BIKE IQ TECHNICAL WHITEPAPER · V1.0 · APRIL 2026

Bike IQ — Cadence Whitepaper Cadence Without a Sensor

How Bike IQ Estimates Pedaling Rhythm from the Phone in Your Pocket

A BIKE IQ TECHNICAL WHITEPAPER

Summary

Bike IQ has developed a method for estimating cycling cadence, the number of pedal revolutions per minute, using only the motion sensors already built into any modern iPhone. No additional hardware is required. No crank sensor, no shoe pod, no pairing, no extra battery to charge.

Across a validation dataset of real rides collected during the research program, on a diverse set of bikes in a variety of real-world conditions, the system reports a mean absolute error of **3.45 RPM**, with **86%** of estimates falling within 5 RPM of a reference Bluetooth cadence sensor and **77%** within 3 RPM. These numbers are strong enough to make cadence a first-class training metric in the app for riders who would never otherwise buy a dedicated sensor, and the system is honest about the cases where it does not yet match hardware precision.

This document describes the problem, our approach at a high level, how we validated it, what it does not do well today, and where the research is heading next.

Why Cadence Matters

Cadence is one of the most important numbers in cycling training. It correlates with muscular fatigue, aerobic efficiency, pedaling economy, and technique. Coaches prescribe cadence targets the way runners prescribe pace. A rider who learns to spin at 90 RPM on the flats and hold a steady 75 on climbs will produce more power with less fatigue than a rider who grinds at 55 for an entire ride.

And yet, the vast majority of cyclists never measure their cadence. A dedicated Bluetooth cadence sensor costs money, needs installing on the crank arm or the shoe, requires its own battery maintenance, and has to be paired through another menu in another app before every ride. For most cyclists the friction is too high, so cadence stays invisible and training stays blind.

A cadence estimate that works on any bike, with any setup, using only the phone the rider already carries, removes that friction entirely. That is what this research set out to deliver.

Why It Is Hard

At first glance, estimating cadence from a phone sounds like a textbook problem. Pedaling is periodic. Phones have accelerometers and gyroscopes. Finding periodicity in a signal is high school math.

In practice, nearly everything conspires against you.

The signal is not where you expect. A phone mounted on the bike does not see the rider's legs directly. What it sees is the very small side-to-side sway of the frame as the rider transfers weight from one pedal to the other, plus vibration propagating through the fork. A phone carried in a jersey pocket sees the rider's hip motion instead. The signal takes a different shape on every bike, every rider, and every carry location.

Road vibration is periodic too. Tires hitting pavement cracks, expansion joints, and surface texture produce pulses squarely in the same frequency band as pedaling. A naive frequency detector happily reports "pedaling" while the rider is coasting down a hill.

Harmonics are real. Pedaling has a natural second harmonic, because there are two legs and two pedal strokes per revolution. On many riders the second harmonic is stronger than the fundamental. A detector that is not careful will lock onto twice the actual cadence and confidently report nonsense.

You cannot calibrate. Every rider has a different bike, a different pedaling style, different shoes, different tire pressure, and a different typical carry position. There is no opportunity to collect training data from a new user before giving them a useful number on their very first ride.

Everything else that happens on a bike also generates motion. Wind, potholes, shifting, braking, standing sprints, turning, traffic, rough surfaces. Each one produces signal in the same band as the thing you are trying to measure.

The result is that a naive frequency detector, applied to raw phone motion, produces cadence estimates that are wrong more often than they are right. The difficult and interesting work is the second-order problem: telling true pedaling apart from everything else that happens to look periodic.

Our Approach

Bike IQ’s cadence estimator combines several independent sources of information and a temporal model that holds them together. Each source on its own would be insufficient. The combination is what produces a usable result.

Motion Signal Analysis

The phone’s motion sensors are sampled continuously while a ride is in progress. The raw streams are filtered to attenuate drift and low-frequency motion, then recombined to emphasize the components most likely to carry pedaling information. A frequency-domain analysis of short overlapping time windows produces a set of candidate cadences, each with an associated confidence score. The confidence score, not just the candidate itself, is load-bearing throughout the rest of the pipeline.

Drivetrain Physics

This is where Bike IQ’s approach diverges from most of what is described in the signal-processing literature.

