# PHENOTYPING SLEEP FROM SMART-RING TRAJECTORIES: UNSUPERVISED DISCOVERY OF FOUR LONGITUDINAL SLEEP PATTERNS IN A 456-USER INDIAN COHORT

**Subhabrata Paul[1], Dr. Madhurina Saha[2], Manosij Sengupta[3], Ankit Saha[4]**
[1]Co-Founder, CIO, CareNow Healthcare Pvt Ltd, Kolkata, India.
[2]MBBS, General Physician, Kolkata, India.
[3]Co-Founder, COO, CareNow Healthcare Pvt Ltd, Kolkata, India.
[4]Business Operations Specialist, CareNow Healthcare Pvt Ltd, Kolkata, India.
**Email:**[1]subhabrata@careplix.com,[2]madhu.saha.7@gmail.com,[3]manosij@careplix.com,[4]ankit@careplix.com

## ABSTRACT

Traditional sleep statistics condense weeks' worth of data into one night's value, flushing the organization of restorative across days. Unsupervised clustering was applied to longitudinal sleep-stage trajectories derived from a smart ring (Loop Ring, CarePlix) worn continuously by 456 adults in India over four months between Jan 2026 and Apr 2026. After cleaning, we retained 6,006 nights and limited discovery of phenotypes to the 75 subjects with ≥25 valid nights. Using K-means clustering on an 11-dimensional trajectory feature vector, four phenotypes were identified: Circadian Anchor (28%), Persistent Drifter (32%), Weekend Compensator (24%), and Restless Sleeper (16%). Only 25.2% achieved the adult recommendation of ≥7 h per night while 49.4% of nights were <6 hours. The classic textbook "rebound" sawtooth pattern failed to cluster discretely; rebound events were small and occasional (≈1% of nights). The phenotypic unit of importance for digital health is not the arithmetic mean but rather the geometry of a month of sleep.

**Keywords:** Digital Chronobiology, Smart Ring Wearable, Sleep Staging, Unsupervised Clustering, Time-Series Phenotyping, Sleep Regularity, Social Jetlag, India.

## INTRODUCTION

Adults sleep, on average, 6.5–7.5 h a night. That number is not wrong — but it is the wrong unit. Sleep is enacted across weeks, shaped by occupation, family, and the working week. Two people whose 30-day means are identical can have radically different sleep biology: one sleeps a steady 6.4 h every night; the other crashes through five-hour weekdays and recovers with nine-hour Saturdays. The mean does not see the difference. The body does.

Smart-ring wearables have made longitudinal stage-resolved sleep estimation broadly available. These devices are not polysomnographic-grade but they are dense, and density unlocks a different question: what does the shape of a person's month of sleep tell us, beyond what the average tells us?

We treat each user's nightly sequence as a single trajectory and cluster the trajectories. The hypothesis we begin with is a four-phenotype framework — a stable Anchor, a debt-and-rebound Sawtooth, a weekend-shifting Social Jetlagger, and a fragmentation-dominated Restless type. The point of running an unsupervised algorithm on real data is to test whether such typology survives contact with the wild. We report what did and did not.

*www.ajmrhs.com*
eISSN: 2583-7761

Date of Received: 10-06-2026
Date Acceptance: 19-06-2026
Date of Publication: 18-07-2026

Two contributions follow. First, an interpretable feature pipeline that turns nightly stage outputs into a small trajectory vector deployable on any wearable platform. Second, the application of this pipeline to an Indian working-age cohort — a population conspicuously under-represented in published wearable-sleep research — and a baseline that is, frankly, more sobering than the global headlines: half of all valid nights here fell below six hours.

## METHODS

### A. Dataset and pre-processing

The raw export covered 456 users with one or more recorded nights between Jan 2026 and Apr 2026, totalling 9,069 user-night records. Each record stored REM, deep, light, and awake durations (ms), an activity-day date, and a sync timestamp. Stage-level epoch data were not available; we worked with nightly totals. Records with zero total sleep (n = 3,011) were dropped as non-wear. Nights with total sleep below 2 h or above 14 h (n = 42 above the upper bound) were excluded as physiologically implausible — almost certainly device-on-table artefacts. Stage durations exceeding 5 h of REM or 5 h of deep were rejected on the same grounds. The cleaned dataset retained 6,006 nights from 456 users.

