



BIKE IQ TECHNICAL WHITEPAPER · V1.0 · APRIL 2026

Bike IQ — Grade Estimation Whitepaper

Knowing How Steep the Road Is

How Bike IQ Estimates Road Grade in Real Time From a Phone Barometer

A BIKE IQ TECHNICAL WHITEPAPER

Summary

Bike IQ produces a real-time estimate of the steepness of the road the rider is currently on, expressed as a percent grade, using the barometric pressure sensor inside the iPhone, the GPS speed signal, and a sensor-fusion pipeline that turns those raw inputs into a usable number. Grade is one of the most fundamental signals in cycling and one of the hardest to measure cleanly from a moving bike.

The system delivers a grade signal that powers Bike IQ's power estimation, drives climb detection and the climb-ahead UI, and shows up directly to the rider as the steepness readout during the ride. It is accurate enough that the power estimate built on top of it agrees with dedicated power meters within a few watts on average. There are conditions where the underlying physics imposes hard limits on what is possible, and we are open about those.

Why This Matters

Grade affects almost everything a cyclist cares about. It determines how hard a section of road is, how much power the rider has to put out to maintain a given speed, how to pace a climb, and how to interpret a personal record. A 30 minute effort up a 6% climb and a 30 minute effort on a flat road are not comparable training stimuli, and the only thing distinguishing them in the data is the grade.

Beyond power, grade matters for route planning, climb classification, and the kind of moment-to-moment readouts a cyclist looks at on a screen during a ride. A grade signal that is wrong, or that lags by several seconds, makes every downstream metric wrong with it.

Why It Is Hard

Grade looks like a simple measurement. The road has a slope. Measure it. In practice, every approach has a fundamental problem that has to be engineered around.

GPS altitude is not good enough. GPS gives latitude and longitude with reasonable accuracy, but vertical position from GPS alone is many times noisier than horizontal position. Grade computed from GPS altitude differences over short distances is dominated by GPS noise, not real terrain. This rules out the most obvious approach.

The barometer is the right tool, but it has hardware lag. Modern phones include a barometric pressure sensor that, indoors and at rest, can resolve elevation changes of a fraction of a meter. The sensor is the right physical principle for grade measurement on a moving bike. It also has a built-in response time of a few seconds, because the sensor cavity has to equalize pressure with the outside world. That response time is invisible at rest but becomes critical at speed: at 40 kilometers per hour, a few seconds of lag is many tens of meters of road.

The lag is not constant. The same sensor lags more at high speed than at low speed, because the airflow around the phone affects the pressure equalization in non-trivial ways. This is a real physical effect documented in the aviation literature for pitot-static systems. It means a fixed lag compensation is not enough.

Wind generates pressure noise. A 25 kilometer per hour crosswind produces measurable Bernoulli pressure variation around the bike that the sensor sees as if it were altitude change. On a windy day this can produce false grade readings of several percent over short windows.

Pedaling vibration affects the signal too. Hard pedaling produces motion that, depending on how the phone is mounted, can register on the accelerometer-based corrections that some grade pipelines use. Older approaches that combined barometer and accelerometer naively had a tendency to freeze the grade output during hard climbing efforts, which is exactly when accurate grade matters most.

Climbs and descents are physically asymmetric. A bike going uphill has reasonably symmetric airflow over the phone. A bike descending at high speed creates pressure conditions where the sensor reads slightly differently. This produces real differences in measured grade between the up and down directions of an out-and-back ride that the algorithm has to either correct or accept.

Ground truth is hard to come by. Unlike speed (GPS-truthed), cadence (Bluetooth sensor truthed), and power (meter truthed), grade has no easily available reference signal on a real bike ride. The closest reference is the high-resolution digital elevation model of the road, which is itself imperfect, and post-processed power meter analysis, which inherits all the same difficulties in reverse.

Our Approach

Bike IQ's grade estimator is the result of an extensive research program covering hundreds of rides, dozens of algorithm families, and tens of thousands of automated parameter sweeps. The current production system is a multi-stage filter that takes raw barometric pressure and GPS speed and produces a usable grade signal in real time, with several layers of safeguards against the failure modes described above.

Barometer as the Primary Source

The barometric pressure sensor is the dominant input. Atmospheric pressure changes predictably with altitude (roughly 12 pascals per meter at sea level), and the sensor in a modern iPhone resolves changes well below that. After conversion to altitude, the rate of altitude change combined with the rider's speed yields the road grade.

Rate-Clamped Smoothing

The raw barometric altitude is rate-clamped to reject outliers caused by wind gusts, going through doors, riding past tall buildings, and weather changes. Realistic vertical speeds on a bike are bounded; samples that imply impossible vertical speeds are rejected before they reach the rest of the pipeline.

Speed-Adaptive Smoothing

The amount of smoothing applied to the altitude signal is not constant. It adapts to the rider's speed in a non-obvious way: more responsive at low speed and more smoothed at high speed. This is the opposite of what intuition would suggest based on signal-to-noise reasoning, but it is the right answer because grade matters most for the physics at low speed (climbing, where gravity dominates power) and least at high speed (descending or flat, where aerodynamic drag dominates instead).

Real-Time State Estimation

The smoothed altitude is combined with the speed signal in a multi-state filter that tracks altitude, vertical rate, and speed jointly, rather than treating them as independent. The filter knows that grade is the relationship between vertical rate and forward speed, and that none of those quantities can change discontinuously, so it produces a grade output that is consistent with all three signals together.

