Enhancing Food Tracking Practice with LLMs: Towards Goal-Oriented Reflection

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

Abstract

Current personal informatics systems often rely on visualizations, which makes it difficult to accommodate users' evolving and personalized health goals. We aim to enhance food tracking systems by integrating LLMs to better support goal-oriented reflection. We propose a photo-based food journaling approach that leverages LLMs' context-aware reasoning and personalized feedback generation to dynamically present relevant tracking metrics, deliver goal-aligned insights, and identify historical eating patterns to support long-term reflection. Drawing on an analysis of popular food tracking apps, we outline three key design strategies and present a system workflow from photo input to personalized feedback. We have started prototyping using Figma and GPT-4o. Future work will include empirical studies to evaluate how LLMs can support reflection in food tracking as users' goals evolve.

CCS Concepts

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

Keywords

Personal informatics, Large Language models, Food tracking

ACM Reference Format:
Weijun Li and Daniel A. Epstein. 2025. Enhancing Food Tracking Practice with LLMs: Towards Goal-Oriented Reflection. In *Proceedings of (CHI '25 Workshop on Envisioning the Future of Interactive Health).* ACM, New York, NY, USA, 4 pages.

1 Introduction

Personal Informatics (PI) systems can support gaining self-knowledge and take steps toward achieving their personal goals [6, 10]. These systems, primarily focused on health, are widely used in domains such as physical activity [10, 11], food journaling [13, 18], mental health [1]. They enable users to collect data, receive feedback, and adjust their behaviors to align with their personal goals.

In the practice of using PI systems, reflection plays a crucial role, helping users make sense of their raw data and apply insights to change their behavior. Most existing PI systems (such as Apple Health, Fitbit, and MyFitnessPal) depend on data visualization, such as rule-based trend graphs, statistical summaries, and progress tracking to help users reflect their behaviors. Research shows [2] that reflection does not happen automatically and requires active guidance. However, current PI systems often assume that visualization alone fosters meaningful reflection, making reflection highly dependent on user motivation [3], which might reduce its effectiveness. Some tools provide brief explanations below visualizations to help users interpret data, but these general insights are not customized to individual behaviors [2]. While many PI systems provide goal-based feedback, their tracking goals, feedback mechanisms, and data presentation are typically system-defined and rule-based, limiting users' ability to adapt them to their evolving or unique needs.

Large language models (LLMs) offer a potential solution to enhance the reflection mechanisms of current PI systems by enabling context-aware reasoning, adaptive feedback generation, and personalized insight delivery. Emerging research has started exploring how LLMs can be integrated into PI systems. For instance, GPT-Coach [8] leverages LLMs' conversational flexibility and multimodal reasoning for personalized health interventions. Similarly, Vital Insight [12] combines interactive data visualization and LLM-driven inference to assist experts in iterative sensemaking. Narrating Fitness [20] leverages LLMs' ability to generate qualitative descriptions of fitness data. While these approaches show LLMs' capability to provide adaptive feedback and explain data, current explorations are still primarily focused on generated qualitative narratives or expert-oriented interactive systems. Visualization-based interactions are primarily designed to support expert analysis rather than everyday self-tracking for general users, despite the prevalence of visualization-oriented feedback in PI applications for this audience. Therefore, we aim to explore how LLMs can enhance existing PI systems to better support reflection, particularly goal-oriented reflection in self-tracking.

We focus on food journaling as our research domain to enhance current PI system. Food journaling plays a crucial role in health and well-being, and it is a widely studied domain in personal informatics. Prior research has identified diverse goals for food tracking, including weight loss [15], managing chronic diseases [4], identifying intolerances [9], making healthier food choices [7]. Additionally, food tracking involves diverse data modalities [19], including manual text entry, nutritional databases, image recognition, and contextual factors. Unlike static health metrics such as step counts or heart rate, dietary behaviors are more dynamic and personalized, with the reflection process evolving alongside users' goals and contexts. Therefore, we base our design on photo-based food journaling, as photos offer a flexible and rich medium for capturing meals, providing diverse contextual information that reflects various aspects of eating behaviors [5]. We aim to support goal-oriented food reflection by interpreting meal photos and providing dynamic,

CHI '25 Workshop on Envisioning the Future of Interactive Health, April 27th, 2025, Yokohama, Japan

