

ORIGINAL RESEARCH

Open access

Smart Food Expiration Monitoring Through OCR and Cloud-Integrated Image Recognition

João Silva^{1*}, Pedro Costa¹, Ana Beatriz²

Abstract

Food waste is a critical concern for both developing and advanced economies alike. Digital technologies hold promise for helping to alleviate this problem, and our investigation sets out to deliver a workable prototype employing cloud-based image classification alongside tailored OCR techniques. The product of our work is an economical yet high-performing mobile app prototype that establishes a foundation for future inquiry. We drew on sophisticated principles of application integration—encompassing mobile frameworks, Firebase machine learning modules, and OCR procedures—to signal how near food items are to spoiling. Where no printed date exists, the system generates an indicative shelf-life projection built on prudent category-based assumptions. These projections do not constitute safety determinations and are not a substitute for manufacturer date markings or national food-safety regulations. The outcomes lend our paper unmistakable originality and scholarly value, boosting the operational efficiency of storage facilities and retail food outlets. The ramifications of our work are technical, economic, and societal in nature.

Keywords On-Cloud Image Classification, OCR and machine learning, Food expiration date determination, Systems' integration, Firebase machine learning

*Correspondence:

João Silva
joao.silva@ufrj.br

¹ Department of Food Safety and Shelf-Life Monitoring, Faculty of Agriculture, Federal University of Rio de Janeiro, Rio de Janeiro, Brazil

² Department of OCR-Enabled Smart Labeling, Faculty of Agricultural Sciences, University of Campinas, Campinas, Brazil

Introduction

Food waste stands as a formidable sustainability issue across Europe, carrying environmental, social, and economic repercussions. Data issued by Eurostat in 2024 for the 2022 measurement period [1] places the overall mass of food waste generated along the European Union's entire supply network at roughly 59.2 million tonnes. Breaking this down, it amounts to about 132 kg of discarded food for each inhabitant of the EU, a sobering statistic that reveals the severity of the predicament. Private households are singled out as the predominant producers, generating 54% of all waste, which corresponds roughly to 72 kg of food thrown away per person. It follows logically that confronting this challenge can originate at the personal level, where individuals possess the capacity to implement meaningful shifts in everyday buying and preservation routines. Prompted by intensifying anxieties over food waste, several digital interventions have surfaced to encourage more sustainable patterns of consumption. Too Good To Go [2] and Munch's Bonapp [3] represent exchange-oriented ecosystems that team up with eateries, bakeries, and grocery outlets to put forward steep price reductions on goods approaching their use-by dates, thus facilitating eleventh-hour salvaging transactions. Olio, by contrast, advances a person-to-person sharing framework, permitting citizens and neighborhood businesses to pass along surplus provisions to those nearby at zero charge, thereby consolidating locality-focused redistribution [4]. NoWaste

pursues yet another route by centering on domestic stock control, enabling individuals to catalogue foodstuffs and obtain automated nudges before they spoil [5].

While these instruments contribute valuably to consciousness-raising and habit transformation, a great many of them are still heavily contingent on hand-entered data or affiliate cooperation. More to the point, they zero in overwhelmingly on pre-marked spoilage dates without the faculty to visually discern the food article itself or to forecast a usability timeframe should the expiry marking be gone or indecipherable. Hence, a persistent requirement remains for a consolidated, consumer-oriented remedy that can autonomously pinpoint food items, pull out any legible dates, and extrapolate a defensible consumption interval when labels cannot be read.

Within the scholarly corpus, multiple examinations and inquiries have addressed food waste governance, with earlier outcomes directed at shopper education [6], statutory and regulatory measures [7], ecological footprint [8], policy angles [9], particular facilitators and hindrances [10], unified methodologies [11], and fiscal fallout [12].

Numerous contemporary food administration apps deploy optical character recognition (OCR) engines to decipher expiry markings printed on merchandise wrappings. Although this tactic functions adequately for crisp, unobstructed labels, it frequently proves deficient under authentic conditions where use-by dates may be abraded, blotched, bleached, or fragmentary owing to inferior printing or rough treatment [13]. Certain edible goods—fresh produce and baked items, for example—routinely carry no date stamps at all, most notably when sold unwrapped or destined for swift intake. Under these circumstances, visual determination of the product rather than formal textual decoding must take precedence [14].

