

ECG and Blood Glucose: A Meta-Analysis of the Signal

Ben Letchford

29 June 2026

ABSTRACT

A structured synthesis of the evidence linking ECG features to blood glucose, from QT changes to wearable machine-learning models.

A weak physiological echo is still a signal

Whatever else it is, the ECG is not a glucose sensor, and the field gets into trouble the moment it forgets that.

A glucose meter measures chemistry directly. A continuous glucose monitor reads an electrochemical reaction in the fluid between your cells. An ECG does neither. It records the heart's electrical activity at the skin. There is no glucose molecule in that trace, no spectral line reading 6.1 mmol/L, no hidden blood sugar channel waiting for the right decoder. The quantity we want is simply not the quantity we are recording.

The idea survives anyway, and not because it is wishful. It is indirect. Glucose changes the body, and the body changes the heart. A falling glucose level triggers counterregulatory physiology. A rising one shifts autonomic tone, repolarization, vascular state, and electrolyte handling. Diabetes sits underneath all of it and slowly rewires autonomic control and cardiac risk. None of that is glucose, but all of it leaves marks the ECG can pick up: in heart rate, in heart-rate variability, in the QT interval, in T-wave shape, in ST and PR features, and in the higher-order patterns a neural network can pull out of raw beats. The heart is a poor witness to glucose, but it is not a blank one.

That changes the question worth asking. Forget whether the ECG can measure glucose the way a meter does, because it can't. Ask instead how much glucose-related information actually reaches the ECG when glucose moves, and how much of that information is still there once you test the model on someone it has never seen.

My own read, up front so you can argue with the rest of it: the correlation is biologically real, not a statistical artifact, and it is clearest during hypoglycemia, acute glucose challenge, and some hyperglycemic states. It is also not strong or specific enough to trust on its own, and most of the field's overclaiming lives in the gap between those two facts. The role that holds up for now is a supporting one. An alarm, a risk score, a bit of trend context, one input among several in a wearable. Not the number you dose insulin against.

1 What kind of meta-analysis this can honestly be

You cannot pool this literature into a single effect size, and any paper that tries should make you suspicious. The studies are barely measuring the same thing. One quantifies QTc drift during clamp-induced hypoglycemia. The next classifies hypoglycemic episodes as events. Another classifies hyperglycemia. Another regresses continuous glucose against the waveform. The reference standard wanders too, from plasma glucose to capillary to interstitial CGM, and one recent line of work splits fasting blood glucose from HbA1c entirely, on the correct grounds that the two live on different time scales and should not be averaged together ([Azam & Singh, 2026](#)).

The populations differ as much as the endpoints: healthy volunteers, children with type 1 diabetes, adults with type 1, adults with type 2, hospital-grade protocols, nocturnal ambulatory recordings, oral glucose tolerance tests, scraped public datasets. Validation ranges from a subject-specific model that knows one person cold to a subject-independent network that has to generalize to people it never trained on. Average across all of that and you get a number with three decimal places and no meaning.

So this is a structured synthesis rather than a forest plot. I have stacked the evidence by causal layer, working from physiology up to product:

- the physiological changes the ECG shows during glycemic excursions;
- the hand-engineered features built to read glucose state off those changes;
- the deep-learning and wearable models that learn the morphology themselves;
- and the validation limits that explain why a decade of strong retrospective scores still has not produced a clinical ECG glucose monitor.

The layering matters because a QT study and a neural-network study can both be true while answering completely different questions, and a lot of the confusion in this field comes from mistaking that for agreement.

2 The signal is strongest for hypoglycemia

Hypoglycemia detection is the oldest coherent branch of the field, and it is also where the physiology is clearest. Falling glucose is an emergency, and the body treats it as one. Sympathoadrenal activation kicks in, potassium shifts between compartments, and ventricular repolarization bends under the strain. Those pathways run close enough to the heart's electrical behavior to show up on the trace, which is why the ECG reads glucose best at exactly the moment the body is most alarmed by it.

