NutriSight: Event-Triggered Calorie Tracking and Allergen Risk Detection with Smart Eyewear

Zihao Huang*
2413989@mail.nankai.edu.cn
College of Software, Nankai University
Tianjin, China

Tianyi Zhang*
2412227@mail.nankai.edu.cn
College of Software, Nankai University
Tianjin, China

Xiaoya Tang*
2413126@mail.nankai.edu.cn
College of Software, Nankai University
Tianjin, China

Ziyuan Lin
2410936@mail.nankai.edu.cn
College of Cryptology and Cyber Science, Nankai University
Tianjin, China

Huize Chao†
chz24@mails.tsinghua.edu.cn
Department of Electronic Engineering, Tsinghua University
Beijing, China

Abstract

Dietary monitoring is most useful when it captures intake as it happens, yet phone diaries ask people to remember, frame, and enter food after a decision has already been made. A formative inquiry with 215 questionnaire responses and 12 interviews found that mobile logging is often interrupted during snacks, busy work, and social meals, while first-person sensing also raises privacy concerns. We present NutriSight, a user-prioritized hybrid event-triggering system for smart eyewear. Rokid Glasses RV101 supplies a live preview, low-rate scene frames, and high-resolution JPEG captures to a phone-centered pipeline. Users can select manual, automatic, or hybrid sensing: automatic visual candidates can initiate meal capture, explicit App commands and gestures can intervene, and a fist gesture ends an eating event and suppresses pending escalation. Package and ingredient scanning remains explicitly user initiated. The phone routes confirmed captures to food recognition or package OCR, structures evidence, and presents results for confirmation rather than writing directly to a health record. A successful three-path demonstration covers automatic meal capture, manual package scanning, and gesture-prioritized hybrid control. Existing component checks of 30 food images and 36 package images provide limited support for the downstream analysis path; they are not trigger-performance results. We contribute the hybrid trigger policy, its bounded event state machine, and a concrete evaluation protocol for event quality, sensing cost, and user control.

*These authors contributed equally to this work.
†Corresponding author.

UbiComp Companion '26, Shanghai, China

ACM ISBN 979-8-4007-2533-3/2026/10
https://doi.org/10.1145/3798063.3839148

CCS Concepts

• **Human-centered computing** → **Ubiquitous and mobile computing systems and tools**; • **Applied computing** → *Health informatics.*

Keywords

smart eyewear, dietary monitoring, event-triggered sensing, eating detection, calorie tracking, allergen risk, first-person vision

ACM Reference Format:
Zihao Huang, Tianyi Zhang, Xiaoya Tang, Ziyuan Lin, and Huize Chao. 2026. NutriSight: Event-Triggered Calorie Tracking and Allergen Risk Detection with Smart Eyewear. In *Companion of the 2026 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp Companion '26), October 11–15, 2026, Shanghai, China.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3798063.3839148

1 Introduction

"The most profound technologies are those that disappear," wrote Weiser when describing computing woven into everyday life [10]. Dietary support has not yet reached that condition. Mobile food diaries reduce part of the input burden, but users still need to open an application, frame the relevant view, estimate what was consumed, and correct the record later [3, 5]. These actions are easy to skip during a brief snack, a busy workday, or a social meal. A first-person system is useful only if it reduces this activation cost without turning everyday life into retained video.

Our formative inquiry provides direct motivation for this design. We collected 215 questionnaire responses and conducted 12 semi-structured interviews about diet logging, package decisions, proactive prompts, and first-person privacy. Seventy-eight percent of respondents had tried mobile diet applications, but reported persistence of no more than three weeks. More than 85% of described interruptions occurred while snacking, being busy at work, or eating socially. At the same time, 45% expressed concern about continuous first-person recording. These findings do not establish a

Table 1: Phone diaries and user-prioritized smart-eyewear sensing.

Dimension	Phone diary	NutriSight
Activation	User opens, frames, and submits	Automatic candidate or explicit command
Best-fit moment	Retrospective review and correction	Hands-busy snacks, shared meals, label checks
Main cost	Missed or delayed entries	Energy, privacy, and false capture risk
Control	Explicit by default	Automatic sensing with gesture override

population-level behavioral effect; they identify the decision window in which a user-controlled, bounded capture mechanism may be valuable.