Because Bike IQ already knows the rider’s current speed from GPS, and already knows the rider’s bike configuration (chainring sizes, cassette ratios, wheel size, tire configuration), it can compute, for any hypothetical gear the rider might be in, the cadence that gear must produce at the current speed. That physical relationship is rigid. If the rider is in a 50 by 17 at 32 kilometers per hour, the cadence is determined, not guessed.

This gives the estimator a physics-based prior on which cadences are actually possible right now. Instead of searching a wide frequency range blindly, it evaluates the motion signal against a small set of candidates that the bike’s own geometry says are feasible. This is the single largest contributor to the system’s real-world accuracy, and it is what distinguishes Bike IQ’s cadence estimate from a generic smartphone step counter applied to cycling.

Temporal State Tracking

Cadence and gear selection do not jump around at random. A rider stays in a gear for many pedal strokes at a time, and shifts are discrete events with a predictable cost. By tracking the most likely gear over time, the system smooths out moment-to-moment noise in the motion signal and resists the harmonic flipping that defeats simpler methods. When the rider does shift, the tracker is set up to accept the new state quickly without getting stuck on the previous one.

Adaptive Personalization

As a ride progresses, the system learns each rider's personal cadence range and refines its internal parameters accordingly. A rider who naturally spins at 95 RPM on flat ground and a rider who grinds at 65 RPM get different treatment within the same algorithm, without either of them ever having to configure anything.

Self-Calibration

Small errors in the rider's configured tire size would otherwise shift every gear-implied cadence by a constant percentage. The system continuously refines its estimate of the rider's true effective wheel circumference from moments when the motion signal is unambiguous, and corrects for that error automatically within the first minute of a ride.

Independent Coasting Detection

A separate pipeline, distinct from the cadence estimate itself, decides whether the rider is pedaling at all. This separation is deliberate. A system that is confident but wrong about cadence is worse than a system that stays silent. By making the coasting decision on its own terms, and only reporting a cadence when the coasting detector confirms the rider is actually pedaling, the app avoids the most embarrassing failure mode, which is a rider coasting down a hill while the screen shows 85 RPM.

None of These Are Individually Novel

Each of the components above has been studied, in some form, in the academic literature or in other consumer products. What is novel is the combination, tuned specifically for the real-world case of a pocketable phone on a moving bike, across many bikes and many riders, with no per-user calibration and no hardware beyond the device itself.

How We Validated It

We validated the system using diagnostic rides collected during the research program. Each diagnostic ride contains the raw phone motion streams, the GPS track, the bike configuration, and, where available, a ground-truth cadence feed from a Bluetooth cadence sensor paired to the phone for the same ride. The Bluetooth feed is treated as truth for the purposes of evaluation.

The current validation set:

- 45 diagnostic rides collected during the research program
- 22 of those rides have a paired Bluetooth cadence sensor providing ground truth
- A diverse set of drivetrains, from 1x7 commuter setups to 2x11 road racing setups
- Multiple rider body types, pedaling styles, and carry locations
- Mixed conditions including urban riding, climbing, and open road
- Rides collected over several weeks from multiple testers

For every short time window of every ride, the estimator's output is compared point-by-point against the reference. The headline results:

Metric	Result
Mean absolute error	**3.45 RPM**
Within 3 RPM of reference	**77.3%**
Within 5 RPM of reference	**86.0%**
Large errors (greater than 10 RPM)	7.8%

A note on the reference itself. Bluetooth cadence sensors are not a perfect truth. Most consumer cadence sensors average over several pedal revolutions before reporting a number, which means they lag the true instantaneous cadence by a second or two, particularly at the beginning and end of coasting intervals. A small fraction of the "error" in our numbers is actually our estimate tracking the instantaneous truth more closely than the reference does. For validation purposes we treat the reference as correct anyway, which means the reported numbers are a slight underestimate of the system's real accuracy.

What It Does Not Do Well Yet

Bike IQ believes in being straightforward about the limits of its research. A rider using cadence estimation in the app today should understand the following.