The first clue was the QT interval. In 1998, Eckert and Agardh put healthy men into experimental hypoglycaemia and watched the QT interval lengthen ([Eckert & Agardh, 1998](#)). Robinson and colleagues followed with an insulin clamp and confirmed QTc changes during hypoglycemia, alongside the expected counterregulatory

and electrolyte effects (Robinson et al., 2003). Koivikko and colleagues then complicated the tidy QT story in a useful way. In type 1 diabetes and healthy controls alike, controlled hypoglycemia did not confine itself to a single number. T-wave amplitude, T-wave area, the QRS-T angle, and the geometry of the T-wave loop all moved as well, which means the signal was spread across the whole shape of repolarization rather than sitting in QTc alone (Koivikko et al., 2008).

This is worth dwelling on. A wearable that watches only QTc is reading one descriptor, and not even a clean one, since QTc depends on its own correction formula, on heart rate, and on whether you are standing or lying down. The ECG’s response to glucose is not a single value, and an algorithm that treats it as one risks measuring its own assumptions.

The harder question is whether any of this survives outside the clamp lab. Some of it does. Lee and colleagues used simultaneous Holter ECG and continuous glucose monitoring and caught QT prolongation accompanying hypoglycemia in insulin-treated type 1 and type 2 diabetes, in ordinary conditions rather than under a protocol (Lee et al., 2016). That is the test most clean lab effects fail. People sleep, move, digest, get stressed, and wear imperfect sensors, and crisp physiology tends to blur. Here it held up well enough to be seen, though not well enough to retire the glucose meter.

The same biology that produces the signal also explains its weakness. The people at greatest risk of severe hypoglycemia are often the ones whose hearts can no longer raise much of a response. Lipponen and colleagues found that repolarization-based detection worked well in healthy subjects and in otherwise-healthy type 1 diabetes, then weakened in type 1 diabetes with complications, where the autonomic response was already blunted (Lipponen et al., 2011). That is the central paradox of the whole approach. The signal is the body’s reaction to glucose, and the body’s reaction is exactly what long-standing diabetes erodes, so the detector tends to work best on the patients who need it least.

3 Hyperglycemia changes the ECG too, but less cleanly

Run glucose the other way and the signal gets quieter and harder to trust. Acute hyperglycaemia can prolong QTc even in healthy subjects, which Marfella and colleagues showed with a glucose infusion, so the electrophysiology is real and not just borrowed from the hypoglycemia story (Marfella et al., 2000). There is a foundation here. It is just a narrower one.

The trap is treating hyperglycemia as hypoglycemia in reverse. It isn’t. High glucose comes in too many forms to share one ECG signature: the post-meal spike, the fasting plateau, the slow grind of an elevated HbA1c, and the conditions that travel with it, like illness, dehydration, neuropathy, and cardiovascular disease. Some of those states change HRV, some change repolarization, and most sit behind a long list of confounders: age, medication, fitness, sleep, caffeine, stress, infection, and how long the person has had diabetes. Hypoglycemia gives you one alarmed body to read.

Hyperglycemia gives you a crowd, and the ECG has to pick out the relevant face.

The cleaner approach is to pin glucose down with an acute challenge. Tobore and colleagues did this with 16 participants and a 75 g oral glucose tolerance test, and found glucose-associated patterns across HRV, heart rate, PRQ, QT, and ST segments (Tobore et al., 2019). Ha and colleagues went further and handled a problem most studies skip. Working with 30 adults, an OGTT, Holter ECG, and interstitial glucose monitoring, they measured each participant’s HRV response delay and aligned to it before scoring. With the timing matched up, the HRV-glucose association tightened and acute hyperglycemia detection reached about 0.89 accuracy and AUROC under temporal validation (Ha et al., 2026).

That delay is easy to skip past, but it quietly governs the whole field. The ECG does not move at the same instant as glucose. Glucose rises first, and the heart follows on a lag that varies from person to person. Ignore the lag and you smear a real effect into noise and conclude the physiology is weak. Estimate the lag from the whole dataset and you let the model use future information to align the past, which produces a good-looking score built on leakage. The lag is both the key to the signal and an easy way to fool yourself, and a serious study has to deal with both sides of that.

