signals [1].

These approaches are well-supported in recent literature, particularly in systems emphasizing continuous health tracking, stress detection, and personalized monitoring frameworks [2].

2.2.2 Probabilistic and Bayesian Models

Probabilistic methods provide structured ways to incorporate uncertainty, temporal relationships, and prior knowledge, making them suitable for sequential and behavioral modeling tasks in wearable health systems [1].

- **Hidden Markov Models (HMMs)** and **Dynamic Bayesian Networks** are frequently used for modeling transitions in physical activity states, especially when labeled data is sparse or noisy.
- **Conditional Random Fields (CRFs)** offer a discriminative framework for sequence labeling, and have been applied in activity recognition pipelines that rely on accelerometer and gyroscope data [2].
- **Bayesian neural networks** integrate deep learning with probabilistic reasoning, making predictions more robust in uncertain or variable sensor conditions.
- **The Wake–Sleep algorithm**, often applied in Helmholtz Machines, supports generative modeling of wearable signals and enables unsupervised learning when labels are unavailable.

These models are particularly useful in domains requiring explainability, structured temporal analysis, or probabilistic inference—areas increasingly relevant in personalized healthcare.

2.2.3 Large Language Models and Few-Shot Learners

Recent research has begun exploring how transformer-based architectures and large language models (LLMs) can be adapted to process wearable time-series data. These approaches are still emerging but offer promising capabilities for contextual understanding and label-efficient learning [3].

- Early explorations show that LLMs such as GPT-4 and LLaMA can be prompted to analyze structured wearable logs, infer activity sequences, and predict user behavior with minimal supervision [3].
- Few-shot learning techniques have demonstrated that even limited physiological datasets can be effectively leveraged by large models to classify conditions such as cardiac anomalies or metabolic imbalances.
- Hybrid systems combining transformers with natural language input—like user diaries or symptom notes—enable richer, multimodal analysis for applications such as mental health tracking or sleep-behavior correlation.

This class of models remains experimental but represents a shift toward more general-purpose, adaptable AI frameworks in wearable health research.

2.3 Summary

This chapter provided a comprehensive exploration of how machine learning is applied to wearable data in health analytics. We began by analyzing the unique characteristics of wearable data, including its temporal continuity, person-specific variability, and susceptibility to sensor noise. Addressing these complexities requires careful preprocessing, feature extraction, and context-aware modeling.

We then categorized prominent machine learning techniques into three core groups: deep learning architectures, probabilistic and Bayesian models, and emerging large language model frameworks. Each category supports different aspects of wearable data analysis—from capturing sequential dependencies and detecting anomalies to generating personalized feedback with minimal supervision. These methods are increasingly being used in health applications such as activity recognition, sleep classification, cardiovascular monitoring, and adaptive health scoring.

In addition, we emphasized evolving trends such as multi-modal sensor fusion, model personalization, and self-supervised learning—all of which are vital for scaling health insights across diverse users. These developments align with the goals of this thesis, which seeks to compute personalized health scores and generate tailored recommendations based on continuous, real-world data.

With this conceptual foundation established, the following chapter moves from theory to implementation. It outlines the data acquisition and preprocessing pipeline used to collect, organize, and prepare wearable sensor data for use in the health scoring system developed as part of this work.

Chapter 3

Wearable Data Collection and Preprocessing

The reliability and usefulness of any data-driven health analysis system fundamentally depend on the quality and structure of the input data. In the context of wearable-based health monitoring, data is gathered from various sources such as smartwatches, fitness trackers, mobile apps, and cloud platforms. This chapter describes the end-to-end process of data acquisition and preprocessing, beginning from raw collection to the generation of a clean, structured weekly summary per user. It also introduces the role of MongoDB as the core database technology for flexible and scalable storage.

3.1 Nature of Wearable Health Data

Wearable devices continuously collect a rich stream of physiological and behavioral metrics that reflect an individual's daily health patterns. These include:

- **Sleep data:** total sleep duration, light sleep, deep sleep, REM sleep, sleep onset time, wake periods.
- **Cardiovascular data:** resting heart rate, heart rate variability, peak heart rate, stress index.
- **Fitness data:** daily steps, activity minutes (moderate, intense), calories burned, distance traveled.
- **Nutrition data:** calories consumed, macronutrient breakdown (carbs, proteins, fats), hydration levels.
- **Metabolic indicators:** weight, BMI, basal metabolic rate, body fat percentage, VO max (if supported).