### B. Cohort selection and feature engineering

Phenotype discovery requires enough nights per user to estimate within-person variability. We retained users with ≥25 valid nights, yielding a 75-user cohort with a median of 34 nights per user (mean 36.8). All clustering was conducted on this cohort; the wider 456-user, 6,006-night dataset is used only for population descriptives.

Each cohort user was reduced to an 11-dimensional feature vector. Duration features: mean, SD, and CV of total sleep. Architecture features: mean REM% and deep%. Fragmentation features: mean and SD of awake/(awake+TST). Timing irregularity: circular variance of sync-time-of-day on the unit circle (1 − R). Trajectory shape: lag-1 Pearson autocorrelation of nightly duration and a sawtooth-rebound rate (fraction of nights exceeding the user's mean by >1 SD after ≥2 consecutive nights below mean − 0.5 SD). A weekend-Δ feature captured the Saturday/Sunday minus weekday duration difference; in the South-Asian context we treated Friday-night and Saturday-night sleeps as the weekend pair.

### C. Clustering

Features were Z-standardised across the 75-user cohort. We ran k-means with k from 2 to 7 (50 random initialisations each) and selected k by combining silhouette inspection with theoretical alignment to the four-phenotype framework. Silhouette peaked at k = 2 (0.197) and was relatively flat between k = 3 and k = 7; the value at k = 4 was 0.157. We selected k = 4: the silhouette curve is shallow, there is no compelling internal-validity argument for any particular k, and we wanted to test the four-phenotype hypothesis directly. Modest silhouette is the rule, not the exception, in mixed-population wearable data, and we report it transparently rather than over-claim. The final solution used 200 random initialisations. Phenotype names were assigned post-hoc from the Z-score profile of each centroid.

## RESULTS

### A. Population baseline

The 6,006-night cleaned dataset is summarised in Fig. 1. Mean total sleep was 5.91 ± 1.69 h, median 6 h; 49.4% of all valid nights fell below six hours and only 25.2% met or exceeded the ≥7-hour adult recommendation. Mean REM proportion was 24.9%, mean deep proportion 30.3% and mean fragmentation 8.7%.

Two observations follow. First, this cohort is meaningfully under-slept by recommendation standards, and the under-sleep is the population median — not an extreme tail. Second, the 30.3% deep-sleep figure exceeds typical polysomnographic norms (13–23% of TST). This is a finding about the algorithm, not about Indian biology: consumer rings, including Loop Ring's current staging model, tend to over-allocate to the deep class. We therefore treat absolute deep% as a within-device metric and rely on relative differences across phenotypes for inference.

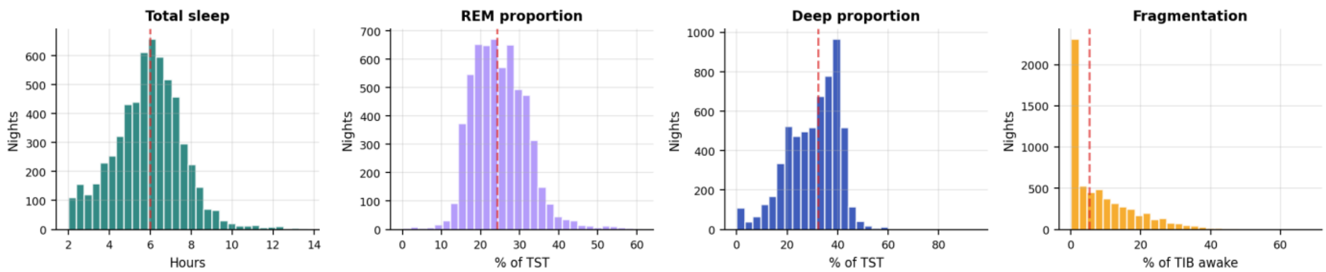


**Fig. 1.** Per-Night Distributions across All 6,006 Cleaned Nights from 456 Users. Median Lines Marked. Note The Strong Left-Skew Of Total Sleep Duration And The Long Tail Of Fragmentation.

