

LifeTrack: An Integrated Statistical and Probabilistic Framework for Anomaly Detection and Behavioral Prediction in Wearable Health Data

Basavaraj Bhairawadagi, Dr. Basavaraj S Pol

Department of Computer Science and Engineering, R L Jalappa Institute of Technology, Doddaballapur,
Karnataka, India basavarajspol@rljit.in

Abstract— Wearable devices continuously generate activity, sleep, heart-rate, calorie, intensity, and other behavioral measurements. These observations are frequently distributed across heterogeneous files with different schemas, timestamps, and sampling granularities, making longitudinal interpretation difficult. This paper presents LifeTrack, an integrated statistical and probabilistic framework for transforming multi-source wearable-health data into interpretable behavioral information. The framework combines Gaussian z-score analysis for deviation detection, Shannon entropy for behavioral variability, a Behavioral Consistency Index (BCI) for composite lifestyle stability, and Markov-chain modeling for short-term behavioral-state prediction. Up to eighteen Fitbit-style datasets can be processed through a unified temporal preprocessing pipeline. The system creates a daily analytical representation, identifies unusual patterns, quantifies routine irregularity, classifies behavioral conditions into Healthy, Moderate, and Unstable states, and estimates probabilistic state transitions. Optional ARIMA forecasting is incorporated for continuous BCI trend analysis. The framework is intended for research and self-monitoring rather than medical diagnosis. By connecting complementary statistical, information-theoretic, and probabilistic methods, LifeTrack provides a transparent pathway from raw wearable observations to personalized behavioral analytics.

Index Terms— Wearable Health Data, Anomaly Detection, Gaussian Model, Z-Score, Shannon Entropy, Behavioral Consistency Index, Markov Chain, ARIMA, Lifestyle Analytics, Health Monitoring.

I. INTRODUCTION

Wearable technologies have changed the way everyday health-related behavior can be observed. Fitness trackers and smartwatches can continuously record steps, calories, heart rate, activity intensity, sedentary duration, sleep characteristics, and other physiological indicators. These measurements create a longitudinal record that can be analyzed for changes in routine, deviations from normal behavior, and relationships among lifestyle variables. The resulting data can support self-monitoring and research into patterns that may be difficult to recognize from isolated observations.

The availability of data, however, does not automatically produce useful insight. Wearable records are commonly distributed across multiple files and tables, and different sources may use different schemas, sampling intervals, timestamp conventions, and units. Daily summaries may coexist with hourly and minute-level activity records, while sleep and heart-rate measurements may follow separate recording structures. Missing values can represent either a genuine absence of an event or a failure of the device to record it. These issues make direct comparison and longitudinal modeling difficult.

Most consumer-oriented monitoring systems emphasize descriptive dashboards that display steps, calories, sleep duration, goals, and historical trends. Such dashboards are valuable for awareness, but they provide limited analytical support for determining whether an observation is unusual for a particular user, whether behavioral variability is increasing, or how a current behavioral state may evolve over subsequent days. A useful research framework therefore needs to combine data preparation with methods that quantify deviation, variability, consistency, and temporal progression.

LifeTrack is proposed as an integrated statistical and probabilistic framework for this purpose. Instead of depending on a single machine-learning model, the framework connects several interpretable analytical methods around a common daily representation of wearable observations. Gaussian z-score analysis is used to identify measurements that deviate from a reference distribution. Shannon entropy is used to describe the variability of behavioral observations. A Behavioral Consistency Index combines selected lifestyle indicators into a normalized

score, while a Markov-chain model represents transitions among discrete behavioral states. An optional ARIMA module provides continuous BCI trend forecasting.

A major motivation for combining these techniques is interpretability. An end-to-end black-box predictor may provide a final label without clearly showing why the label was generated. LifeTrack instead exposes intermediate quantities such as standardized deviations, entropy values, BCI scores, state labels, and transition probabilities. This makes the framework easier to inspect, debug, compare, and communicate in academic and exploratory settings.

The framework is designed for research, educational experimentation, and personal self-monitoring. It should not be interpreted as a clinical diagnostic system. Medical applications would require independent validation, appropriate data governance, clinically relevant ground truth, and evaluation under the intended conditions of use. The proposed system therefore focuses on the analytical problem of converting heterogeneous wearable observations into understandable behavioral information.

