THE ARISING

The Space Between Heartbeats

HRV & RECOVERY



NOVEMBER 2026

The Space Between Heartbeats

HRV, recovery, and what our devices can-and cannot-tell us about living well

The Arising | Living Health Magazines · November 2026 · John Griffin, MD

The heart is not a metronome. Even at rest, the interval between one beat and the next varies slightly. We usually cannot feel the difference. A watch, ring, chest strap or smart bed may detect it, and an algorithm may turn those differences into a number called **heart rate variability**, or HRV.

That number has become a daily ritual for millions of people. They wake up, check whether their HRV rose or fell, and decide whether they are ready for a workout, a difficult workday or an easier morning. There is something appealing about the idea: a quiet biological signal that might help us listen to our bodies more carefully.

But the useful question is not whether a higher number looks reassuring. It is whether the number tells us something reliable-and whether using it changes what happens next.

The answer is interesting, incomplete, and more useful than the slogans on either side.

What the interval actually measures

An electrocardiogram, or ECG, records the heart’s electrical activity. From a clean tracing, researchers can identify successive normal heartbeats and calculate the variation in the intervals between them. Two common measures are **SDNN**, the standard deviation of normal-to-normal intervals, and **RMSSD**, the root mean square of successive differences between adjacent normal intervals. In everyday terms, SDNN describes how spread out the beat intervals are across a recording; RMSSD emphasizes changes from one beat to the next. Both are expressed in milliseconds, or thousandths of a second. Because they summarize different things, compare like with like.

A ring or wristwatch generally does not measure the heart's electrical R-to-R intervals continuously. Instead, a light-based sensor uses a method called photoplethysmography, or PPG, to detect changes in blood volume as each pulse arrives. The timing between those pulses produces **pulse-rate variability (PRV)**, which can approximate HRV under favorable conditions. Motion, temperature, blood flow, sensor fit, missed beats, and changes in the time a pulse takes to reach the sensor can alter the result.

This is why the way the measurement is taken matters as much as the number. Five minutes seated quietly in the morning, readings selected from the parts of a whole night with a clear signal, and a brief interval during deep sleep are three different experiments. A value of 42 milliseconds on one platform is not necessarily equivalent to 42 on another.

A 2026 systematic review and meta-analysis reviewed 43 studies and combined numerical results from 10 unique studies. Under selected resting or controlled conditions, pooled standardized absolute error was 0.188 for RMSSD (95% CI 0.066-0.309) and 0.134 for SDNN (95% CI 0.014-0.255). These are standardized error scores, not percentages or milliseconds: lower scores mean closer agreement with the reference measurement. The 95% confidence intervals (CI) show the uncertainty around each average estimate under the analysis's assumptions. They do not describe the range of errors an individual wearer should expect. Those numbers describe measurement agreement, not improved health outcomes. The authors specifically warned against generalizing them to sleep, exercise, stress or unrestricted daily life, or treating the measurements as interchangeable.

A free-living validation study recorded 25 healthy volunteers over two five-day periods. Errors were generally lower during sleep; activity-specific analyses of the Corsano device found greater error during movement than rest. In a separate controlled validation study, 66 participants wearing a particular wrist device showed strong agreement for some measures, including SDNN, but weaker agreement for others. These are useful demonstrations of technical feasibility, not permission to assume every device is accurate in every setting.

Five devices, five interpretations

Oura calculates overnight HRV from five-minute samples and reports an average and maximum. Its Readiness Score also incorporates sleep, resting heart rate, temperature, activity and personal history. The company emphasizes comparisons with a person's own baseline. Oura HRV documentation · Readiness methodology.

WHOOP calculates sleep-based RMSSD and incorporates it into a proprietary Recovery Score. The score is an interpretation of several inputs, not a laboratory measure of biological recovery. WHOOP HRV · WHOOP Recovery.

Garmin uses overnight HRV and a multiweek personal baseline; HRV Status can feed Training Readiness and Training Status on compatible devices. Garmin explicitly notes that higher-than-usual HRV is not always better. Garmin HRV Status.

Apple Watch contributes heart and other sensor measurements to Apple Health and overnight Vitals. Its algorithms and available features depend on device, software, country and regulatory status. Apple's estimated VO_2 max describes how much oxygen the body can use during hard exercise, a measure of aerobic fitness. It is calculated from sensor data rather than measured in a laboratory exercise test with breathing equipment. Apple Watch health features · Cardio fitness.

Sleep Number measures nightly average heart rate, HRV and breathing rate in compatible smart beds. It also reports a SleepIQ score, sleep trends and bed exits. The bed's averages are not real-time clinical monitoring. Sleep Number biosignals.

Google/Fitbit products add another family of HRV, sleep and readiness estimates. Their exact features vary by model; the manufacturer's device-by-device data inventory is more useful than assuming every Fitbit measures everything.

