

Zero-Shot Heart Rate Variability Forecasting from Consumer Wearables Using Time Series Foundation Models

Luukas Peräkylä[†] , Fahad Sohrab[†] , Ville Hautamäki^{*} , Merja Heinäniemi^{*} , Sui Huang^{*} ,
Pekka Abrahamsson[†] 

[†]Tampere University, Tampere, Finland

^{*}University of Eastern Finland, Finland

Abstract

Short-term Heart Rate Variability (HRV) forecasting could provide clinicians with actionable lead time for detecting autonomic dysfunction and adverse cardiac events. Consumer wearable devices generate fragmented, artifact-rich HRV signals that challenge conventional forecasting approaches. In this study, we evaluated the forecasting ability of three Time Series Foundation Models (TSFMs), TimesFM, Chronos, and MOIRAI, against traditional baselines (Mean, Exponential Smoothing, and Exponentially Weighted Moving Average) on real-world wearable data collected from 49 healthy individuals. To address data fragmentation, we introduce a variability-preserving imputation method that augments linear interpolation with locally adaptive stochastic noise, retaining physiological dynamics essential for accurate forecasting. The results show that TSFMs outperformed all baselines without fine-tuning, achieving average Mean Absolute Scaled Error (MASE) between 0.81 and 0.87 across TSFMs and both context lengths (32 and 64 time steps), with Chronos and TimesFM as the top models, though MOIRAI showed limited gains over baselines. With up to a 2-hour forecast horizon, the results establish a baseline for TSFMs' performance on a real-world dataset, highlighting domain-specific fine-tuning as a promising direction for clinical deployment.

1. Introduction

Health care wearables, from smartwatches to biosensors, enable continuous and real-time physiological monitoring, tracking heart rate, sleep, and stress among other metrics, supporting personalized treatments, and can be used for the detection of diseases, such as atrial fibrillation [1]. A promising direction within this space is the tracking of Heart Rate Variability (HRV) as a non-invasive tool to improve personal health outcomes [2]. HRV reflects autonomic nervous system function and cardiovascu-

lar health, with higher HRV associated with better physiological adaptability and lower HRV indicating elevated risk [3, 4]. Accurate short-term HRV forecasting could therefore enable early alerting of autonomic dysfunction and adverse cardiac conditions, providing clinicians and monitoring systems with additional time for intervention [3].

The clinical value of HRV as a predictive signal is well established. Classifiers evaluating extracted HRV features have been shown to anticipate acute clinical events prior to symptom onset, such as paroxysmal atrial fibrillation [5] and ventricular fibrillation [6]. Using non-linear HRV features such as multiscale entropy, machine learning classifiers have successfully distinguished asymptomatic patients in the preclinical stages of cardiac autonomic neuropathy from healthy individuals [7]. Previous work has also validated deep learning on wearable data for detecting a broader range of medical conditions, including diabetes, high cholesterol, high blood pressure, and sleep apnea [8].

However, these studies largely rely on high-fidelity, continuous clinical electrocardiogram (ECG) recordings rather than consumer-grade wearable signals. Translating HRV-based prediction to remote monitoring settings introduces significant challenges: data is inherently irregular and noisy due to device charging, motion artifacts, and non-wear periods, while limited continuous recording windows constrain the available context. Robust wearable-derived HRV can successfully predict major adverse cardiovascular events [9], and continuous photoplethysmography (PPG) measurements allow for real-time ischemic risk assessment, though mitigating signal noise remains a hurdle for these task-specific models [10].

To overcome the limitations of task-specific, single-modality approaches, generalized Time Series Forecasting Models (TSFMs) provide an architecture explicitly designed to process heterogeneous, multivariate physiological signals generated by modern consumer devices [11]. Models such as TimesFM [12], Chronos [13], and MOIRAI [14] are designed for robust zero-shot generalization across diverse domains without task-specific fine-




Figure 1: Data preprocessing pipeline for HRV time series, showing segmentation, gap handling, and imputation steps used to generate the final dataset for modelling.

tuning, making them natural candidates for wearable HRV forecasting. This study investigates whether these models can accurately forecast wearable HRV without fine-tuning, how context and horizon length affect forecasting accuracy, and whether they outperform classical statistical baselines on real-world data.