The principal contributions of this work are: (1) a unified preprocessing strategy for heterogeneous wearable-health datasets; (2) transparent Gaussian anomaly scoring; (3) entropy-based measurement of behavioral variability; (4) an interpretable Behavioral Consistency Index; (5) three-state behavioral representation; (6) Markov-chain prediction and stochastic simulation; (7) optional ARIMA forecasting for continuous trends; and (8) an interactive analysis workflow for visualization and export.

II. LITERATURE REVIEW

Gaussian models provide an accessible statistical representation for observations that deviate from a reference distribution. In wearable analytics, a standardized score can provide a direct explanation for why a daily value is unusual. The advantage is transparency, although validity depends on the distributional characteristics of the selected feature.

Information entropy provides another perspective. Shannon's formulation measures uncertainty in a probability distribution and can be adapted to describe variability in behavioral observations. A point anomaly may be absent even when a user's routine becomes increasingly irregular; entropy can capture this broader pattern.

Markov models are useful for sequential data because they represent the probability of moving from one state to another. In a behavioral context, discrete states can summarize daily conditions and provide a compact representation of temporal dynamics. Monte-Carlo simulation can extend the transition model to generate multiple possible trajectories.

ARIMA models are established tools for forecasting temporal numerical series. Their role in LifeTrack is complementary to the Markov model: ARIMA estimates the future direction of a continuous BCI sequence, while the Markov component focuses on discrete behavioral states.

The reviewed concepts motivate an integrated framework rather than a single predictive algorithm. LifeTrack combines these methods around a common daily wearable representation and preserves the interpretability of each analytical stage.

III. PROPOSED LIFETRACK FRAMEWORK

LifeTrack accepts up to eighteen Fitbit-style datasets, including daily activity, hourly activity, minute-level activity, sleep, heart-rate, intensity, and weight information. Each source is processed independently before being converted to a standardized analytical representation. This reduces the risk that differences in source schemas propagate into the modeling stage.

The preprocessing module detects date or timestamp fields, normalizes temporal formats, standardizes feature names, handles missing values, aligns observations, and aggregates fine-grained records. The resulting daily table acts as the common input for anomaly analysis, entropy estimation, BCI calculation, state classification, and prediction.

The complete workflow is organized as data acquisition, temporal preprocessing, feature engineering, Gaussian anomaly analysis, entropy estimation, BCI calculation, state classification, Markov prediction, optional ARIMA forecasting, and visualization/export.




Fig. 1. Integrated LifeTrack processing architecture.

A. Gaussian Anomaly Detection

For an observation x , the standardized deviation is calculated from its reference mean μ and standard deviation σ . This converts a raw measurement into a scale indicating how far it lies from the historical center.

$$z = (x - \mu) / \sigma \quad (1)$$

A practical threshold such as $|z| > 2$ can be used to flag unusual observations, although thresholds should be calibrated to the selected dataset. Positive and negative scores indicate the direction of the deviation.

B. Shannon Entropy

For a discrete distribution with probabilities $p(x_i)$, entropy is calculated as:

$$H(X) = - \sum p(x_i) \log p(x_i) \quad (2)$$

Higher entropy indicates a more distributed or variable pattern, while lower entropy represents a more concentrated distribution.

Entropy therefore complements z-score analysis by summarizing behavioral variability instead of only identifying individual extreme observations.

C. Behavioral Consistency Index

The BCI converts multiple behavioral measurements into one interpretable score. For metric M_i , goal G_i , and weight w_i , the weighted goal-normalized formulation is:

$$BCI = [\sum w_i \cdot \min(1, M_i/G_i) / \sum w_i] \times 100 \quad (3)$$

The normalization prevents an individual metric from contributing beyond its selected target. The resulting score can summarize daily consistency and can subsequently be mapped to behavioral states.

D. Behavioral-State Construction

Daily behavior is represented using Healthy, Moderate, and Unstable categories. These categories are analytical states rather than clinical diagnoses. Their thresholds should be selected and validated according to the dataset, study objective, and intended population.

E. Markov State Prediction

The transition probability from state i to state j is:

$$P_{ij} = P(S_{t+1} = j | S_t = i) \quad (4)$$

Observed state sequences are converted into a transition matrix. The next state can be selected according to the largest transition probability, while

F. ARIMA Trend Forecasting

ARIMA is an optional extension for the continuous BCI series. It can model temporal dependence and forecast future BCI values, complementing the discrete Markov state representation.

IV. SYSTEM IMPLEMENTATION

The framework can be implemented using Python-based data and visualization libraries. Pandas and NumPy support data loading, cleaning, aggregation, merging, and numerical operations. SciPy can support statistical calculations, Plotly can provide interactive visualizations, and Statsmodels can be used for optional ARIMA forecasting. Streamlit provides the interactive application layer.

