

Deficit Loop: Closed-Loop Adaptive Calorie Targeting with Per-User AI Agents over Messaging

Xiaomeng Li · NanoRhino, Houston, TX · xiaomengli@nanorhino.com · October 2026

Abstract

Digital weight-loss tools prescribe daily calorie targets open-loop: a prediction equation estimates energy expenditure, a deficit is subtracted, and the number is rarely checked against what actually happens to the individual. We describe the *deficit loop*, a closed-loop calorie-target controller deployed in production inside an AI coaching service delivered entirely over consumer messaging (SMS and WhatsApp). Each user is served by an isolated, stateful AI agent; a deterministic controller re-estimates the user's energy expenditure each week from their own logged meals and weigh-ins, and re-derives the target under layered physiological safety guardrails. The central design problem is that the measurement channel is systematically dirty — self-reported intake carries large, person-specific under-reporting bias — so the controller maintains two expenditure estimates in deliberately different units and uses their relationship to *route* between opposite interventions: adjust the target, or repair the measurement. We prove a fixed-point property of the control law: with an absolute (not multiplicative) deficit applied in the user's own logging units, the steady-state true deficit equals the intended deficit regardless of stable reporting bias. We report engineering constraints of operating hundreds of long-lived stateful agents over carrier messaging (latency, delivery, safety filtering, structural separation between a deterministic controller and a generative messaging layer), and field observations from seven weeks of closed-loop operation: measured expenditure runs materially below both consumer-calculator and textbook-formula predictions (median ≈ 300 kcal/day below the plan's textbook formula and ≈ 500 below the consumer-calculator recommendation), and on 28-day forward weight-change prediction the loop nearly eliminates the formula baseline's systematic bias (-0.02 vs -0.15 kg/week) while improving median absolute error (0.26 vs 0.36 kg/week). We discuss why the dominant harm of expenditure over-estimation in an LLM-mediated coach is *behavioral* — the model confidently urges the user to eat more — and how the architecture contains it. This is an observational systems report, not a clinical-efficacy claim; limitations and ethics are discussed throughout.

Keywords: closed-loop control, energy balance, conversational agents, LLM agents, mHealth, weight management, self-report bias

1. Introduction

Every consumer calorie tool begins the same way: height, weight, age, and an activity guess go into a prediction equation (typically Mifflin-St Jeor [1]), an intended deficit is subtracted, and the user receives a daily target. The number is a population estimate wearing the costume of a personal fact. Individual expenditure varies widely around the prediction, and — more corrosively — the *measurement channel* through which any correction would have to flow is itself biased: against doubly-labeled-water ground truth, self-reported energy intake typically under-runs true intake by roughly 10–45% depending on instrument and population, approaching 50% in diet-resistant cohorts, and the under-reporting is stable and person-specific rather than random [2–4].

The consequence is a specific, repeatable failure. A user eats to their prescribed number, the scale stops moving, and the system — if it reacts at all — reacts with the only move it has: lower the number. But a stall has two indistinguishable causes demanding *opposite* responses. Either the target is genuinely too high (the prediction over-estimated expenditure), in which case lowering it is correct; or a material share of intake never reaches the log, in which case the target may be physiologically fine and lowering it punishes a measurement problem with a dietary restriction — an intervention that is at best useless and at worst harmful.

We treat weight coaching as a control problem with a dirty sensor. This paper describes the **deficit loop**, a closed-loop calorie-target controller that has been operating in production since late July 2026 inside NanoRhino, an AI nutrition-coaching service delivered entirely over consumer messaging. The system holds, for each user, two energy-expenditure estimates in deliberately different units — one back-solved from the user’s own logs (inheriting their reporting bias by construction), one predicted from anthropometrics (physiological) — and uses the relationship between them not as a target and not as a bound, but as a **router** that decides, per user per week, whether this is a *target problem* or a *measurement problem*.