**Table 1.** Cluster Centroids on Key Trajectory Features (Cohort N = 75).

| Phenotype | n (%) | Mean dur (h) | SD dur (h) | Deep % | REM % | Frag % | Wkd-Δ (h) | ACF(1) |
|---|---|---|---|---|---|---|---|---|
| Circadian Anchor | 21 (28%) | 6.38 | 1.04 | 35.5% | 22.2% | 4.5% | -0.19 | -0.099 |
| Persistent Drifter | 24 (32%) | 5.59 | 1.25 | 34.0% | 22.6% | 10.1% | +0.05 | +0.131 |
| Weekend Compensator | 18 (24%) | 6.46 | 1.47 | 24.4% | 28.4% | 7.0% | +0.47 | +0.019 |
| Restless Sleeper | 12 (16%) | 4.68 | 1.32 | 24.2% | 28.6% | 18.3% | +0.01 | +0.003 |

### B. The four phenotypes

Cluster sizes, centroids, and signature features are summarised in Table I. Each phenotype's Z-scored signature on every feature is shown in Fig. 2.

The Circadian Anchor (n = 21, 28%) sleeps 6.38 h per night with within-person SD of only 1.04 h — CV of 0.165, the lowest in the cohort. Fragmentation is also lowest at 4.5% of bed-time. The weekend-Δ is mildly negative (−0.19 h), suggesting the weekday schedule is the optimum and weekend timing reflects social commitments rather than recovery imperative. The closest analogue in this dataset to a textbook good sleeper.

The Persistent Drifter (n = 24, 32%) is the largest cluster and the most subtle. Mean duration is 5.59 h — below median, but not dramatically — and within-person variability is unremarkable. What distinguishes this cluster is the lag-1 autocorrelation of +0.131, the highest in the cohort: a short night is followed by another short night, with relatively few abrupt transitions. There is no weekend rebound (Wkd-Δ ≈ 0). This is the phenotype the population mean conceals best — chronic, persistent restriction without compensatory recovery.

The Weekend Compensator (n = 18, 24%) maps closely to the construct of "social jetlag." Mean duration 6.46 h — equal to the Anchor — but SD is 1.47 h, well above cohort mean. The weekend-Δ of +0.47 h is the largest in the cohort: ≈28 minutes per night more on Saturday/Sunday than on weekdays. REM proportion is 28.4% (vs cohort 24.9%) and deep is 24.4%, consistent with later sleep timings preserving REM (concentrated in the second half of the night) at the expense of slow-wave sleep.

The Restless Sleeper (n = 12, 16%) is the smallest and most clinically concerning cluster. Mean duration is 4.68 h — below five — and the fragmentation index is 18.3%, more than four times the Anchor's 4.5%. Within-person variability is high (CV = 0.282). At an individual level, exemplars in this cluster show single nights with 48–65% of bed-time spent awake. The data alone cannot distinguish primary insomnia from a high-arousal lifestyle, but the signal flags users who would benefit from clinician follow-up.

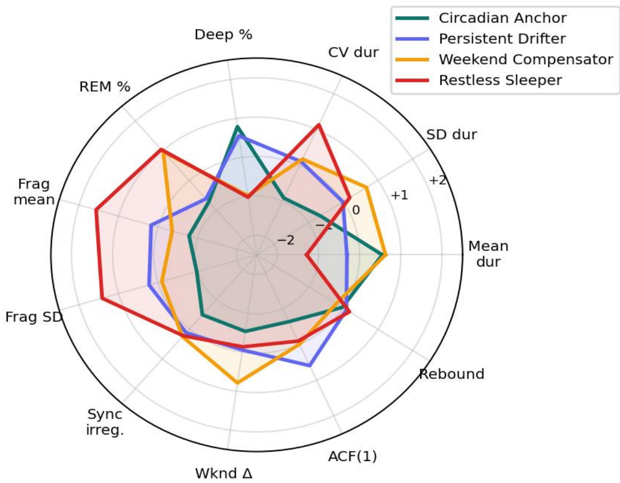


**Fig. 2.** Z-Scored Phenotype Signatures vs. Cohort Means.

### C. Trajectory exemplars and population view

Fig. 3 shows the first 30 valid nights of one representative user from each phenotype, with stage composition stacked under the duration line. The visual signatures are immediately legible: the Anchor's flat, narrow band; the Compensator's recurring peaks; the Drifter's slow undulation; the Restless Sleeper's compressed and choppy pattern with conspicuous nights below 3 hours. Fig. 4 stacks all 75 cohort users into a single heatmap, sorted by phenotype: horizontal banding by phenotype is visible, with Anchors forming a band of consistent mid-tone teal pixels and Restless Sleepers forming a band where the pixels darken and become noisier.