These data points may be captured at different frequencies—from minute-level heart rate monitoring to daily logs of meals or weight—and require careful handling to align and integrate meaningfully.

3.2 Data Storage Using MongoDB

To manage the heterogeneity and volume of wearable data, MongoDB—a document-based NoSQL database—has been used as the central data store. MongoDB offers a flexible, scalable environment that supports nested documents, dynamic schemas, and fast query capabilities. These properties make it highly suitable for storing time-series data with diverse and irregular formats.

MongoDB Collections and Schema Design

The key collections in the system include:

- `profile_details`: stores basic demographic and static health-related information for each user (e.g., age, gender, weight, goals).
- `profile_preferences`: stores user-defined settings and thresholds for personalized scoring (e.g., sleep targets, activity sensitivity).
- `body_matters`: stores daily logs and time-series data from wearable devices. This includes structured subdocuments for sleep, nutrition, cardio, fitness, and metabolism, indexed by date and profile_id.

Each document is tagged with a unique `profile_id`, allowing aggregation and linkage of multiple data sources per user. The structure is intentionally modular, facilitating independent processing of each health domain.

3.3 Data Collection Methods

The data stored in MongoDB is collected through the following mechanisms:

- **Wearable device APIs:** Platforms like Fitbit [7], Garmin, or Apple HealthKit offer secure APIs to extract health data after user authorization. OAuth2 protocols are used to manage access tokens.
- **Mobile app integrations:** In some cases, third-party health apps (e.g., MyFitnessPal for nutrition) provide supplementary data via REST APIs.
- **Manual logs and CSV uploads:** For users without automated sync, data can be uploaded in standardized CSV formats and parsed into JSON before insertion.

A data ingestion module periodically fetches or receives data from these sources, parses it into a unified schema, and inserts it into MongoDB. Basic validation and timestamp standardization are applied during this stage.

3.4 Data Preprocessing Pipeline

Once raw data is stored, it undergoes several preprocessing steps to ensure consistency, reliability, and readiness for downstream analysis. The goal is to transform daily logs into clean, structured records that can be aggregated into weekly summaries per user.

3.4.1 Timestamp Normalization and Alignment

Wearable data can originate from users across different time zones, with different formats (e.g., Unix timestamps, ISO strings). All timestamps are converted to a standard UTC-based format and aligned to calendar days. This ensures that sleep recorded at 2 AM is assigned to the correct night, and that daily totals are computed consistently.

3.4.2 Missing and Incomplete Data Handling

Wearable data often contains missing entries due to device non-usage, sync failures, or battery outages. Handling strategies include:

- **Imputation:** For short-term missing values (e.g., heart rate gaps of a few minutes), linear interpolation is applied.
- **Forward filling:** For metrics like weight or hydration, the last known value is propagated until a new value appears.
- **Masking or exclusion:** If key metrics like total sleep are entirely missing for a day, that day is flagged and excluded from scoring.

3.4.3 Resampling and Aggregation

High-frequency data (e.g., minute-wise heart rate or steps) is resampled to daily aggregates using mean, min, max, and sum operations. Activity data is bucketed by intensity level, and sleep stages are aggregated into total durations. Daily summaries are computed and stored per user per day.

3.4.4 Derived Feature Computation

Beyond raw values, domain-specific features are derived to enrich the dataset:

- Sleep efficiency = (Total sleep time) / (Time in bed)
- Active-to-sedentary ratio = (Active minutes) / (Sedentary minutes)
- Nutrient balance score = Deviation from ideal macronutrient split
- Resting HR trend = 3-day moving average of resting heart rate

These features capture trends, balance, and behavior patterns more effectively than raw values.

3.4.5 Normalization and Standardization

To support user-independent scoring in the next stage, numerical features are scaled. Options include:

- **Min-max normalization** to a [0, 1] range for scoring
- **Z-score normalization** when comparing a user's value to population trends

Normalization ensures consistency in interpreting values across users with different baselines.