Today's technological landscape stays disjointed: portable OCR can lift expiration dates solely when they remain whole, whereas slimmed-down CNNs (Convolutional Neural Networks) can categorize unlabeled provisions [15]. Still, barely any user-facing platforms fuse these dual faculties to approximate a believable usage window whenever a date is absent or illegible. Even though frameworks like Firebase ML Kit already bundle cloud-based text recognition and image-tagging interfaces, neither scholarly proofs of concept nor business apps have put to the test an all-encompassing pipeline that pinpoints the product, cross-checks any captured data, and yields a consumption interval live before dispatching notifications to a cloud-mirrored stock list. Our investigation thus endeavors to bridge a vital technological deficit by weaving these pieces into a seamless, handset-compatible sequence.

The current inquiry sets out to architect and empirically field-test an Android application that cuts down the manual drudgery tied to residential food surveillance. The suggested scheme autonomously registers goods through lean, cloud-driven image identification, harvests any perceptible expiration dates via the OCR facility built into Firebase ML Kit, and interprets a plausible storage life in instances devoid of a printed date by leaning on class-specific heuristics housed in a local information repository. It likewise pushes instant alerts while keeping the stock database mirrored in the cloud via Firebase Firestore and Cloud Messaging. The central ambition is to furnish a smartphone-optimized flow that supplies on-the-spot insight, keeps user effort to a minimum, and bears the potential to support domestic food-waste abatement by sharpening attentiveness and prompt usage. The contribution delivered here is architectural rather than algorithmic. We put forward a consolidated mobile pipeline that marries cloud-driven OCR, cloud-driven image annotation, and a regulation-guided shelf-life substitute to trim down human input without compromising real-time swiftness. When the apparatus guesstimates an expiry point in the absence of a scannable marking, the output should be regarded as a directional indicator. The true safe-to-eat horizon hinges on the user's keeping conditions, thermal oversight, packaging soundness, and manufacturer-specific makeups. Owing to the broad heterogeneity in packing conventions and producer protocols, we do not assert that the conjectured date is exact or uniformly applicable.

To ascertain the practicability of our mobile design, the present inquiry tackles five investigative questions: (Question 1) With what precision can a lightweight, cloud-reliant image-classification engine distinguish comestibles inside naturalistic, user-shot snapshots? (Question 2) How proficient is Firebase ML Kit at fishing out spoilage dates from wrappers under assorted realistic circumstances such as glare, curved surfaces, and scuffed typography? (Question 3) With what dependability can the framework extrapolate a reasonable keeping period when no printed stamp comes to light,

employing the categorically identified food type as a stand-in? (Question 4) What is the end-to-end efficacy of the holistic approach, taking in lag, resource draw, and felt user-friendliness on regular mid-tier handsets running contemporaneously? (Question 5) To what extent is an automated scheme capable of upholding higher user involvement and exhibiting the potential to spur more planet-friendly food-utilization customs in comparison with hand-operated or bar-coded solutions?

The balance of this paper unfolds as follows: the subsequent section surveys allied scholarship and constructs the investigatory hypotheses; the following section explicates the chosen research method, spanning data aggregation, model architecture, and assessment criteria; the next section lays out and deliberates on the experimental outcomes against antecedent investigations, spotlights the technical merits, and sketches conceivable societal and monetary repercussions. Finally, the concluding section wraps up with pivotal discoveries, constraints, and paths forward for subsequent exploration, anchored in the study findings yielded by this scientific undertaking.

Related Work

The shift toward digital tools has proven to be an exceptionally powerful means of reducing food waste within households, and mobile applications, currently the most ubiquitous digital interface, dominate interventions aimed at individuals by enabling surplus provisions to be shared person-to-person, offering time-sensitive discounts on items close to spoiling, and allowing on-phone capture of use-by dates for seamless kitchen stock management [16].

Mobile food-waste applications

Within this framing, the present subsection surveys applications already in circulation that target domestic food discards. It delves into machine-learning approaches—spanning on-device OCR, streamlined CNNs, and Firebase ML Kit APIs—that make possible a smooth, entirely hands-off food-logging sequence.