These differences are not trivial. The devices may be asking related questions using different sensors, windows, filtering rules and mathematical summaries. A change in hardware or software can change the observed baseline without changing the person.

Can HRV help us train better?

This is where the evidence begins to move from measurement into action.

In a 2026 randomized clinical trial, 48 patients recovering from a sudden problem with blood supply to the heart, such as a heart attack (acute coronary syndrome) were assigned to three months of HRV-guided exercise or standard exercise training. Both groups trained three times weekly. In the HRV-guided group, a five-minute morning measurement of beat-to-beat variation (RMSSD) influenced whether a high-intensity or active-recovery session was prescribed.

Peak oxygen uptake, a measure of aerobic fitness during hard exercise, improved by **2.1 mL/kg/min** in the HRV-guided group and **1.9 mL/kg/min** in the standard group. The main comparison did not provide clear evidence that one program improved this measure more than the other (interaction P = 0.794). A P value describes how unusual a result would be under a statistical model assuming no difference; it is not the probability that a treatment works. A secondary measure favored HRV guidance (P = 0.043), despite less total training. That measure tracked oxygen use at the point where breathing begins to rise more quickly as effort increases, adjusted for body mass other than fat. It can offer another view of exercise capacity, but it was not the trial’s main result. The trial was small and did not establish fewer heart attacks, hospitalizations or deaths.

Peak oxygen uptake improved in both groups, and this small trial did not detect a difference between them. That does not prove the programs are equally effective, or that HRV guidance is no worse than standard training by an acceptable amount. The lower training load and greater improvement in the secondary ventilatory-threshold measure warrant further study; they do not establish that HRV guidance improves clinical outcomes.

The larger clinical goal remains physical function: being able to walk, climb stairs, recover after illness and participate in ordinary life. HRV may eventually help us adjust the route to those outcomes. It is not itself the destination.

Can we train HRV itself?

HRV biofeedback lets people watch a live heart-rate signal while practicing breathing, with the aim of learning to influence the body’s responses. It is different from passively checking a recovery score.

A 2026 trial that randomly assigned people to real biofeedback or simulated feedback enrolled 47 healthy adults. After 24 days of practice, the HRV-biofeedback group reported improvements in positive feelings and depression symptoms compared with the simulated-feedback group, but the investigators did **not** find corresponding changes in the measures of the nervous system's automatic control of bodily functions at rest, during stress or recovery. That disconnect is scientifically valuable: subjective benefit need not depend on a large change in the physiological signal being trained.

A separate randomized study of 37 emergency and ICU workers found better resilience scores after four weeks of HRV biofeedback compared with a waitlist (effect size d = 0.79). This score describes the size of the difference relative to the variation in participants' scores; it does not mean a 79% improvement. The results did not clearly separate the groups on stress, depression or burnout. People waiting for treatment know they are not receiving it, so expectations and the attention given to the treatment group could influence the comparison. The sample was small.

A systematic review of 18 remote-biofeedback studies involving a reported 1,352 participants overall found an average reduction in depression symptoms in its 10-study depression analysis (Hedges g approximately −0.41). This is another standardized measure of difference, not a 41% reduction; the minus sign indicates fewer symptoms. Results varied substantially between studies. Most studies had weaknesses that could distort their results. That supports further work on accessible interventions, but the average across different populations and programs cannot be assumed to predict an individual result.

A person can practice slow, comfortable breathing without buying a new device. The clinical question is whether a structured program improves symptoms or function in a particular population. The number on the screen may be helpful feedback; it should not become the sole measure of success.

Why sleep matters-and why sleep scores do not settle the question

Heart rate, respiration and the nervous system's automatic activity vary through the night. Sleep loss, alcohol, illness, travel, intense training, some medications and menstrual-cycle changes can all shift HRV. A wearable's overnight value is therefore best interpreted as a context-rich trend, not a verdict about personal discipline.

Consider a person who wakes with a lower-than-usual HRV after a short night. The sensible question is whether the pattern reflects sleep debt, a brewing illness, heavy exertion or a recording error. It is not whether the person has suddenly become less healthy than yesterday.

Sleep duration and regularity have independent relevance to health and daytime function. A proprietary sleep score can be a convenient summary, but it does not diagnose insomnia, sleep apnea or a circadian disorder. Those conditions require symptoms, history and sometimes formal testing.

Smart beds add a useful perspective because they require nothing on the wrist. They can reveal trends in time in bed, restlessness, bed exits and nighttime body signals. Yet the absence of a wrist sensor does not make a mattress a clinical sleep laboratory.

Cardiovascular health: a signal, not a treatment target

Problems with automatic nervous-system control, heart and blood-vessel disease, abnormal heart rhythms, medications and aging can influence HRV. Researchers have studied whether HRV helps predict the course of several illnesses. Predicting what may happen does not show that changing HRV will improve it.