3.4.6 Weekly Summary Generation

All cleaned and processed data is finally grouped by user and week (e.g., Monday to Sunday). The weekly summary includes:

- Daily values for each domain
- Weekly averages, trends, and min/max markers
- Flags for missing or anomalous data

This weekly summary becomes the core input to the scoring and recommendation engine described in the next chapter.

3.5 Challenges and Design Decisions

Throughout the data collection and preprocessing stages, several challenges emerged:

- **Device heterogeneity:** Different wearables report metrics using different names, units, and resolution. Wrapper scripts were built to standardize field mappings.
- **Sparse and irregular data:** Some users provide dense time-series data; others log values manually. Adaptive preprocessing pipelines were designed to handle both.
- **Outlier detection:** Implausible values (e.g., 2 hours of sleep or 500 bpm heart rate) are detected and corrected or discarded.
- **Scalability:** MongoDB aggregation pipelines were optimized with indexes and projections to support fast weekly rollups even with growing datasets.

3.6 Summary

This chapter outlined the foundational data infrastructure for the personalized health report system. It described how diverse wearable data—such as sleep, activity, heart rate, nutrition, and metabolic inputs—is collected, validated, cleaned, and aggregated into consistent weekly summaries for each user. Using MongoDB's flexible schema and Python-based preprocessing routines, the system handles missing data, aligns time-series inputs, and standardizes multi-domain features. This structured and modular data output ensures that downstream components receive reliable and ready-to-use inputs.

The preprocessing pipeline serves as the critical bridge between raw sensor logs and health intelligence. It addresses real-world challenges such as irregular logging, device inconsistencies, and temporal misalignment. These efforts make the subsequent computation of health scores both accurate and robust.

In the next chapter, we build on this preprocessed dataset to design and implement a comprehensive ML framework for personalized health feedback. This includes defining domain-wise scoring rules, integrating natural language generation for user-specific advice, and developing a complete report generation system that translates raw data into interpretable and actionable health insights.

Chapter 4

A New ML Framework for Personalised Health Advice From Wearable Data

4.1 Introduction

The proliferation of wearable health devices has enabled the continuous monitoring of vital signs, lifestyle behaviors, and physiological parameters in a non-invasive and user-friendly manner. However, while the data collected is rich and multidimensional, most users struggle to interpret raw metrics such as step count, resting heart rate, or sleep stages in a meaningful way. This disconnect between data availability and actionable health insight presents a significant gap in the current landscape of digital health systems.

Existing platforms often provide basic visualizations or generalized tips, lacking context-sensitive interpretation or personalized recommendations. To truly empower users in managing their health, it is essential to convert wearable data into structured, interpretable scores and generate individualized advice that aligns with their behaviors and goals. This requires an integrated framework that combines data preprocessing, health scoring logic, and intelligent feedback generation.

The primary objective of this chapter is to introduce a new machine learning–inspired framework that addresses this gap. The system transforms weekly wearable data into modular health scores across different domains—sleep, cardiovascular activity, fitness, nutrition, and metabolism—and then uses language generation models to produce personalized summaries and advice. The framework is designed to be modular, interpretable, and extensible, allowing seamless integration of new data streams or scoring methods.

This approach not only bridges the gap between raw sensor data and actionable feedback but also supports early detection of unhealthy trends and motivates sustained behavior change. The remainder of the chapter details the proposed system design, health scoring methodology, and the integration of natural language feedback mechanisms, followed by experimental results and discussion.

4.2 Proposed Methodology

The proposed methodology centers on transforming raw, multi-source wearable data into comprehensible and actionable health insights. This is achieved through a structured pipeline that includes data ingestion, preprocessing, scoring, interpretation, and recommendation generation. The system is modular and extensible, capable of handling various types of health metrics, adapting to different user profiles, and scaling across multiple use cases. This section describes the overall system architecture and the core computational logic that drives health score generation and personalized advice.

4.2.1 System Design and Architecture

The framework follows a five-stage pipeline that begins with weekly wearable data and ends with a personalized health report:

1. **Data Ingestion and Storage:** Raw data from multiple wearable sources is collected and stored in a MongoDB database, organized by user profile and health domains (sleep, cardio, fitness, nutrition, and metabolism). This stage was covered in detail in the previous chapter.