2. Methodology

This section describes the dataset, preprocessing pipeline, and experimental setup used to evaluate the forecasting models. We first outline the characteristics of the wearable HRV data and the imputation strategy used to handle missingness, followed by the models and evaluation protocol.

2.1. Data and Imputation

We evaluated models on a public real-world HRV dataset comprising wearable-based measurements from 49 healthy individuals collected over four weeks using Samsung Galaxy Watch Active 2 devices [15]. The device uses Photoplethysmography (PPG), a common but noise-susceptible method for digital biomarker measurement [2]. Among HRV time-domain metrics, the Root Mean Square of Successive Differences (RMSSD) robustly captures short-term parasympathetic activity and is particularly suited for short, PPG-derived wearable recordings [3, 4]. Only pre-filtered (missingness score ≤ 0.35) RMSSD data was used, computed from 5-minute PPG segments.

The dataset reflects the highly fragmented nature of real-world consumer wearable data. After applying a 15-minute gap threshold, participants contributed a mean of 109 continuous segments each, with longest continuous recordings averaging 4.4 hours and a maximum of approximately 14.75 hours, driven by nightly charging cycles, motion artifacts, and non-wear periods [15]. The chosen context lengths of 32 and 64 steps (2.7 and 5.3 hours) reflect these real-world recording limitations rather than arbitrary design choices.

To generate regularly sampled time series suitable for the models, smaller gaps were imputed. Standard imputation methods (linear interpolation, KNN, cubic splines, and forward-fill) can produce unrealistically smooth sequences that affect both metrics and model learning. To address this, we augmented linear interpolation with locally adaptive stochastic noise:

$$\hat{x}_t = x_t^{\text{linear}} + \epsilon_t. \quad (1)$$

Here $\hat{x}_t$ is the imputed value at time t , x_t^{linear} is the linearly interpolated value, and ϵ_t is zero-mean Gaussian noise with standard deviation $\sigma_{\text{local}}(1 + |\beta|)$. Local standard deviation σ_{local} is computed from a 5-step window, β represents the local slope, and noise is clipped at $3\sigma_{\text{local}}$ to prevent outliers. The multiplicative term $(1 + |\beta|)$ adaptively modulates noise variance based on local slope magnitude, preserving physiological variability structure essential for accurate forecasting. The preprocessing pipeline is illustrated in Figure 1.




Figure 2: MASE across forecast horizons for context length 32 (left) and 64 (right). The dashed line denotes the naive forecasting threshold (MASE = 1.0). ETS exceeds the plot range at longer horizons in context 32.

2.2. Models and Experimental Setup

Three foundation models were evaluated without fine-tuning: TimesFM-1.0-200m [12], Chronos-bolt-base [13], and MOIRAI-1.0-R-base [14]. These were benchmarked against traditional baselines: Mean, Exponentially Weighted Moving Average (EWMA, $\alpha=0.3$) [16], and Exponential Smoothing (ETS) with additive damped trend. Two context lengths were tested: 32 steps (2.7 hours) and 64 steps (5.3 hours), each step representing a 5-minute RMSSD average. Forecast horizons ranged from 1 to 24 steps (up to 2 hours ahead). Probabilistic model outputs (Chronos, MOIRAI) were reduced to the median for point-forecast comparability with TimesFM. Evaluation metrics were Mean Absolute Error (MAE), Root Mean Square Error (RMSE) [16], and Mean Absolute Scaled Error (MASE) [17].

As this is a zero-shot evaluation, no training or fine-tuning was performed on any subset of this dataset; the entire dataset serves as the evaluation set for all foundation models. The classical baselines require no training: EWMA uses a fixed smoothing factor of $\alpha = 0.3$ chosen a priori. ETS smoothing parameters are fitted independently on each context window with no shared training across subjects, ensuring no cross-subject data leakage.