Our contributions:

1. **A production architecture for per-user stateful AI coaching agents** over carrier messaging: isolation, layered memory, proactive scheduling, and a strict structural separation between a deterministic numeric controller and a generative messaging layer (§3).
2. **The deficit-loop controller** (§4): dual-unit estimation, misreporting-aware routing, confidence-weighted target generation, layered physiological guardrails, and a provable bias-invariance fixed point of the absolute-offset control law.
3. **Engineering field notes** (§5) from operating this fleet: latency envelopes, carrier delivery reality, outbound safety filtering, and cost structure — a regime (hundreds of long-lived stateful agents, weeks-to-months horizons) largely absent from the current LLM-agent literature.

4. **Field observations** (§6) from the first seven weeks of closed-loop operation, reported with their selection caveats: the gap between measured and predicted expenditure, forward-prediction accuracy against a formula baseline, guardrail activation profile, and a behavioral finding — the dominant harm of an over-estimated target in an LLM-mediated coach is that the model *tells the user to eat more*.

We are explicit about what this paper is not: it is an observational report on a deployed system, with a self-selected user base and no control group. Efficacy claims require the randomized or at least controlled comparison we outline as future work (§8).

2. Related Work

Energy-balance modeling and self-report error. Hall et al. [5] established quantitative dynamic models of human energy balance, replacing the folk “3,500 kcal per pound” static rule. The intake-balance method [6] back-solves energy intake and expenditure from observed weight change and validated the approach against doubly-labeled water (DLW) to within ≈ 40 kcal/day over two-year horizons. Most recently, a predictive equation derived from 6,497 DLW measurements enables detection of erroneous self-reported energy intake at screening time [7], and time-weighted TDEE averaging has been proposed for estimating intake during active weight loss [8]. These works validate the two premises our controller is built on — expenditure can be back-solved from weight + intake data, and self-report error is large, systematic, and *detectable* — but both operate as offline measurement/screening methods. Our contribution is to move the misreporting judgment **inside an online controller**, per user per week, where it selects between opposite *interventions* rather than flagging bad data: and with a twist — we *deliberately keep* the back-solved estimate in the user’s biased logging units rather than correcting it (§4.1), because a target derived in the user’s own units is attainable in the user’s own units; the bias cancels at the fixed point (§4.4).

Adaptive interventions and control-theoretic behavior change. The just-in-time adaptive intervention (JITAI) literature [9,10] frames digital health as sequential decision-making over momentary states, and control-systems formulations of behavioral intervention design [11] have long argued for feedback over open-loop prescriptions; recent deployments optimize intervention *timing and prompting* with reinforcement learning under real-world constraints [12,13]. Our system is complementary and unusual within this frame: the controlled variable is not prompt timing but the **prescription itself** (the daily calorie target), the plant model is physiological energy balance rather than a learned response model, and the central inference is not “is the user receptive now?” but “is the sensor lying, and in which direction?” The SmartLoss pilot trial [14] remains the closest clinical antecedent — model-based adherence feedback from weigh-in data (−9.4% vs −0.6% body weight against control over 12 weeks) — but its recommendations

flowed through human counselors; the loop was not autonomous. Consumer tools such as MacroFactor recalculate targets from weight trends for self-directed trackers; the user runs the tool, and honest logging is assumed rather than adjudicated.

LLM coaches for diet and behavior change. A fast-growing literature evaluates LLMs as health coaches: LLM-generated weight-loss coaching messages whose revised versions were rated comparably to human coaches' [15]; GPTCoach's tool-using, motivational-interviewing physical-activity coach [16]; Bloom, an LLM-augmented physical-activity application evaluated in a four-week controlled field study [17]; an agentic workflow for barrier-driven personalized nutrition coaching [18]; and a first randomized evaluation of LLM-delivered diet coaching in the wild, which found no consistent real-world benefit despite strong intrinsic performance [19]. Meta-analytic evidence on pre-LLM chatbots shows modest lifestyle-behavior effects [20,21]. A consistent theme — LLMs are fluent but unreliable at sustained numeric bookkeeping — is a design *premise* of our architecture: every number is computed by a deterministic controller, and the generative layer's numeric vocabulary is restricted by construction (§3.4, §4.5). To our knowledge no published system closes the loop from a user's own longitudinal logs and weigh-ins back onto the prescribed target through an LLM-mediated channel with structural non-disclosure of estimator internals.