   ```
   from pymongo import MongoClient
   client = MongoClient('mongodb://localhost:27017/')
   db = client['health_db']
   collection = db['wearable_data']
   user_data = collection.find({"profile_id": profile_id})
   ```

2. **Preprocessing and Feature Extraction:** Daily-level data is cleaned, resampled, and normalized. Derived features such as sleep efficiency, active-to-sedentary ratio, or nutrient balance are computed and grouped into a weekly summary per user.

   ```
   # Example: Sleep efficiency calculation
   sleep_efficiency = total_minutes_asleep / total_minutes_in_bed
   ```

3. **Health Scoring Engine:** Each health domain has a modular scoring function that evaluates the weekly data based on thresholds, trends, and variability. The outputs are normalized scores between 0 and 100, along with brief interpretation tags (e.g., "good recovery," "low cardio activity").

   ```
   # Example: Scoring function
   if 7 <= avg_sleep_hours <= 9:
       sleep_score = 100
   elif avg_sleep_hours < 5:
       sleep_score = 50
   else:
       sleep_score = 70
   ```

4. **Language-Based Recommendation Generator:** The weekly summary and scores are formatted as a prompt and sent to a large language model (LLM) API. The model generates personalized health advice based on domain-specific cues, recent trends, and user goals.

```
prompt = f"""
Weekly Summary:
Sleep Score: {sleep_score} - {sleep_tag}
Cardio Score: {cardio_score} - {cardio_tag}
...
Generate feedback for a user who wants to improve sleep and fitness.
"""
response = openai.ChatCompletion.create(
    model="gpt-4",
    messages=[{"role": "user", "content": prompt}]
)
```

5. **Report Generator:** The final stage involves formatting the scores, feedback, and visualizations (e.g., pie charts, score circles) into a PDF report using ReportLab or similar libraries.

```
from reportlab.lib.pagesizes import A4
from reportlab.pdfgen import canvas
c = canvas.Canvas("report.pdf", pagesize=A4)
c.setFont("Helvetica-Bold", 18)
c.drawString(100, 800, "Weekly Health Report")
c.save()
```

The system architecture is designed to support flexibility in health domains, allowing for independent improvements or tuning in one module (e.g., sleep scoring) without affecting others. The separation of data, logic, and presentation layers ensures robustness and ease of maintenance.

4.2.2 Methods to Compute Health Score

A key innovation of the framework lies in the computation of modular health scores that are both personalized and interpretable. The scoring engine evaluates user data across five major domains:

a. Sleep Score

This score is computed based on sleep duration, consistency, sleep efficiency, and stage balance (REM/deep/light). Ideal sleep ranges are defined per age group, and deviations are penalized based on magnitude. Sleep fragmentation and variability across days also reduce the score. Example components include:

- Optimal duration range (e.g., 7–9 hours)
- Sleep efficiency $= \frac{\text{time asleep}}{\text{time in bed}}$
- Sleep regularity (std. deviation of onset times)

b. Cardiovascular Score

This module evaluates resting heart rate (RHR), heart rate variability (HRV), and stress indices across the week. Scores are normalized based on evidence-based thresholds and personal trends:

- Low RHR and high HRV typically indicate better cardiovascular health.
- Trends (e.g., rising RHR) are flagged even if values are within acceptable range.

c. Fitness Score

The fitness score is computed using step count, active minutes (moderate and intense), and calories burned. Activity thresholds are adapted from WHO recommendations, and scores adjust for age and gender:

- 150+ active minutes/week is a benchmark for good fitness.
- Sedentary behavior (low step count) reduces score even if calories are high.

d. Nutrition Score

This score evaluates calorie intake, macronutrient balance, and hydration consistency. The system flags both underconsumption and overconsumption. Macronutrient ratios are compared against ideal targets (e.g., 50% carbs, 20–30% proteins, 20–30% fats).

e. Metabolic Score

Based on weight, BMI, BMR, and trends in body composition (if available). Stability and gradual improvement are rewarded, while sudden spikes or drops reduce the score. A user-specific baseline is used for better interpretation.

Score Normalization and Aggregation Each domain score is independently normalized to a 0–100 scale. Weights can be assigned to different domains depending on user preferences or clinical priority. A final overall health score is computed as:

$$\text{Health Score} = \sum_{i=1}^{5} w_i \cdot s_i$$

where s_i is the domain score and w_i is its weight.

