

Supporting The Many Lives of Personal Data with Rebite: LLM-Powered Goal-Directed Framing in Food Journaling

Weijun Li
weijunl3@uci.edu
University of California, Irvine
Irvine, California, USA

Daniel A. Epstein
epstein@ics.uci.edu
University of California, Irvine
Irvine, California, USA

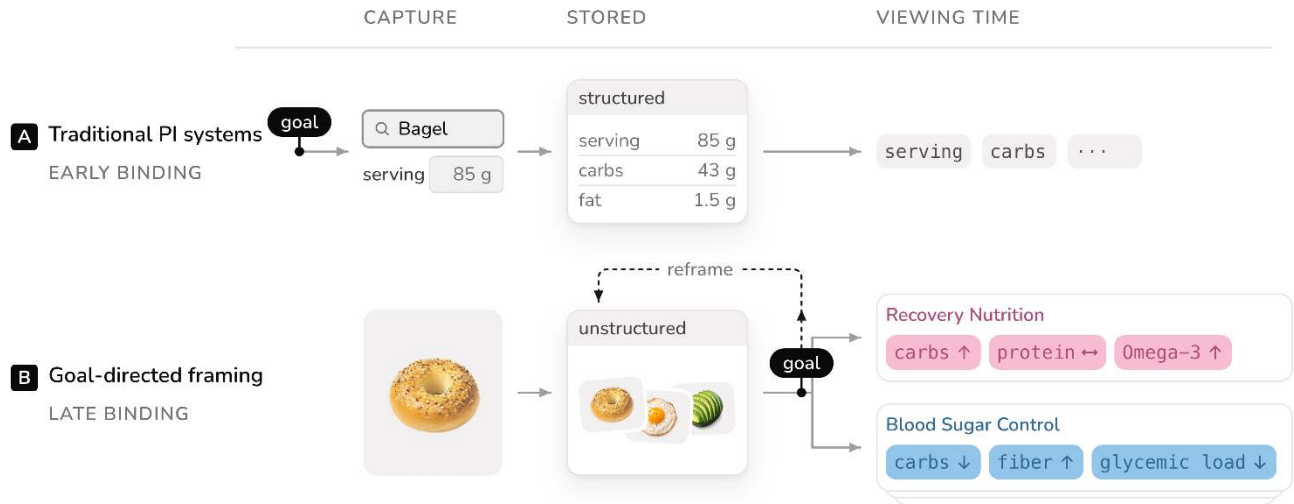


Figure 1: (A) Traditional Personal Informatics systems fix the metrics when data are logged, so a goal adopted later relies on the same structured data. (B) Our approach of Goal-directed framing keeps data as unstructured and applies the goal at viewing, so one record can be interpreted and reframed to support more than one goal.

Abstract

People’s health and tracking goals frequently change, but most personal informatics systems struggle to adapt, leading people to abandon their data and start over. We propose *goal-directed framing*, an approach that repositions goals within personal informatics systems. Instead of fixing the meaning of data at capture time, the approach frames the collected data through the current goal and reframes it whenever the goal changes. We realize this in *Rebite*, a photo-based food journaling system that uses LLMs to read unstructured meal photos and produce goal-directed feedback. In a one-week deployment with 21 participants managing multiple dietary goals, we find that goal-directed framing shaped how participants engaged with their goals. Translating a goal into metrics helped them see what it meant in practice, confirming existing priorities, surfacing what they overlooked, and revealing where the metrics fell short. When goals changed, seeing past meals reframed under the new goal exposed overlaps and conflicts, prompting participants to negotiate trade-offs and refine priorities. We discuss

how goal-directed framing both supports and complicates reflection as goals change, and offer design implications for personal informatics systems to support evolving goals.

CCS Concepts

• **Human-centered computing** → **Interactive systems and tools**; **Interaction design**.

Keywords

Personal informatics, Large Language models, Food tracking, Goal

ACM Reference Format:

Weijun Li and Daniel A. Epstein. 2026. Supporting The Many Lives of Personal Data with Rebite: LLM-Powered Goal-Directed Framing in Food Journaling. In *The 39th Annual ACM Symposium on User Interface Software and Technology (UIST '26)*, November 02–05, 2026, Detroit, MI, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3830398.3830692>

1 Introduction

Personal informatics research has long argued that aligning the data people collect with their goals supports reflection and action [23, 42, 60]. Yet, goals change. People revise, refine, expand, or abandon their goals as routines, health needs, and priorities evolve [23, 26, 27, 61]. Goals may be translated from qualitative aspirations into concrete, quantitative targets [54], and may also multiply, evolve, and shift in priority over time [1]. When the system’s metrics no

longer match the user’s evolved goal, people often suspend tracking or switch to new tools, leaving past data behind [14, 26, 27, 45]. How to build PI systems that adapt when goals change remains largely unexplored.

Why do current PI systems fail to support goal change? Goals evolve in at least two ways. They may gradually become clearer as users learn what matters [1, 54], or shift as priorities change [20, 27]. In both cases, the interpretation of existing data should also change, including which metrics are relevant, at what level of detail, and how to present them for reflection. Yet, most current systems do not fully adapt. They either hard-code which metrics to present [51] or allow users to configure tracking only prospectively [37, 61]. As a result, the meaning of collected data remains tied to decisions made at logging time. When goals change, users often cannot reinterpret their existing data through the new goal because the information needed for that interpretation was never preserved or made available.

We propose **goal-directed framing**, an approach that repositions goals within personal informatics systems where goals shape interpretation at the time of viewing data. The approach preserves rich, unstructured data (such as photos) and, when the user states a goal, collaboratively translates it into relevant metrics and structures feedback through that goal’s lens. When the goal changes, the system *reframes*, by reinterpreting all past data under the new goal, so the same data acquires new meaning without re-logging. We use large language models (LLMs) to realize this, as they can reason over unstructured inputs and generate goal-directed feedback [62].

We realize goal-directed framing in **Rebite**, a photo-based food journaling system (Figure 2). Food journaling is a useful testbed for this approach because people pursue a range of dietary goals that change over time [20, 29], from weight loss [16] and managing chronic conditions like diabetes [50] to identifying food triggers [35] and eating more mindfully [25]. Meal photos carry rich nutritional detail across many dimensions (e.g., calories, macronutrients), and interpreting it under a specific goal can require expertise that most users do not have [15, 49]. Structured support that connects meal photos to the user’s current goal is therefore important for meaningful reflection.

To understand how people experience goal-directed framing in practice, we deployed Rebite with 21 participants managing multiple dietary goals in a one-week study with a mid-study goal switch. We find that the goal-to-metric mapping helps participants make a vague goal concrete, confirming what they already care about, surfacing what they overlook, and revealing where the metric space falls short. When goals change, reframing past meals under the new goal exposes overlaps and conflicts between goals, prompting participants to negotiate trade-offs and reconsider what their goals mean in practice. We discuss how goal-to-metric mapping enables reframing, how reframing brings both benefits and risks, and what it means for personal data to stay meaningful as goals change. We make three contributions.

- *An approach for PI systems with evolving goals.* We propose goal-directed framing, which repositions the goal within personal informatics systems by having the system take on the role of translating the current goal into relevant metrics,

framing preserved data through that goal’s lens, and reframing it when the goal changes, enabling personal data to “live many lives.”

- *Design and implementation of Rebite.* A photo-based food journaling system that realizes goal-directed framing, letting the same collected data be reflected upon under evolving goals without new data entry.
- *Empirical evidence from a deployment study ($N = 21$).* We illustrate how reframing past data under new goals can make cross-goal relationships visible, including overlaps that validated past efforts, conflicts that prompted trade-off negotiation, and breakdowns that revealed when users preferred to start fresh instead of revisiting.

2 Background and Related Work

To motivate our approach, we first review *goals in personal informatics* and examine *when interpretations become fixed in existing systems*. We then discuss *LLMs for personal data* as creating new opportunities for flexible interpretation, and finally situate our work within the context of *food journaling*.

2.1 Goals in Personal Informatics

Goals are central to what and how people track. Prior work shows that goals guide what people record and how they reflect on their data [42, 43, 60]. In lived practice, however, goals evolve with everyday routines, life transitions, and health needs [23, 26, 27, 61], and people revise, expand, or abandon them accordingly.

People progressively sharpen goals from abstract needs to quantifiable targets [54], select among multiple problems, simplify them into specific goals, and iteratively readjust [1], and refine goals through reflection [41]. The direction and pace of change are shaped by personal, contextual, and social factors [20, 58]. A recent scoping review argues that PI systems should be designed to support goal adjustment over time [21].

Prior systems have pursued alignment between goals and data through flexible configuration of tracking fields [37], supporting selection of apps that match tracking needs [40], and presenting data in ways that reflect stated goals [28, 61]. Specifically, Schroeder et al. [60] proposed goal-directed self-tracking as a framework in which designs centered around an individual’s goals support tracking exactly and only the data that individual needs. Migraine-Tracker [61] provides empirical evidence for this approach, aligning data collection and presentation with clinical goals to support reflection and decision making. However, none of these systems reinterpret previously collected data when goals evolve. A growing body of work has called for PI systems that support goal adjustment [20, 21, 26, 27], but reinterpreting past data under a new goal has so far been done by hand where it has been done at all, and how it shapes people’s engagement with their goals has not been studied. We address this gap by framing already-collected data through the goals a person already holds and moves between, complementing prior work on forming and refining goals [21, 41, 54, 60].

2.2 When Does Interpretation Get Locked In?

Current PI systems bind interpretation early, in one of two ways, and as a result struggle to reinterpret past data when goals change.