LLM agents in production. The emerging agent-systems literature has begun to characterize stateful long-horizon agent workloads, notably the memory-system taxonomy and fleet-scale cost analysis of [22], alongside surveys of agent memory mechanisms [23]. That literature is benchmark-driven; long-lived *per-user* agents — each carrying months of personal state, operating over carrier messaging, with proactive (not purely reactive) behavior — remain scarcely documented from the field. We report this regime's engineering constraints first-hand (§5), including the memory-consolidation cost asymmetries that [22] predicts, observed at production scale.

What is deliberately out of scope here. The specific numeric thresholds, window schedules, and rate-limit constants of the production controller, and the estimator's guard logic, are not disclosed in this paper. The architecture and method are described to the depth of the system's US patent filing (patent pending); operational constants remain proprietary. We believe the contribution — the dual-unit routing formulation and its fixed-point property — survives this redaction intact.

3. System Architecture

3.1 TOPOLOGY: ONE GATEWAY, ISOLATED PER-USER AGENTS

NanoRhino runs a single message gateway that terminates consumer channels (SMS via a telephony provider; WhatsApp) and routes each inbound message to that user's **isolated agent workspace**. Isolation is physical, not logical: each user's agent has its own working directory, context documents, data stores, and scheduled jobs. No cross-user state exists at the agent layer; an agent cannot read, or be prompted into reading, another user's data. The fleet currently spans hundreds of user workspaces.

Each agent comprises (i) a generative language model that composes user-facing messages, (ii) a set of markdown context documents injected per turn (persona, user profile, plan, coaching-focus directives), and (iii) transactional stores: meal log (including photo-derived entries), weight log, plan store, engagement state, and run history.

3.2 LAYERED MEMORY

Message-native coaching lives or dies on memory: the medium implies continuity ("my coach knows me"), and a model alone forgets. Each agent maintains three memory strata — a short-term rolling store, a medium-term consolidated document, and a long-term document — with a daily consolidation pass that promotes durable facts upward. The design lesson after a year of operation: memory *writes* must be conservative and audited, because a wrong "fact" about a user compounds silently across every subsequent conversation.

3.3 PROACTIVITY WITHOUT SPAM

A coach that only answers is a food database. Agents schedule proactive touches — meal reminders, weigh-in prompts, weekly reports — through a per-user scheduler under a global anti-burst policy and an engagement lifecycle (active → pause → recall → silent) that decays contact pressure for unresponsive users and periodically attempts recall. Every proactive pipeline passes a pre-send check against recent user activity ("already logged today" suppresses the nudge).

3.4 THE GENERATIVE LAYER IS FENCED

Two structural rules govern the LLM layer. **Arithmetic is never delegated to the model:** every number a user sees (targets, macros, paces, dates) is computed deterministically and handed to the model as data. **Outbound messages pass a layered filter** (rule-based screen, then a

small-model classifier, then an exclusion list) that blocks internal reasoning, meta-commentary, and controller internals from ever reaching a phone. §4.5 extends this fencing to the controller's own estimates.

4. The Deficit Loop

4.1 DUAL-UNIT ESTIMATION

The controller maintains two expenditure estimates per user:

Apparent expenditure (`TDEE_apparent`, *logging units*) is back-solved from the user's own data over an adaptive multi-week observation window: mean logged daily intake minus the energy-storage rate implied by the regression slope of body weight (times an energy density of tissue change — the $\approx 7,700$ kcal/kg convention [24], with the dynamic-model correction of [25] as the acknowledged refinement). Because logged intake carries the user's stable multiplicative reporting bias k (logged = $k \times$ true), the estimate inherits that bias **by construction**. This is deliberate: a target derived from it is attainable *in the units the user actually logs in*.

Predicted expenditure (`TDEE_predicted`, *physiological units*) comes from anthropometric prediction times an activity factor. It never sets the displayed target; it exclusively sets the **safety floor** and the **deficit ceiling**.

