Blood: source-bounded personal anxiety scoring from glucose and wearable signals

Alan N. Pham[1]

[1]AO Labs, Worcester, MA, USA

Blood is a private personal health record that turns automatic CONTOUR NEXT ONE glucose readings and phone-derived health signals into a one-decimal stabilization score, one readable health read, six aligned graphs, and one equal-weight current-readings snapshot for glucose, heart rate, estimated-or-source HRV, sleep, and steps. The anxiety estimate is not a diagnosis and does not infer mental state from a model trained on other people. It is a deterministic, source-bounded index calculated from current uploaded glucose, heart rate, true HRV when available, fresh-enough sleep recovery, daily step state, small rate-of-change terms from recent glucose, HR, and HRV history, and recent personal HR/HRV baselines. Estimated HRV is displayed as a proxy recovery estimate and can matter when it dips against recent history, but its absolute value is not treated like a true RMSSD threshold. When true RMSSD and proxy HRV overlap on prior days, later proxy HRV values are calibrated against Alan's own true-RMSSD/proxy ratio before they enter the graph or score. The health read is a source-bounded assessment written as one natural clinical read: one good sign, one biggest watchout, and one concrete food, water, or movement move to improve the day; diagnosis boundaries, rolling-window basis lines, and recalculation mechanics stay in the paper, references, API, or diagnostics instead of appearing as tiny secondary text under the main read. Sleep stays visible as graph history and bounded score context, but old sleep is not turned into a current instruction. Heart-rate samples are retained with a separate per-type history limit so dense partial syncs cannot crowd older HR days out of the summary; the public HR graph uses five-minute median buckets while preserving the latest raw HR point. Steps are collapsed to one latest daily total per Eastern date, matching the Samsung Health day-total model, while short-range views draw HRV, sleep, and step lines through adjacent context points so 24h behaves as a zoom into the longer trend. The anxiety graph samples across the full available stored input range instead of keeping only the newest dense cluster. Low steps before morning are displayed as early-day context rather than as an

anxiety watchout. A protected Disregard action can remove a selected glucose measurement from public calculations without erasing the protected audit record. The system is therefore a transparent self-monitoring surface: every score, health read, current reading, and graph can be reconstructed from source values, source history, recent baselines, and the rule table.

1 Introduction

Many personal health apps show glucose, sleep, heart rate, or activity as separate traces. That separation is technically clean, but it leaves the user to decide which signal explains the current body state. Blood was built for the opposite interface problem: a current glucose reading or wearable signal should not become another small dashboard burden. It should become one clear score, then the source graphs that make the score auditable.

The score is intentionally called an anxiety estimate rather than an anxiety diagnosis. Blood does not classify psychiatric anxiety, does not use population diagnosis labels, and does not claim that one biomarker proves a mental state. The implementation uses physiologic signals only as personal stabilization evidence. Low blood glucose can produce adrenergic symptoms such as fast heartbeat, shaking, sweating, nervousness, or anxiety (*1*). HRV standards define variability from beat-to-beat RR intervals and distinguish true HRV measurement from looser heart-rate summaries (*2*). Reduced HRV is associated with anxiety disorders in meta-analytic literature, but the effect is group-level and not sufficient for diagnosing an individual state (*3*). Blood therefore treats HRV as one recovery signal among several, not as a standalone mental-health label.

The data path is also bounded. Android Health Connect exposes categories such as activity, sleep, vitals, and blood glucose, Samsung Health can share selected Samsung Health data through Health Connect, and Health Connect defines a distinct RMSSD HRV record type when a source actually exposes it (*4–6*). In the current Blood deployment, glucose reaches the server from a Blood Bridge path connected to the CONTOUR NEXT ONE meter, while heart rate, sleep, and steps come from Health Connect records written by the phone or wearable. The bridge reads Health Connect records in pages, and the interface separates health-upload time from Samsung/Health Connect source time, so a recent bridge upload cannot make an older shared heart-rate sample look current. When a source does not provide true HRV, Blood computes a labeled proxy from heart-rate samples only when the sample geometry is strong enough to make the estimate useful. If a glucose measurement is wrong, accidental, or not useful for the visible record, the website can mark that one reading as disregarded through a protected edit-key action that is separate from the bridge ingest token. Disregarded readings remain available in protected export/history, but they are removed from the public summary, latest glucose, glucose graph, anxiety reconstruction, time-pattern detector, and visible readings table.