Many systems hard-code a single frame at design time. Calorie-focused food apps [70] foreground energy budgets, and photo-based nutritional analysis [55] extracts a fixed metric set when data enters. Moore et al. [51] call this the “personal informatics analysis gap”, where tools “baking-in” analysis workflows, leaving users unable to ask questions the designers did not anticipate. Other systems give users control over what to track, but apply it only going forward. OmniTrack [37] lets users compose custom fields, MigraineTracker [61] lets patients select and adjust tracking routines, yet reconfiguration reaches only future data, and previously logged entries are not re-assessed. When a goal changes and the collected data no longer supports that goal, people often abandon the tool or start over [24, 26]. A growing line of food journaling work preserves richer, open-ended data to support reflection [48, 64], but does not reinterpret it when goals change.

A nascent body of work points the other way, suggesting that reinterpretation is both desired and possible. Ng et al. [53] found that people bound periods of disruption in their data, such as a pregnancy, and interpret data inside and outside those boundaries differently, and Jung et al. [34] showed that letting users manipulate their own tracked values can trigger new reflection. Design work has begun exploring open-ended tracking that resists fixing meaning at design time [6, 69].

The goal has long been seen as central to what and how people track [42, 60], and prior systems let it shape what is collected and how it is shown. What no system does is re-apply it to data already in hand. The goal is bound once, prospectively, and never turned back on the record. We call this missing capability *goal-directed framing*, in which the system translates a goal into relevant metrics and frames collected data through that goal’s lens, then *reframes* it whenever the goal changes, so past data stays relevant instead of locked in its original frame.

This capability has overlap with strategies in interactive visualization, where user actions can transform views on data. For example, dynamic queries that update a visualization when the user adjusts a parameter [2], and reactive visualization grammars that propagate a signal change through a dataflow graph to every dependent mark [59]. In this way, the goal acts as a parameter, yielding different interpretations of an existing dataset. However, the parallel falls short in that visualization typically assumes the dataset is already captured, and is being interactively analyzed. In personal informatics, since goals emerge and change and future goals are hard to anticipate, it is instead unclear at collection time what to record and which dimensions will later matter. People also revisit and reanalyze their data in different ways than traditional visualization. PI practice favors smaller, more frequent reflection situated in everyday life, not a sit-down session with a dashboard [8]. Supporting this kind of everyday reflection as goals change is thus a challenge PI needs to explore, and one our work takes up.

2.3 LLMs and Personal Informatics

Recent work shows that LLMs can enhance personal informatics along two dimensions. For interpretability, LLMs have turned raw metrics into natural-language narratives to support reflection [66], elicited experiential detail like emotions to complement numeric tracking [36], and paired narrative summaries with visual timelines

to help experts question and validate AI inferences [44]. For personalization, LLMs have translated goals into actionable IF-THEN plans [62], generated context-sensitive journaling prompts from behavioral sensing data [52], grounded activity coaching in wearable data [33], and let users choose among generated prompts and request summaries on demand [65].

These systems show that LLMs can personalize feedback and generate interpretive narratives around tracked data. Yet LLMs alone are not sufficient. Chopra et al. [11] found that people querying generative AI about their health data had to keep scaffolding and refining their prompts to get goal-relevant answers, and abandoned queries when answers stayed too generic. Goal alignment therefore belongs in the system, not in the user’s prompting. Prior work also suggests how to structure the output. Sensecape [67] shows that organizing LLM output into levels of abstraction helps users navigate complex information, and Graphologue [32] shows that turning responses into interactive diagrams makes them more interpretable than plain text. This informs how Rebite presents goal-directed feedback in layers.

2.4 Food Journaling

Food journaling is a widely studied domain in personal informatics [23], and the focus of our implementation of goal-directed framing. People use food journaling to pursue a wide range of goals, including weight loss [16], managing chronic diseases [13, 50], and identifying intolerances [35]. The practice also serves qualitative purposes, such as reflecting on the sensory and contextual aspects of meals [72] and building awareness through lightweight, mindful engagement [25], and social purposes such as self-presentation [12], social learning [45], and family support [47].

Researchers have explored multiple input modalities for food journaling. Calorie logs have been the traditional approach, leveraging database lookups of text descriptions [70] or barcodes. Photo-based journaling has been suggested as an alternative to balance entry difficulty [15], with a long line of research seeking to derive nutrient estimates from meal images, whether by crowdsourcing [55] or by automated recognition [3, 68]. Other strategies include conversational journaling [48, 64] and multi-device routines [63]. Each modality carries trade-offs. Text entry is tedious and often incomplete [5, 16], image-based approaches reduce effort but struggle with portion estimation [3, 68], and semi-automated approaches raise privacy concerns [46].

Food tracking goals frequently change. Life transitions such as moving or changing jobs reshape food routines and priorities [29]. What people track also differs across goals. An IBS patient follows symptom–food relationships [13, 49], someone managing type 2 diabetes needs personalized targets around carbohydrate intake [50], and someone pursuing weight loss focuses on calories and macronutrients [16]. Luo et al. [49] showed that different conditions require distinct tracking items and representations. Such customization, however, is typically set up once and does not support transitions between evolving goals. Cordeiro et al. [15] showed a decade ago that the value of photo capture varies with the journaler’s goal, and argued that journals should support a variety of goals. What is missing is not the idea but a system that performs the reinterpretation, and evidence for what happens when it does.

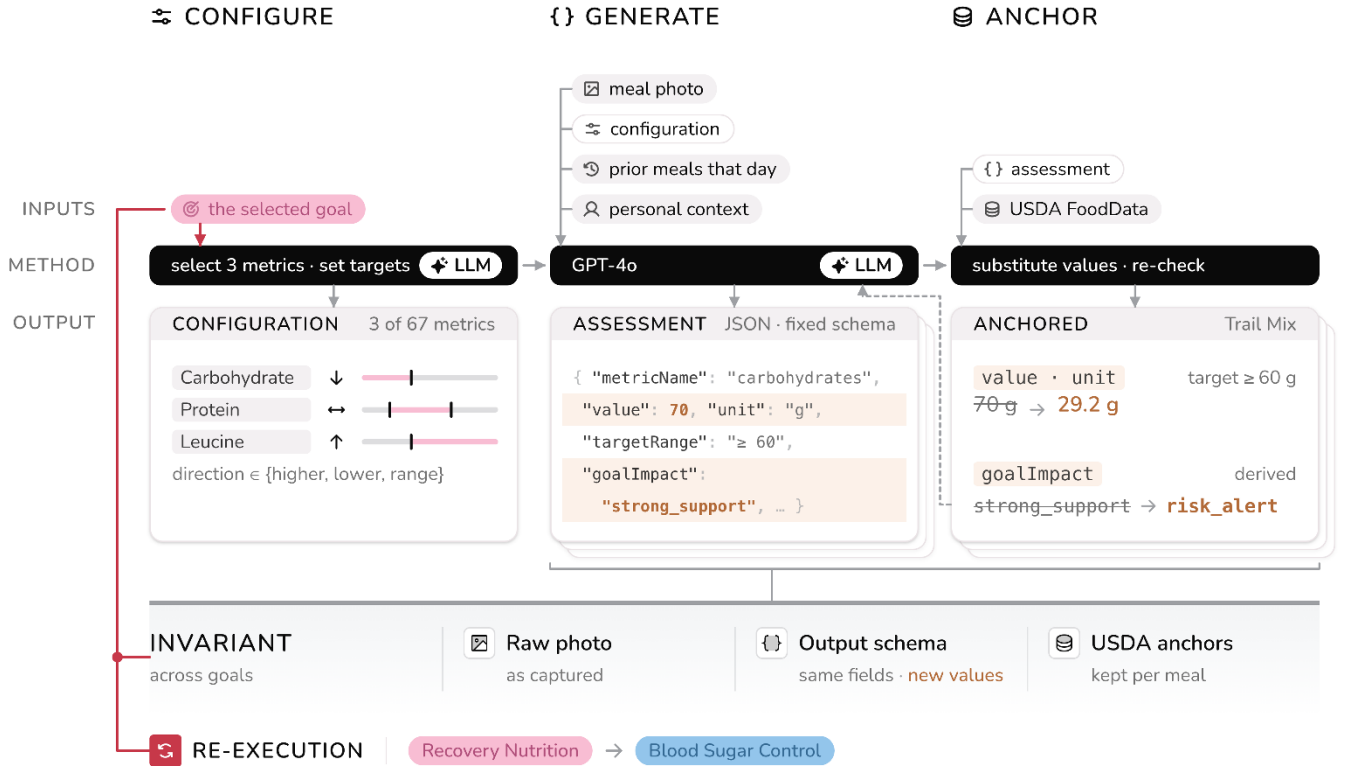


Figure 2: The interpretation pipeline that enables goal-directed framing. Based on a selected goal, the LLM creates a configuration of target metrics. When new data (meal photos) are logged, this configuration is passed alongside context to the LLM to generate a goal-specific assessment. This is then anchored against a food database. This pipeline can be re-executed on existing logs when a person's goal changes.

3 Goal-Directed Framing: An Approach for Personal Informatics Systems with Evolving Goals

As articulated in Section 2.2, current PI systems lock interpretation because they bind meaning at design or setup time. To address this, we propose goal-directed framing. Instead of users conforming to the system's metrics, *systems can serve the user's evolving intention*. Figure 1 contrasts this with current PI systems.

Realizing this model poses three design challenges, each derived from limitations of the strategies in Section 2.2.

D1. Translate goals into trackable configurations. Goals range from vague aspirations to specific targets [11, 54], so the system can propose metrics for a stated goal and let users refine them as their intent sharpens.