To ground the technology-oriented analysis, **Table 1** draws together the hallmark properties of four mobile offerings situated at the leading edge of residential food-waste curtailment: Too Good To Go [2], Olio [4], Bonapp by Munch [3], and NoWaste [5]. The parallel breakdown isolates each solution's foundational waste-minimization rationale, dominant user-journey design, and supporting mobile mechanisms, thereby casting light on both the heterogeneity of present-day digital formulas and the common footing in core handset features that our own OCR and CNN-based proposition means to build upon.

Table 1. Comparative features of mobile applications for household-level food-waste reduction.

Mobile applications	Waste-reduction method	Technologies employed (mobile)
Too good to go	A discount platform that offers “Surprise Bags” of extra food from nearby cafés, bakeries, and grocery stores at lower prices.	Uses map-based geolocation to browse nearby deals, an in-app reservation and payment system, push notifications, and an impact-tracking dashboard [17].
Olio	A community-based sharing app where people (and volunteers) exchange surplus food and household goods within local areas.	Supports photo-based listings using a smartphone camera, in-app chat for arranging pickups, distance-based filtering, and location-triggered alerts/notifications [18].
Bonapp by Munch	A regional discount marketplace providing unsold food items at 50–70% reduced prices shortly before expiration.	Includes geo-tagged discount offers, in-app purchasing, and scheduled pickup time slots [19].

NoWaste	A home food management app designed to reduce waste by tracking food items and sending expiration reminders.	Features barcode scanning for quick item entry, cloud synchronization across devices, expiration-date organization, and push notifications [5].
----------------	--	---

Source: Author's compilation based on the references listed below.

Every platform intervenes at a different junction of the food-loss chain, calling upon specific handset capacities such as location services linked with in-app checkout, direct messaging between users, or barcode and OCR decoding. The contrastive overview reveals where incumbent tools deliver real utility and where vital competencies remain unrealized, particularly the automatic gleaning of stamped expiry markings and CNN-powered discernment of loose, unpackaged edibles, reinforcing the case for the unified architecture put forward in this investigation.

Looking beyond third-party marketplaces, leading grocery conglomerates throughout Europe have also enmeshed food waste mitigation protocols into their branded mobile storefronts. Outfits including Lidl [20], Kaufland [21], Carrefour [22], and PENNY [23] deploy their apps to issue digital rebate coupons, inform patrons of "final-hour" markdowns on goods nearing expiration, and give entry to surplus lots via mechanized price-reduction infrastructures. In several scenarios, such as Carrefour and PENNY, these measures are amplified through tie-ups with specialist services like Too Good To Go or Bonapp.eco, which redistribute unsold merchandise through instantaneous, geo-tagged push alerts [24, 25]. These amalgamated strategies rest on mobile technologies—barcode lookup, loyalty scheme utilization, and stock-sensitive pricing logic—to sharpen inventory turnover and pare away wastage directly at the point of sale.

Technologies used in food-waste reduction applications

Contemporary mobile platforms engineered to curb food spoilage marshal an array of sophisticated techniques to mechanize pantry surveillance and minimize decay. These tools habitually embed image-driven recognition and code-deciphering utilities so that consumers can record food goods with negligible exertion merely by snapping photographs or scanning barcodes.

Such innovations simplify the task of keeping tabs on a dwelling's edible reserves and dispatch preemptive warnings right before short-lived items go bad. By swapping tedious hand-entry workflows with self-acting detection routines, these apps aid in ensuring that provisions are consumed in a timely fashion, drastically cutting the probability that forgotten supplies end up trashed [26].

Optical Character Recognition (OCR) methods are acquiring ever-greater weight in mobile apps devoted to food waste abatement because of their knack for automating the retrieval of pivotal particulars from merchandise wrappers, most crucially use-by stamps. By lessening dependence on keystroke input, OCR notably reduces operator fatigue while sharpening the precision and rapidity of household stock-keeping [27]. OCR refers to the transmutation of pixel-based text depictions into character streams readable by machines through a multi-step sequence spanning image acquisition, preparatory cleanup (such as thresholding, speckle elimination, and deskewing), singling out of distinct letterforms, characteristic mapping, and concluding symbol assignment [28].