Interpretability Each score is accompanied by a textual tag or flag (e.g., "excellent sleep", "low hydration", "improving fitness") based on predefined thresholds and recent trends. These help guide the LLM-based recommendation in the next stage.

4.2.3 Methods for Report Generation

Language Model for Feedback Generation

The structured weekly summary and scores are converted into a prompt and sent to a large language model (LLM) API (e.g., OpenAI). The prompt includes:

- Weekly average scores and key trends
- Tags for each domain (e.g., "low cardio", "balanced nutrition")
- Optional user preferences or goals (e.g., focus on stress management)

The LLM returns personalized, easy-to-read recommendations such as:

"Your sleep patterns were consistent this week, but you averaged only 6.5 hours per night. Consider winding down 30 minutes earlier. Cardio activity was low—try incorporating 20-minute brisk walks three times this week."

PDF Report Generation

Using the ReportLab library, the system formats:

- Domain scores as circular visual indicators
- A sleep pie chart (e.g., REM, light, deep proportions)
- LLM-generated recommendation section
- User profile and weekly summary table

```
from reportlab.graphics.charts.piecharts import Pie
pie = Pie()
pie.data = [30, 45, 25]  # Deep, Light, REM
pie.labels = ['Deep', 'Light', 'REM']
drawing.add(pie)
```

Modularity and Extensibility The architecture allows new health modules (e.g., mental wellness, blood glucose) to be added with minimal changes. Likewise, alternate scoring formulas or new model prompts can be tested and integrated independently.

4.3 Experimental Results and Discussion

To evaluate the performance and usability of the proposed ML framework for health analysis from wearable data, a set of experiments was conducted on real-world data across multiple user profiles and time periods. The primary aim was to assess the robustness, interpretability, and personalization of the system through health score generation, trend tracking, and feedback generation. This section presents experimental cases, pipeline behavior, generated scores, system response to incomplete data, and sample output excerpts.

4.3.1 Experimental Setup and Evaluation Design

The experimental evaluation of the proposed ML-based framework for personalized health insights from wearable data was carried out in a realistic and comprehensive manner. The primary aim was to assess the system's effectiveness in interpreting multi-dimensional health data, computing meaningful scores, and generating personalized recommendations that reflect an individual's lifestyle patterns. The experimental design focused not only on technical performance but also on interpretability, user relevance, and resilience to noisy or incomplete data.

Data Collection and Participant Diversity For the evaluation, data was collected from a diverse pool of volunteer participants over multiple weeks. Participants included adults across a range of age groups (20 to 45 years), varying occupational and lifestyle backgrounds (e.g., students, working professionals, recreational athletes), and differing levels of health awareness. The goal was to ensure that the system could generalize well across different health behaviors, activity levels, and usage patterns of wearable devices.

Each participant's data was gathered using commercially available wearable devices such as Fitbit, Garmin, or Apple Watch, as well as complementary health tracking apps. The data spanned several health domains, including:

- **Sleep metrics:** total sleep duration, time in REM and deep sleep, number of interruptions, sleep onset and wake times.
- **Cardiovascular indicators:** resting heart rate (RHR), heart rate variability (HRV), and stress scores.
- **Fitness activity:** steps taken per day, active minutes, sedentary duration, and calories burned.
- **Nutritional data:** daily calorie intake, macronutrient composition, and hydration logs.
- **Metabolic readings:** weight, body mass index (BMI), and basal metabolic rate (BMR).

The collected data was anonymized and stored securely in a centralized MongoDB database, ensuring that records were appropriately linked by a unique profile ID but devoid of personally identifiable information.

System Evaluation Goals The core evaluation goals of the experiment were fourfold:

- To validate the system's ability to accurately summarize raw wearable data into meaningful health scores.
- To assess whether the personalized recommendations generated by the language model aligned with user patterns and were perceived as helpful.
- To observe the system's robustness in handling real-world data irregularities, such as missing days, device sync failures, or inconsistent logging habits.
- To evaluate the overall usability and interpretability of the final health report, including visual design, explanation clarity, and personalization.