3. Results

Tables 1 and 2 summarise average forecast performance across all horizons. Chronos-bolt-base achieves the lowest MAE, RMSE, and MASE at context length 32, while TimesFM-200m achieves the best MAE and MASE at context length 64. Increasing context length from 32 to 64 improves performance for all foundation models, with

TimesFM gaining the largest improvement and surpassing Chronos at the longer context. Among traditional baselines, EWMA remains competitive at short horizons but is consistently outperformed by foundation models for longer predictions. MOIRAI did not significantly outperform the Mean baseline at either context length ($p=0.953$ and $p=0.794$ respectively), which may reflect sensitivity to the short, highly fragmented segment lengths characteristic of consumer wearable data, where its patch-based architecture may not have sufficient continuous context to operate effectively.

Table 1: Average performance for context length 32, over all forecast horizons (1–24).

Model	MAE	RMSE	MASE
Mean	16.74	<u>26.04</u>	0.87
EWMA	16.19	26.18	0.84
ETS	18.37	34.80	0.95
MOIRAI	16.73	26.68	0.87
TimesFM	<u>16.16</u>	26.55	0.84
Chronos	15.74	25.64	0.82

Table 2: Average performance for context length 64, over all forecast horizons (1–24).

Model	MAE	RMSE	MASE
Mean	15.99	25.80	0.85
EWMA	15.77	26.34	0.84
ETS	16.15	27.88	0.86
MOIRAI	15.94	26.34	0.85
TimesFM	15.32	26.41	0.81
Chronos	<u>15.44</u>	<u>26.24</u>	<u>0.82</u>

Figure 2 illustrates performance trends across all horizons, confirming that foundation models maintain their

advantage throughout and remain below the MASE 1.0 threshold relative to naive forecasting. Forecasting accuracy decreased only modestly with longer horizons, likely because the maximum horizon of 24 steps remains short relative to HRV’s chaotic dynamics, and because averaging error metrics over more forecast points reduces the impact of individual outliers.

Paired t-tests (Table 3) confirm that differences between foundation models and traditional baselines are statistically significant for most comparisons. Chronos outperforms EWMA across both contexts ($p < 0.001$). Differences between TimesFM and Chronos are significant only at context 32, with the two models becoming comparable at context 64 as TimesFM leverages the longer history more effectively.

Table 3: Statistical significance of key model comparisons (paired t-test, $n=24$ horizons, $\alpha=0.05$). Δ MAE is used as the test statistic; negative values indicate Model 1 is better.

Model 1	Model 2	Context 32		Context 64	
		Δ MAE	p-val	Δ MAE	p-val
TimesFM	Chronos	0.42	< 0.001*	-0.12	0.057
TimesFM	EWMA	-0.03	0.450	-0.46	< 0.001*
TimesFM	Mean	-0.58	< 0.001*	-0.67	< 0.001*
Chronos	EWMA	-0.45	< 0.001*	-0.33	< 0.001*
Chronos	Mean	-0.99	< 0.001*	-0.55	0.001*
MOIRAI	EWMA	0.54	< 0.001*	0.17	0.121
MOIRAI	Mean	-0.01	0.953	-0.04	0.794
EWMA	Mean	-0.54	< 0.001*	-0.21	0.141

4. Conclusions

HRV is a valuable non-invasive indicator of autonomic function and cardiovascular health, increasingly accessible through consumer wearable devices. In this study, we evaluated the zero-shot forecasting capability of TSFMs on real-world wearable HRV data, benchmarking them against traditional statistical models. Our results demonstrate that foundation models outperform traditional baselines, with Chronos-bolt-base excelling at shorter contexts and TimesFM-200m at longer contexts. Importantly, foundation models maintain superior accuracy at longer forecast horizons, enabling earlier prediction of adverse autonomic events and providing greater lead time for clinical intervention. These results suggest Chronos and TimesFM are viable candidates for deployment in consumer cardiac monitoring pipelines with domain-specific fine-tuning, while MOIRAI’s sensitivity to data fragmentation warrants further investigation. A key limitation is that the evaluation was restricted to healthy individuals. The primary remaining step is validation on patient cohorts with known cardiac conditions, which this study’s baseline performance on healthy individuals is designed to support.