Handset-anchored remedies fold in OCR engines like Tesseract or Google ML Kit Text Recognition to comb product labels for expiration dates—an operation capable of swiftly lifting temporal data even from tiny or washed-out print and refreshing a user's kitchen ledger with barely any handholding [27].

Research-grade demonstrators have posted extraordinarily high correctness rates for self-acting expiry-date interpretation, having topped 99% recognition fidelity with a purpose-built OCR pipeline (drawing on image preparation and deep learning) targeting spoilage stamps on beverage containers [29], and a separate inquiry reported around 97.8% exactness by refining an OCR architecture with a convolutional neural network attention block [27]. These indicators suggest that OCR, above all when paired with deep learning, can constitute a trustworthy instrument for pinpointing expiration dates.

An operational caveat is that OCR exactness falls off when picture fidelity is undermined by weak lighting, meagre contrast, or smeared lettering—states that impede trustworthy text fragmentation and classification [30]. Follow-on experimental scholarship verifies that directed preprocessing, like adaptive thresholding and contrast amplification, can partially offset these faults, yet a residue of susceptibility lingers, especially in casual images snapped on smartphones [31]. All the same, baking OCR into food-waste-reduction applications stays highly advantageous: the self-acting logging of spoilage dates keeps domestic larders current and fires off well-timed prompts that urge occupants to finish goods before decomposition sets in, all while bypassing the drudgery of hand-typed entries [32].

Where OCR paves the way for text-anchored harvesting of expiry markers, it sidesteps the puzzle of spotting unpackaged or label-devoid provisions—think fresh produce or leftover meals. For this niche, visual taxonomy via Convolutional Neural Networks (CNNs) presents a counterpart remedy, adept at singling out food classes straight from imagery [33].

A convolutional neural network (CNN) is a deep feed-forward construct that takes an incoming picture as a multi-channel tensor and subjects it to successive rounds of learned convolutions, non-linear activation mappings, and spatial condensation operations, consequently manufacturing layered feature representations that evolve from rudimentary boundary sensors to sophisticated whole-object encodings [34]. Up-to-date assessments attest that mobile-adapted CNNs reach top-1 accuracy marks north of 90% on granular food-recognition test suites, even under erratic illumination and cluttered backdrops [35].

The tandem of OCR for text harvesting and CNNs for image taxonomy now operates briskly on handsets thanks to hardware acceleration and stripped-down model configurations. Notwithstanding such progress, programmers often wrestle with disjointed tooling ecosystems and daunting onboarding demands when attempting to unite disparate vision modules in live settings.

Firestore ML Kit supplies an all-in-one scaffold that caters to both OCR and CNN-grounded operations. Through its local and cloud-resident components, it facilitates real-time character decoding and image tagging for object naming, all while calibrating execution for mobile constraints. Firestore ML Kit is a multi-surface mobile SDK crafted by Google that gathers svelte [36], ready-made computer-vision and natural-language frameworks together with a TensorFlow Lite distribution pipe and over-the-air model-refresh protocol, thus letting creators embed deep-learning deduction that executes offline with sub-millisecond turnaround and sparse source code inside Android and iOS environments [37].

These levers can energize central app modules for orchestrating residential food caches, tracking expirations, and flagging unwrapped edibles. Recent literature stresses that such capabilities (e.g., automated decay-date nudges and stock surveillance) are indispensable for reining in household-level food squandering [38].

Firestore ML Kit's Text Recognition API permits on-the-fly pulling of printed words from snapshots right on portable gadgets, guided by a deep learning-orchestrated workflow incorporating image conditioning, glyph separation, and language-attuned deciphering. Standing apart from earlier OCR packages like Tesseract, ML Kit is optimized for handheld ecosystems and enables effortless model upgrades through Firestore, doing away with full-application redeployment needs and rendering it more nimble and thrifty with resources [39].