Smart eyewear is therefore a complement to, rather than a replacement for, phone diaries. A phone remains better for detailed correction, profile editing, and retrospective review. Eyewear offers a hands-free first-person view when the user is already looking at a meal or a package, especially for short snacks, shared meals, and pre-consumption label checks. The tradeoff is additional hardware, energy use, accidental capture, and bystander privacy. NutriSight addresses this tradeoff with selectable sensing modes and an explicit user-priority policy.

NutriSight treats monitoring as a sequence of bounded dietary events. Rokid Glasses RV101 provides Camera2 frames; the paired phone discovers the device, displays a WebRTC preview, receives low-rate JPEG scene frames, and requests high-resolution captures. The phone-side state machine supports manual, automatic, and hybrid modes. Automatic visual candidates can initiate a meal event, while App commands and hand gestures express user intent. An open-palm gesture starts an eating state and a fist ends it; explicit user control takes precedence over an automatic candidate. Package analysis remains explicitly initiated to avoid turning a passing view of a shelf into a stored product record.

The captured event is routed to food recognition or package OCR, converted into source-linked fields, and shown for review. Deterministic rules perform arithmetic and profile matching; the language layer structures or explains evidence but does not decide that a product is safe. This paper makes three contributions:

- a user-prioritized hybrid trigger policy that combines automatic visual candidates, App commands, and open-palm/fist gesture control;
- an RV101-to-phone event pipeline that bounds capture, routes meals and packages separately, and preserves reviewable evidence; and
- a three-path feasibility demonstration and evaluation protocol for event quality, sensing cost, and user control.

2 Related Work and Research Gap

Image-assisted dietary assessment has moved from manual photo review toward automated food recognition, portion estimation, and nutrient inference. Im2Calories and Nutrition5k illustrate the progress and the remaining dependence on visible geometry, reference scale, and suitable training data [5, 7]. Single-view energy estimation is particularly fragile for mixed dishes, occlusion, and missing scale [4]. These limitations matter in first-person views, where the wearer may look at only part of a shared dish and may not hold the head still.

Wearable intake monitoring addresses a different part of the problem: identifying that eating is occurring. Reviews cover acoustic, inertial, proximity, and body-worn sensing approaches [8]; FitByte illustrates the value of multimodal eyeglass sensing for intake-related activity [2]. NutriSight does not assume that such signals are available on every device. Instead, it uses the implemented RV101 visual stream to nominate low-cost candidates and reserves high-resolution capture for a bounded event. Wearable cameras provide rich context but also introduce bystander, retention, and social-acceptability concerns; work on privacy-preserved egocentric captioning shows one way to reduce downstream exposure to raw imagery [6].

NutriSight connects these two lines of work. Eating detection alone does not yield calories or allergens, while food recognition alone assumes that the correct moment and task have already been selected. Our focus is the transition between them: escalation from a low-cost cue to a bounded visual event, followed by task-specific interpretation and longitudinal recording.

3 Design Requirements

The project is organized around two research questions. **RQ1** asks when a smart-eyewear system should escalate from low-cost context sensing to image capture. **RQ2** asks how a confirmed event should be interpreted so that calorie and allergen outputs retain the meaning of their evidence source. Four requirements follow from these questions.

R1: cover the decision window without equating monitoring with recording. Package risk is useful before consumption, while a meal record is useful during or shortly after eating. The system therefore needs more than a post-meal camera button, but it should not retain all-day first-person video. Triggering and retention are separate decisions: a transient cue can start analysis without automatically becoming a stored record.

R2: preserve task-specific evidence. A meal frame, a nutrition panel, and an ingredient list answer different questions. The system must route them before inference, preserve the source crop, and avoid converting a missing field into a negative finding. A food-name candidate may support nutrition lookup; only visible label text or explicit user knowledge can support a stronger allergen statement.

Table 2: Evidence sources and the claims they can support.