References

- [1] Huhn S, Axt M, et al. The impact of wearable technologies in health research: Scoping review. *JMIR mHealth and uHealth* 2022;10(1):e34384.
- [2] Li K, Cardoso C, et al. Heart rate variability measurement through a smart wearable device. *Int J Environ Res Public Health* 2023;20(24):7146.
- [3] Shaffer F, Ginsberg JP. An overview of heart rate variability metrics and norms. *Front Public Health* 2017;5:258.
- [4] Task Force of the ESC and NASPE. Heart rate variability: Standards of measurement, physiological interpretation, and clinical use. *Circulation* 1996;93(5):1043–1065.
- [5] Maier C, Bauch M, Dickhaus H. Screening and prediction of paroxysmal atrial fibrillation by analysis of heart rate variability parameters. In *Computers in Cardiology* 2001. Vol. 28 (Cat. No. 01CH37287). IEEE, 2001; 129–132.
- [6] Farhadi Sedehi J, Jafarnia Dabanloo N, Attarodi G, Eslamizadeh M. Prediction of ventricular fibrillation from hrv signals using combination of genetic algorithm and neural networks. In *2017 Computing in Cardiology (CinC)*. IEEE, 2017; .
- [7] Cornforth D, Tarvainen MP, Jelinek HF. Automated selection of measures of heart rate variability for detection of early cardiac autonomic neuropathy. In *2014 Computing in Cardiology Conference (CinC)*. IEEE, 2014; 93–96.
- [8] Ballinger B, Hsieh J, Singh A, et al. Deepheart: semi-supervised sequence learning for cardiovascular risk prediction. In *Proc. AAAI*. 2018; 253–261.
- [9] Orini M, Flores JL, Chaturvedi N, Hughes AD. Wearable-derived long-term heart rate variability predicts major adverse cardiovascular events in middle aged individuals without previous cardiovascular disease. In *2023 Computing in Cardiology (CinC)*. IEEE, 2023; .
- [10] Tasmurzayev N, Amangeldy B, Imankulov T, Imanbek B, Postolache OA, Konysbekova A. A wearable iot-based measurement system for real-time cardiovascular risk prediction using heart rate variability. *Eng* 2025;6(10):259.
- [11] Luo Y, Chen Y, Salekin A, Rahman T. Toward foundation model for multivariate wearable sensing of physiological signals. *arXiv preprint arXiv241209758* 2024;.
- [12] Das A, Kong W, Sen R, Zhou Y. A decoder-only foundation model for time-series forecasting. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*. PMLR, 2024; 10148–10167.
- [13] Ansari AF, Stella L, et al. Chronos: Learning the language of time series. *Trans Mach Learn Res TMLR* 2024;.
- [14] Woo G, Liu C, Kumar A, et al. Unified training of universal time series forecasting transformers. In *Proc. 41st Int. Conf. Mach. Learn. (ICML)*, volume 235 of *PMLR*. 2024; 53411–53435.
- [15] Baigutanova A, Park S, Constantinides M, et al. A continuous real-world dataset comprising wearable-based heart rate variability alongside sleep diaries. *Sci Data* 2025; 12:1474.
- [16] Holt CC. Forecasting seasonals and trends by exponentially weighted moving averages. *Int J Forecast* 2004;20(1):5–10.
- [17] Hyndman RJ, Koehler AB. Another look at measures of forecast accuracy. *International Journal of Forecasting* 2006;22:679–688.

Address for correspondence:

Fahad Sohrab (fahad.sohrab@tuni.fi)
P.O. Box 553, FI-33014, Tampere Finland