Feature	Existing Apps	Potential Enhancements with LLMs
Nutritional & Calorie Analysis	Fixed calorie and macronutrient breakdowns	Dynamically adjusts presentation based on user goals
Energy Balance (Deficit/Surplus)	Static calculations	Calculations and context-aware insights
Macronutrient & Micronutrient Trends	Displays macro/micronutrient	Highlights key focus based on goal focus
Meal Timing & Distribution	Logs meal times	Detects meal timing pattern
Mindful Eating Support	Predefined templates or reminders	Generates personalized reflection prompts
Behavioral Triggers & Contextual Factors	User input	Analyze social setting based on photos
Personalized Feedback	System-defined, generic recommendations	Personal insights
Long-Term Trend Analysis	Historical data in charts and tables	Links them to goals
Food History & Consumption Patterns	Basic meal history logs	Detects pattern

Figure 1: Existing food tracking apps include a variety of features designed to support reflection which have the potential to be enhanced by LLMs. Numerical summaries of tracked data can be enhanced by dynamic presentation of tracking metrics, which account for user goals. User-specific trends can be summarized through context-aware and adaptive feedback, extending existing approaches. Trends are typically summarized through static charts, which can be improved with LLMs through historical pattern recognition for goal alignment.

personalized feedback that aligns with users' specific needs and goals. In this poster, we describe our initial design and implementation efforts, with the intent of discussing our plans in more detail at the workshop.

2 Design and Proposed Approach

To inform our design, we analyzed widely used food tracking apps, including general-purpose loggers (Lose It!, MyFitnessPal, MyNetDiary), nutrient-focused tools (Cronometer), and mindful eating apps (Ate). These apps support reflection through calorie and macronutrient tracking, weight trend analysis, meal timing insights, and goal progress monitoring, primarily using numerical summaries, visual analytics, and automated food recommendations. Some, like Ate, incorporate behavioral reflection by encouraging mindful eating and logging meal-related motivations. While these reflective features are useful, prior work has also pointed to their limitations. For example, a qualitative study on MyFitnessPal [16]shows that users are frustrated with the lack of personalized feedback and limited customization. They prefer quick access to relevant data without navigating complex menus and want tracking systems to adapt to changing goals instead of following fixed rules.

2.1 Key Design Strategies

Based on our analysis and prior research, we summarized the key reflection functions of existing food tracking systems in Figure 1 and identify potential areas where LLMs can enhance reflection. We propose three key design strategies. **(1)** Dynamic presentation of tracking metrics ensures that users receive the most relevant metrics and insights based on their goals. As shown in Figure 1, conventional food journaling systems present fixed calorie and macronutrient breakdowns, meal timing, energy balance, which often require users to manually interpret and prioritize information based on their needs. Our approach prioritizes goal-specific indicators and presents different key metrics accordingly based on user goals. For example, if a user focuses on muscle gain, the system will highlight protein intake and recovery-related nutrients. By intelligently surfacing the most relevant data, the system could reduce cognitive load and provides a more personalized reflection experience. **(2)** Context-aware and adaptive feedback improves how users engage with their data. Existing apps mainly rely on static charts and predefined feedback, and these systems often present information in a rigid format. Our approach addresses this limitation by dynamically adjusting both UI presentation and textual explanations to provide more actionable insights. Instead of displaying uniform numerical summaries, the interface visually highlights trends that are most relevant to the user's current goals. It then generates tailored feedback, offering specific strategies to adjust future choices. In this way, we improve reflection while keeping the existing visualization system. **(3)** Historical pattern recognition for goal alignment bridges single meal tracking with long-term dietary trends. While conventional apps provide historical data in charts and tables, users must manually identify patterns and trends from the raw numbers. Our approach automates this process by detecting habits. Rather than simply displaying past intake data, the system will proactively evaluate how these patterns align with or deviate from users' goal.

2.2 Prototype Implementation Plans

Technically, as shown in Figure 2, we plan to implement our system using a structured pipeline that processes meal photos, analyzes