4 The engineered-feature era

Before deep learning took over the field, researchers had to decide by hand which parts of the ECG were worth measuring. Each chosen feature was a small bet about where glucose information might be hiding.

The early HypoMon work bet on heart rate and corrected QT. In a 24-child study, Nguyen, Ghevondian, and Jones found both rising during naturally occurring nocturnal hypoglycemia, not induced episodes, and used a Bayesian neural network to reconstruct glucose profiles with significant correlation on a held-out test set (Nguyen et al., 2009). The question was older still. A 1995 note in Diabetes Care had already asked whether an ECG recording could detect hypoglycemia at all (Heinemann et al., 1995). The field has been circling the same idea for three decades.

The list of candidate features grew from there. Nuryani, Ling, and Nguyen used swarm-tuned support vector machines for hypoglycemia detection (Nuryani et al., 2012). Lipponen and colleagues used repolarization features, including QT and T-wave flattening, to flag hypoglycemic events in clamp data (Lipponen et al., 2011). Nguyen, Su, and Nguyen used a wider set at once, including heart rate, corrected QT, PR, corrected RT, and corrected TpTe, applied to both hypo- and hyperglycemia in type 1 diabetes, and reported that the parameters moved in opposite directions for low versus high glucose (Nguyen et al., 2012).

Underneath the different methods, this era left two findings that pull against each other.

The encouraging one is that the glucose information really is in there. Different groups, using different descriptors, populations, and protocols, kept finding it. That

consistency is hard to dismiss as coincidence.

The sobering one is that none of those features belong to glucose. Heart rate is not glucose-specific, and neither are HRV, QT, or T-wave shape. Each of them also responds to the protocol, the person, the sensor, the sleep stage, the meal schedule, and the disease phenotype. So a classifier can post a strong score by learning one of those instead, unless the validation is deliberately built to remove those shortcuts. That problem only gets worse in the deep-learning era, where the shortcuts become more tempting and much harder to spot.

5 The deep-learning era

Deep learning changed the central question from which feature to trust to whether a model can find the relevant morphology itself, including the parts nobody thought to name. It can. The same flexibility is also what makes these results so easy to fool yourself with.

Porumb and colleagues built personalized detectors for nocturnal hypoglycemia from short raw ECG segments recorded by wearables over 14 days in healthy people (Porumb et al., 2020). The word personalized is doing real work. They argued that people differ enough from one another that a single cohort-level detector had been unreliable, so they did not claim to have found one universal ECG-to-glucose mapping. The claim was narrower and sturdier: a given person’s heart has its own signature, and a model trained on that person can read it.

Cordeiro, Karimian, and Park went the other way, toward scale. On ECG from 1119 subjects, their hyperglycemia classifier reported 94.53% AUC, 87.57% sensitivity, and 85.04% specificity (Cordeiro et al., 2021). That result matters because it points across people rather than within one, which suggests the hyperglycemia signal is more than one person’s baseline drift relabeled as prediction.

Then the field reached for the harder target of a continuous glucose number rather than a yes-or-no, and the limits showed. Fella Arbi and colleagues attempted continuous glucose estimation from ECG using three patients and eight days of the D1NAMO dataset, with a CNN to segment the beats and regression on the extracted parameters. The best model performed well within patients, but a three-person sample makes it a proof of concept rather than a usable estimator (Fella Arbi et al., 2023). Zhang and colleagues returned to the same D1NAMO data in 2026 with a more mature design, a multi-expert SC-ResNet that splits segments by glucose state, extracts and fuses deep features, and hands a random forest the final prediction, and reported improved RMSE, MARD, and Clarke error grid results (Zhang et al., 2026). The architecture is genuinely more sophisticated than earlier pipelines. The open question is the same one every paper in this section inherits: how much of the performance carries over to a new person, a longer time window, a different device, and an ordinary chaotic day.