2 Results

2.1 A single score remains reconstructable from source values

Blood renders the anxiety estimate first, then a best-effort instability window, and then six aligned graphs for anxiety, glucose, heart rate, HRV, sleep, and steps. The placement matters because the score summarizes the latest uploaded source state plus recent change, the pattern window names the time block that has carried or is currently carrying more actionable source instability, and the graphs are the audit trail. The anxiety graph is reconstructed from the full available stored glucose and health input range, with sampling used only to keep the browser payload bounded. The display value is computed by:

$$S = \text{round}_{0.1}\left(\min\left(10, \max\left(1, 2R\right)\right)\right)$$

where R is the raw additive score and S is the visible 1–10 score. The raw score starts at $R = 2.2$. Source-specific factors then add or subtract the values shown in Table 1. Positive values raise the anxiety estimate. Negative values lower it.

Table 1. Deterministic anxiety-score factors in the current Blood deployment. The final visible score is $2R$, clamped to 1–10 and rounded to one decimal place.

Signal	Rule	Raw change
Glucose	<70 mg/dL	+1.00
Glucose	>180 mg/dL	+0.90
Glucose	<82 or >140 mg/dL	+0.35
Glucose	82–140 mg/dL	-0.15
Heart rate	≥100 bpm	+0.85
Heart rate	≥85 bpm	+0.40
Heart rate	55–75 bpm	-0.15
Source HRV	<25 ms	+0.80
Source HRV	<40 ms	+0.45
Source HRV	≥65 ms	-0.25
Estimated HRV	shown as estimate; absolute value does not trigger threshold weight	0.00
Sleep	<300 minutes asleep	+0.90
Sleep	<360 minutes asleep	+0.45
Sleep	≥420 minutes asleep	-0.25
Steps	<1,500 daily steps	+0.35
Steps	<4,000 daily steps	+0.15
Steps	≥8,000 daily steps	-0.25
Glucose dynamics	drop of at least 25 mg/dL near the low edge, or fast downward rate	+0.30 to +0.55
Glucose dynamics	rise of at least 30 mg/dL near the high edge, or fast upward rate	+0.25 to +0.50
Heart-rate dynamics	rise of at least 15 bpm into the raised/high range	+0.25 to +0.45
HRV dynamics	drop of at least 10 ms and 22% into the low range	+0.25 to +0.45
Heart-rate baseline	current HR at least 12 bpm above recent Blood baseline	+0.25
HRV baseline	current HRV at least 8 ms and 22% below recent Blood baseline	+0.28 to +0.35

This additive design is deliberately small. A glucose crash, a fast glucose change, low fresh sleep recovery, high heart rate, fast HR rise, low HRV, or HRV dip can move the score, but no single signal silently becomes the whole answer. Stable signals can offset a stressed signal. Missing or stale signals do not create fake precision; they simply do not contribute to the current action state.

2.2 The live score has a full audit trail

Table 2 reconstructs the live Blood score checked on June 27, 2026 at 6:13 PM ET. The latest score is not a black-box output. It is the arithmetic sum of five visible source states.

Table 2. Live score reconstruction from the Blood summary API on June 28, 2026 after estimated-HRV absolute thresholds were removed.

Term	Source state	Raw contribution
Baseline	Personal neutral starting point	+2.20
Glucose	159 mg/dL, near the high edge	+0.35
Heart rate	64 bpm, inside the calm band	-0.15
Estimated HRV	23 ms estimated HRV, shown as proxy estimate	0.00
Sleep	7.3 h sleep, at least 420 minutes	-0.25
Steps	3,419 daily steps, light	+0.15
Glucose dynamics	Recent rise from 94 to 159 mg/dL	+0.25
Raw total	$2.20 + 0.35 - 0.15 + 0.00 - 0.25 + 0.15 + 0.25$	2.55
Visible score	2×2.55, clamped and rounded	5.1/10

The label associated with the score is also deterministic: scores at or below 2.5 are labeled low, scores through 4.5 steady, scores through 6.5 watch, scores through 8 high, and higher scores very high. The live 4.4/10 score is therefore labeled steady.