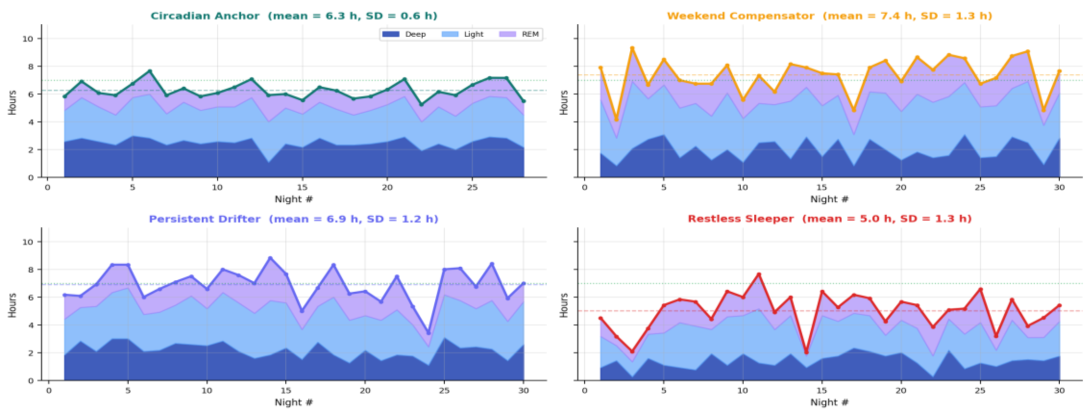


**Fig. 3.** Sleep-Stage Trajectories of Representative Users from Each Phenotype across Their First 30 Valid Nights. Dashed Line: User-Specific Mean. Dotted Green Line: 7-Hour Adult Recommendation.

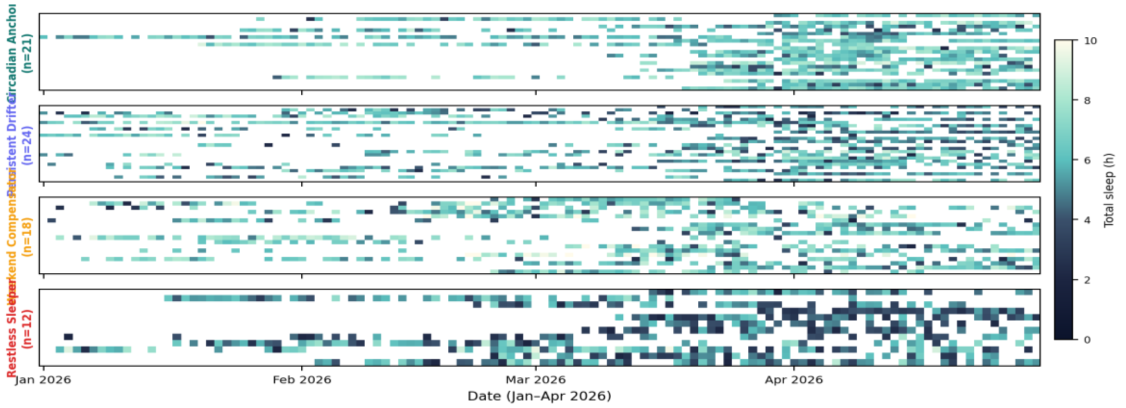


**Fig. 4.** Population Heatmap of Total Sleep Duration. Each Row Is One User, Each Column One Calendar Date In The 4-Month Window. Colour Encodes Total Sleep (Dark = Low, Light = High). White Cells Indicate No Recording On That Date.

## DISCUSSION

### A. Where the rebound went

The phenotype that did not appear is as instructive as the four that did. Our framework anticipated a sawtooth Rebounder — short weekdays, recovery weekend, cycle restarts. In the data, dimensional rebound is rare (cohort-mean rebound rate ≈1%) and discrete clustering does not isolate Rebounders as a population. What we see instead is the Persistent Drifter: chronic, autocorrelated under-sleep without compensatory recovery. This is, if anything, more concerning. The Rebounder at least gets a recovery dose of slow-wave and REM sleep; the Drifter does not.

Why? One hypothesis is contextual: the Indian working week, particularly in metropolitan and urban settings, runs Monday-through-Saturday for a substantial fraction of the workforce. If only one day per week is genuinely off-duty — and even that day carries family or domestic obligations — the social pressure to recover sleep is reduced, and the rebound mechanism never fires. The Weekend Compensator, who does fire, is essentially a person whose work schedule is closer to the Western Saturday-and-Sunday-off norm. This is testable in a cohort with confirmed work-week metadata.