D2. Reframe past data when goals change. Current systems interpret data once and apply new configurations only to future entries [37, 61], so a goal change has to send the system back over everything already stored. Reinterpretation only means something if the underlying facts hold still. If the recorded values themselves shift with the goal, a user cannot tell a new reading of the same meal from a new account of it.

D3. Bridge lightweight input to structured, goal-directed feedback. Users should not have to specify metrics or label data. The system extracts goal-directed structure from unstructured input and

presents it so that the goal connects to concrete assessments and evidence.

Because LLMs can reason over unstructured data and generate structured interpretations on demand, we propose an LLM-powered approach to address these challenges. The user logs unstructured inputs and states a goal, and the system does the mapping, interpretation and re-interpretation. Figure 2 shows this as we implement it in Rebite.

Goal-directed framing could extend to domains where rich raw data can be preserved and reinterpreted as goals change. Exercise, sleep, and mental health are candidates, where sensor streams carry enough signal to be read through different goals. We instantiate and evaluate it in food journaling for the reasons given in Section 1.

4 The Rebite System

Rebite is our photo-based food journaling system that realizes goal-directed framing. Users log meals through photos and state a goal in everyday language. The system then translates the goal into trackable metrics (D1), interprets each meal against those metrics and re-interprets all stored meals when goals change (D2), and presents structured, four-layer feedback that moves from goal context to concrete food evidence (D3).

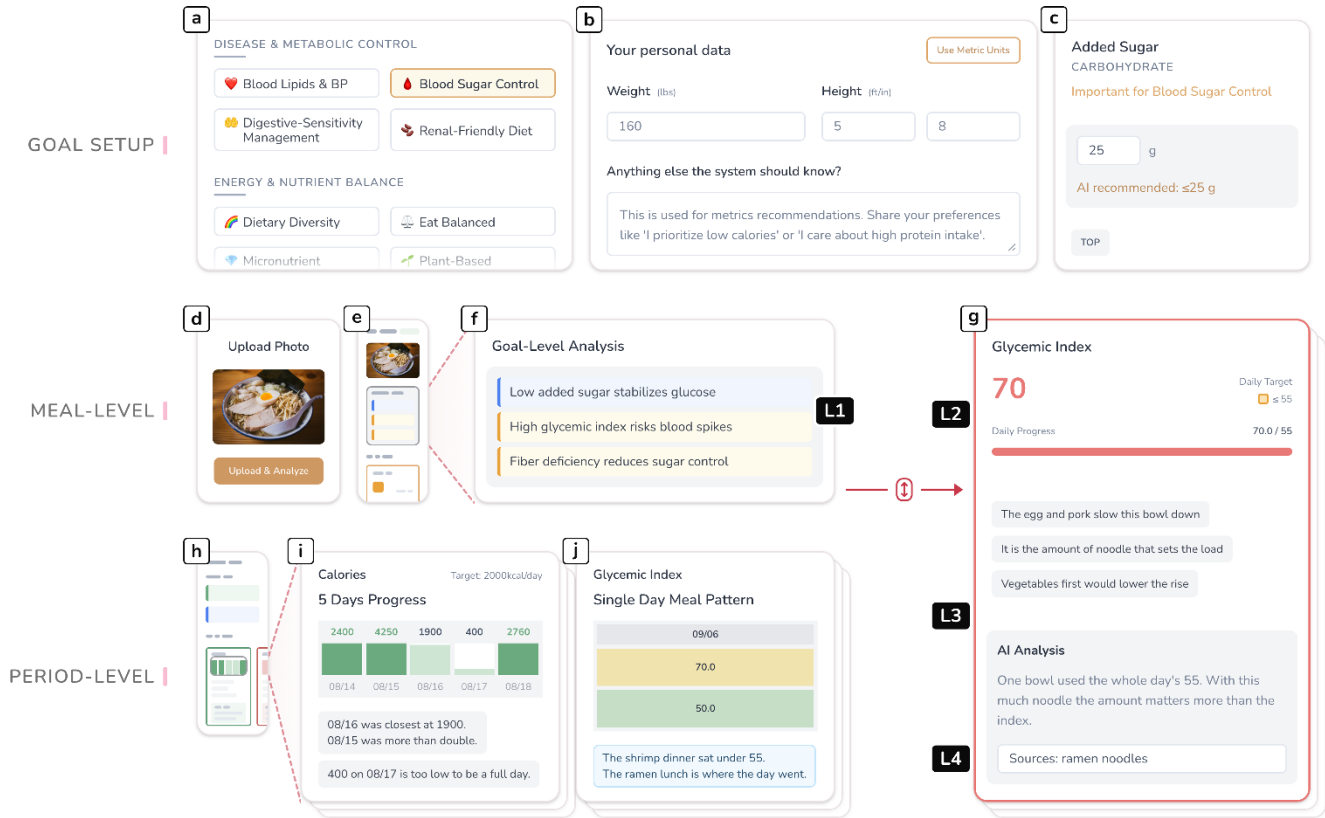


Figure 3: Goal-directed framing instantiated in Rebite. To set up their goal, a person (a) selects from options and (b) provides personal context, resulting in an (c) AI-recommended configuration of metrics. When (d) new meal photos are logged, (e-f) they are summarized at the meal level with first layer L1 analysis of how they support the goal. (g) A person can scroll down for a metric card with L2 comparing progress against their target, L3 AI interpretation, and L4 specific food evidence. (h) Data are summarized at a period-level, with visualizations for metrics which should be (i) increased or decreased or (j) should stay within a range. When goals change, each interpretation layer reframes existing data while interaction remains the same.

4.1 Goal Setup

Users begin by selecting a goal in everyday language (e.g., “lose weight,” “recover from training”) and optionally providing personal context such as height, weight, or health conditions (Figure 3a–b). The system then translates these into a concrete configuration of which metrics to track, in which direction (increase/decrease), and with what thresholds. Users can inspect the resulting configuration (Figure 3c) and adjust metrics or targets as their understanding deepens.

This translation is supported by a goal–metric framework that we compiled from published nutrition and public-health guidance (full mappings are provided in the supplementary material). The framework pairs 67 nutritional metrics with 19 goals, assigning each goal three priority metrics and several supporting ones, together with a direction and a starting threshold. For example, *Lose Weight* prioritizes calories, protein, and fiber, whereas *Blood Sugar Control* prioritizes carbohydrate dose, fiber, and glycemic load. The mappings draw on published nutrition guidelines and supporting evidence from the literature [4, 22, 31, 57, 71]. Switching goals automatically produces a new configuration that reframes the same food records.

The framework provides a default configuration for each goal. Users can add, remove, or retarget metrics to better reflect their own priorities. They can also maintain multiple goals and switch among them, revisiting the same food records through different goal-directed framings.

4.2 Goal-Directed Interpretation

Rebite does not bind a meal to a single interpretation when it is logged. Instead, the same meal can be framed through the user’s *current* goal configuration whenever feedback is requested.

Meal logging. Users capture meals by uploading a photo (Figure 3d). After reviewing the system’s analysis, users can refine the estimates by adding ingredient descriptions with quantities (e.g., “100 g pasta with 120 g ground beef”), correcting misidentified items, or adjusting portions. Rebite then interprets each meal through three stages, where the output of each feeds the next (Figure 2).

Configure. The system sends the user’s goal to the model together with the goal–metric framework (Section 4.1). The model selects three priority metrics from the framework’s library and personalizes each target’s direction and threshold (the supplementary material gives an abridged version of all three pipeline prompts),

producing a concrete configuration of which metrics to evaluate, in which direction, and with what thresholds. This configuration shapes all downstream interpretation.

Generate. The system constructs a prompt from the meal photo, the configuration from ‘Configure’, prior meals from the day, and the user’s personal context. This prompt is sent to a VLM, GPT-4o in our implementation, which returns a structured JSON response following a fixed output schema. The response includes food identification, nutritional estimates for all tracked metrics, per-metric assessments relative to the goal, and goal-level tags. Because the configuration determines which metrics the LLM evaluates, the same photo produces different assessments under different goals.

Anchor. The system looks up the foods the model identified in USDA FoodData Central and recomputes every tracked metric the database covers from its composition data and the model’s estimated weights. What stays estimated is which foods are present and how much of each. Metrics the database does not carry, such as glycemic load and FODMAP content, stay model estimates, so goals resting on them keep the estimation error anchoring removes elsewhere. If the correction changes a metric’s status relative to its target, the corrected values go back to the model, which re-generates that assessment.

When a user changes goals, ‘Configure’ produces a new configuration, and ‘Generate’ and ‘Anchor’ re-execute on every stored meal. Both calls conform to a fixed output schema regardless of the active goal, so reframing requires no schema migration.

4.3 Goal-Directed Feedback

The interface is designed around the principle that *the structure of feedback stays constant while its content reframes*. This gives users a stable visual language for interpreting goal-directed assessments, even as the underlying metrics, colors, and narratives change when goals change.

Meal-level feedback. Each meal’s feedback is organized as four layers that move from goal context to factual evidence (Figure 3e–g). **Goal-level tags** color-code each nutritional dimension by its impact on the current goal, separating dimensions that support it from those that carry risk or ask to be monitored (Figure 3f, L1). When the goal changes, the same meal’s tags change color, making reframing immediately visible. **Metric cards** display the anchored value against the target (Figure 3g, L2), with a breakdown showing how this meal sits within the day. A status indicator shows how the metric performs relative to the goal. **Goal-directed interpretations** explain what the metric value means for the user’s specific goal (Figure 3g, L3). **Food evidence** grounds each interpretation in specific ingredients from the meal (Figure 3g, L4).