Source	Supports	Does not establish
Meal frames	food candidates, visible components, coarse portion cues	exact weight or complete ingredients
Package label	per-serving energy, ingredient and allergen fields	amount actually consumed
User input	serving correction, ingredient confirmation, event boundary	independent nutrition truth
Profile	declared allergens, goals, and limits	diagnosis or inferred disease

R3: keep quantities revisable. Calorie tracking compounds errors over time. A wrong serving estimate should remain linked to its assumption so that a user can change the mass or serving count without repeating recognition. Event records therefore store inputs, derived values, and confirmation state separately.

R4: reserve interruption for consequential uncertainty. Asking about every candidate recreates the burden of manual logging. NutriSight should remain quiet when no dietary event is likely, ask for a new view when a label is unreadable, and request serving confirmation only when the value changes the retained record. The phone is suited to detailed correction; eyewear audio is reserved for short questions and time-sensitive risk.

4 Dietary Event Model and System Overview

NutriSight treats a dietary event as the smallest retained unit of monitoring. An event has an onset and end state, a trigger source, selected evidence, an interpretation path, and a confirmation state. Monitoring over a day is the aggregation of these events; continuous video is neither the analytical unit nor the intended stored artifact.

Figure 1 shows the complete RV101-to-phone path. RV101 provides first-person Camera2 capture; the paired phone manages local discovery, WebRTC preview, mode selection, gesture state, review, and routing. Low-rate JPEG scene frames support automatic meal candidates, while high-resolution JPEG captures provide evidence for analysis. Compute-intensive recognition and parsing run after capture. Deterministic code performs arithmetic, schema checks, and profile matching; a language model may structure OCR text or phrase a result, but it does not invent thresholds or declare a product safe.

5 When to Recognize: Hybrid Event Triggering

The central design decision is not whether to keep a camera on, but which signal is allowed to request a bounded capture.

Table 3: Trigger sources and control responsibilities.

Source	Primary role	Boundary
Low-rate scene frame	nominate a meal candidate	does not prove consumption
App command	request capture or select a mode	may be accidental; user can cancel
Open palm	enter eating state	does not write a health record
Fist	end or cancel eating state	suppresses pending automatic capture
Package command	mark a label for OCR	required before package retention

NutriSight exposes manual, automatic, and hybrid modes so that users can choose the desired balance between activation cost and sensing control. Table 3 separates the implemented trigger sources and their responsibilities.

5.1 Three User-Selectable Modes

In *manual mode*, the user controls capture through the phone App or an explicit gesture. This mode minimizes automatic sensing and is appropriate when the wearer wants predictable capture. In *automatic mode*, the phone receives low-rate JPEG scene frames from RV101 and evaluates them for a stable meal context. A stable candidate requests a high-resolution JPEG, after which the normal food-analysis path and user review apply. Automatic mode does not autonomously scan packages: label and ingredient capture always requires an explicit package command.

In *hybrid mode*, automatic meal candidates run in parallel with user controls. An open-palm gesture can immediately enter an eating state, while a fist ends the state and cancels pending automatic escalation. The App's immediate-capture command and mode changes also override a candidate. Thus automation can reduce the cost of remembering to start a meal record, but it cannot overrule a clear user stop or write a record without confirmation.

5.2 State Machine and Priority

Figure 2 summarizes the hybrid event flow. After selecting manual, automatic, or hybrid sensing, the user-trigger lane and automatic-vision lane run singly or in parallel and converge at candidate arbitration. The phone applies stability checks and gives explicit user actions priority before entering *Active Eating*. A high-resolution JPEG is then acquired in *Capture*, routed to recognition or package OCR, and shown in *User Review*. Package capture remains explicitly user initiated. Confirmation, correction, ignore, or deletion closes the event and returns the system to *Idle*.