### B. Why averages mislead

The Drifter is the cleanest illustration. Their mean duration of 5.59 h sits within the cohort middle and would not flag any consumer-grade dashboard threshold. Their autocorrelation, however, is the highest in the cohort: the user is in a stable, self-reinforcing low-sleep regime. From an intervention perspective, these users require a different prompt than the Compensator — who needs to redistribute the same total sleep more evenly — and a fundamentally different prompt than the Restless Sleeper, who needs clinical evaluation.

### C. Considerations in Interpretation

This study is based on real-world wearable data at scale, so results should be understood in the context of current-generation devices and models.

In real-world data, behaviours do not form perfectly separate groups. Instead, patterns exist on a spectrum. The four sleep types were chosen because they are easy to interpret and align well with known sleep behaviours, making them useful for real applications.

The dataset represents a working-age population over four months. While longer studies can add more depth, this analysis already shows how continuous wearable data can reveal meaningful and actionable sleep patterns at scale.

### D. Implications

This framework makes sleep analysis more practical and useful in real-world settings by focusing on patterns instead of just averages. By identifying distinct sleep types, it enables more relevant and personalised recommendations for each user, improving engagement and outcomes. It can also be applied at scale in wellness or healthcare programs to spot individuals who may need attention and support early action.

Overall, it helps turn wearable sleep data into clear, actionable insights that can drive better health decisions.

## CONCLUSION

Sleep is a behaviour expressed across weeks, not a quantity expressed per night. Applied to four months of smart-ring data from 456 Indian adults, an unsupervised trajectory-clustering pipeline produced a four-phenotype solution — Anchor, Drifter, Compensator, Restless — that is clinically interpretable, statistically defensible within the modest separation typical of real-world wearable data, and actionable from a digital-health product perspective. The geometry of a month of sleep should sit alongside the nightly average in any future generation of wearable sleep dashboards.

## REFERENCES

1. A. J. K. Phillips et al., "Irregular sleep/wake patterns are associated with poorer academic

performance and delayed circadian and sleep/wake timing," *Sci. Rep.*, vol. 7, 3216, 2017.

2. D. J. Buysse, "Sleep health: can we define it? Does it matter?," *Sleep*, vol. 37, no. 1, pp. 9–17, 2014.
3. J. Paparrizos and L. Gravano, "k-Shape: Efficient and accurate clustering of time series," in *Proc. ACM SIGMOD*, 2015, pp. 1855–1870.
4. T. Roenneberg, K. V. Allebrandt, M. Merrow, and C. Vetter, "Social jetlag and obesity," *Curr. Biol.*, vol. 22, no. 10, pp. 939–943, 2012.
5. N. F. Watson et al., "Recommended amount of sleep for a healthy adult," *Sleep*, vol. 38, no. 6, pp. 843–844, 2015.
6. M. Altini and H. Kinnunen, "The promise of sleep: a multi-sensor approach for accurate sleep stage detection using the Oura ring," *Sensors*, vol. 21, no. 13, p. 4302, 2021.
7. A. S. P. Lim et al., "Quantification of the fragmentation of rest–activity patterns," *Sleep*, vol. 34, no. 11, pp. 1569–1581, 2011.
8. M. M. Ohayon, M. A. Carskadon, C. Guilleminault, and M. V. Vitiello, "Meta-analysis of quantitative sleep parameters from childhood to old age," *Sleep*, vol. 27, no. 7, pp. 1255–1273, 2004.
9. E. Keogh and C. A. Ratanamahatana, "Exact indexing of dynamic time warping," *Knowl. Inf. Syst.*, vol. 7, no. 3, pp. 358–386, 2005.

**How to cite this article:** Subhabrata Paul, Dr. Madhurina Saha, Manosij Sengupta, Ankit Saha, PHENOTYPING SLEEP FROM SMART-RING TRAJECTORIES: UNSUPERVISED DISCOVERY OF FOUR LONGITUDINAL SLEEP PATTERNS IN A 456-USER INDIAN COHORT, Asian J. Med. Res. Health Sci., 2026; 4 (2):2322-2326.

**Source of Support:** Nil, Conflicts of Interest: None declared.