The two are never mixed: comparing a logging-unit target against a physiological bound (e.g., BMR), or adding measured exercise energy to the apparent estimate (which already absorbs habitual exercise), is the category of unit error the architecture makes structurally impossible.

Sensing is defensive: same-day weigh-ins collapse to one reading; unlogged days are imputed at the user's own window mean rather than counted as zero (zero would fabricate an enormous spurious deficit); the observation window grows with user history between a minimum and maximum length; and a battery of safety interlocks (eating-disorder flags, medical directives, minimum age, opt-out, pause states) is evaluated *before* any estimate exists to be acted on.

4.2 THE ROUTER: TARGET PROBLEM OR MEASUREMENT PROBLEM?

The dimensionless index $\hat{k} = \text{TDEE_apparent} / \text{TDEE_predicted}$ estimates logging completeness. The router maps each weekly run to one of five branches; **only one branch ever moves the target**:

- **control** ($\hat{k}$ within a middle band, sufficient weigh-ins): normal closed-loop operation — re-derive the target (§4.3).

- **coverage** ($\hat{k}$ well below the band): a large share of intake never reaches the log; *calories are not this user's bottleneck*. The target is frozen, a logging-calibration interaction is initiated, and target-reduction messaging is affirmatively suppressed.
- **verify** ($\hat{k}$ above the band): usually a short-window artifact (water-mass swings, sparse atypical logs); hold and lengthen the window.
- **reject** ($\hat{k}$ implausible): frozen, silent, flagged for human review — contradictory data is a data problem, not a coaching moment.
- **low-confidence** (no computable estimate / too few weigh-ins): hold; at most a single weigh-in invitation.

The router does not predict who will stall; it decides **how to respond to a stall that has already happened**. (A stall-precursor predictor we evaluated during development topped out near precision 0.45 and was discarded; the routing formulation is what survived contact with production data.) Production $\hat{k}$ spread is wide — decile span roughly 0.3 to 1.5 (production snapshot 2026-07-29, $n = 40$: p10 0.317, median 0.859, p90 1.543; the same snapshot appears in the provisional disclosure package) — direct evidence that the routing problem is real: a single-response controller would mis-treat a material fraction of users, and the users it would most confidently restrict (lowest $\hat{k}$) are precisely the ones who need measurement repair, not restriction.

A per-user **phase machine** (calibrate / loss / hold / maintain / pause) gates what “success” means, so a user who *chose* maintenance is never misread as plateaued and intervened on. The controller autonomously performs exactly one transition (calibrate → loss, when confidence first suffices); maintenance and pause are entered only by the coach or an operator.

4.3 CONFIDENCE-WEIGHTED TARGET GENERATION AND GUARDRAILS

Within the control branch, the target-generating estimate shrinks toward the physiological prior as data weakens: $\text{TDEE_used} = w \cdot \text{TDEE_apparent} + (1 - w) \cdot \text{TDEE_predicted}$, where the confidence weight w is a product of clamped monotone data-sufficiency gates (logging coverage, weigh-in count, weigh-in span) — multiplicative because the estimate is only as good as its *weakest* evidence stream. Below a confidence gate the target freezes entirely.

The intended deficit derives from a *rate* intent (fraction of body weight per week, clamped into a safe band on both write and read), applied as an **absolute offset in logging units** (§4.4). A layered guardrail pipeline then applies, in order: a physiological **safety floor** (from the predicted estimate only — and stricter for the autonomous loop than the same system permits a human operator: *a loop that lowers targets on its own gets less rope than a human does*); a **deficit ceiling** in physiological units; an **activity-first switch** (when the achievable deficit falls below a meaningful threshold, the plan's lever changes from calories to movement, and deficit/goal-date

language is suppressed — a sub-noise deficit is not a lever worth burdening the user with); an **asymmetric rate limiter** (downward moves capped tighter than upward: lowering is the risky direction); and **floor rescue** (correction toward the floor bypasses the rate limiter — a safety rule outranks a comfort rule).

Every user-facing promise (expected pace, projected completion date) derives from the *settled* target, never the rate-limited transient. Deriving a promise from a mid-walk transient once produced an arithmetically-true pace of 0.03 kg/week and a goal date years out — the single most churn-inducing statement such a system can make.