The priority order is deterministic: a fist or cancel command has highest priority, followed by an App capture or mode command, an open-palm start, and finally an automatic visual candidate. This ordering makes user intent observable and testable rather than treating all triggers as equal votes. Consecutive consistent scene results stabilize

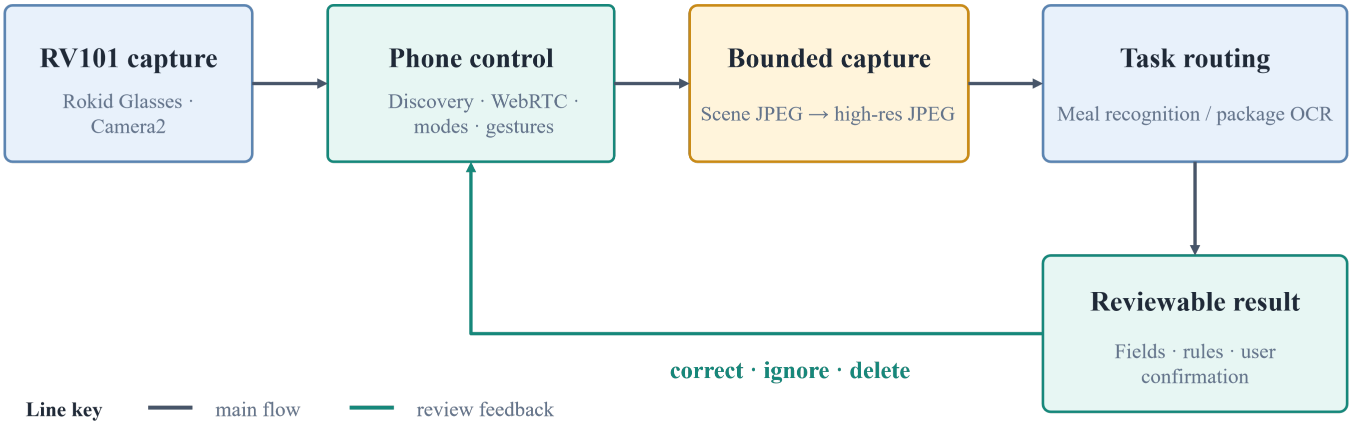

Figure 1: RV101-to-phone dietary-event path. The phone owns mode selection, gesture interpretation, and user review; the eyewear provides first-person capture and transport.

transitions, while brightness, blur, and duplicate checks reject unusable captures. A short in-memory buffer may preserve the moments immediately before confirmation; rejected candidates are not retained as dietary records [9].

5.3 Package Capture and User Intent

Package analysis is intentionally explicit. The user selects package mode or issues a package command, the phone requests a high-resolution capture, and the OCR path retains the source crop with the parsed fields. This prevents a walk past a store shelf from becoming dozens of product records. If glare, curvature, transparent folds, or small text make the ingredient field unreadable, the system requests recapture or human inspection and never emits a "no allergen" conclusion.

5.4 Key-Frame Selection

Selecting one sharp image is not enough for every event. For a meal, the selector should balance visual quality with temporal diversity: an early overview can show the plate, while a later frame may reveal a side dish or a changed serving. For a package, the useful set may contain the front, ingredient list, and nutrition panel. NutriSight therefore ranks frames by blur, exposure, duplicate similarity, and task coverage, then retains a small set of complementary views. This selection occurs before cloud upload when device resources permit. The event stores why each view was retained, allowing a later error to be traced to missing capture, poor OCR, or interpretation.

6 How to Recognize: Two Evidence Paths

6.1 Meals: From Visible Food to an Energy Range

The meal path selects representative frames and returns open-vocabulary food candidates with visible components. Multiple frames are merged at the event level so that a later side dish is added rather than treated as a new meal. Ambiguous mixed dishes remain editable candidates. Portion inference first looks for a known plate or packaged object as a scale anchor; a hand landmark can provide a weaker secondary cue. Because a monocular image does not reveal hidden depth or the amount actually eaten, the interface asks the user to confirm a serving class or approximate mass when the estimate affects the record.

For n recognized foods, the basic calculation is

$$E_{meal} = \sum_{i=1}^{n} W_i \frac{K_i}{100}, \tag{1}$$

where W_i is confirmed or estimated mass in grams and K_i is energy per 100 grams from a nutrition source. When W_i is uncertain, NutriSight stores an interval rather than a precise value. Visual appearance can raise an ingredient question, but it cannot prove the absence of an allergen in a cooked meal.