Period-level feedback. At the period level (e.g., a week), additive metrics (e.g., calories) use bar charts scaled against the metric’s target (Figure 3i), while non-additive metrics (e.g., glycemic index) use heatmap grids (Figure 3j). Each claim lists the meals behind it and whether each one supports the goal, so a claim can be checked against the entries it came from. When goals change, the active metrics and visualization layout reframe (D3).

4.4 Implementation

Rebite extends the open-source *Clean Slate* calorie tracker (Apache License 2.0)¹ and is implemented as a web application using React and a GraphQL backend with a PostgreSQL database. The interpretation pipeline uses the OpenAI API to make calls to GPT-4o with JSON-constrained output. The pipeline can be adapted to other LLMs that support structured output. Before deploying Rebite with users, we ran a technical evaluation of how much reframing varies across goals, what factual anchoring contributes, and how much anchoring changes the values the model returns from a photo. The method and results are in the supplementary material.

5 User Study

We conducted a mixed-method deployment study to understand how people experience goal-directed framing in practice, complemented by a pre/post survey for quantitative reference. The study protocol was reviewed and approved by our university’s IRB, and all participants gave informed consent. We organize our investigation around two questions.

RQ1. *How do users make sense of goal-directed feedback?* When the system maps a vague goal into concrete metrics (D1) and interprets meals through that lens (D3), how do users experience the resulting feedback?

RQ2. *What happens when goals change and past data is reframed?* When the system re-maps and re-interprets stored data under a new goal (D2), what forms of reflection emerge, and where does the mechanism break down?

5.1 Method

Over one week, participants tracked meals under an initial goal for four days, then switched to a different goal mid-study. Because we were interested in how participants **perceived and made sense of** reframed feedback and not in measuring behavior change, we recruited people who already held multiple dietary goals and asked them to switch between existing priorities.

Figure 4 summarizes the schedule. Participants completed a pre-study survey about their prior tracking practices, then joined a 20-minute onboarding session where they set up Rebite on their phones and selected an initial goal (Goal A). Over seven days, they logged meals with photos and reviewed feedback daily. After four days, participants joined a mid-study interview (~40 minutes) where they reviewed feedback under Goal A, switched to a different goal (Goal B), and discussed reinterpretations of their logged meals under the new goal. At the end of the study, participants completed an after-use survey and a 60-minute reflection interview. The full procedure is in the supplementary material.

5.2 Participants

We recruited 25 adults in the United States who owned a smartphone, had previously tracked meals for at least two weeks, and were currently pursuing multiple diet-related goals. The multiple-goal requirement kept the switch on existing priorities and off invented ones [21, 26, 54], and the prior-tracking requirement kept the study on goal-directed framing and off the novelty of food

¹<https://github.com/successible/cleanslate>



Figure 4: Study timeline (N=21). Participants logged meals under Goal A, then switched to Goal B following the mid-study interview.

journaling. Because food tracking can be sensitive in the context of eating disorders, recruitment materials emphasized a range of dietary goals (e.g., muscle gain, anti-inflammatory eating, weight management). We recruited through flyers and diet-focused Discord communities.

Four participants withdrew during the study (citing usability issues and busyness), leaving $N=21$ (13F, 8M; ages 21–49). Each held two or three concurrent goals, 49 in all, spanning 14 of the framework’s 19 goals. Weight loss ($n=14$) and muscle gain ($n=10$) were the most common, followed by anti-inflammatory eating and recovery nutrition ($n=4$ each), balanced eating and blood sugar control ($n=3$ each), and plant-based adequacy, reduced fat intake and glycogen reload ($n=2$ each). Gut health, mental clarity, blood lipids and blood pressure, micronutrient optimization, and digestive sensitivity were each held by one participant. Participants received an \$80 Amazon gift card on completion. Per-participant demographics, prior tracking experience, and goals are tabulated in the supplementary material.

5.3 Analysis

5.3.1 Survey Analysis. We administered a pre/post 11-item survey of 6 of the 9 Technology-Supported Reflection Inventory (TSRI) items [7], covering its Exploration and Insight dimensions, plus 5 goal-related items (Q7–Q11). We excluded its third dimension, Comparison, whose items concern comparison with other people. We reworded the TSRI items for the food domain and administered them on a 5-point scale, so we report them as TSRI-derived measures and not as the validated instrument. Pre and post used the

same items, so the comparison is internal to that adaptation. Pre-study responses referred to participants’ prior tools and post-study responses to Rebite. We compared scores with paired two-tailed t -tests, Cohen’s d_z , and Holm–Bonferroni correction. The full instrument is in the supplementary material.

5.3.2 Qualitative Analysis. Transcripts from mid-study interviews, final interviews, and daily mini-surveys were coded in ATLAS.ti, combining deductive codes from the research questions and system features with inductive codes for emergent practices (e.g., goal conflicts surfaced by the system). Following reflexive thematic analysis [9, 10], we emphasized depth over inter-coder reliability and kept analytic memos.

5.3.3 Integration. We used survey trends to locate where reflection scores changed and the qualitative data to explain how, cross-checking emergent themes against survey patterns to separate individual cases from broader trends.

5.4 Study Limitations

Although goals often evolve over months [20], we chose a one-week deployment to balance burden and engagement. The switch drew on goals participants already held, but the study scheduled it instead of letting it arise on its own, so our findings capture how a first reframing lands and not how reframing behaves across repeated, self-initiated changes or once novelty fades. The survey compares Rebite against participants’ recalled impressions of prior tools, which is subject to recall and novelty effects and is not a controlled baseline. We therefore read it as secondary evidence that reflection changed, and rest the findings themselves on the qualitative accounts. Participants’ goals clustered around common themes such as weight loss and muscle gain, leaving out rarer or culturally specific ones. Rebite relies on VLM interpretation, which our technical evaluation shows is only coarsely accurate from photos alone (supplementary material), and estimates remain imperfect for portion sizes and hidden ingredients. Our goal-metric framework emphasizes common health metrics and is an artifact we authored from published guidance, not a clinically validated instrument reviewed by a dietitian or clinician, so Rebite’s feedback is not medical advice. Finally, we report what participants said about reframing, not how much they logged.

6 Study Results

We organize findings around RQ1 and RQ2, using survey results as a quantitative reference and qualitative evidence to show how goal-directed framing shaped engagement with goals.

6.1 Statistical Analysis of Pre–Post Survey Scores

Figure 5 summarizes the pre–post survey results. The strongest effects appear in the goal-related reflection items (Q7–Q11), where participants scored significantly higher with Rebite than in their recalled ratings of prior tools on all five items ($d_z = 0.75$ – 1.10).

TSRI-derived subscales. The Exploration subscale (Q1–Q3) showed a small increase that did not reach significance ($\Delta M = 0.27$, $p = .057$, adjusted $p = .078$, $d_z = 0.44$). The Insight subscale




Figure 5: Participants indicated that Rebite assisted with goal-related reflection more than their prior food journaling tools. No significant difference was observed around exploration and insight components of reflection. Significance levels are based on corrected p values.

(Q4–Q6) showed a small increase in the same direction ($\Delta M = 0.41$, $p = .039$, $d_z = 0.48$), though it did not survive correction for multiple comparisons (adjusted $p = .078$).

Goal-Related Reflection Items (Q7–Q11). Participants reported significantly higher awareness of the active dietary goal when using Rebite (Q7, $M_{prior} = 4.00$, $M_{Rebite} = 4.71$, $t(20) = 4.18$, $p < .001$, $d_z = 0.91$). They also more frequently reported reflecting on whether individual meals fit their current goal (Q8, $\Delta M = 0.81$, $t(20) = 4.25$, $p < .001$, $d_z = 0.93$), as well as noticing consistent or inconsistent patterns across meals (Q9, $\Delta M = 1.10$, $t(20) = 5.04$, $p < .001$, $d_z = 1.10$). Participants also more often rethought what it means to achieve a dietary goal (Q10, $\Delta M = 0.81$, $t(20) = 3.44$, $p = .003$, $d_z = 0.75$). Finally, they reported reflecting more frequently on the goals they were pursuing beyond specific eating behaviors (Q11, $\Delta M = 0.86$, $t(20) = 3.70$, $p = .001$, $d_z = 0.81$). The p values reported per item are uncorrected; all five remain significant after Holm-Bonferroni correction across the seven tests (adjusted $p \leq .008$). Participants therefore reported engaging with goals in a wider range of ways with Rebite than with their prior tools.

6.2 Making Sense of Goal-Directed Feedback (RQ1)

Users arrived with vague goals, such as wanting to “lose weight” or “recover from training”. The system mapped each into a concrete configuration with specific metrics. This mapping offered a foundation for reframing. Because the system determined what to measure, it could later re-determine when goals changed. Three patterns emerged. Participants confirmed what they already understood, expanded their goals, and found mismatches between stated metrics and lived experience.




Figure 6: Examples of metrics and analysis surfaced to participants based on their goals. P1’s goal of adhering to a plant-based diet surfaced monitoring iron and B12, while P6 monitors carbohydrates to support their weight loss goal.

6.2.1 Confirming existing understanding. When the system surfaced metrics that participants already cared about, it established trust in the translation mechanism (Q7–Q9). P1, a plant-based eater, saw B12 and iron flagged as below target (Figure 6, top). She said, “B12 and iron are exactly what I always worry about ... it felt like the system understood me.” P6 (weight loss) saw carbohydrates highlighted (Figure 6, bottom), and P2 (muscle gain) focused on protein. Participants who trusted the system’s initial translation were readier to attend to its re-translation when goals later changed.