2.3 HRV is source-preferred and proxy-labeled

The HRV value used in the anxiety score is true HRV only when a source HRV/RMSSD record exists. Blood preserves that value and labels its basis as source RMSSD. When a true HRV record is absent, Blood estimates HRV from heart-rate samples and labels the result with `ms_est`. The estimator converts heart rate to an RR interval using:

$$RR_i = \frac{60000}{\mathrm{bpm}_i}$$

It then computes an RMSSD-like proxy over adjacent RR differences:

$$\widehat{HRV} = \sqrt{\frac{1}{n}\sum_{i=2}^{n+1}(RR_i - RR_{i-1})^2}$$

The current estimator rejects weak windows before producing a value. It deduplicates heart-rate

records by timestamp, trims the edges of sleep sessions, prefers points inside sleep windows when the trimmed segment has enough samples, otherwise searches resting windows, and evaluates 30-minute candidate windows. A sleep candidate must contain at least 24 points, at least 20 accepted adjacent pairs, a duration of at least 22 minutes, median sample gap no higher than 1.75 minutes, coverage at least 86%, median heart rate no higher than 88 bpm, heart-rate standard deviation no higher than 5.5 bpm, median adjacent heart-rate step no higher than 2.8 bpm, 90th-percentile adjacent heart-rate step no higher than 6.5 bpm, and 90th-percentile adjacent RR difference no higher than 135 ms. A fallback resting candidate is stricter: at least 30 points, 26 accepted adjacent pairs, median sample gap no higher than 1.5 minutes, coverage at least 90%, and lower heart-rate and RR-difference noise thresholds. Blood removes noisy adjacent pairs with tighter median-absolute-deviation gates before computing the RMSSD-like value. It then selects up to 12 independent or low-overlap candidate windows, requires at least three selected windows and 75 accepted pairs, and reports the pair-count-weighted median value across the selected windows. If true source RMSSD and proxy HRV exist on at least two overlapping source dates, Blood computes the median true/proxy ratio and calibrates later proxy estimates within bounded limits.

This is stronger than a day-wide heart-rate swing, but it is still not the same as a chest-strap or ECG-grade beat-to-beat HRV record. The deployed API therefore carries `basis`, `quality`, `confidence`, sample count, accepted-pair count, rejected-pair count, median gap, coverage ratio, selected-window count, window spread, window length, calibration factor, calibration sample count, and the uncalibrated value when calibration is active. The value remains an estimate unless Health Connect or another source exposes true RMSSD; its source timestamp is the end of the selected sleep/rest window, not a live beat-interval feed.

3 Discussion

Blood's anxiety score is best understood as a personal stabilization index. It is useful when it reduces the work of deciding which current body signal deserves attention. It is not useful if it is treated as a medical diagnosis, a universal anxiety scale, or proof that one biomarker explains a mood state.

The main design choice is transparency. The score can be computed by hand from the table, but the page should not make Alan do that arithmetic or interpret generator mechanics. The front page therefore shows the score first, then one health read in simple English, then the graphs as the source audit. The health read answers one question: what should Alan do with this data right now? It starts