4.4 THE FIXED-POINT PROPERTY: BIAS INVARIANCE OF THE ABSOLUTE OFFSET

Model the user's stable reporting bias as $k > 0$ (logged intake $L = k \cdot I$ for true intake I). Weight is measured directly, so the storage rate is unbiased: $\text{stored} \approx I - E_{\text{true}}$. Then

$$\text{TDEE}_{\text{apparent}} = \text{mean}(L) - \text{stored} = k \cdot I - (I - E_{\text{true}}) = E_{\text{true}} - (1 - k) \cdot I.$$

With the absolute-offset law $T = \text{TDEE}_{\text{apparent}} - D$ and a user who complies *in logging units* (logs $L' = T$, truly eats $I' = T/k$), the loop re-estimates each period. At the fixed point the loop reproduces itself: $k \cdot I^* = E_{\text{true}} - (1 - k) \cdot I^* - D$, hence

$I^* = E_{\text{true}} - D$: the steady-state true deficit equals the intended deficit exactly, for every $k > 0$.

The estimation error (food not logged) and the execution error (food eaten beyond the logged target) are the same quantity with opposite signs; the absolute offset cancels them at the fixed point. Two comparison laws lack the property. An **open-loop physiological target** compared against biased logs leaves a $k = 0.7$ user “compliant” while in true surplus ($\approx +370$ kcal/day for $E = 2,200$, $D = 400$) — the user gains weight while obeying, and a single-response loop then ratchets the target toward the floor while the actual problem, measurement, is never addressed. A **multiplicative law** $T = (1 - d) \cdot \text{TDEE}_{\text{apparent}}$ under-delivers the true deficit whenever $k < 1$ ($k = 0.7$, $d = 0.2$ delivers 14.9% against 20% intended).

The property assumes k stable within the window; drift is handled by window lengthening and confidence reduction, and residual error is measured (§4.6). The safety floor is unaffected: it binds in physiological units regardless of k , so bias invariance can never drive a user below physiological limits.

4.5 SPEAKING RARELY, AND STRUCTURAL NON-DISCLOSURE

The loop runs weekly and therefore must speak rarely (target: most runs silent). Message selection follows a strict priority order: an open *adherence gap* outranks any target change (a user logging far above their standing target does not have a target problem — and this signal needs zero weigh-ins, so it is available from day 3); branch messages (a coverage calibration invitation must not mention any new target); silence on no-change and on sub-threshold change — which deliberately outranks floor announcements (an 11-kcal re-derivation is not news even at the floor); a newly-binding floor is announce-and-confirm, stating the honest achievable pace and promising nothing faster; larger moves tier from a digest line up to a standalone announcement that must give a plain-language cause, never a formula.

Finally, non-disclosure is structural: the controller emits a restricted payload enumerating the *only* numeric values the generative layer may surface. `TDEE_apparent`, the blend, $\hat{k}$, and confidence internals are absent from the payload — excluded by construction, not by instruction — because for a heavily under-logging user the apparent figure can read on the order of 900 kcal/day, and presenting that as “your metabolism” would be both false and harmful. Alongside the payload, a branch-dependent *coaching-stance directive* is idempotently injected into the agent’s context (e.g., during a coverage week: “the target is intentionally frozen; coach the log, not the intake; do not suggest eating less”).

4.6 FALSIFIABILITY

Every run records the predicted weight-change rate; a later run backfills the realized rate and the resulting calibration error. Median absolute calibration error is the controller’s single core accountability metric — without it the loop would be exactly as unfalsifiable as the open-loop formula it replaces. One measurement lesson: the *one-week* calibration error is dominated by scale noise (production medians 0.78 kg/wk for the loop’s own estimate and 0.71 for the formula baseline) and must not be used as an acceptance metric; the 28-day forward horizon is where signal emerges (§6.2).