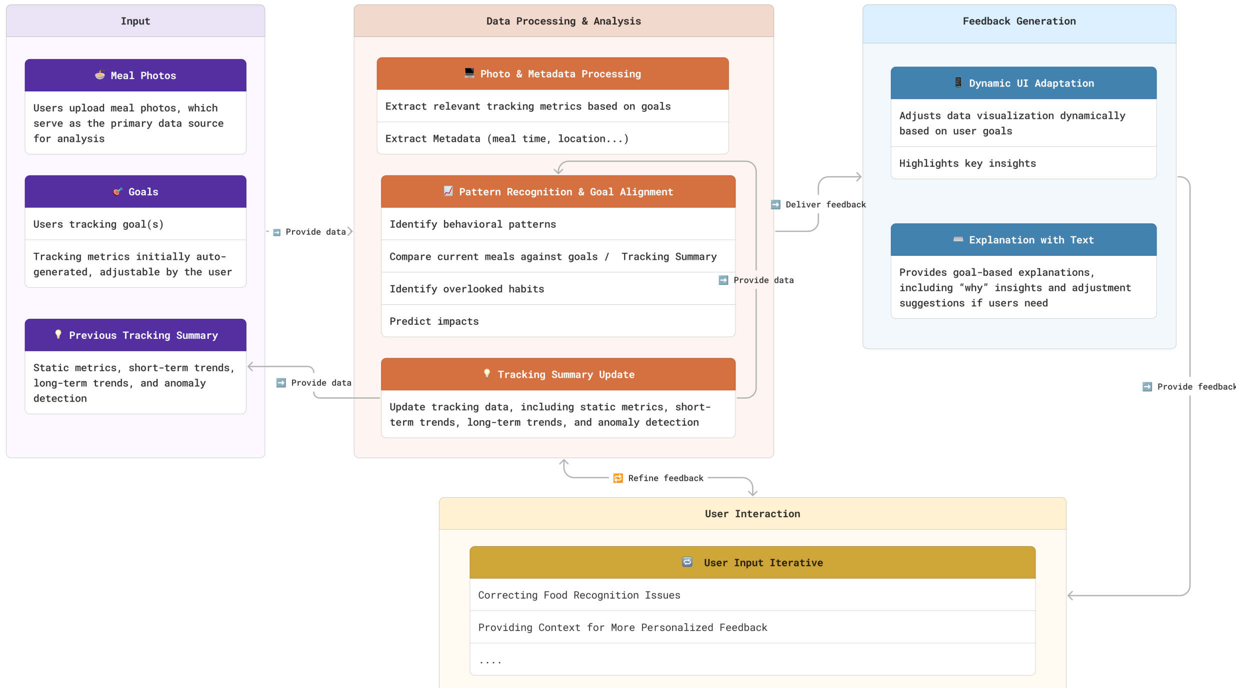

Figure 2: Goal-oriented food tracking pipeline: Our system will follow a structured process, starting with Input, where users upload a meal photo and specify their tracking goals. A tracking history summary is referenced to provide additional context. In Data Processing & Analysis stage, the system extracts food items, nutritional data, and metadata such as meal time and location. It then compares the meal to past behaviors, detecting patterns, overlooked habits, and goal alignment. In Feedback Generation, the system dynamically adjusts the UI and textual explanations to emphasize relevant insights. Finally, User Interaction & Iterative Refinement allows users to correct recognition errors, add contextual details, and refine preferences, ensuring that feedback remains personalized and goal-aligned.

data, generates feedback, and iteratively refines insights to support user goals. We outline a proposed workflow using a single meal upload as an example, demonstrating how feedback will dynamically adapt based on user-defined goals.

The process begins with **Input**, where users upload a meal photo. The system will reference predefined tracking goals, such as weight loss, muscle gain, or blood sugar control, to determine relevant tracking metrics. Additionally, a tracking history summary will provide behavioral trends and past meal data to enrich feedback with historical context. In the **Data Processing & Analysis** stage, the system will extract food items and identify relevant nutritional metrics from the meal photo. Metadata, such as meal time and location, will be incorporated to enhance contextual understanding. The system will then compare the meal against the user's tracking goals and history, identifying behavioral patterns, detecting overlooked habits, and recognizing long-term trends. It will also predict the meal's potential impact on broader dietary objectives. The **Feedback Generation** stage will adapt the user interface dynamically to highlight key insights. Depending on the goal, the system will emphasize different aspects—for example, flagging high-GI foods for users managing blood sugar or highlighting protein-rich meals for those focused on muscle gain. Beyond visualization, the system will provide goal-based textual explanations, detailing why certain meal components matter and suggesting actionable dietary adjustments. Finally, in the **User Interaction & Iterative Refinement** stage, users will be able to fine-tune the system's responses. This includes correcting recognition errors, adding contextual details (such as whether a meal was post-workout or eaten in a social setting), and updating preferences to refine future feedback.

We have begun to formalize our design in Figma, and have implemented stages of the prototype using ZotGPT, a version of OpenAI's GPT-4o which keeps dialogs confidential [17, 21]. We plan for our final implementation to be an iOS application, following the standard of mobile food journals as the primary device of input and reflection [5, 14, 18].