with one concrete food, water, or movement recommendation, then explains why that move helps, then names the strongest source-backed good signal. Current metrics, recurring patterns, spikes, and dips can feed those sentences, but they must be resolved into that simple read instead of a scattered metric paragraph or role-prefix labels such as `Good sign:`, `Biggest watchout:`, or `Best move:`. The read uses one normal-weight, uncolored text flow; the page does not split the sentence into green and red phrase highlights. Provenance, diagnosis boundaries, rolling-window basis lines, and recalculation mechanics belong in references, this paper, API/debug detail, or diagnostics unless the boundary is the one primary message. The anxiety graph uses the same formula as the top score and the same x-axis as the source graphs, so changes in the score can be compared directly with glucose, HR, HRV, sleep, and steps. The chart renderer separates source measurement time from plotting endpoint: current latest-upload values use the current upload endpoint for x-position so the last current point lands on the right edge, while the label and tooltip still show the true source measurement time. Short-range views such as 24h keep the selected time window but use adjacent context points for all six graph lines when visible points exist, so the view reads like a zoom into the seven-day trend rather than a single dot or cut-off fragment. Heart-rate history is retained per type and the public HR graph is bucketed at five minutes, so a dense recent sync can update the current point without erasing the previous days from the graph. The current-readings snapshot below the graphs gives glucose, HR, HRV, sleep, and steps equal visual weight. Sleep sessions are reduced to one daily asleep/duration/awake/stage total per source sleep date, so a later nap cannot replace the main sleep value. Step records are reduced to one latest daily total per Eastern date, so repeated Health Connect uploads cannot multiply a day into a fake 24-hour count; low step totals before morning remain context because the day has not had enough time to accumulate ordinary movement. A row-level Disregard control in the latest-readings table removes a selected glucose point from the visible and computed record while preserving the protected raw record; re-uploading the same bridge reading keeps that disregard flag. The same structure also keeps recommendations bounded. Blood selects easy actions from the strongest positive actionable factor, including rate-of-change factors, not from a generated personality profile. If glucose is too low, near the low edge, or dropping quickly, Blood names the glucose value or drop and asks for a carb-plus-protein snack and water. If glucose is too high, near the high edge, or rising quickly, it asks for protein/fiber with carbs, water, and an easy walk. If heart rate is too high, raised, or rising quickly, it names the bpm value or rise and asks for water, a protein/fiber snack, and an easy walk. If estimated HRV is merely low as a proxy, Blood says estimated HRV looks normal; if HRV drops meaningfully against recent history, it can become a recovery-trend watchout. If daily steps are low or light after the

day has started, it names the step count and asks for water, normal meals, and an easy walk. Sleep can affect the score only when the source sleep record is still fresh enough to describe current recovery context, but sleep is not selected as the visible recommendation source and stale sleep is left as graph history.

The strongest limitation is HRV. HRV is defined from beat-to-beat intervals, and sampled heart-rate records are a compressed derivative of that signal (*2*). Blood reduces that problem by selecting dense sleep/rest windows, calibrating proxy days against any overlapping true RMSSD days, and labeling the proxy, but the result remains a proxy. The second limitation is source freshness. Android WorkManager can run durable background work, but periodic work has a system-governed schedule and a 15-minute minimum repeat interval; exact execution time still depends on system constraints and optimization (*7*). A watch can monitor heart rate continuously while the health bridge has not uploaded recently, or while Samsung Health has not exposed the newest shared sample through Health Connect. Blood treats those as separate stale states and shows the bridge request, bridge check-in, health-upload time, and sample timestamp rather than inferring a newer value. A direct Samsung Health current-feed path would require Samsung's proprietary Health Data SDK or a separate watch sensor bridge, not only the current Health Connect phone bridge. The third limitation is symptom calibration. The current weights are personal, deterministic, and inspectable; they are not learned from a clinical dataset and have not been validated against symptom diaries. A future version should store Alan's own reported anxiety, caffeine, meal, medication, illness, and stressor context when he chooses to provide it, then evaluate whether the fixed weights should be changed.