6.2.2 Expanding vague goals into concrete dimensions. The system also surfaced metrics that participants would not have chosen themselves, making vague goals more concrete (Q10). P21 came in with the goal of “lose weight”, meaning “eat fewer calories”. The system added fiber, which she initially accepted out of curiosity, but later found genuinely useful. She said, “I didn’t think much about fiber, but then I noticed it actually kept me full for longer.” P5’s “anti-inflammatory eating” became concrete through polyphenols. She explained, “I had heard the word before, but never linked it to daily food ... when I go shopping I also check for it now.” Because the system supplies domain knowledge the user lacks, a new goal supplies different knowledge, so reframing produces structurally different feedback from the same data.

However, not all expansions stuck. P4 removed a spice-related metric she found irrelevant, and P1 found fiber unnecessary given her plant-based diet. These cases show the value of letting users adjust the configuration.

6.2.3 Mismatches as a boundary condition for reframing. Some goals exceeded the available metric space (Q10–Q11). P19 (recovery nutrition) noted, “Recovery is also your mood, your training load,

your rest ... eating the right foods doesn't automatically mean I've recovered well." P20 (digestive management) felt that FODMAPs, lactose, and fiber missed his experience's complexity. Where the metric set could not represent a goal, reframing under that goal had little to work with.

6.3 What Happened When Goals Changed (RQ2)

When participants changed their goal, the system re-mapped it into a new configuration and re-interpreted the same stored meals, with different metrics selected, different assessments assigned, and different interpretations generated. This appeared to make overlaps and conflicts between goals visible through the same structured feedback, prompting participants to negotiate trade-offs, refine priorities, or reject reframing altogether (Q10–Q11).

6.3.1 Finding overlap between goals. Some participants found that Rebite helped them see overlap and continuity across goals, showing that the same metrics, foods, or routines could benefit them in more than one way. This gave them a sense that their past efforts were still valuable and reduced the feeling of having to "start over".

Some metrics supported multiple goals at once, showing participants that their efforts could serve more than one purpose. For example, P19 explained that when the system highlighted protein intake under his recovery goal, the same *metric card* also showed benefits for blood sugar stability. He reflected, *"When I pushed up protein for recovery, it actually turned out to be good for blood sugar too ... eat[ing] more meals like this can work for both."* Before using Rebite, he had focused almost entirely on recovery nutrition, but reframing his logs revealed overlaps across goals.

Others described how the same food could be reinterpreted under different goals. P7, for instance, first saw avocado highlighted under her weight-loss goal as a lower-calorie fat that supported satiety. Later, under her fat-related goal, the *interpretation* praised the same avocado for its unsaturated fat content. She recalled, *"I always thought of avocado as just filling ... the system showed it as good for fat quality too."* In some cases, feedback layered affirmations with constructive contrasts. P2, focused on muscle gain, saw her protein shake affirmed as "supports muscle gain," but when she later switched to her balanced-eating goal, the same shake was flagged as "low diversity." She reflected, *"I started using protein powders with added vegetable powders ... it's an easy way to get protein and still make the meal more balanced."*

Some participants also carried forward old strategies that remained useful under new goals. P4's low-sugar meals, first adopted for weight loss, later *"also improved my inflammation score."* Reframing also turned past logs into a resource for learning. P10 explained how old logs expanded in meaning when he added new goals, connecting diet to mental clarity as well as muscle gain. He said, *"I always knew protein was for muscle, but I never thought it connected to mental clarity too."*

6.3.2 Surfacing conflict between goals. When some users switched goals and metrics, Rebite surfaced how the same food could be beneficial for one goal while problematic for the other. Seen through the same feedback format, these contradictions led participants to conclude that a meal has no single, stable reading. P12 described

this when she logged salmon. Under *Lose Weight* it was praised for protein, while under *Blood Lipids & BP* it highlighted omega-3 but also sodium concerns. She reflected, *"I realized the healthy food can be both good and bad."*

Rebite also surfaced conflicts across daily priorities. P19 focused on endurance cycling, and described how carbohydrates appeared in opposite ways. After logging a pasta dinner post-ride, the *interpretation* under his *Recovery* goal explained, "Carbohydrates in this meal help restore glycogen after endurance training." But under *Blood Sugar Control*, the feedback warned, "Large carb intake may cause unstable blood sugar." He reflected, *"After training I needed carbs, but for blood sugar it flagged the same carbs as a problem."*

6.3.3 Revising and refining goals. As previously described, Rebite made tensions between goals visible by showing different metrics and interpretations under different goals. Instead of treating meals as simply "good" or "bad," participants used these conflicts as opportunities to reflect at different levels. They sometimes made small compromises, changed priorities depending on context, reordered goal hierarchies, and redefined what health meant to them.

Some participants adjusted specific choices without abandoning either goal. P17 talked about red bean soup. It felt like a light dessert that worked fine for her weight-loss plan, but it carried enough sugar to concern her blood sugar goal. *"[Rebite]'s tag shows I should reduce the portion ... Red bean soup is always really low in calories, [it helps me] control my appetite. I just had half a bowl after dinner and put the rest in the fridge for the next day."* Here, P17 found a middle ground so she could look after both weight and blood sugar.

Other participants refined their goals by switching priorities based on the situation. Facing the carbohydrate conflict described above, P19 created a rule and explained, *"After training I need carbs, but if I'm not training I don't want them to spike my glucose. So I just switch depending on the day."* Similarly, P13 acknowledged limits in pursuing everything at once. He said, *"I had to accept that I can't maximize everything. Sometimes I lean one [goal], sometimes the other. ... Right now, [my] weight matters more to me."* These reflections illustrate how surfaced conflicts may prompt participants to explicitly negotiate trade-offs and refine priorities depending on daily context.

In some cases, Rebite prompted participants to reconsider the hierarchy of their goals over time. P12 reflected, *"I thought weight loss was important, but compared to long-term heart health it's actually less important. I know I can't maximize every goal at once. For me, long-term heart health comes first, then weight loss, and finally anti-inflammation."* For her, the value of the system was in making conflicts visible without resolving them, enabling her to articulate a new order of priorities. For others, reframing reshaped what health meant in practice. P17 reflected, *"Back then I only cared about calories, if the number was low, I felt good. But looking back, I saw I was always low on protein and fiber. That made me realize eating less wasn't the same as eating well."*

6.3.4 Resetting after goal changes. When participants changed their goals in Rebite, the system resurfaced and reinterpreted past logs. In interviews, some participants said they would rather reset, discarding old logs and starting fresh with their new goal, revealing a tension with the reframing premise that past data should always

be reframed. For some, a goal switch created a clear break with the past.

For these participants, old logs felt tied to a finished phase and no longer useful for current decisions. P13 explained, *“If my goal has changed, I don’t care what I ate in the past, it’s irrelevant now.”* Yet, some showed mixed feelings. P8 wanted a fresh start but was still curious, *“When I set a new goal, I want a fresh start. But I’ll eat my favorite food anyway, it’s my comfort food. I still want to see how the system explains it under the new goal.”*

Another way participants explained this was as a kind of fresh start ritual, a symbolic act that marked a boundary between “former me” and “current me.” P3 said switching goals was *“like opening a new one ... [like] I’m the current me that’s really trying to lose weight ... [it makes me feel] more committed.”* At the same time, participants also described resetting as a way to protect their emotions and motivation. Part of what made resetting appealing was that reinterpretation could feel like retroactive judgment, where choices that had been reasonable under a previous goal were now recast against the new goal’s standard. Some experienced this as a verdict on their past behavior. P13 explained, *“It felt like something I did before was being judged as wrong ... like the system was telling me I hadn’t done it right.”*

7 Discussion

Our study suggests that goal-directed framing can support reflection through a two-step process. The system’s goal-to-metric mapping established a concrete, shared language between user and system, which then enabled reframing, where re-interpreting past data under a new goal surfaced overlaps, conflicts, and trade-offs across goals. We discuss how mapping enables reflection, what reframing offers and costs as an interaction, and what it means for data to “live many lives.”

7.1 Enabling Reframing: How Goal-to-Metric Mapping Lays the Foundation

People often start with high-level goals and need help turning them into specific, trackable metrics [21]. The mapping (D1) supported exactly this, translating a vague goal into a shared language that the system uses to interpret meals and participants use to understand what their goal means in practice. That shared understanding enabled reframing to support reflection. Because participants already grasped their old goal in concrete terms, when the system re-interpreted the same meals under a new goal, they could compare the two readings and see where goals overlapped or conflicted. This enabled them to read the change as a shift between goals they understood, not an arbitrary system revision.

However, participants noted that this understanding was only as broad as the knowledge base behind the mapping. Some participants emphasized that nutrition-only mappings missed factors like training load or stress, which limited what reframing could surface. Future systems could extend the mapping to incorporate signals beyond nutrition [44, 52].

Mapping and reframing also reposition the role of the goal. Because Rebite can re-map whenever the goal changes, users do not have to fix their goal in advance [37, 61], and they can revise it in response to what reframing reveals instead of abandoning tracking

when it no longer fits [14, 26]. In our study, participants negotiated, reordered, and sometimes redefined their goals across goal switches (§6.3). This suggests that goals can become *outcomes of reflection* instead of fixed preconditions, and that past data stays available to revisit and reflect on as goals change.

7.2 Designing Reframing: Benefits, Risks, and Implications

LLMs make it technically feasible to re-interpret unstructured personal data under different goals at scale, but feasibility alone does not determine what the interaction should look like. Our deployment points to three implications.