This is why the most useful deep-learning papers are not the ones with the highest score but the ones that show what the model is using. Porumb and colleagues

visualized which regions the model relied on, including T-wave and ST-interval contributions, so the result could be checked against physiology (Porumb et al., 2020). Ha and colleagues paired saliency analysis with their delay alignment to tie the model’s decisions back to plausible ECG and HRV features (Ha et al., 2026). Without that link, an ECG glucose model can look accurate while quietly learning a shortcut, right up until it meets someone new.

6 Multimodal wearables make the problem more honest

If the ECG is a weak, nonspecific proxy, the better move may be to stop forcing it to stand alone and give it some context.

Dave and colleagues paired ECG with an accelerometer over two weeks in five healthy participants. Their best regression model detected both hypo- and hyperglycemic excursions better with ECG plus accelerometer than with ECG alone, reaching roughly mid-to-high 70% sensitivity and specificity for both excursion types (Dave et al., 2022).

Next to the 90-something AUCs in the previous section that looks unimpressive, but it is one of the more honest results in this area. Activity, sleep, and posture all change HRV, and exercise moves glucose and the ECG at the same time, which is exactly the kind of confound that makes a lone ECG misleading. An accelerometer is a cheap way to tell “the heart rate rose because glucose fell” apart from “the heart rate rose because the person walked upstairs.” The single-sensor papers post the bigger numbers, but this is the more realistic framing.

This is probably where the ECG belongs in an actual device. Not as the only sensor, but as one input in a context model that can also include movement, skin temperature, sleep state, time since the last meal, medication timing, the previous glucose reading, and the occasional real glucose measurement to keep the whole thing calibrated.

7 Why validation is the whole story

The most important paper here is not the one with the deepest network or the highest score. It is the one that strips away the usual conveniences and checks what is left. That is Azam and Singh’s 2026 re-evaluation of HRV biomarkers for glucose sensing (Azam & Singh, 2026), and it reads like a corrective to the rest of the field.

They took 43 male type 2 diabetes patients, kept HbA1c and fasting blood glucose as separate targets instead of merging them into one convenient measure of “glucose,” and validated strictly. They used leave-one-subject-out cross-validation, with feature selection and standardization done inside each fold so the model never saw the test subject in any form. Under that setup the associations were still real, but the accuracy was not clinically useful. The best results were modest: an R^2 of 0.222 for HbA1c and 0.086 for fasting blood glucose, with mean absolute errors of 1.18 percentage points and 2.27 mmol/L (Azam & Singh, 2026). A 2.27 mmol/L

error is not a glucose monitor.

The value of the paper is the standard it sets. It lays out the questions a serious ECG glucose study has to answer before its score means anything:

- Are subjects separated between training and testing?
- Are days separated, or can the model memorize a person’s local baseline?
- Are feature selection, scaling, lag estimation, and calibration done inside the training fold only?
- Is the target acute glucose, fasting glucose, interstitial CGM, capillary glucose, plasma glucose, or HbA1c?
- Is performance reported as event detection, regression error, Clarke error grid, MARD, alarm latency, false alarm burden, or all of them?
- Does the model survive sleep, meals, exercise, stress, illness, medications, arrhythmias, autonomic neuropathy, sensor movement, and different ECG hardware?

Without those controls, an accuracy figure can be perfectly real for one dataset and useless everywhere else.

8 What the evidence says when you read it together

Read end to end, the literature does not support the headline claim that an ECG can replace glucose measurement. It never comes close.

What it does support is quieter and more interesting. The ECG carries measurable information about glycemic state, because a glucose excursion drags autonomic and cardiac electrophysiology along with it. The information is real, but it is indirect, partial, and easily drowned out.