The system is organized into input, preprocessing, analytics, prediction, visualization, and export components. This modular implementation allows each analytical component to be tested independently and makes future substitution of algorithms practical.

Layer	Implementation	Role
Input	CSV upload and multi-file ingestion	Collect wearable records
Preprocessing	Pandas, date parsing, schema mapping	Align heterogeneous sources
Analytics	NumPy, SciPy, custom formulas	Compute anomaly, entropy and BCI
Prediction	Markov matrix and optional ARIMA	Estimate future states/trends
Interface	Streamlit and Plotly	Interactive analysis

TABLE I. LIFETRACK SOFTWARE ARCHITECTURE.

The interface is intended to expose intermediate results rather than only a final state. Typical outputs include daily measurements, anomaly flags, entropy values, BCI values, state labels, transition probabilities, forecast trajectories, and downloadable analytical tables.

V. EXPERIMENTAL DESIGN AND TESTING

The project material describes functional testing of the complete pipeline. The tests focus on file ingestion, temporal normalization, anomaly behavior, entropy and BCI calculations, Markov transitions, user-interface interaction, and export. These tests establish functional feasibility rather than clinical performance.

Test	Scenario	Expected Result
TC01	Mixed wearable CSV files	Successful parsing and daily merging
TC02	Different timestamp formats	Dates normalized correctly
TC03	Injected outlier values	Gaussian method flags unusual records
TC04	Behavioral variability data	Entropy and BCI remain valid
TC05	Three-state sequence	3×3 transition matrix generated
TC06	Interactive dashboard	Charts, controls and export work

TABLE II. FUNCTIONAL TESTING OF THE PROPOSED SYSTEM.

Functional testing alone does not establish generalization or clinical validity. A stronger experimental study should use a larger longitudinal cohort, predefined anomaly labels, baseline methods, cross-validation where appropriate, sensitivity and specificity analysis, calibration, and uncertainty reporting.

VI. RESULTS AND DISCUSSION

The project results indicate that LifeTrack can transform multiple wearable-data sources into a unified daily analytical table and calculate statistical and probabilistic indicators. The integration provides complementary interpretations: z-scores identify unusual individual measurements, entropy captures variability across behavioral observations, and BCI summarizes multiple lifestyle dimensions.

The resulting state representation adds temporal context. Instead of treating every day independently, the Markov component considers how a behavioral condition follows another condition. This enables the system to estimate the most likely next state and simulate alternative short-term trajectories.

One important advantage is interpretability. A user can inspect the feature contributing to an anomaly, the standardized deviation, entropy, BCI, state category, and transition probabilities. This makes the system appropriate for educational and exploratory research settings where understanding the analytical path is as important as prediction.

The framework also has limitations. Wearable variables may be skewed, multimodal, or affected by device-specific measurement behavior, which can reduce the suitability of Gaussian assumptions. Entropy estimates depend on how continuous values are discretized. BCI results depend on goal values and weights. Markov predictions assume that historical transition behavior remains informative for future observations.

VII. COMPARATIVE ANALYSIS

Capability	Basic Dashboard	LifeTrack
Multi-source wearable data	Limited	Integrated
Historical deviation	Trend only	Z-score anomaly
Behavioral variability	Basic charts	Entropy
Composite stability	Goal display	BCI
Discrete state prediction	Usually absent	Markov chain
Continuous trend forecast	Limited	Optional ARIMA
Interpretability	High	High

TABLE III. COMPARISON OF BASIC MONITORING AND LIFETRACK ANALYTICS.

The comparison highlights that the proposed framework is not simply a visualization layer. Its distinction is the combination of historical deviation, variability, composite scoring, and sequential prediction. These capabilities coexist in one workflow while retaining relatively simple mathematical explanations.

VIII. APPLICATIONS

LifeTrack can be applied to personal activity monitoring, sleep-pattern exploration, behavioral variability analysis, and longitudinal wearable-data research. In activity monitoring, unusual changes in steps or intensity can be detected relative to historical behavior. In sleep analysis, duration and related

measurements can be examined for deviations and variability.

The framework can also be used as an academic platform for demonstrating statistical modeling, entropy, Markov chains, time-series forecasting, data preprocessing, and interactive visualization. Its modular structure allows researchers to substitute alternative anomaly detectors or forecasting methods and compare results.