6.2 Packages: From Text to Declared Risk

The package path locates label regions, corrects orientation where possible, and applies OCR. A parser separates product name, serving size, energy, ingredient text, and declared allergen statements while retaining the source crop. Ingredient aliases then map to canonical allergens in the user's declared profile. Package energy is computed as $E_{pkg} = sK_s$, where K_s is label energy per serving and s is the consumed serving count confirmed by the user.

Package evidence is usually stronger than meal imagery for allergens, but it still has failure modes. A missing ingredient field is not evidence of safety. OCR on curved cups, glossy bags, transparent folds, and small text can omit precisely the line that matters. NutriSight therefore distinguishes *label-confirmed*, *visually suspected*, and *unknown* allergen states. Only the first supports a direct profile match; the others request inspection or recapture.

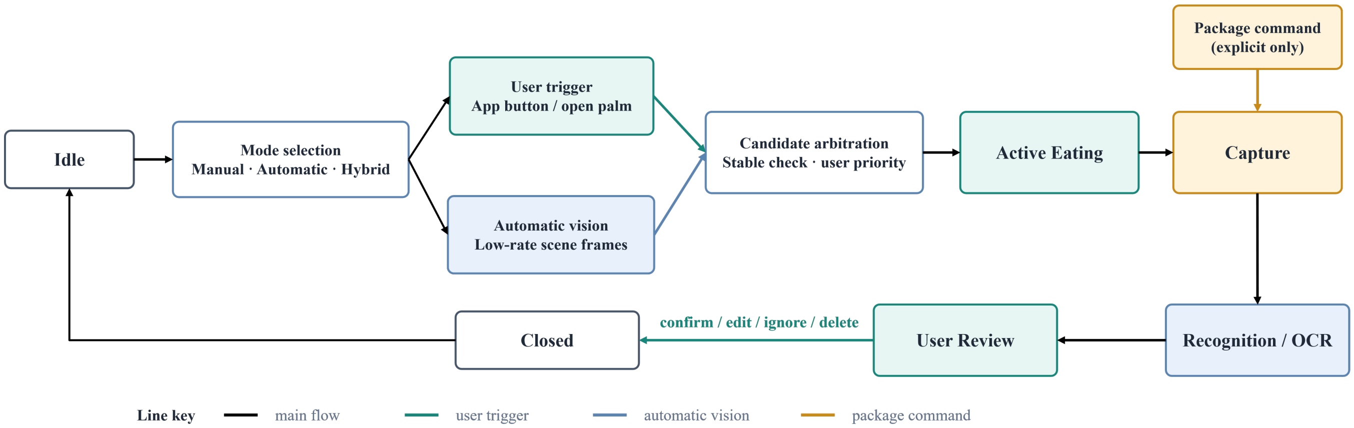

Figure 2: User-prioritized hybrid event flow. Manual, automatic, and hybrid modes enable one or both trigger lanes, which converge at candidate arbitration before the shared capture, analysis, and review path. Package capture is explicitly user initiated.

6.3 From Structured Evidence to Feedback

Both paths emit the same event schema: timestamp, trigger, source image or crop, parsed fields, confidence notes, energy value or range, allergen state, and confirmation status. Schema validation and arithmetic precede feedback. Deterministic rules accumulate daily energy and compare explicit label/profile matches. A language layer may convert the structured result into a concise sentence, but it is downstream of the calculation. Following human-AI interaction guidance, correction and dismissal are ordinary actions rather than hidden error recovery [1].

The phone is the review surface for serving corrections, history, and profile editing. Eyewear audio is reserved for short, timely messages, such as asking whether a meal event should be retained or warning that a label explicitly names a declared allergen. This division keeps detailed inspection off the small wearable interface while preserving intervention at the decision window.

7 Claim-Aware Uncertainty and Record Semantics

A single confidence score would hide where an event can fail. NutriSight instead keeps four kinds of status: whether a dietary event occurred, whether the retained views are adequate, whether the food or label fields were interpreted, and whether the consumed quantity is known. These statuses support different claims and should not be multiplied into an opaque probability. A stable scene candidate, for example, increases confidence that a capture is worth requesting but says nothing about food identity or consumption amount. Likewise, clear ingredient text can support an allergen warning even when the consumed amount remains unknown.