4 Materials and Methods

4.1 System inputs

The Blood server stores glucose readings and health metrics separately. Glucose records include measured time, capture time, source, reading identifier, mg/dL value, relation to meal when available, and specimen source. The current automatic glucose path is the CONTOUR NEXT ONE meter bridge. Health metrics include heart rate, HRV when exposed, sleep sessions, and step windows. Blood Bridge requests HRV permission when available, but missing HRV permission does not block heart-rate, sleep, or step uploads. After `Connect Blood` receives the required Bluetooth and Health Connect permissions, the open activity reads a recent seven-day window and performs the first meter and health upload directly; it posts the accepted or blocked result to the bridge-state endpoint before leaving later freshness to Android scheduling. A CONTOUR scan timeout is recorded as a meter-source result rather than a cancellation, so a missing or sleeping meter cannot abort steps, heart-rate, or sleep upload. The recurring background worker uses the same paged Health Connect read so dense Samsung heart-rate records can reach the Blood API in upload-sized batches. App open, permission grant, boot, app update, and public-site fresh-run requests also queue background work, and the bridge maintains a rolling invisible one-time worker in addition to the durable periodic worker. After a CONTOUR meter reading, opening the public Blood page records a fresh-sync request; the bridge uploads stored meter records first, then Health Connect records, when Android allows the invisible worker to run. Fresh-run requests append a follow-up immediate run when one is already active instead of replacing or cancelling the old run or dropping the new request, and worker cancellation is not reported as a completed no-data sync. The public summary endpoint records a fresh-sync request and returns bridge request/check-in state; the bridge posts check-ins after both the direct first upload and automatic workers. Server writes merge health metrics by stable metric identifier; readback uses per-type limits, not one shared row cap, so dense heart-rate records cannot evict older sleep, step, HRV, or HR history from the summary. Bridge and fallback ingest endpoints require the Blood ingest token. The protected `DELETE /api/blood/readings/{readingId}` route marks one glucose reading as disregarded with the Blood edit key. The summary builder filters disregarded readings before computing latest glucose, graph points, range statistics, dynamic glucose factors, anxiety score, instability patterns, and the visible readings table.

4.2 Score implementation

The current anxiety score is implemented by `estimateAnxietyState` in `server.js`. The function receives the latest glucose, heart-rate, HRV, fresh-enough sleep, daily-step value, dynamic factors, and recent personal baselines from the summary builder. Dynamic factors are computed by `estimateDynamicFactorsAtTime`: it compares the current glucose point with a prior point 20 minutes to six hours earlier, the current heart-rate point with a prior point 20 minutes to three hours earlier, and the current HRV point with a prior point two hours to 14 days earlier. It only emits a positive factor when the direction and magnitude are large enough to support a concrete food, water, or movement response. Baseline factors add smaller weights when current heart rate is at least 12 bpm above the recent Blood baseline, or current HRV is at least 8 ms and 22% below its recent Blood baseline. Sleep is filtered out before the current score action path when its source end time is older than the freshness window. The function initializes `raw = 2.2`, applies the factor table, computes `score = clamp(1, 10, raw * 2)`, and rounds to one decimal place. It returns `score`, `scale`, `label`, `raw`, `factors`, `dynamics`, `suggestion`, `condition`, and the note "Source-bounded personal health read, not diagnosis."

The anxiety trend is implemented by `estimateAnxietyTrend`. It walks stored glucose readings and health-trend samples from the beginning of available input history, reconstructs score points from the most recent source values available at each timestamp, recomputes the dynamic factor leading into that timestamp, and requires at least two source inputs before keeping a historical point. It appends the latest score as the newest visible point when a score exists. If the reconstructed history is dense, the server samples across the full range instead of slicing off the beginning, with a default cap of 5,000 anxiety points for browser payload size. This makes the top score and the anxiety chart share one formula while keeping the graph tied to source records rather than generated filler.

The time-pattern surface is implemented by `estimateInstabilityPatterns`. It takes recent glucose readings plus health trends, assigns each source record to an Eastern time block, scores each record, and reports the block with the highest recent average instability when a visible actionable reason is available. Dense heart-rate samples are reduced to hourly median buckets so a watch sampling frequently cannot dominate the pattern only by producing more rows. The reported pattern contains the current block, predicted block, source counts, observations, directional condition text with values, a simple-detail sentence, and a basis line that states whether the result is a 45-day pattern or a best-effort current-upload signal. The top read turns this pattern into the one biggest watchout sentence and keeps

the basis line out of the first-read surface. The visible detail names actionable source states such as "HR too high (114 bpm)", "glucose near high edge (159 mg/dL)", or "HRV dip" instead of a vague aggregate count of samples that were high, low, short, light, or raised. When a clear repeated pattern is not yet available, Blood still names the best current abnormal spike, dip, raised, low, or light signal it can infer and pairs it with a concrete food, water, or movement action. Sleep-only history remains on the sleep graph and does not produce a visible unstable-window instruction. Because the summary endpoint recomputes this object from stored source records, the pattern changes automatically when additional glucose, HR, HRV, sleep, or step data changes the distribution.