Decoupled Display and Physics

The grade number shown to the rider on the screen and the grade number used internally for power estimation are processed slightly differently. The display value is biased toward smoothness, because riders find rapidly twitching readouts distracting. The power-input value is biased toward responsiveness, because every second of grade lag during a climb costs measurable power accuracy. The same underlying barometer and filter feed both, with different output smoothing on each.

Continuous Research Beyond the Production Pipeline

The production pipeline above is the current shipping system. Several research efforts continue beyond it. Inertial-sensor fusion (using the accelerometer's ability to detect tilt directly) can in principle compensate for barometer lag and is producing promising results in offline experiments, particularly at high speed. Machine-learned correction of the barometer's transfer function, trained on rides where ground truth power is available, has demonstrated significant additional accuracy gains in post-processing. Both are described in the roadmap section.

How We Validated It

Validating a grade estimator without a direct ground truth requires care. Bike IQ's primary validation method is to use the grade estimate as an input to the power estimator, and then compare the resulting power output against a paired dedicated power meter. The physics equation that produces power from grade, speed, mass, and aerodynamics is well known. If the grade estimate is correct, the resulting power matches the meter. If the grade estimate is wrong, the power error reveals it.

This indirect validation has the property that it weights grade errors by how much they actually affect what the rider cares about. A small grade error on a flat road at high speed costs few watts; the same error on a climb at low speed costs many. The power-MAE figure of merit captures this asymmetry naturally.

The current validation set:

- Over a hundred research rides recorded with a paired dedicated power meter
- Multiple riders, multiple bikes, multiple terrain types

- Both shipped production data and held-out research data
- Repeated direction-asymmetric rides (out-and-back climbs) for symmetry checks

Across this validation set, the production grade pipeline drives a power estimate that agrees with dedicated power meters within a few watts of mean bias and roughly twelve percent of typical error at meaningful time windows. Ongoing research approaches reduce the grade-related component of that error substantially in offline tests.

What It Does Not Do Well Yet

Lag at high speed. The barometer's physical response time imposes a floor on how quickly the grade estimate can react to a real terrain change. At 40 kilometers per hour, the lag is on the order of seconds, and is most noticeable at the moment a flat section transitions into a climb or descent. The display catches up quickly but not instantly.

Descent grade is harder than climb grade. Several physical effects (asymmetric airflow over the phone, pressure changes from braking, the way pedaling effort is distributed differently when the bike is going down) combine to make descent grade systematically harder. Production accuracy on descents is meaningfully worse than on climbs, and we are honest about this in the validation numbers.

Strong wind days. On days with sustained winds above roughly 25 kilometers per hour, the barometric signal carries enough wind-induced pressure noise that grade accuracy degrades visibly. The rate-clamping filter prevents the worst spikes, but the steady-state noise floor is higher.

Very short transients. Speed bumps, sudden potholes, and short ramps shorter than the filter's response time get smoothed out. This is intentional (the alternative is a constantly twitching readout) but it does mean Bike IQ's grade is the macroscopic road grade, not a centimeter-by-centimeter terrain profile.

The Roadmap

Inertial-sensor fusion for lag compensation. The accelerometer on the phone can sense tilt directly, with no lag. The challenge is that the accelerometer also senses every bump, every pedal stroke, and every body movement, and naively combining it with the barometer makes things worse rather than better. Adaptive fusion that detects when the inertial signal is trustworthy (phone stable on the bike, smooth pedaling) and dynamically blends it with the barometer signal is producing dramatic lag reductions in offline experiments. Real-world deployment is the next step.

Machine-learned barometer correction. Training a model on rides with ground-truth power to learn the systematic transfer function of the barometer (lag, wind sensitivity, descent asymmetry) reduces the grade-related power error meaningfully in post-processing studies. The roadmap includes deploying this correction in real time, gated on having enough labeled data to ensure cross-rider generalization.

Raw pressure recording. Bike IQ recently started capturing the raw pressure signal at higher fidelity than was previously available. The aviation literature on barometer lag (going back to a 1963 Air Force report) demonstrates that the lag operates fundamentally in pressure space rather than altitude space, and the correction is more accurate when applied at the pressure stage. This data collection is the foundation for the next round of correction algorithms.

Higher-fidelity descent handling. Targeted research on descents specifically is planned, including separate filter parameters for negative-grade conditions and using the inertia term in the power physics as a cross-check on the grade signal.

Cross-ride terrain learning. Persisting a corrected grade profile of frequently-ridden roads across rides removes most of the per-ride uncertainty from those segments. This works particularly well for commute routes and local training loops where the rider passes the same terrain repeatedly.

What This Means for Riders

A cyclist using Bike IQ during a ride sees a steepness readout that is reliable enough to pace climbs by and is internally consistent with the power, speed, and elevation numbers shown alongside it. After the ride, the climb segments and elevation gain match what a paper map of the route would suggest. The power that Bike IQ infers from the grade matches what a dedicated power meter would have read on the same ride within a few percent on aggregate.

Grade is also the single largest hidden contributor to the power estimate. Improvements in the grade estimator flow directly to better power numbers, better climb segments, and better post-ride analysis. This is why grade is one of the most active research areas inside Bike IQ even though it is one of the least visible to most riders.

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.