Table 4 turns uncertainty into an action policy. An event may be retained while one claim is withheld: a package can contribute label-derived energy after serving confirmation

Table 4: Minimum evidence and abstention behavior for user-facing claims.

Claim	Minimum evidence	If absent
Meal energy range	food candidates, portion bound, nutrition source	retain incomplete; exclude from total
Package energy	readable energy and serving fields; consumption confirmed	recapture or ask serving count
Allergen warning	explicit label text matched to declared profile	mark unknown; request inspection
"No risk" message	not supported by the present evidence model	never emit

even if its ingredient panel remains unreadable, but its allergen state stays unknown. Conversely, a clear allergen match should be delivered before consumption even if calorie fields have not been parsed. This claim-specific policy avoids an all-or-nothing record and makes the cost of missing evidence visible to the user.

Corrections propagate through the structured record rather than overwriting its history. Each derived quantity stores its source and assumptions. Changing a meal portion updates the corresponding W_i and all dependent summaries; replacing an OCR crop reruns only the fields derived from that crop. For m retained events with energy intervals $[E_j^-, E_j^+]$, the daily summary remains an interval,

$$[E_{day}^-, E_{day}^+] = \left[\sum_{j=1}^{m} E_j^-, \sum_{j=1}^{m} E_j^+\right], \quad (2)$$

until a user correction narrows one or more terms. Trigger confidence is stored as provenance but does not numerically shrink a calorie estimate; uncertainty about whether an event happened and uncertainty about its quantity are presented separately. This record semantics gives later evaluation a concrete target: not only whether a model label was correct, but whether the final event contained enough supported fields to justify each user-facing output.

Table 5: End-to-end feasibility demonstration paths.

Path	Trigger	Processing	User gate
Automatic meal	Stable low-rate scene candidate	High-resolution JPEG, food route, structured meal fields	Confirm, edit, or delete
Manual package	App command or package gesture	JPEG, OCR, ingredient/allergen fields	Confirm label and serving
Hybrid override	Automatic candidate plus open palm/fist	Gesture state overrides pending escalation	Fist cancels; user reviews

8 Event Completion and Feedback Timing

In a meal event, a stable scene candidate or open-palm gesture creates one event record rather than a sequence of independent recognitions. The phone requests a high-resolution JPEG, filters it, and routes it to food recognition; later non-duplicate candidates during the same eating state are merged into the record. If no scale anchor is visible, the phone asks for a serving class instead of accepting precise grams. A fist closes the event and cancels pending automatic escalation, so the final review concerns one traceable meal record.

For a package event, capture remains explicit. The user scans front, ingredient, or nutrition panels before eating, and quality checks can ask for another angle when glare, curvature, or small print degrades OCR. Parsed fields stay linked to their source crop. An explicit label/profile allergen match can be spoken before consumption, but energy is added only after serving count is confirmed; a scan is evidence of consideration, not consumption. If ingredient text cannot be recovered, the state remains unknown rather than "no allergen found."

Daily and weekly views aggregate confirmed or retained events while preserving provenance. Estimated meal values remain distinguishable from label-derived values, and corrections update dependent totals without discarding original evidence. This interaction pattern supports timely warnings, lightweight in-meal questions, and more detailed phone-based correction after the event.

9 Prototype Demonstration and Validation Plan

The prototype connects Rokid Glasses RV101 to a phone-centered capture and analysis pipeline. RV101 provides WebRTC preview, low-rate scene JPEGs, and requested high-resolution JPEGs; the phone handles discovery, mode and gesture routing, capture requests, and food or package result cards. Automatic, manual, and hybrid modes were exercised in an end-to-end demonstration.

This demonstration establishes functional integration across RV101 capture, phone transfer, event routing, downstream analysis, and user review. It does not estimate deployment-scale trigger precision, latency, gesture accuracy, camera duty cycle, or behavior change. Existing manual checks are therefore only compact support for the downstream analysis path: 28/30 food Top-1 matches, 27/36 package ingredient fields, and 28/36 package allergen fields.