The firmest ground is hypoglycemia and repolarization, meaning QT, T-wave morphology, and the related ventricular repolarization descriptors that keep reappearing across labs ([Eckert & Agardh, 1998](#); [Robinson et al., 2003](#); [Koivikko et al., 2008](#); [Lee et al., 2016](#)). The best modern case for hyperglycemia rests on large-scale classification and on delay-aware, challenge-controlled HRV work ([Cordeiro et al., 2021](#); [Ha et al., 2026](#)). The sharpest warning comes from subject-independent revalidation, where the association survives but the clinical accuracy does not ([Azam & Singh, 2026](#)). Hold those three together and you have the real shape of the field.

So the honest framing is proxy sensing rather than glucose sensing. The ECG never sees glucose. It sees the cardiovascular response to glucose, mixed in with the response to sleep, movement, stress, and disease. That response is sometimes strong, sometimes delayed, sometimes absent, and usually sharing the trace with something else.

9 Where this should go next

The field does not need another small retrospective model with an impressive within-subject score, because there are already plenty of those. It needs the slower, less exciting work:

- pre-registered protocols, so the hypothesis comes before the result;
- simultaneous ECG, CGM, and reference blood glucose;
- subject-independent and device-independent validation;
- explicit handling of physiological lag;
- alarm latency and false-alarm burden reported rather than buried;
- subgroup analysis for autonomic neuropathy, diabetes type, medications, arrhythmias, sleep, and exercise, since those are the cases where the signal is known to weaken;
- comparison against simple baselines such as time of day, recent meal, movement, heart rate, and previous glucose, which have to be beaten first;
- external replication on data the model’s designers never touched.

It also needs honesty about what is being sold. “This person may be entering a risky glucose state” is a useful statement even though it is not “your glucose is 5.7 mmol/L.” A carefully tuned supporting alarm and a non-invasive continuous glucose monitor are different medical claims, and treating them as the same product is how reasonable engineering turns into a regulatory and clinical problem.

That is the real promise of this literature, and it is worth more than the magic version. Not glucose read off the heartbeat, but a cheap, continuous signal that people already wear, which, combined with context and checked against the occasional real measurement, might warn them before glucose becomes dangerous. The signal is real. The translation is not finished yet.

References

- Azam, M. B., & Singh, S. I. (2026). Re-evaluating heart rate variability biomarkers for glucose sensing: The impact of age normalisation and subject-independent validation. *BMC Medical Informatics and Decision Making*, 26, Article 166. <https://doi.org/10.1186/s12911-026-03455-8>
- Cordeiro, R., Karimian, N., & Park, Y. (2021). Hyperglycemia identification using ECG in deep learning era. *Sensors*, 21(18), Article 6263. <https://doi.org/10.3390/s21186263>
- Coronado-Reyes, O. I., Tellez-Anguiano, A. C., Castro-Pimentel, L. A., & Gutierrez-Gnecchi, J. A. (2025). Non-invasive hyperglycemia detection via electrocardiogram using discrete wavelet transform and machine learning. *Cureus*, 17(3), Article e80548. <https://doi.org/10.7759/cureus.80548>