A further salient innovation is Firestore ML Kit's image labeling API, which makes possible on-device reasoning through an all-purpose model that distinguishes over 400 classes, letting programs independently sort a photograph—say, singling out edible items like apples or carrots—without requiring supplementary cues. This faculty proves especially handy for flagging unwrapped merchandise instantaneously [40]. Firestore also surfaces a cloud-based image labeling pathway (resting on Google Cloud Vision) with wider reach: encompassing 10,000+ categories, per-category Knowledge Graph entity markers for language-neutral disambiguation, and a metered price scheme with a capped gratis band (first 1000 queries/month). This cloud avenue fits cases where broader category breadth or canonical entity references are needed, while taking on the compromises of round-trip lag and per-invocation billing [41].

One more fundamental building block is barcode reading, which lets individuals catalogue packaged groceries by cracking UPC, EAN, or QR patterns right off the merchandise. ML Kit's barcode scanning API handles on-device, real-time decoding of these symbologies and eases automatic pulling of product data such as stock-keeping unit codes or article names [42].

Collectively, these consolidated perceptual technologies—script decoding, object labeling, and symbology reading—make Firebase ML Kit a robust foundation for crafting sophisticated mobile experiences. Encapsulating complex neural inference workloads behind straightforward API surfaces permits engineering teams to devote attention to shaping compelling interaction designs rather than grappling with raw vision-processing plumbing. This convergence demonstrates particular utility in solutions centered on food waste mitigation, where self-operating record-keeping, categorical sorting, and spoilage monitoring stand as prerequisites for catalyzing habit change within households.

Building upon the deficiencies noted across the present-day milieu of mobile interventions tackling domestic food surplus, this inquiry puts forward the construction of a unified handheld platform that exploits the synergistic pairing of Optical Character Recognition (OCR) and Convolutional Neural Networks (CNNs), delivered through Firebase ML Kit. Although multiple incumbent products depend on disjointed rollouts of either barcode lookup, manually typed entries, or narrow OCR feature sets, virtually none supply a seamless ecosystem that handles the entirety of food stock administration in a fashion that marries simplicity with operational effectiveness.

The principal research proposition animating this endeavor contends that crafting a mobile tool interweaving OCR-grounded text harvesting and CNN-driven visual taxonomy through Firebase ML Kit will materially elevate the autonomy, fidelity, and accessibility of domestic pantry tracking when contrasted against currently deployed mobile utilities that confront these same functional areas only in incomplete or patchwork forms.

This supposition presumes that housing both OCR and CNN faculties within a unified environment will empower end users to register hermetically sealed goods alongside bulk, unwrapped provisions with scarcely any deliberate effort. OCR eases the autonomous pulling of decay timestamps and article descriptors from packaging surfaces, whereas CNN-fueled picture sorting simultaneously unlocks the identification of fresh produce, baked items, and stored leftovers that habitually carry no scannable identifiers or printed markings whatsoever. Firebase ML Kit steps in as the architectural catalyst, offering developers a harmonized substrate that hosts both streams of machine intelligence within an optimized, portable-device-ready enclosure, assuring seamless throughput and near-instantaneous inferencing executed entirely on the handset.

The amalgamation of these technical strands is projected to cut back keyboard intervention, strengthen the trustworthiness of consumable logging, and yield quicker, sharper warnings concerning impending spoilage events. This dynamic, in turn, harbors the capacity to arm consumers with the insight needed to reach more considered judgments around edible usage and to pare back avoidable domestic discarding. The proposition will anchor the conceptual bedrock for the architecture, the build phase, and the subsequent end-user assessment cycles.

Materials and Methods

The methodology maps out the procedural and engineering scaffolding brought to bear during the conception and realization of the mobile software. It records the interaction scenarios, the constituent technology palette, the overarching structural scheme, and the decisional branching that buttressed the UI/UX configuration and brought into relief the consonance between the functional demands and the enacted remedies. This back-and-forth trajectory shepherded the shift from rough mockups to the finished deliverable, guaranteeing room for growth, solidity, and a steady experiential thread.

Use cases

To press onward through the next developmental increments, we spelled out precisely what the application needed to deliver as perceived through the eyes of its intended operators. Rendering the use cases explicit staked out the behavioral perimeter and the yardsticks for conformance, consequently laying the substratum for the code architecture and for the sculpting of the interface and interaction paradigm. The software encompasses six primary use cases, as illustrated in **Figure 1**.