There is no established clinical rule that a consumer should raise their nightly HRV above a universal threshold to prevent a heart attack. Atrial fibrillation, an irregular heart rhythm, and frequent extra beats can also make standard HRV summaries difficult to interpret. If a device detects an irregular rhythm or a person has a racing or fluttering heart, fainting, chest discomfort or other concerning symptoms, the appropriate next step is clinical evaluation-not attempting to optimize a recovery score. Chest discomfort, fainting, or severe shortness of breath may require urgent medical assessment.

For blood pressure, a validated upper-arm cuff provides actionable measurements in a way a generic recovery score does not. For diabetes, continuous glucose monitoring can improve blood-sugar control in selected populations, as reflected in the 2026 ADA Standards of Care. Those examples show what evidence-linked measurement looks like: a measurement tied to a decision, a treatment and a meaningful outcome.

A practical way to use the number

If you already own a device, a low-cost experiment is to observe your own trend for several weeks without treating each day as a command. Use the same device, comparable recording conditions and a stable time window. Note sleep, exercise, illness, alcohol, medication changes and how you actually feel. Pay attention to patterns that recur, not one-day excursions.

If you train, keep the fundamentals visible: sustainable activity, strength work, adequate recovery and symptoms. An unusual HRV trend may invite an easier day or a closer look at sleep; it should not automatically override a rehabilitation plan or medical advice.

If you do not own a device, you have not missed the foundation of good recovery. Regular sleep opportunity, appropriate physical activity, nutrition, social connection and evidence-based care remain available without a subscription.

And if tracking creates anxiety or compulsive checking, the cost of the information may exceed its benefit. A measurement is useful when it improves understanding or action. It is less useful when it turns ordinary physiological variation into a daily judgment.

Three real-world reading scenarios

A runner after a difficult week. The runner has slept poorly, traveled, and finished several demanding sessions. The wearable shows a lower-than-usual HRV and a higher resting heart rate. That pattern may be consistent with accumulated strain, but it does not identify a single cause. The useful decision is to look at the training plan, perceived effort, symptoms, recent illness and sleep opportunity together. An easier day may be sensible because of the whole picture-not because a red icon has issued an order. If the pattern repeats, it can become a useful observation to discuss with a coach or clinician.

An older adult recovering from hospitalization. This person wants to return to walking outdoors, climbing stairs and managing daily activities. A wearable may report steps, resting heart rate and nighttime HRV. None of those numbers substitutes for a gait assessment, a chair-rise test, strength examination, medication review or a physical therapist's evaluation. If HRV is included in future rehabilitation research, the primary outcomes should be mobility, falls, symptoms, participation and independence. A model that predicts a low recovery score without improving any of those outcomes has not yet earned a clinical role.

A caregiver with fragmented sleep. A caregiver wakes repeatedly to help someone at home. Their watch shows reduced sleep duration and changing HRV. A recommendation to optimize recovery without recognizing the caregiving circumstances is not particularly helpful. The more useful response may involve respite, a safer nighttime care plan, social support or an evaluation of the care recipient's sleep and symptoms. Here the most important variable may not be a biosignal at all. It may be the availability of another person to help.

These scenarios are illustrative, not reports of actual patients. They show why context belongs beside every number and why the right next action is sometimes outside the device.

What a useful explanation should include

If an app recommends a change, you should be able to see which measurement prompted it, how it compares with your usual range and what evidence supports the advice. A software update can change the way a score is calculated. That should be made clear, so a new calculation is not mistaken for a change in your health.

An AI assistant may help organize readings and explain the study behind a recommendation. It should not diagnose illness from a score or automatically change treatment. The real test is whether its explanations help people and clinicians make safer, clearer decisions without creating more worry or work.

What the next generation should prove

The most promising future is not a more elaborate color-coded score. It is a validated way to answer specific questions: Can we identify an athlete training harder than they can recover from before injury? Can we tailor rehabilitation to preserve function with less burden? Can an older adult recover from illness without losing independence? Can a patient with cardiovascular disease receive a safer, more effective exercise prescription?

To answer those questions, trials must compare a defined HRV-guided decision rule with a credible alternative, measure whether people follow the plan and whether it causes harm, and follow outcomes people care about. They should include older adults, women, patients with chronic disease and the people whose sensors work least reliably. The research should also distinguish an improvement in HRV from an improvement in health.

There is a quiet kind of promise in the space between heartbeats. It is not that a sensor can tell us everything about ourselves. It is that carefully measured physiology, joined to good science and humane judgment, may help us make a few better decisions-and preserve more of the life we want to live.

Go deeper with The Arising

- Sleep & circadian health
- Fitness & movement
- Cardiovascular health
- Wearables & health measurement
- Public Atlas

Illustration: original conceptual art, not a clinical measurement or trial figure.

Sources and further reading

Primary studies and manufacturer documentation are linked beside the relevant claims. This issue is educational and does not provide individual diagnosis or treatment.