7.2.1 Reframing surfaces cross-goal structure that users cannot see on their own. Reframing surfaced the complexity of pursuing multiple, sometimes conflicting goals. Overlaps reduced the sense of “starting over,” while conflicts pushed participants to negotiate trade-offs. Because LLMs can re-assess unstructured data under different goals, these cross-goal relationships can become visible without requiring users to hold multiple framings in memory. This suggests that LLM-powered reframing may extend beyond accommodating new goals, potentially making the *relationships between goals* an explicit object of reflection.

In our study of Rebite, these cross-goal relationships surfaced sequentially, where users saw them by reframing under each goal in turn, never side by side. This fits how PI is often used in everyday life, favoring brief single-goal glances over deliberate multi-goal comparison [8]. Still, moments like weighing trade-offs across goals might benefit from a direct cross-goal view. Designing a tool that provides this sort of view in everyday reflection, without burdening it, could be valuable future work.

7.2.2 Reframing should support negotiating trade-offs between goals. Beyond observing cross-goal relationships, participants actively negotiated trade-offs between their goals. They adjusted portions and reordered which goal took priority, sometimes introducing context-dependent rules such as prioritizing recovery on training days and blood sugar on rest days. This negotiation ran on concrete, explained conflict. The four-layer feedback let participants see *which* metric conflicted, *why*, and *which food* caused it, and seeing the same meal assessed under both goals put the trade-off directly in front of them. These natural-language interpretations gave richer anchors than a numeric dashboard alone.

Negotiation therefore depends on the conflict staying visible. Blending multiple goals into a single score would settle the disagreement for participants and remove the very material they reasoned with (§6.3). Systems should instead keep the conflict intact and help people act on it themselves, for instance by letting them record the context-dependent rules they arrive at, without resolving the trade-off for them.

7.2.3 Users need control over whether and how much to reframe. Reframing also caused discomfort for some. Some participants preferred to reset entirely, treating previous logs as a closed chapter (e.g., happy abandonment [14, 24]). Others felt uneasy when feedback recast previously “healthy” foods as problematic. Because re-interpreting all stored data at once is technically cheap, a single

goal change can deliver far more re-assessment than the user asked to see.

This discomfort shows that reframing does not merely re-describe past behavior. It re-evaluates that behavior against a *new* standard. The same content points two ways. As advice about the next meal it feels actionable, but applied to meals already eaten it can read as a verdict, and some participants experienced it as being judged for past choices (§6.3). For some, that is why they preferred to reset to escape feelings of judgment.

Our instantiation of goal-directed reframing in Rebite left users with a largely binary choice of accepting the new assessment or resetting and walking away. Neither option lets a user agree with part of a re-assessment while pushing back on another, or contest a specific verdict without abandoning the broader reframing. A first step towards resolving this tension is giving users control over the *scope* of reframing by indicating which data to re-interpret, at what granularity, and whether to engage at all. Beyond scope, users would also benefit from ways to negotiate with the re-assessment itself. Prior HCI work has made system output adjustable instead of take-it-or-leave-it [39] and let users revise their own input after the fact [30, 38]. Translated to reframing, a re-assessment could allow users to accept some reinterpretations, defer or reject others, or annotate why a past choice was reasonable under its original goal, so they keep what reframing offers while staying in control of their own data.

7.3 Generalizing Reframing: When Data Lives Many Lives

7.3.1 The same data can mean many things. Goal-directed framing broadens what PI systems can do with already-collected data. By keeping raw data and deferring interpretation, the same meal can be framed as recovery fuel, a weight-management risk, or a contribution to dietary diversity without new data entry. Prior critiques have shown how tracking tools can reproduce narrow narratives, such as equating health with thinness [19, 56]. Keeping data open to reinterpretation may help resist these rigidities by treating health as situational and changeable.

However, LLMs also carry risks of amplifying harmful narratives. Biased outputs may reinforce restrictive dieting or pro-eating disorder tendencies, tied to cultural biases embedded in the models themselves. Anchoring metric selection to published guidance constrains this only partly. Our framework is authored, not clinically validated, and fixing the metric set does not fix the language the model uses about it. Whether goal plurality widens the range of values a system can express, or only gives a biased model more surfaces to speak from, is an open question. Reinterpreting the same data under different goals also shows that no assessment is neutral. The system’s interpretation, the user’s own understanding of the meal, and the values built into the active goal can diverge, and reframing surfaces those differences for the user to weigh.

7.3.2 The cost of deferring interpretation. This “capture now, interpret later” stance is double-edged. Decoupling capture from any single goal lowers the barrier to logging and keeps data useful when priorities change, which is much of what makes reframing possible in the first place. Yet it sits uneasily with a long-standing view in personal informatics that *why* people record shapes *what* they

record [26]. If the “why” is deferred indefinitely, the motivation that sustains careful logging may erode, and capture may drift toward data hoarding. How to keep open-ended logging worthwhile is a question future work should examine. Automated and semi-automated collection through wearables and passive sensing makes preserving rich data increasingly feasible, which only sharpens this question, as preserving reusable data may be necessary but not sufficient. Future systems may therefore need to support goal *discovery* and refinement [20], letting users explore data and goals together so that each decision improves the other over time.

7.3.3 Beyond food journaling. Unclear and changing goals are not unique to food journaling [20]. The approach may generalize to domains that satisfy three conditions. Raw data supports multiple interpretations (e.g., sleep sensor data encodes latency, efficiency, stages, and HRV), a structured knowledge base can map goals to relevant dimensions (e.g., clinical guidelines for sleep hygiene vs. athletic recovery), and factual estimates can be separated from goal-directed assessments. Prior work has shown that personalized systems can adapt to individual users in domains like sleep [18] and conversational goal tracking [17]. Extending to a new domain would likely require building a domain-specific mapping layer, though the rest of the pipeline could be reused.

8 Conclusion

We introduce goal-directed framing, in which the goal acts as an active lens over preserved data, and realize it in Rebite. Reframing past meals under a new goal surfaced overlaps and conflicts that prompted users to negotiate trade-offs and refine priorities, though some preferred to start fresh. The approach could extend to domains such as exercise, sleep, and mental health.

Acknowledgments

This work was partially supported by the National Science Foundation under award IIS-2237389. Generative AI tools were used for grammar and sentence editing, advice on shortening, and some development assistance.

References

- [1] Elena Agapie, Patricia A. Areán, Gary Hsieh, and Sean A. Munson. 2022. A Longitudinal Goal Setting Model for Addressing Complex Personal Problems in Mental Health. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 451 (2022), 29 pages. doi:10.1145/3555160
- [2] Christopher Ahlberg, Christopher Williamson, and Ben Shneiderman. 1992. Dynamic Queries for Information Exploration: An Implementation and Evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '92)*. ACM, 619–626. doi:10.1145/142750.143054
- [3] Dario Allegra, Sebastiano Battiato, Alessandro Ortis, Salvatore Urso, and Riccardo Polosa. 2020. A review on food recognition technology for health applications. *Health psychology research* 8, 3 (2020), 9297. doi:10.4081/hpr.2020.9297
- [4] American Diabetes Association Professional Practice Committee. 2024. Facilitating Positive Health Behaviors and Well-being to Improve Health Outcomes: Standards of Care in Diabetes—2024. *Diabetes Care* 47, Supplement 1 (2024), S77–S110. doi:10.2337/dc24-S005 Erratum in: *Diabetes Care* 47, 4 (2024), 761–762. https://doi.org/10.2337/dc24-er04.
- [5] Adrienne H Andrew, Gaetano Borriello, and James Fogarty. 2013. Simplifying mobile phone food diaries. In *2013 7th International Conference on Pervasive Computing Technologies for Healthcare and Workshops*. IEEE, 260–263. doi:10.4108/icst.pervasivehealth.2013.252101
- [6] Matthew Barker-Canler, Daniel Gooch, Janet van der Linden, and Marian Petre. 2024. Flexible Minimalist Self-Tracking to Support Individual Reflection. *ACM Transactions on Computer-Human Interaction* 31, 3 (2024), 35 pages. doi:10.1145/3660339