Figure 1. Use Case Diagram for the Expiry Application (Source: Author's own elaboration (processed diagram based on the defined use cases)).

- Account creation and sign-in: Operators instantiate profiles, prove their identity through email–password pairings, uphold guarded interactive sessions, terminate access, and reestablish forgotten credentials.
- Add product by image capture and automatic identification: The individual frames a snapshot of the merchandise; thereupon, the engine classifies the article and plucks the consumption deadline from its encasing. Whenever the timestamp proves absent or indecipherable, the engine extrapolates a tentative estimate leaning on the article's taxonomic grouping. The individual inspects and potentially adjusts the proposed fields, and after ratification, the engine commits a canonical entry bearing designation, category, deadline, and derived count of days left.

- Add product through manual entry: The individual types out a title, designates a category and unit count, and picks a termination date by way of a calendar selector. Every supplied datum undergoes verification and is persisted within a standardized data mold aligned with the camera-based intake avenue.
- Inventory with "time to expiration": The program surfaces the live holdings and renders a day-level countdown tied to each line item. It permits arrangement by imminence of spoilage and exposes tags such as "near expiration" to streamline prioritization of what gets consumed next.
- Search and filtering: The holdings register supports constriction and reordering along category and title axes. Unstructured textual search is made available. The roster recomputes dynamically, allowing nimble traversal even amid voluminous stockpiles [43].
- Group messaging: Dwellers under one roof can exchange communiqués inside purpose-built conversational spaces backed by permanent archival and temporally sorted threading, enabling orchestration without ever leaving the software environment [44].

These six use cases deliver an exhaustive portrayal of operator intentions and externally observable effects. They demarcate the edge of system responsibility, disentangle the obligations of every actor category, and transmute operator expectations into confirmable acceptance touchstones. They further cement lineage connecting requirements all the way through structural modules, persistent storage layouts, and evaluation vignettes.

Application flowchart

This part disentangles the start-to-finish workflows that realize the six use cases in practice. Every workflow names its initiating spark, the ordered set of undertakings, the payload deposited into durable storage, and the externally perceptible manifestations within the interface and experiential layer. The workflows are articulated at a platform-agnostic pitch so as to stay broadly applicable while remaining auditable according to the acceptance touchstones.

The identity lifecycle initiates the moment an operator fires up the software and elects enrollment or credential entry. Throughout enrollment, the newcomer furnishes personalized attributes; the engine cross-checks them, spawns the identity entry, and configures a shielded interactive session. During credential entry, the operator proffers authentication tokens; the engine corroborates them, rehydrates the interactive context, and populates the personal dossier. Every one of these dimensions is made visible in **Figure 2**.




Figure 2. Application flowchart for the expiry management system (Source: Authors' own elaboration, based on the defined use cases).

Commodity intake follows two forking avenues that ultimately join upon a singular representational format. Along the photographic avenue, the operator triggers a scan, frames one shot, the engine determines what the commodity is, and surfaces or approximates the spoilage cutoff, whereupon the operator reviews and optionally fine-tunes the proposed entries. Along the hand-driven avenue, the operator keys in title, taxonomic bucket, count, and consumption cutoff. Under both sets of circumstances, the platform scrutinizes the supplied entries, inscribes a harmonized line item accompanied by a derived interval until decay, and propagates the change to the stock interface.

Stock tabulation executes perpetually for the authenticated operator or residential collective. The engine computes the residual lifespan for every tracked article, attaches state classifiers governed by preset bust points, and depicts the lineup complete with countdown indicators, graphic cues, and a sort order keyed to spoilage imminence.

Search and filtering empower narrowing of focus; the operator activates category or designation sifters and optionally supplies an open-ended query text. The platform tests these constraints against the live stock set, projects a narrowed and sequenced subset, and reflows the view without perceptible delay; rescinding the sieve controls reinstates the unfiltered catalog.

Cooperative dynamics hinge upon shared messaging and timely prods. Operators enter residential conversational rooms, author notes, and encounter temporally sequenced dialog chains that stay navigable from the stock vantage point. Simultaneously, the engine gauges items against notification schedules, pushes alerts bearing a link to the relevant entry particulars, and journals each dispatched nudge to circumvent repeats.