Observed component failures define recovery behavior. Food errors were label-granularity failures on fruit assortments; package failures clustered around glare, low contrast, curved or transparent material, and dense small text. An unreadable ingredient field requests recapture or human inspection and never becomes a negative allergen finding. The event record preserves source crop, parsed fields, trigger source, confidence notes, and confirmation state so future trigger errors can be separated from OCR or interpretation errors.

The next evaluation should annotate continuous RV101 sequences with meal onset, pauses, package inspection, and confounds such as talking, walking, head turns, drinking, and non-dietary hand movements. It should compare automatic, manual, and hybrid modes using event precision/recall, false camera activations per hour, onset latency, end-time error, camera-on duty cycle, user cancellation, and confirmation rate. A second stage should measure portion and calorie error against weighed foods, then report event completeness and missed allergen alerts. Participant- and meal-level splits should avoid sequence leakage; logs should include camera-on time, uploaded bytes, prompts, and user corrections.

10 Conclusion

NutriSight treats smart-eyewear dietary monitoring as a user-prioritized event-triggering problem. Its RV101-to-phone pipeline supports automatic meal candidates, explicit package capture, and gesture or App control within one bounded state machine. Automatic sensing lowers the activation cost for high-frequency diet logging, while fist termination, explicit confirmation, and claim-aware abstention keep automation from silently creating health records. The present evidence demonstrates the integrated paths and downstream recognition components; the proposed evaluation protocol now makes trigger quality, sensing cost, and user control measurable without presenting them as already established clinical or behavioral outcomes.

References

[1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13. doi:10.1145/3290605.3300233

[2] Abdelkareem Bedri, Diana Li, Rushil Khurana, Kunal Bhuwalka, and Mayank Goel. 2020. FitByte: Automatic Diet Monitoring in Unconstrained Situations Using Multimodal Sensing on Eyeglasses. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12. doi:10.1145/3313831.3376869

[3] Carol J. Boushey, Mark Spoden, Frank M. Delp, Quing Zhu, Robert Bosch, E. J. Ahmad, Y. Shvetsov, B. DeLany, D. A. Kerr, and E. J. McClung. 2017. New Mobile Methods for Dietary Assessment. *Proceedings of the Nutrition Society* 76, 3 (2017), 283–294.

[4] Shaojie Fang, Zhe Shao, Ruoyu Mao, Chang Fu, Edward J. Delp, and Fengqing Zhu. 2018. Single-view Food Portion Estimation: Learning Image-to-Energy Mappings Using Generative Adversarial Networks. In *Proceedings of the IEEE International Conference on Image Processing*. 251–255.

[5] Austin Meyers, Nick Johnston, Vivek Rathod, Anoop Korattikara, Alex Gorban, Nathan Silberman, Sergio Guadarrama, George Papandreou, Jonathan Huang, and Kevin P. Murphy. 2015. Im2Calories: Towards an Automated Mobile Vision Food Diary. In *Proceedings of the IEEE International Conference on Computer Vision*. 1233–1241.

[6] Jianing Qiu, Frank P.-W. Lo, Xiao Gu, Modou L. Jobarteh, Wenyan Jia, Tom Baranowski, Matilda Steiner-Asiedu, Alex K. Anderson, Megan A. McCrory, Edward Sazonov, Mingui Sun, Gary Frost, and Benny Lo. 2023. Egocentric Image Captioning for Privacy-Preserved Passive Dietary Intake Monitoring. *IEEE Transactions on Cybernetics* (2023). doi:10.1109/TCYB.2023.3243999

[7] Quinn Thames, Arjun Karpur, William Norris, Fan Xia, Livia Panait, Tobias Weyand, and John Sim. 2021. Nutrition5k: Towards Automatic Nutritional Understanding of Generic Food. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8903–8911.

[8] Tri Vu, Feng Lin, Nabil Alshurafa, and Wenyao Xu. 2017. Wearable Food Intake Monitoring Technologies: A Comprehensive Review. *Computers* 6, 1 (2017), 4. doi:10.3390/computers6010004

[9] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. doi:10.1109/TIP.2003.819861

[10] Mark Weiser. 1991. The Computer for the 21st Century. *Scientific American* 265, 3 (1991), 94–104. doi:10.1038/scientificamerican0991-94