- [7] Marit Bentvelzen, Jasmin Niess, Mikołaj P. Woźniak, and Paweł W. Woźniak. 2021. The Development and Validation of the Technology-Supported Reflection Inventory. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 366, 8 pages. doi:10.1145/3411764.3445673
- [8] Tanja Blascheck, Frank Bentley, Eun Kyoung Choe, Tom Horak, and Petra Isenberger. 2021. Characterizing Glanceable Visualizations: From Perception to Behavior Change. In *Mobile Data Visualization*, Bongshin Lee, Raimund Dachsel, Petra Isenberger, and Eun Kyoung Choe (Eds.). Chapman and Hall/CRC, 151–176. doi:10.1201/9781003090823-5
- [9] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101. doi:10.1191/1478088706qp0630a
- [10] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (2019), 589–597. doi:10.1080/2159767X.2019.1628806
- [11] Shaan Chopra, Katherine Juarez, James Fogarty, and Sean A. Munson. 2025. Engagements with Generative AI and Personal Health Informatics: Opportunities for Planning, Tracking, Reflecting, and Acting around Personal Health Data. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 9, 3, Article 75 (Sept. 2025), 33 pages. doi:10.1145/3749503
- [12] Chia-Fang Chung, Elena Agapie, Jessica Schroeder, Sonali Mishra, James Fogarty, and Sean A. Munson. 2017. When Personal Tracking Becomes Social: Examining the Use of Instagram for Healthy Eating. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 1674–1687. doi:10.1145/3025453.3025747
- [13] Chia-Fang Chung, Qiaosi Wang, Jessica Schroeder, Allison Cole, Jasmine Zia, James Fogarty, and Sean A. Munson. 2019. Identifying and Planning for Individualized Change: Patient-Provider Collaboration Using Lightweight Food Diaries in Healthy Eating and Irritable Bowel Syndrome. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 1, Article 7 (March 2019), 27 pages. doi:10.1145/3314394
- [14] James Clawson, Jessica A. Pater, Andrew D. Miller, Elizabeth D. Mynatt, and Lena Mamykina. 2015. No longer wearing: investigating the abandonment of personal health-tracking technologies on craigslist. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Osaka, Japan) (UbiComp '15). Association for Computing Machinery, New York, NY, USA, 647–658. doi:10.1145/2750858.2807554
- [15] Felicia Cordeiro, Elizabeth Bales, Erin Cherry, and James Fogarty. 2015. Rethinking the Mobile Food Journal: Exploring Opportunities for Lightweight Photo-Based Capture. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 3207–3216. doi:10.1145/2702123.2702154
- [16] Felicia Cordeiro, Daniel A. Epstein, Edison Thomaz, Elizabeth Bales, Arvind K. Jagannathan, Gregory D. Abowd, and James Fogarty. 2015. Barriers and Negative Nudges: Exploring Challenges in Food Journaling. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 1159–1162. doi:10.1145/2702123.2702155
- [17] Adam J. Coscia, Shunan Guo, Eunye Koh, and Alex Endert. 2025. OnGoal: Tracking and Visualizing Conversational Goals in Multi-Turn Dialogue with Large Language Models. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology* (UIST '25). ACM, New York, NY, USA, 18 pages. doi:10.1145/3746059.3747746
- [18] Nediya Daskalova, Danaë Metaxa-Kakavouli, Adrienne Tran, Nicole Nugent, Julie Boergers, John McGeary, and Jeff Huang. 2016. SleepCoach: A Personalized Automated Self-Experimentation System for Sleep Recommendations. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology* (UIST '16). ACM, 347–358. doi:10.1145/2984511.2984534
- [19] Elizabeth V. Eikey and Madhu C. Reddy. 2017. "It's Definitely Been a Journey": A Qualitative Study on How Women with Eating Disorders Use Weight Loss Apps. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 642–654. doi:10.1145/3025453.3025591
- [20] Tina Ekhtiari, Armağan Karahanoglu, Rûben Gouveia, and Geke D. S. Ludden. 2025. Changing Health Goals with Personal Informatics. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference* (DIS '25). Association for Computing Machinery, New York, NY, USA, 2867–2883. doi:10.1145/3715336.3735813
- [21] Tina Ekhtiari, Armağan Karahanoglu, Rûben Gouveia, and Geke Ludden. 2023. Goals for Goal Setting: A Scoping Review on Personal Informatics. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) (DIS '23). Association for Computing Machinery, New York, NY, USA, 2625–2641. doi:10.1145/3563657.3596087
- [22] Julia A. Ello-Martin, Liane S. Roe, Jenny H. Ledikwe, Amanda M. Beach, and Barbara J. Rolls. 2007. Dietary energy density in the treatment of obesity: a year-long trial comparing 2 weight-loss diets. *The American Journal of Clinical Nutrition* 85, 6 (2007), 1465–1477. doi:10.1093/ajcn/85.6.1465
- [23] Daniel A. Epstein, Clara Caldeira, Mayara Costa Figueiredo, Xi Lu, Lucas M. Silva, Lucretia Williams, Jong Ho Lee, Qingyang Li, Simran Ahuja, Qiuer Chen, Payam Dowlatyari, Craig Hilby, Sazeda Sultana, Elizabeth V. Eikey, and Yunan Chen. 2020. Mapping and Taking Stock of the Personal Informatics Literature. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 4, Article 126 (Dec. 2020), 38 pages. doi:10.1145/3432231
- [24] Daniel A. Epstein, Monica Caraway, Chuck Johnston, An Ping, James Fogarty, and Sean A. Munson. 2016. Beyond Abandonment to Next Steps: Understanding and Designing for Life after Personal Informatics Tool Use. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 1109–1113. doi:10.1145/2858036.2858045
- [25] Daniel A. Epstein, Felicia Cordeiro, James Fogarty, Gary Hsieh, and Sean A. Munson. 2016. Crumbs: Lightweight Daily Food Challenges to Promote Engagement and Mindfulness. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 5632–5644. doi:10.1145/2858036.2858044
- [26] Daniel A. Epstein, An Ping, James Fogarty, and Sean A. Munson. 2015. A lived informatics model of personal informatics. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Osaka, Japan) (UbiComp '15). Association for Computing Machinery, New York, NY, USA, 731–742. doi:10.1145/2750858.2804250
- [27] Catrin Feron, Tina Ekhtiari, and Ruben Gouveia. 2022. Transitions in Personal Informatics: Investigating Self-Tracking During Moments of Change. In *Adjunct Proceedings of the 2022 Nordic Human-Computer Interaction Conference* (Aarhus, Denmark) (NordiCHI '22 Adjunct). Association for Computing Machinery, New York, NY, USA, Article 43, 5 pages. doi:10.1145/3547522.3547686
- [28] Rûben Gouveia and Daniel A. Epstein. 2023. This Watchface Fits with my Tattoos: Investigating Customisation Needs and Preferences in Personal Tracking. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 327, 15 pages. doi:10.1145/3544548.3580955
- [29] Seung Wan Ha, Novia Nurain, Elena Agapie, and Chia-Fang Chung. 2025. Preparing and Experiencing Food During Life Events: Implications for Technology Supporting Social and Value Changes. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 354, 14 pages. doi:10.1145/3706598.3713183
- [30] Michael J. Hoefer, Raegan Rychecky, Max Gong, and Stephen Volda. 2025. Tell-Time: An AI-Augmented Calendar with a Voice Interface for Collecting Time-Use Data. In *Proceedings of the 30th International Conference on Intelligent User Interfaces* (IUI '25). ACM, 1366–1380. doi:10.1145/3708359.3712116
- [31] Institute of Medicine, Food and Nutrition Board. 2006. *Dietary Reference Intakes: The Essential Guide to Nutrient Requirements*. The National Academies Press, Washington, DC. doi:10.17226/11537
- [32] Peiling Jiang, Jude Rayan, Steven P. Dow, and Haijun Xia. 2023. Graphologue: Exploring Large Language Model Responses with Interactive Diagrams. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (UIST '23). ACM, 3:1–3:20. doi:10.1145/3586183.3606737
- [33] Matthew Jörke, Shardul Sapkota, Lyndsea Warkenthien, Niklas Vainio, Paul Schmiedmayer, Emma Brunskill, and James A. Landay. 2025. GPTCoach: Towards LLM-Based Physical Activity Coaching. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). ACM, 46 pages. doi:10.1145/3706598.3713819
- [34] Yeohyun Jung, Kwangyoung Lee, and Hwajung Hong. 2025. Thinking Outside the Data Box: Investigating the Potential of Data Manipulation for Self-Reflection on Personal Data. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference* (DIS '25). Association for Computing Machinery, New York, NY, USA, 2189–2207. doi:10.1145/3715336.3735840
- [35] Ravi Karkar, Jessica Schroeder, Daniel A. Epstein, Laura R. Pina, Jeffrey Scofield, James Fogarty, Julie A. Kientz, Sean A. Munson, Roger Vilardaga, and Jasmine Zia. 2017. TummyTrials: A Feasibility Study of Using Self-Experimentation to Detect Individualized Food Triggers. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 6850–6863. doi:10.1145/3025453.3025480
- [36] Taewan Kim, Seolyeong Bae, Hyun Ah Kim, Su-Woo Lee, Hwajung Hong, Chanmo Yang, and Young-Ho Kim. 2024. MindfulDiary: Harnessing Large Language Model to Support Psychiatric Patients' Journaling. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 701, 20 pages. doi:10.1145/3613904.3642937
- [37] Young-Ho Kim, Jae Ho Jeon, Bongshin Lee, Eun Kyoung Choe, and Jinwook Seo. 2017. OmniTrack: A Flexible Self-Tracking Approach Leveraging Semi-Automated Tracking. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 67 (Sept. 2017), 28 pages. doi:10.1145/3130930