Informational fidelity and cross-view coherence are safeguarded at every stage; every mutation undergoes validation ahead of durable commitment, screen representations stay locked to the backend truth across gadget boundaries, and a forensic trail logs the acting party, temporal marker, and mutation class. Colliding alterations are unwound via a predetermined reconciliation rule, and the settled condition is surfaced to the affected party.

In summation, these workflows tether operator ambitions to deterministic state transformations and externally legible responses, certifying stable derivation of inferred properties and dependable cooperative functionality, all while furnishing a base for verification, observability, and forward extensibility.

User interface and user experience design in figma

Built upon the previously codified interaction scenarios and the full-chain operational blueprint, an exhaustive screen-level and experiential specification was produced within the Figma design environment [45]. The driving goal consisted of reinterpreting functional demands into unambiguous views, uniform manipulation idioms, and a modular stylistic toolkit, thereby enabling the engineering stage to advance against reduced interpretive drift and concrete ergonomic benchmarks.

The visual language settles on an understated, content-led atmosphere deploying gentle layered panels, generously curved geometries, and a muted tonal spectrum within which a signature green shade carries the brand's visual identity while semantically tinted tags broadcast the urgency tied to decaying provisions through graduated colour shifts. The ordering of information weights the immediate job and what demands attention next heaviest, whereas supporting particulars materialize via stepwise exposition so as to hold perceptual strain at a minimum.

Navigation is structured around a persistent lower tab strip that exposes the software's primary zones, centred on a landing hub that blends situation summaries, one-tap shortcuts, and the entries of greatest concern. Data-entry constructs rely on live field checking, transparent labelling, and situation-specific prompts, while operations with destructive ramifications mandate unambiguous ratification. Reactions from the platform arrive as compact transient banners, overlay dialogues, and condition indicators, with subtle motion flourishes that confirm actions without intrusive animation, as reflected in **Figure 3**.




Figure 3. Representative application screens designed in Figma (Source: Author's own elaboration; mockups derived from the use cases and the application flowchart).

The Figma artefact draws together the interaction scenarios and the operational flowchart into a unified screen ensemble possessing visual constancy, broad accessibility, and straightforward extensibility. The design-token underpinning and interchangeable widget catalogue promise ocular and interactive steadiness from one view to the next, while the clickable walkthroughs manifest well-defined goal paths spanning initiation to resolution. By fixing precedence hierarchies, system acknowledgements, and conditional presentations in meticulous detail, this reference work prunes ambiguity during coding and shortens revision cycles amid validation rounds. What emerges is a steady launchpad for construction, ergonomic appraisal, and prospective extensions that preserve a verifiable thread linking original stipulations to delivered functionalities.

System architecture

The platform follows an end-user-device-centric, function-backed, occurrence-activated structural pattern, as diagrammed in **Figure 4**. Application logic is housed on the Android front-end, expressed in Java and bound to the Firebase Android SDK. Repository duties, credential oversight, pictorial processing, and push-based dispatches fall to maintaining Firebase and Google Cloud service surfaces [46].




Figure 4. System architecture and end-to-end data flow (Source: Authors' own elaboration).

The complete processing thread ignites when an operator proves their identity through Firebase Authentication [47], after which point the Android shell captures a photographic frame and transfers it into Cloud Storage. The object-finalize callback launches the `analyzeProductImages` Cloud Function, which in turn queries Cloud Vision, obtains categorical annotations tagged with certainty ratings together with OCR-derived character sequences, and undertakes standardization together with contextual augmentation.

The routine subsequently commits a normalized merchandise entry into Firestore, packing the standardized label, taxonomic bucket, unit count, spoilage timestamp or taxonomy-informed approximation, calculated residual days, condition marker, and origin trace. The front-end upholds an open onSnapshot listener, meaning Firestore streams back a live mutation that portrays the classified item along with its temporal tally inside the stock panel. Whenever threshold crossings or matching events materialize—for instance, a near-decay state or a freshly posted message—the service tier pushes a Firebase Cloud Messaging alert toward the individual's connected endpoints.