Evaluation Strategy Each participant was monitored for a period of 3 to 6 weeks, and a weekly summary was computed for each individual. These summaries were generated automatically by the data pipeline described in Chapter 3, and were then passed through the scoring and recommendation engine introduced in this chapter.

During the evaluation period, the following strategies were employed:

- **Trend observation:** The evolution of each user's health scores over multiple weeks was tracked to understand consistency, improvement, or decline in specific domains.

- **Feedback realism:** The natural-language advice generated each week was reviewed for appropriateness, personalization, and motivational tone.
- **Score-aid alignment:** The consistency between numerical scores and generated advice was checked—for instance, whether low sleep scores led to appropriate suggestions on sleep hygiene or stress relief.
- **User interpretation:** Informal interviews and feedback were gathered from participants to determine how understandable, actionable, and relatable they found their reports.

Infrastructure and Workflow The data processing and evaluation workflow was hosted on a standard desktop computing setup, with weekly automated generation of reports. The database was continuously updated with incoming data, and weekly cron jobs triggered scoring, prompt generation, and PDF creation. Reports were shared with users for feedback and observation, though no direct health interventions were conducted.

To ensure fairness and consistency, all evaluations were conducted using a fixed version of the scoring rules and prompt templates. Minor changes in data density or wearable brand differences were accounted for by normalization procedures, ensuring comparability across users.

Nature of the Dataset The dataset resulting from this evaluation setup was highly representative of real-world conditions. It included:

- Users with complete logs and disciplined tracking behavior.
- Users with inconsistent or sparse data entries (e.g., nutrition missing for some weeks).
- Profiles with clear patterns of stress, sleep deprivation, or sudden lifestyle changes (e.g., travel, illness).

This heterogeneity allowed the system to be evaluated not only in ideal conditions but also in edge cases, which is essential for validating a health feedback system meant for widespread adoption.

Evaluation Constraints and Ethics Given the personalized nature of health data, ethical considerations were paramount. Data collection was voluntary, and users had full visibility into how their data was used. The system was presented as a prototype health assistant and not a medical diagnostic tool. Participants were informed that the advice was automatically generated and meant for lifestyle awareness, not for clinical decision-making.

4.4 Summary

In this chapter, we introduced and implemented a comprehensive machine learning framework designed to convert wearable data into personalized health intelligence. The system architecture integrates modular components for data ingestion, domain-wise scoring, personalized recommendation generation, and automated report creation. Each user's weekly data—processed through the pipeline described in the previous chapter—is analyzed across five core health dimensions: sleep, cardiovascular activity, fitness, nutrition, and metabolic stability.

A hybrid scoring approach was adopted, combining expert-informed rules with trend-based logic to compute interpretable weekly health scores. To make the insights actionable and user-friendly, a large language model was used to generate natural language recommendations tailored to each user's behavior. Finally, all elements were compiled into a structured PDF report containing visual summaries, scores, and motivational advice.

This framework not only demonstrates how diverse sensor data can be transformed into meaningful feedback but also sets a foundation for dynamic personalization and long-term health engagement. It provides a scalable solution for digital health reporting that balances technical accuracy with user accessibility.

The next chapter will discuss key findings, real-world observations from diverse user profiles, strengths of the approach, and its current limitations. It will reflect on the interpretability, modularity, and user-centric design of the system. Following that, we outline several future directions for advancing this work, including real-time feedback, integration of mental health data, healthcare system connectivity, and ethical considerations for AI-driven health support.

Chapter 5

Conclusion and Future Direction

The widespread adoption of wearable devices has significantly transformed personal health tracking by providing continuous, real-time monitoring of physiological and behavioral indicators. Despite this progress, many users still face difficulty in interpreting the complex data generated by these devices. This thesis addresses that gap by introducing a complete system that translates raw sensor streams into personalized health assessments and actionable guidance.

The framework developed in this work offers a structured approach to wearable data interpretation. By integrating scoring logic, data processing techniques, and natural language generation, the system produces weekly health summaries that combine numerical insights with narrative feedback. Designed with modularity and extensibility in mind, the platform accommodates a range of health domains—including sleep, fitness, cardio, nutrition, and metabolic signals—allowing for holistic and personalized evaluation.