4.3 HRV estimator

The HRV trend function first keeps any source HRV records and marks them as source RMSSD. It then calculates proxy HRV only for dates without source HRV. Heart-rate records are grouped by date, deduplicated by `measuredAt`, sorted by time, and converted to RR intervals. Sleep-session edges are trimmed before candidate selection. Candidate windows are continuous 30-minute spans with no accepted sample gap above two minutes. Windows with too few samples, too few accepted adjacent pairs, short duration, sparse coverage, high median heart rate, high heart-rate standard deviation, high adjacent heart-rate steps, high RR-difference noise, or unstable selected-window spread are rejected. The selected estimate is the pair-count-weighted median RMSSD-like value from up to 12 independent or low-overlap candidates. When at least two source RMSSD days overlap proxy-estimate days, the server uses the median true/proxy ratio to calibrate later proxy estimates and exposes the factor, sample count, spread, and uncalibrated value in the API and chart tooltip.

5 Acknowledgements

The paper records the current Blood implementation and Alan Pham's requested interface contract: the anxiety score belongs at the top of the page, formatting must remain compact and consistent, the six aligned graphs follow immediately, current latest-upload points align to the shared right-edge endpoint, chart labels carry source timestamps, all six short-range graph lines preserve line continuity from adjacent context points, the score uses both endpoint values and recent source movement, the pattern text names concrete metric states and values even under thin data, sleep is a daily source-date total rather than the latest same-date session, steps are daily totals rather than summed repeated packets, the header names health-upload freshness versus Samsung/Health Connect source freshness, and HRV

must be labeled according to the strongest source basis available.

6 Funding

No external funding is associated with this software paper.

7 Author contributions

Alan N. Pham defined the application need, source preferences, scoring placement, source-freshness requirement, and automatic-bridge requirements. Codex implemented the current paper and checked it against the deployed Blood source.

8 Competing interests

The author declares no competing interests.

9 Data availability

The live public summary exposes current compact Blood state at the Blood summary API. Private write tokens, raw export tokens, and private device data are not public.

10 Code availability

The implementation is in the Blood application source, including `server.js`, `app.js`, and the Android bridge connector. Public routes are `https://blood.aolabs.io/` and the fallback mirror `https://aolabs.io/blood/`.

11 Additional information

This paper is a single-user system record. It should not be used for diagnosis, treatment, or medication decisions. The anxiety estimate is a personal software index for stabilizing attention around current physiologic signals.

References

1. Centers for Disease Control and Prevention, *Low Blood Sugar (Hypoglycemia)* (2026; https://www.cdc.gov/diabetes/about/low-blood-sugar-hypoglycemia.html).

2. Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology, Heart rate variability: standards of measurement, physiological interpretation and clinical use. *Circulation* **93**, 1043–1065 (1996).

3. J. A. Chalmers, D. S. Quintana, M. J.-A. Abbott, A. H. Kemp, Anxiety disorders are associated with reduced heart rate variability: a meta-analysis. *Frontiers in Psychiatry* **5**, 80 (2014).

4. Android Developers, *Health Connect data types* (2026; https://developer.android.com/health-and-fitness/health-connect/data-types).

5. Android Developers, *HeartRateVariabilityRmssdRecord* (2026; https://developer.android.com/reference/android/health/connect/datatypes/HeartRateVariabilityRmssdRecord).

6. Samsung Developer, *Accessing Samsung Health Data through Health Connect* (2026; https://developer.samsung.com/health/blog/en/accessing-samsung-health-data-through-health-connect).

7. Android Developers, *Define work requests* (2026; https://developer.android.com/develop/background-work/background-tasks/persistent/getting-started/define-work).