Low cadences are harder. Riders who pedal slowly, roughly below one revolution per second, are underrepresented in our validation data and present a fundamental ambiguity. At low cadences, the signal's second harmonic sits in the same range as a healthy cadence, and without enough independent evidence the system can report twice the true value. Our adaptive personalization catches this for most riders, but not for every rider on every ride.

Unusual mounting and carry locations. The majority of our validation rides were collected with common mounting options, including on-the-bike mounting and jersey pocket carry. Riders who keep the phone in a frame bag, stem bag, handlebar bag, or backpack present a weaker and differently shaped signal that our current validation set does not yet cover adequately.

Heavy vibration can mask coasting. On rough roads, chip seal, or cobbles, the vibration reaching the phone can briefly resemble pedaling. The independent coasting detector handles most of these cases, but a small percentage of coasting windows still register as low-cadence pedaling in the most hostile conditions.

Tire size errors take a moment to correct. The self-calibration routine needs roughly half a minute of clear pedaling at a known gear before it has real confidence in the rider's effective wheel circumference. In the first minute of a ride with a significantly mis-configured tire size, accuracy is temporarily lower than the steady-state numbers above.

Rider population. The current validation dataset reflects a specific population. Bike IQ's beta testers skew toward experienced road cyclists on well-maintained bikes. We do not yet have good data on e-bikes, fixed-gear and track bikes, BMX, children's bikes, or recumbents, and we do not claim accuracy numbers for those populations until we have validated them.

The Roadmap

The following work is either already underway or planned for upcoming Bike IQ releases.

Broader validation. The research program is expanding to cover a wider cross-section of bikes, riders, and conditions. The validation dataset is expected to grow substantially over the coming months. A larger and more diverse dataset both sharpens the accuracy numbers we can honestly report and surfaces edge cases we have not yet seen.

Better coasting detection. The current independent coasting detector handles common cases well but has known blind spots on rough surfaces and in high-vibration conditions. A learningbased coasting classifier is on the research roadmap, gated on having enough labeled data from enough riders to train a model that generalizes rather than overfits.

Higher-rate motion data. A higher sampling rate on the underlying motion data would unlock entirely new families of techniques that are out of reach at current rates. Bike IQ plans to characterize both the accuracy benefit and the battery cost of running at a higher rate, and to make a data-driven decision about whether the tradeoff is worth it for real-time use.

Cross-ride personalization. Today, the adaptive personalization logic starts from scratch at the beginning of every ride and takes a minute or two to settle. Persisting the learned profile across rides will make the first minute of every ride as accurate as the rest of the ride.

Low-cadence robustness. A targeted research sprint is planned to push the lower limit of the useful cadence range well below where it sits today. The working hypothesis is that a combination of longer analysis windows, tighter drivetrain constraints, and cross-ride learning can resolve the low-cadence harmonic ambiguity that currently defeats the system on a subset of rides.

Publication. Bike IQ intends to share further technical detail with the academic and industry research community over time, in venues where the work can be properly reviewed.

What This Means for Riders

A cyclist who opens Bike IQ on a bike ride gets a usable cadence number without buying, installing, pairing, or charging anything beyond the phone they already carry. That number is accurate enough for pacing, training, and post-ride analysis on the large majority of rides, and the app is honest about when it does not know.

This is not yet a perfect replacement for a dedicated cadence sensor on every bike and every rider in every condition. We are open about the cases where purpose-built hardware still has an edge. What it represents is a meaningful reduction in the friction of measuring one of the most important metrics in cycling training, delivered to a large population of riders who would otherwise never see the number at all.

Cadence is the first metric Bike IQ has attacked with this combination of signal analysis, physics-informed priors, and temporal tracking. The same approach is being applied to several other measurements inside the app, and we expect the results to tell a similar story.

About Bike IQ

Bike IQ is an iOS cycling app focused on turning the phone into a legitimate cycling computer for riders who want serious training metrics without buying serious training hardware. The app is built by a small independent team and ships regular research-driven updates to its beta and production users.

For press inquiries or technical follow-up, please contact the Bike IQ team at contact@bikeiq.app.

This whitepaper describes research in active development. The systems and methods described herein are the subject of pending U.S. patent application(s), including U.S. Provisional Application No. 64/040,848 filed April 16, 2026. Numbers reported above reflect the state of the system as of the document date and may improve in subsequent releases.