The system's architecture has been tested across varied user profiles and demonstrated robustness even when faced with incomplete or noisy data. Health scores help users identify trends, while the recommendation engine, powered by language models, transforms insights into engaging, human-readable advice. This dual output—quantitative and narrative—enables more intuitive self-reflection and supports long-term behavioral change.

This chapter concludes the work by outlining its core contributions and presenting forward-looking perspectives. While the current system already offers a solid foundation for personalized health feedback, there are opportunities to improve and expand it. Enhancing the depth of analysis, incorporating machine learning-based personalization, and extending support for new types of wearable signals represent promising directions for future development.

The sections that follow will detail the specific contributions of the thesis and propose enhancements that can further advance the field of intelligent health monitoring.

5.1 Major Contributions

This thesis introduces an integrated, interpretable, and extensible framework for converting raw data from wearable devices into individualized health insights. It bridges the gap between continuous health monitoring and user comprehension through a modular, end-to-end system. The major contributions of this research are outlined below:

- **Development of a Complete End-to-End Interpretation Pipeline**
  A key contribution is the creation of a full-stack system that connects data ac-

quisition, processing, scoring, and reporting. Unlike partial tools that handle only data logging or visualization, this pipeline integrates database access, preprocessing, structured scoring, and personalized natural language feedback [6]. This comprehensive architecture enables seamless interpretation of real-world sensor data into user-facing outputs.

- **Structured Scoring Framework Across Multiple Health Domains**
  The system introduces a domain-wise scoring mechanism that evaluates weekly health behavior across five core areas: sleep, cardiovascular health, fitness, nutrition, and metabolic balance. Each domain score is based on thresholds, variability, and behavioral trends. The scoring design is modular, making it easy to extend or adapt to new metrics, and also allows for a cumulative health score to support comparative tracking.

- **Natural Language Feedback via Language Model Integration**
  This work incorporates a large language model to translate structured summaries into conversational health feedback. The model generates weekly advice in plain, motivational language, helping users interpret trends without requiring technical knowledge. This bridges the gap between numerical analytics and practical, behavior-oriented guidance.

- **Robustness in the Presence of Incomplete or Noisy Data**
  Real-world sensor data is often inconsistent. This system is designed to be tolerant of missing values and irregular logging. Scores are computed only when data sufficiency is met, and the feedback mechanism adapts to acknowledge gaps without discouraging the user. This resilience makes the framework suitable for deployment in everyday, imperfect usage scenarios.

- **Automated Weekly Health Report Generation in PDF Format**
  Another significant contribution is the creation of an automated report generation tool. The system compiles health scores, trend analyses, and language-based recommendations into a cohesive, visually clean PDF format. Features like circular score indicators and sleep distribution charts make the reports easy to understand and suitable for long-term tracking.

- **Validation Across Diverse Real-World User Profiles**
  The system was tested using real users with varying health routines and device usage patterns. It consistently produced reliable summaries and feedback, with users reporting the insights as clear, encouraging, and relevant. These results validate the framework's adaptability to individual variation.

- **Research Contribution to Wearable Data Interpretation**
  From a research perspective, this work advances the field of health informatics by demonstrating a practical, modular framework for transforming raw sensor data into personalized feedback. It lays a foundation for future enhancements such as ML-based anomaly detection, adaptive scoring models, or behavioral recommendation engines.

- **Extensible and Future-Ready Design Philosophy**
  The system architecture is built for extensibility. Each health domain is modular

and can be updated or expanded—for example, to include mental health or respiratory monitoring. Scoring rules and feedback prompts can be customized for different populations, preferences, or clinical integration.

Collectively, these contributions demonstrate that the proposed system is more than a proof-of-concept—it is a scalable, user-aware solution that advances how wearable health data is interpreted and used. It provides a meaningful step toward accessible, personalized digital health systems powered by AI.

5.2 Future Directions

While the system presented in this thesis effectively translates wearable data into personalized health insights, it also lays the groundwork for a wide range of future advancements. As wearable technology continues to evolve and user expectations grow, there are many opportunities to enhance the system's capabilities, increase its adaptability, and broaden its applicability. This section outlines several key directions for future research and development, centered around intelligence, personalization, scalability, integration, and behavioral impact.