- Dave, D., Vyas, K., Branan, K., McKay, S., DeSalvo, D. J., Gutierrez-Osuna, R., Cote, G. L., & Erraguntla, M. (2022). Detection of hypoglycemia and hyperglycemia using noninvasive wearable sensors: Electrocardiograms and accelerometry. *Journal of Diabetes Science and Technology*, 18(2), 351-362. <https://doi.org/10.1177/19322968221116393>
- Eckert, B., & Agardh, C.-D. (1998). Hypoglycaemia leads to an increased QT interval in normal men. *Clinical Physiology*, 18(6), 570-575. <https://doi.org/10.1046/j.1365-2281.1998.00138.x>
- Fellah Arbi, K., Soulimane, S., Saffih, F., Bechar, M. A., & Azzoug, O. (2023). Blood glucose estimation based on ECG signal. *Physical and Engineering Sciences in Medicine*, 46(1), 255-264. <https://doi.org/10.1007/s13246-022-01214-3>
- Ha, J., Hwang, H. B., Kim, H., Lee, S., Lee, J., Park, J. H., Lee, J., & Kim, I. Y. (2026). Noninvasive detection of acute hyperglycemia using signal from wearable ECG sensors considering individual HRV response delays to glucose. *Biosensors*, 16(5), Article 251. <https://doi.org/10.3390/bios16050251>
- Heinemann, L., Meinhold, J., & Kunze, W. (1995). Hypoglycemia detection by ECG recording? *Diabetes Care*, 18(1), 139-140. <https://doi.org/10.2337/diacare.18.1.139>
- Koivikko, M. L., Karsikas, M., Salmela, P. I., Tapanainen, J. S., Ruukonen, A., Seppanen, T., Huikuri, H. V., & Perkiomaki, J. S. (2008). Effects of controlled hypoglycaemia on cardiac repolarisation in patients with type 1 diabetes. *Diabetologia*, 51(3), 426-435. <https://doi.org/10.1007/s00125-007-0902-y>
- Lee, A. S., Brooks, B. A., Simmons, L., Kilborn, M. J., Wong, J., Twigg, S. M., & Yue, D. K. (2016). Hypoglycaemia and QT interval prolongation: Detection by simultaneous Holter and continuous glucose monitoring. *Diabetes Research and Clinical Practice*, 113, 211-214. <https://doi.org/10.1016/j.diabres.2016.01.005>
- Lipponen, J. A., Kemppainen, J., Karjalainen, P. A., Laitinen, T., Mikola, H., Karki, T., & Tarvainen, M. P. (2011). Hypoglycemia detection based on cardiac repolarization features. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 4697-4700. <https://doi.org/10.1109/IEMBS.2011.6091163>
- Marfella, R., De Angelis, L., Siniscalchi, M., Rossi, F., Giugliano, D., & Nappo, F. (2000). The effect of acute hyperglycaemia on QTc duration in healthy man. *Diabetologia*, 43(5), 571-575. <https://doi.org/10.1007/s001250051345>
- Nguyen, H. T., Ghevondian, N., & Jones, T. W. (2009). Real-time detection of nocturnal hypoglycemic episodes using a novel non-invasive hypoglycemia monitor. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 3822-3825. <https://doi.org/10.1109/IEMBS.2009.5335144>
- Nguyen, L. L., Su, S., & Nguyen, H. T. (2012). Identification of hypoglycemia and hyperglycemia in type 1 diabetic patients using ECG parameters. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2716-2719. <https://doi.org/10.1109/EMBC.2012.6346525>
- Nuryani, Ling, S. S. H., & Nguyen, H. T. (2012). Electrocardiographic signals and swarm-based support vector machine for hypoglycemia detection. *Annals of Biomedical Engineering*, 40(4), 934-945. <https://doi.org/10.1007/s10439-011-0446-7>
- Porumb, M., Stranges, S., Pescapè, A., & Pecchia, L. (2020). Precision medicine and artificial intelligence: A pilot study on deep learning for hypoglycemic events detection based on ECG. *Scientific Reports*, 10, Article 170. <https://doi.org/10.1038/s41598-019-56927-5>

- Robinson, R. T. C. E., Harris, N. D., Ireland, R. H., Lindholm, A., & Heller, S. R. (2003). Comparative effect of human soluble insulin and insulin aspart upon hypoglycaemia-induced alterations in cardiac repolarization. *British Journal of Clinical Pharmacology*, 55(3), 246-251. <https://doi.org/10.1046/j.1365-2125.2003.01726.x>
- Tobore, I., Li, J., Kandwal, A., Liu, Y., Nie, Z., & Wang, L. (2019). Statistical and spectral analysis of ECG signal towards achieving non-invasive blood glucose monitoring. *BMC Medical Informatics and Decision Making*, 19(Suppl 6), Article 266. <https://doi.org/10.1186/s12911-019-0959-9>
- Zhang, J., Wang, J., Feng, W., Zareen, S. S., Wu, J., & Luo, C. (2026). Noninvasive blood glucose estimation via ECG: A multi-expert SC-ResNet model. *IEEE Journal of Biomedical and Health Informatics*. <https://doi.org/10.1109/JBHI.2026.3687523>