5. Deployment: Engineering Constraints of a Messaging-Native Agent Fleet

Latency. Coaching over SMS sets a hard conversational envelope; we operate to a $p90 \leq 30$ s reply target. The two dominant costs we found: first-turn initialization of a long-lived agent (plug-in and import overhead costing tens of seconds before any model token is generated) and queue-versus-generation ambiguity in naive latency measurement (measuring from dequeue systematically understates user-experienced latency).

Delivery is not a given. Carrier messaging silently drops a nontrivial fraction of inbound traffic (we have measured losses of several percent); outbound campaigns face carrier-level filtering. Every proactive pipeline therefore needs delivery verification and idempotent retry, and every “user went silent” analysis must first rule out transport loss.

Safety filtering at the outbound edge. With an LLM composing messages to real phones, the outbound filter (§3.4) is load-bearing infrastructure: it has to catch internal monologue, controller internals, and stance leakage across every proactive and reactive path, including cron-generated messages that never passed through a live conversation.

Cost. Long-lived agents accumulate context; per-turn cost is dominated by context assembly rather than generation. Consolidation (§3.2) and context-budget discipline are what keep a months-old user no more expensive than a week-old one.

Human oversight is batched, not in-line. Every apply-mode sweep of the controller is reviewed by an operator before persistence; the reject branch routes to a human queue. The system’s autonomy is bounded by construction: interlocks before estimation, guardrails after, humans above.

6. Field Observations (Seven Weeks of Closed-Loop Operation)

Caveats first. Self-selected users of a paid consumer service; no control group; observation windows differ per user; most cumulative weight loss in the fleet predates the loop’s activation and is not attributable to it. Counts below are from the production fleet snapshot of mid-September 2026.

6.1 THE MEASUREMENT GAP IS LARGE AND ONE-SIDED

Across users with at least one closed-loop control window ($n = 29$), measured expenditure ran below the consumer-calculator recommendation by a median ≈ 511 kcal/day (22 of 24 below) and below the plan’s textbook-formula prediction by a median ≈ 284 kcal/day (24 of 29 below; 13 below by 300+). Framed as pace: for the median user the formula “promises” ≈ 1.76 lb/week at the prescribed intake while the measured reality delivers ≈ 1.23 — the formula over-credits roughly half a pound per week of expected progress. (An earlier public post of ours, from an August snapshot over a broader population with a coarser logs-implied estimator, reported a median formula-vs-logs gap of 564 kcal/day; the control-window figures here are consistent in direction and magnitude, with the narrowing attributable to the restriction to control-eligible windows and the dual-unit estimator.)

6.2 THE LOOP PREDICTS FORWARD BETTER THAN THE FORMULA

On a 28-day forward test (predict weight-change rate at run time; compare against realized rate from subsequent weigh-ins; 57 windows / 19 users), the loop's estimate nearly eliminates the formula's bias (≈ -0.02 vs -0.15 kg/week, the latter equivalent to over-estimating expenditure by ≈ 160 kcal/day) and improves median absolute error (≈ 0.26 vs 0.36 kg/week; per-person first-window 0.20 vs 0.25 ; loop better in 11 of 19).

6.3 GUARDRAILS DO THE QUIET WORK

Over seven weeks: 65 target movements landed (35 down / 30 up; median $|\Delta| = 100$ kcal) — the loop raises targets nearly as often as it lowers them, which no ratchet-down controller can say. Floor rescues lifted a handful of users off unsafely-low inherited targets (largest single correction $\approx +600$ kcal/day); sub-floor detection reclassified roughly **half** of the recently-active fleet as “not a calorie problem” (19 of 36: 12 genuinely eating below their safety floor, whose only admissible move is upward, and 7 whose logs were physically impossible and were routed to measurement repair instead of a target change).

6.4 THE BEHAVIORAL FINDING: OVER-ESTIMATION'S REAL HARM IS WHAT THE AI SAYS

The most consequential observation was not about weight. Scanning seven weeks of outbound coaching for “you still have room to eat” language (3,345 messages), about one in five such nudges (725; 21%, concentrated in 10 users) was sent while the user's standing target exceeded their *measured* expenditure by 150+ kcal/day — the model, arithmetically honest against a wrong number, was confidently urging users to eat more. Closed-loop correction shrank the population in that state from 9 users at registration-target baseline to 6, all of whom sit in frozen (low-confidence) branches where the message policy suppresses intake encouragement. In an LLM-mediated coach, an over-estimated target does not merely slow progress; it *actively generates confident bad advice*. This, more than any pace argument, is why the target must be measured rather than assumed.