- **Incorporating Machine Learning for Data-Driven Scoring**
  Currently, the scoring logic is based on expert-informed thresholds and rules. Future iterations can introduce machine learning models that learn health score mappings from historical user data. These models could detect subtle trends, nonlinear relationships, or seasonality in health behaviors, offering a more dynamic and personalized evaluation.

- **Real-Time Monitoring and Instant Feedback**
  While the system now generates weekly summaries, future versions could offer continuous feedback, delivering timely prompts or micro-recommendations in response to live data trends. Such a feature would enhance real-time behavior correction and promote daily health awareness.

- **Long-Term Personalization and Context-Aware Insights**
  Deeper personalization can be achieved by modeling individual behavior over longer timeframes and incorporating contextual variables such as stress, mood, environment, and lifestyle events. Adaptive baselines, derived from each user's historical patterns, could further improve the relevance of both scores and feedback.

- **Incorporation of Mental and Emotional Well-Being Metrics**
  Expanding beyond physical health, future systems could integrate indicators related to emotional state and mental health—such as stress perception, mood variability, and social interaction—especially as wearables and companion apps begin to collect such data more reliably.

- **Healthcare System Integration and Clinical Use Cases**
  A promising extension involves connecting the system with healthcare providers. Health summaries could be securely shared with doctors, therapists, or fitness experts to support informed consultations. Clinical thresholds and chronic disease markers could also be embedded to flag potential medical risks.

- **Cross-Device Compatibility and Unified Data Integration**
  Supporting multiple devices and platforms is crucial for real-world adoption. Future development could involve syncing data across smartwatches, fitness bands, and health apps using standardized protocols like HL7 or FHIR. This would ensure more comprehensive and consistent data coverage.

- **Population Analytics and Public Health Applications**
  Aggregated, anonymized user data could be analyzed to uncover population-level health patterns. These insights could inform public health initiatives, detect early signs of community health issues, or create regional wellness benchmarks across demographics or seasons.

- **Goal-Oriented and Customizable Recommendations**
  Future systems could allow users to set specific wellness goals—such as increasing stamina, improving sleep, or reducing stress—and receive adaptive guidance tailored to those objectives. Goal tracking and dynamic feedback loops would enhance user engagement.

- **Explainable AI and Ethical Design in Feedback Systems**
  As AI becomes more integral to personal health, ensuring transparency and trust becomes critical. Incorporating explainable AI (XAI) techniques would allow users to understand how scores are derived or why certain advice is given, fostering confidence in the system.

- **Evaluating Long-Term Impact and Behavior Change**
  A vital direction for future research lies in assessing the lasting effects of such systems. Studies could evaluate whether regular feedback leads to sustained behavior change, better health outcomes, or reduced healthcare utilization over time.

In summary, this thesis provides a strong foundation for building intelligent health guidance systems grounded in wearable data. By embracing machine learning, enhancing personalization, expanding health domains, and supporting ethical, user-centric design, future work can further empower individuals to take charge of their health in a meaningful, scalable, and holistic manner.

Bibliography

[1] Singh, S., Gupta, D., and Kaur, M. (2022). *Machine Learning for Healthcare Wearable Devices: The Big Picture*. Journal of Healthcare Engineering, 2022, Article ID 4653923.

[2] Dey, N., Ashour, A. S., and Balas, V. E. (2021). *Leveraging Machine Learning for Personalized Wearable Biomedical Devices: A Review*. Computers in Biology and Medicine, 132, 104336.

[3] Rusu, A., Walton, M. K., and Shakib, S. (2022). *A Scoping Review of the Use of Machine Learning in Health Economics and Outcomes Research: Part 1—Data From Wearable Devices*. Value in Health, 25(12), 2094–2103.

[4] Lee, J., Kim, H., and Lee, S. (2020). *Wearables in medicine: Current and future applications*. Journal of Medical Systems, 44(12), Springer.

[5] Shashikumar, S. P., et al. (2021). *Deep learning models for health prediction using wearable sensor data*. npj Digital Medicine, 4(1), 1–11.

[6] Krishnan, R., Ghosh, D., and Ramanathan, V. (2019). *A personalized health dashboard using wearable data*. Proceedings of the IEEE International Conference on Biomedical Informatics.

[7] Fitbit Inc. (2022). *Fitbit Web API Documentation*. Retrieved from: `https://dev.fitbit.com/build/reference/web-api/`