3 Future Work

Overall, our work focuses on designing an LLM-enhanced approach to goal-oriented food reflection, integrating real-time insights with long-term patterns to support a continuous reflection process. Future work will build on this design, implementing the system and conducting user studies to evaluate how LLMs support goal-oriented reflection. This workshop will provide valuable insights that will help us refine our design by identifying potential improvements, addressing limitations, supporting that the system effectively supports goal-oriented reflection in food tracking.

References

[1] Phil Adams, Mashfiqui Rabbi, Tauhidur Rahman, Mark Matthews, Amy Voida, Geri Gay, Tanzeem Choudhury, and Stephen Voida. 2014. Towards personal stress informatics: comparing minimally invasive techniques for measuring daily stress in the wild. In *Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare* (Oldenburg, Germany) *(PervasiveHealth '14)*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Brussels, BEL, 72–79. https://doi.org/10.4108/icst.pervasivehealth.2014.254959

[2] Eric P.S. Baumer. 2015. Reflective Informatics: Conceptual Dimensions for Designing Technologies of Reflection. 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, 585–594. https://doi.org/10.1145/2702123.2702234

[3] Janghee Cho, Tian Xu, Abigail Zimmermann-Niefield, and Stephen Voida. 2022. Reflection in Theory and Reflection in Practice: An Exploration of the Gaps in Reflection Support among Personal Informatics Apps. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 142, 23 pages. https://doi.org/10.1145/3491102.3501991

[4] 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. https://doi.org/10.1145/3314394

[5] 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. https://doi.org/10.1145/2702123.2702154

[6] 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. https://doi.org/10.1145/3432231

[7] Celia Framson, Alan R Kristal, Jeannette M Schenk, Alyson J Littman, Steve Zeliadt, and Denise Benitez. 2009. Development and validation of the mindful eating questionnaire. *Journal of the American dietetic Association* 109, 8 (2009), 1439–1444.

[8] Matthew Jörke, Shardul Sapkota, Lyndsea Warkenthien, Niklas Vainio, Paul Schmiedmayer, Emma Brunskill, and James Landay. 2024. Supporting Physical Activity Behavior Change with LLM-Based Conversational Agents. *arXiv preprint arXiv:2405.06061* (2024).

[9] 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. https://doi.org/10.1145/3025453.3025480

[10] 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. https://doi.org/10.1145/1753326.1753409

[11] Ian Li, Anind K. Dey, and Jodi Forlizzi. 2012. Using context to reveal factors that affect physical activity. *ACM Trans. Comput.-Hum. Interact.* 19, 1, Article 7 (May 2012), 21 pages. https://doi.org/10.1145/2147783.2147790

[12] Jiachen Li, Justin Steinberg, Xiwen Li, Akshat Choube, Bingsheng Yao, Dakuo Wang, Elizabeth Mynatt, and Varun Mishra. 2024. Vital Insight: Assisting Experts' Sensemaking Process of Multi-modal Personal Tracking Data Using Visualization and LLM. *arXiv preprint arXiv:2410.14879* (2024).

[13] 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. https://doi.org/10.1145/3479595

[14] Lucas M. Silva and Daniel A. Epstein. 2021. Investigating Preferred Food Description Practices in Digital Food Journaling. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference* (Virtual Event, USA) *(DIS '21)*. Association for Computing Machinery, New York, NY, USA, 589–605. https://doi.org/10.1145/3461778.3462145

[15] 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. https://doi.org/10.1145/2047196.2047198

[16] Elyssia Barrie H. Ong and Briane Paul V. Samson. 2024. Exploring Design Opportunities for Improved Self-motivation in Self-tracking and Health Goal Achievement. *Proc. ACM Hum.-Comput. Interact.* 8, MHCI, Article 256 (Sept. 2024), 17 pages. https://doi.org/10.1145/3676501

[17] OpenAI. 2024. GPT-4o. https://openai.com/research/gpt-4o Accessed: 2025-02-20.

[18] 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. https://doi.org/10.1145/3604256

[19] Lucas M Silva, Elizabeth A Ankrah, Yuqi Huai, and Daniel A Epstein. 2023. Exploring Opportunities for Multimodality and Multiple Devices in Food Journaling. *Proceedings of the ACM on Human-Computer Interaction* 7, MHCI (2023), 1–27.

[20] Konstantin R. Strömel, Stanislas Henry, Tim Johansson, Jasmin Niess, and Paweł 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. https://doi.org/10.1145/3613904.3642032

[21] Irvine University of California. 2025. ZotGPT. https://zotgpt.uci.edu/ Accessed: 2025-02-20.