User-visible reception of target changes was calm: $\sim 90\%$ of users receiving a target-change message continued logging normally within 24 h (192 target-change messages across 39 users); explicit reactions ranged from delight (“Wow... 1925!”) to legitimate challenge (“Why did my calories increase?”) that the stance-carrying agent could answer honestly.

7. Ethics, Safety, and Limitations

The system is a coaching product, not a medical device; it does not diagnose, prescribe, or manage medication. Interlocks exclude minors and users with eating-disorder history flags from the loop entirely; medical directives (e.g., clinician-set refeeding protocols) freeze it; GLP-1 users' dosing remains their clinician's domain, with the loop confined to intake/protein coaching. The floor and activity-first switch exist precisely because an autonomous restriction loop is the dangerous kind of automation; its rope is deliberately shorter than a human operator's. All statistics above are de-identified aggregates; verbatim user quotes appear only under the service's existing written-consent process. The author is the system's designer and a co-founder of the company operating it; this interest is disclosed, and the observational framing (no efficacy claim, no control group, self-selected cohort) should be read accordingly. The user base is currently US-based and English-speaking; generalization is untested.

8. Conclusion and Future Work

Static calorie targets fail in a specific, mechanical way: the prediction is wrong by hundreds of kcal/day for many individuals, the measurement channel that could correct it is biased, and a naive correction loop punishes the wrong failure mode. The deficit loop's answer — two estimates, two unit systems, a router between “fix the target” and “fix the measurement,” an absolute-offset law with provable bias invariance, and physiological guardrails with structural non-disclosure — has now run for seven weeks in production without a safety incident, raising targets almost as often as it lowers them, and predicting forward better than the formula it replaced.

Planned next steps: a controlled evaluation once the cohort supports it ($N \geq 100$ with ≥ 8 -week follow-up; retrospective single-arm against literature baselines, then prospective comparison), and an intake-side extension (the current loop closes on expenditure; the adherence signal suggests the next lever is how targets are *budgeted* across the week).

Acknowledgments and disclosure

The system described is operated by NanoRhino (Link Heart Limited), of which the author is co-founder and CTO; this interest is disclosed in §7. The deficit-loop coaching engine is U.S. patent-pending. Large language models were used as writing and analysis aids in preparing this manuscript; all system descriptions, measurements, and claims were authored and verified by the author.

References

- [1] Mifflin MD, St Jeor ST, Hill LA, Scott BJ, Daugherty SA, Koh YO. A new predictive equation for resting energy expenditure in healthy individuals. *Am J Clin Nutr*. 1990;51(2):241–247.
- [2] Lichtman SW, Pisarska K, Berman ER, et al. Discrepancy between self-reported and actual caloric intake and exercise in obese subjects. *N Engl J Med*. 1992;327(27):1893–1898.
- [3] Trabulsi J, Schoeller DA. Evaluation of dietary assessment instruments against doubly labeled water, a biomarker of habitual energy intake. *Am J Physiol Endocrinol Metab*. 2001;281(5):E891–E899.
- [4] Burrows TL, Ho YY, Rollo ME, Collins CE. Validity of dietary assessment methods when compared to the method of doubly labeled water: a systematic review in adults. *Front Endocrinol*. 2019;10:850.
- [5] Hall KD, Sacks G, Chandramohan D, et al. Quantification of the effect of energy imbalance on bodyweight. *Lancet*. 2011;378(9793):826–837.
- [6] Sanghvi A, Redman LM, Martin CK, Ravussin E, Hall KD. Validation of an inexpensive and accurate mathematical method to measure long-term changes in free-living energy intake. *Am J Clin Nutr*. 2015;102(2):353–358.
- [7] Bajunaid R, Niu C, Hambly C, et al. Predictive equation derived from 6,497 doubly labelled water measurements enables the detection of erroneous self-reported energy intake. *Nat Food*. 2025;6(1):58–71.
- [8] Breit MJ, Pan Z, Ostendorf DM, et al. Using time-weighted averages of total daily energy expenditure to estimate energy intake during a weight loss intervention. *Obesity (Silver Spring)*. 2025;33(11):2093–2102.
- [9] Nahum-Shani I, Smith SN, Spring BJ, et al. Just-in-time adaptive interventions (JITAs) in mobile health: key components and design principles for ongoing health behavior support. *Ann Behav Med*. 2018;52(6):446–462.
- [10] Nahum-Shani I, Murphy SA. Just-in-time adaptive interventions: where are we now and what is next? *Annu Rev Psychol*. 2026;77:679–703.
- [11] Rivera DE, Pew MD, Collins LM. Using engineering control principles to inform the design of adaptive interventions: a conceptual introduction. *Drug Alcohol Depend*. 2007;88(Suppl 2):S31–S40.