For personalized analytics, historical observations can be used as an individual's reference distribution rather than applying only population-level thresholds. This can make indicators more relevant to personal routines, although personalization must be carefully validated.

IX. PRIVACY, RELIABILITY AND ETHICAL CONSIDERATIONS

Wearable-health data may contain sensitive behavioral and physiological information. A practical deployment should minimize unnecessary data retention, protect uploaded and exported files, restrict access, and avoid exposing personally identifiable information in charts or logs.

Reliability is closely linked to data quality. Missing timestamps, sensor dropout, device changes, inconsistent measurement frequencies, and unusual recording periods can influence statistical indicators. Data-quality checks should therefore be presented alongside analytical outputs.

The system should clearly distinguish behavioral analytics from medical conclusions. Labels such as Healthy, Moderate, and Unstable are computational states defined by the model and should not be interpreted as clinical diagnoses.

X. FUTURE RESEARCH

Future work can compare Gaussian z-score detection with robust statistical techniques such as median absolute deviation, interquartile-range methods, and multivariate distance measures. Such comparisons can determine which approaches are most reliable for different wearable-variable distributions.

Adaptive entropy estimation can reduce dependence on fixed discretization rules. Personalized BCI weights can be learned from individual history while preserving interpretability. Hidden Markov models and more expressive sequence-learning architectures can be evaluated for longer temporal dependencies.

Explainable artificial intelligence can identify which features contribute most to predicted state changes.

Real-time streaming can extend the current batch-oriented workflow to continuous monitoring. Privacy-preserving and federated approaches can support multi-user research without unnecessarily centralizing sensitive raw records.

XI. CONCLUSION

This paper presented LifeTrack, an integrated statistical and probabilistic framework for anomaly detection and behavioral prediction using heterogeneous wearable-health data. The system connects data preprocessing, Gaussian z-score analysis, Shannon entropy, a Behavioral Consistency Index, behavioral-state classification, Markov prediction, and optional ARIMA forecasting.

The central contribution is the combination of complementary analytical perspectives in a single interpretable pipeline. Gaussian analysis provides point-wise deviation information, entropy measures broader variability, BCI summarizes multiple behavioral dimensions, and Markov modeling adds sequential prediction. This creates a structured pathway from raw wearable observations to understandable behavioral information.

The framework provides a practical foundation for research and personalized lifestyle analytics. Its modular architecture can be extended with robust anomaly detection, adaptive personalization, explainable machine learning, real-time streaming, multimodal sensing, and privacy-preserving computation. The system should remain an analytical and research tool unless its predictions are independently validated for a specific health or clinical application.

ACKNOWLEDGMENT

The authors sincerely thank the Department of Computer Science and Engineering, R L Jalappa Institute of Technology, Doddaballapur, Karnataka, for providing academic support and research facilities.

REFERENCES

- [1] Y. Wu, Y. Feng, and Z. Lin, "Multivariate time-series classification using Gaussian models and KL-divergence metrics," *Proc. AAAI Conf. Artif. Intell.*, vol. 33, no. 1, pp. 6402–6409, 2019.
- [2] Y. Wang, S. Li, and G. Zhao, "Differential evolution with entropy-based adaptive mutation strategy guided by a Markov chain," *IEEE Access*, vol. 9, pp. 116825–116838, 2021.
- [3] W. Huang, W. Zhang, G. Zeng, and G. Zhou, "A short-term load forecasting model of power system

based on ARMA and Kalman filter,” Energy Rep., vol. 3, pp. 51–56, 2017.

- [4] C. E. Shannon, “A mathematical theory of communication,” Bell Syst. Tech.J., vol. 27, pp. 379–423, 623–656, 1948.
- [5] L. Rabiner, “A tutorial on hidden Markov models and selected applications inspeech recognition,” Proc. IEEE, vol. 77, no. 2, pp. 257–286, 1989.
- [6] J. D. Hamilton, Time Series Analysis. Princeton, NJ, USA: Princeton Univ.Press, 1994.
- [7] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, Time SeriesAnalysis: Forecasting and Control, 4th ed. Hoboken, NJ, USA: Wiley, 2008.
- [8] T. Hastie, R. Tibshirani, and J. Friedman, The Elements of StatisticalLearning, 2nd ed. New York, NY, USA: Springer, 2009.
- [9] L. Breiman, “Random forests,” Mach. Learn., vol. 45, pp. 5–32, 2001.
- [10] A. Hyvärinen, “Independent component analysis: Recent advances,” Philos. Trans. Roy. Soc. A, vol. 371, 2013.