- [38] Ha Le, Veronika Potter, Akshat Choube, Rithika Lakshminarayanan, Varun Mishra, and Stephen Intille. 2025. A Context-Assisted, Semi-Automated Activity Recall Interface Allowing Uncertainty. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 4 (2025), 34 pages. doi:10.1145/3770710
- [39] Jungeun Lee, Sungnam Kim, Minki Cheon, Hyojin Ju, JaeEun Lee, and Inseok Hwang. 2022. SleepGuru: Personalized Sleep Planning System for Real-life Actionability and Negotiability. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. ACM, Article 52, 25 pages. doi:10.1145/3526113.3545709
- [40] Jong Ho Lee, Jessica Schroeder, and Daniel A. Epstein. 2021. Understanding and Supporting Self-Tracking App Selection. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 4, Article 166 (Dec. 2021), 25 pages. doi:10.1145/3494980
- [41] Min Kyung Lee, Junsung Kim, Jodi Forlizzi, and Sara Kiesler. 2015. Personalization revisited: a reflective approach helps people better personalize health services and motivates them to increase physical activity. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (Osaka, Japan) (UbiComp '15)*. Association for Computing Machinery, New York, NY, USA, 743–754. doi:10.1145/2750858.2807552
- [42] Ian Li, Anind Dey, and Jodi Forlizzi. 2010. A stage-based model of personal informatics systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Atlanta, Georgia, USA) (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 557–566. doi:10.1145/1753326.1753409
- [43] Ian Li, Anind K. Dey, and Jodi Forlizzi. 2011. Understanding My Data, Myself: Supporting Self-Reflection with Ubicomp Technologies. In *Proceedings of the 13th International Conference on Ubiquitous Computing (UbiComp '11)*. Association for Computing Machinery, New York, NY, USA, 405–414. doi:10.1145/2030112.2030166
- [44] Jiachen Li, Xiwen Li, Justin Steinberg, Akshat Choube, Bingsheng Yao, Xuhai Xu, Dakuo Wang, Elizabeth Mynatt, and Varun Mishra. 2025. Vital Insight: Assisting Experts' Context-Driven Sensemaking of Multi-modal Personal Tracking Data Using Visualization and Human-in-the-Loop LLM. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 9, 3, Article 101 (Sept. 2025), 37 pages. doi:10.1145/3749508
- [45] Xi Lu, Yunan Chen, and Daniel A. Epstein. 2021. A Model of Socially Sustained Self-Tracking for Food and Diet. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 451 (Oct. 2021), 32 pages. doi:10.1145/3479595
- [46] Xi Lu, Edison Thomaz, and Daniel A. Epstein. 2022. Understanding People's Perceptions of Approaches to Semi-Automated Dietary Monitoring. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 3, Article 129 (Sept. 2022), 27 pages. doi:10.1145/3550288
- [47] Kai Lukoff, Taoxi Li, Yuan Zhuang, and Brian Y. Lim. 2018. TableChat: Mobile Food Journaling to Facilitate Family Support for Healthy Eating. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 114 (Nov. 2018), 28 pages. doi:10.1145/3274383
- [48] Yuhuan Luo, Young-Ho Kim, Bongshin Lee, Naeemul Hassan, and Eun Kyoung Choe. 2021. FoodScrap: Promoting Rich Data Capture and Reflective Food Journaling Through Speech Input. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference (Virtual Event, USA) (DIS '21)*. Association for Computing Machinery, New York, NY, USA, 606–618. doi:10.1145/3461778.3462074
- [49] Yuhuan Luo, Peiyi Liu, and Eun Kyoung Choe. 2019. Co-Designing Food Trackers with Dietitians: Identifying Design Opportunities for Food Tracker Customization. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland UK) (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300822
- [50] Elliot G. Mitchell, Elizabeth M. Heitkemper, Marissa Burgermaster, Matthew E. Levine, Yishen Miao, Maria L. Hwang, Pooja M. Desai, Andrea Cassells, Jonathan N. Tobin, Esteban G. Tabak, David J. Albers, Arlene M. Smaldone, and Lena Mamykina. 2021. From Reflection to Action: Combining Machine Learning with Expert Knowledge for Nutrition Goal Recommendations. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 206, 17 pages. doi:10.1145/3411764.3445555
- [51] Jimmy Moore, Pascal Goffin, Jason Wiese, and Miriah Meyer. 2022. Exploring the Personal Informatics Analysis Gap: "There's a Lot of Bacon". *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 96–106. doi:10.1109/TVCG.2021.3114798
- [52] Subigya Nepal, Arvind Pillai, William Campbell, Talie Massachi, Michael V. Heinz, Ashmita Kunwar, Eunsol Soul Choi, Xuhai Xu, Joanna Kuc, Jeremy F. Huckins, Jason Holden, Sarah M. Preum, Colin Depp, Nicholas Jacobson, Mary P. Czerwinski, Eric Granholm, and Andrew T. Campbell. 2024. MindScape Study: Integrating LLM and Behavioral Sensing for Personalized AI-Driven Journaling Experiences. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 4, Article 186 (Nov. 2024), 44 pages. doi:10.1145/3699761
- [53] Ada Ng, Ashley Marie Walker, Laurie Wakschlag, Nabil Alshurafa, and Madhu Reddy. 2022. Understanding Self-Trackers Data from Bounded Situational Contexts. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference (DIS '22)*. Association for Computing Machinery, New York, NY, USA, 1684–1697. doi:10.1145/3532106.3533498
- [54] Jasmin Niess and Pawel W. Woźniak. 2018. Supporting Meaningful Personal Fitness: the Tracker Goal Evolution Model. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3173745
- [55] Jon Noronha, Eric Hysen, Haoqi Zhang, and Krzysztof Z. Gajos. 2011. Platemate: crowdsourcing nutritional analysis from food photographs. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (Santa Barbara, California, USA) (UIST '11)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/2047196.2047198
- [56] Daisy O'Neill, Max V. Birk, and Regan L. Mandryk. 2024. Unpacking Norms, Narratives, and Nourishment: A Feminist HCI Critique on Food Tracking Technologies. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 514, 20 pages. doi:10.1145/3613904.3642600
- [57] Barbara J. Rolls. 2009. The relationship between dietary energy density and energy intake. *Physiology & Behavior* 97, 5 (2009), 609–615. doi:10.1016/j.physbeh.2009.03.011
- [58] Herman Saksono and Andrea G. Parker. 2024. Socio-Cognitive Framework for Personal Informatics: A Preliminary Framework for Socially-Enabled Health Technologies. *ACM Trans. Comput.-Hum. Interact.* 31, 3, Article 42 (Aug. 2024), 41 pages. doi:10.1145/3674504
- [59] Arvind Satyanarayan, Kanit Wongsuphasawat, and Jeffrey Heer. 2014. Declarative Interaction Design for Data Visualization. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST '14)*. ACM, 669–678. doi:10.1145/2642918.2647360
- [60] Jessica Schroeder, Ravi Karkar, Natalia Murinova, James Fogarty, and Sean A. Munson. 2019. Examining Opportunities for Goal-Directed Self-Tracking to Support Chronic Condition Management. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 4, Article 151 (Dec. 2019), 26 pages. doi:10.1145/3369809
- [61] Yasaman S. Sefidgar, Carla L. Castillo, Shaan Chopra, Liwei Jiang, Tae Jones, Anant Mittal, Hyeoung Ryu, Jessica Schroeder, Allison Cole, Natalia Murinova, Sean A. Munson, and James Fogarty. 2024. MigrateTracker: Examining Patient Experiences with Goal-Directed Self-Tracking for a Chronic Health Condition. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 129, 19 pages. doi:10.1145/3613904.3642075
- [62] Donghoon Shin, Gary Hsieh, and Young-Ho Kim. 2025. PlanFitting: Personalized Exercise Planning with Large Language Model-driven Conversational Agent. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces (CUI '25)*. Association for Computing Machinery, New York, NY, USA, Article 66, 19 pages. doi:10.1145/3719160.3736607
- [63] Lucas M. Silva, Elizabeth A. Ankrah, Yuqi Huai, and Daniel A. Epstein. 2023. Exploring Opportunities for Multimodality and Multiple Devices in Food Journaling. *Proc. ACM Hum.-Comput. Interact.* 7, MHCI, Article 209 (Sept. 2023), 27 pages. doi:10.1145/3604256
- [64] Lucas M. Silva, Xi Lu, Emily X. Liang, and Daniel A. Epstein. 2025. Foody Talk: Exploring Opportunities for Conversational Food Journaling. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1183, 19 pages. doi:10.1145/3706598.3713875
- [65] Inhwa Song, SoHyun Park, Sachin R Pendse, Jessica Lee Schleider, Munmun De Choudhury, and Young-Ho Kim. 2025. ExploreSelf: Fostering User-driven Exploration and Reflection on Personal Challenges with Adaptive Guidance by Large Language Models. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 306, 22 pages. doi:10.1145/3706598.3713883
- [66] Konstantin R. Strömel, Stanislas Henry, Tim Johansson, Jasmin Niess, and Pawel W. Woźniak. 2024. Narrating Fitness: Leveraging Large Language Models for Reflective Fitness Tracker Data Interpretation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 646, 16 pages. doi:10.1145/3613904.3642032
- [67] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. ACM, 1:1–1:18. doi:10.1145/3586183.3606756
- [68] Ghalib Ahmed Tahir and Chu Kiong Loo. 2021. A comprehensive survey of image-based food recognition and volume estimation methods for dietary assessment. *Healthcare* 9, 12 (2021), 1676. doi:10.3390/healthcare9121676
- [69] Romain H. A. Toebosch, Arne Berger, and Carine Lallemand. 2025. Challenging Sensor-based Personal Informatics: A Triadic Framework for Designing Open-ended Self-tracking Devices. In *Proceedings of the 19th International Conference on Tangible, Embedded, and Embodied Interaction (TEI '25)*. Association for Computing Machinery, New York, NY, USA, 15 pages. doi:10.1145/3689050.3704425
- [70] Christopher C Tsai, Gunny Lee, Fred Raab, Gregory J Norman, Timothy Sohn, William G Griswold, and Kevin Patrick. 2007. Usability and feasibility of PMEB: a mobile phone application for monitoring real time caloric balance. *Mobile*

- networks and applications* 12, 2-3 (2007), 173–184. doi:10.1007/s11036-007-0014-4
- [71] U.S. Department of Agriculture and U.S. Department of Health and Human Services. 2020. Dietary Guidelines for Americans, 2020–2025. 9th edition. <https://www.dietaryguidelines.gov/> 9th ed. Available at DietaryGuidelines.gov. Accessed 2025-08-14.
- [72] Yixuan Zhang and Andrea G. Parker. 2020. Eat4Thought: A Design of Food Journaling. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI EA '20*). Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3334480.3383044