- [12] Trella AL, Zhang KW, Jajal H, et al. A deployed online reinforcement learning algorithm in an oral health clinical trial. *Proc AAAI Conf Artif Intell.* 2025;39(28):28792–28800.
- [13] Karine K, Klasnja P, Murphy SA, Marlin BM. Assessing the impact of context inference error and partial observability on RL methods for just-in-time adaptive interventions. *Proc Conf Uncertainty in Artificial Intelligence (UAI)*, PMLR 216, 2023. arXiv:2305.09913.
- [14] Martin CK, Miller AC, Thomas DM, Champagne CM, Han H, Church T. Efficacy of SmartLoss, a smartphone-based weight loss intervention: results from a randomized controlled trial. *Obesity (Silver Spring)*. 2015;23(5):935–942.
- [15] Huang Z, Berry MP, Chwyl C, Hsieh G, Wei J, Forman EM. Comparing large language model AI and human-generated coaching messages for behavioral weight loss. *J Technol Behav Sci.* 2025. doi:10.1007/s41347-025-00491-5. arXiv:2312.04059.
- [16] Jörke M, Sapkota S, Warkenthien L, et al. GPTCoach: towards LLM-based physical activity coaching. *Proc CHI Conf Human Factors in Computing Systems (CHI '25)*, 2025. doi:10.1145/3706598.3713819. arXiv:2405.06061.
- [17] Jörke M, Genç D, Teutschbein V, et al. Bloom: designing for LLM-augmented behavior change interactions. *Proc CHI Conf Human Factors in Computing Systems (CHI '26)*, 2026. doi:10.1145/3772318.3790506. arXiv:2510.05449.
- [18] Yang E, Garcia T, Williams HG, et al. A behavioral science-informed agentic workflow for personalized nutrition coaching: development and validation study. *JMIR Form Res.* 2025;9:e75421.
- [19] Li KJ-H, Balloccu S, Dušek O, Reiter E. When LLMs can't help: real-world evaluation of LLMs in nutrition. *Proc 18th International Natural Language Generation Conference (INLG)*, 2025:753–779. arXiv:2511.20652.
- [20] Singh B, Olds T, Brinsley J, et al. Systematic review and meta-analysis of the effectiveness of chatbots on lifestyle behaviours. *npj Digit Med.* 2023;6:118.
- [21] Noh E, Won J, Jo S, Hahm DH, Lee H. Conversational agents for body weight management: systematic review. *J Med Internet Res.* 2023;25:e42238.
- [22] Omri Y, Gan Z, Broveak Z, et al. Agent memory: characterization and system implications of stateful long-horizon workloads. arXiv:2606.06448, 2026.
- [23] Du P. Memory for autonomous LLM agents: mechanisms, evaluation, and emerging frontiers. arXiv:2603.07670, 2026.

[24] Wishnofsky M. Caloric equivalents of gained or lost weight. *Am J Clin Nutr.* 1958;6(5):542–546.

[25] Hall KD. What is the required energy deficit per unit weight loss? *Int J Obes.* 2008;32(3):573–576.