To synthesize, the advanced managed-service architecture concentrates identification and validation within provisioned cloud utilities while preserving an Android surface that stays lean and perceptually fast. The structure supports elastic demand spikes, homogeneous inference quality irrespective of hardware tier, and near-instantaneous perceptual closure via Firestore listeners and FCM pathways. It likewise remains serviceable and future-amenable, considering that supplementary detectors, policy constraints, or notification rhythms can be fielded within Cloud Functions absent any redeployment of the mobile package.

Application implementation

The software takes shape inside Android Studio, powered by Java, and matches a native shell with a managed serverless Firebase spine. Firebase Authentication steers profile creation and session verification, delivering hardened session endurance, credential restoration, and token-gated ingress that meshes with Firestore protection policies and FCM routing. Visual ingestion taps Firebase ML Kit for swift merchandise identification alongside Text Recognition (OCR) suited to extracting spoilage stamps; where a stamp is absent or impossible to resolve, the program derives one from the article's taxonomic grouping.

The flow path also supplies hand-keyed item registration. Long-lived mirroring and live coherence across the stock ledger and communal discussion feeds rest upon Cloud Firestore, which furnishes sieving and rearrangement along group, moniker, and time-to-expiry vectors through indexed fetch operations and live subscription streams. Firebase Cloud Functions assume server-side reasoning, request body vetting alongside schema normalization, residual-day calculation, reminder queuing, and the fan-out dissemination of group exchanges, whereas Firebase Cloud Messaging issues device alerts flagging goods approaching expiry and new chat contributions. This fusion of on-device inferencing power and an outsourced backend tracks the design intent and skirts any requirement to provision dedicated host machines.

In the paragraphs ahead, we unveil the source-level incarnation of the case where an expiry marker proves absent or indecipherable, and the program forecasts a defensible date extracted from the product's categorical membership. The backend falls back from OCR interpretation onto a regulation-driven category-to-storage-duration lookup and then computes daysLeft plus a status designation before writing a canonicalized entry to Firestore. The Android front-end sustains a live tie and portrays the enriched record as soon as it lands. The subsequent code excerpts capture the definitive fragments across Cloud Functions (JavaScript) and Android (Java) that enact this logic.

First, we expose the event-triggered handler that absorbs a freshly deposited image, calls upon Google Cloud Vision for label identification and textual analysis, executes standardization, and commits canonicalized entries to Firestore.

The handler reacts upon a Storage completion occurrence, retrieves the photograph, invokes labels, physical objects, and OCR from Cloud Vision, standardizes the candidate merchandise label, surfaces or approximates the spoilage marker, calculates daysLeft together with condition indicator, and upserts a normalized entry at `users/{uid}/items/{imgId}`. Next, we unpack the OCR decoding module that searches for domain-customized patterns, reconciles heterogeneous date notations into an ISO instant, and filters out implausible results.

The method aligns with packaging norms by merging numeric templates alongside signal words such as "best before" and "use by," subsequently screening candidates through a requirement for a decodable value situated in the future. The pipeline applies lowercase folding and token canonicalization, traverses multiple date schemas, resolves ambiguous

representations by preferring explicit four-digit years when present, and discards non-credible or historically dated readings.

The backup mapping links food groupings to representative keeping-time spans, supplies a fallback spoilage point, produces daysLeft together with a condition marker, and thereby enables consistent downstream display and alert logic.

The lookup encodes sensible priors for domestic provisions and generates deterministic output when OCR fails to recover an explicit date. The calculated daysLeft and condition markers regularize interface sequencing and notification thresholds across endpoints. The category-informed mapping provides orientational figures exclusively. The application consistently favours an OCR-recovered date over a derived one, permits the operator to revise any field before commit, and flags items as approaching expiry with a conservative margin. Predicted timestamps do not represent food-safety guidance and must not stand in for the producer's use-by or best-before labelling.

To avoid hand-tuned thresholds, our longevity priors were benchmarked against EU advisory literature and presuppose proper home storage (cooling ≤ 5 °C, freezing ≈ -18 °C) [48]. In step with EFSA and the European Commission's date-marking schema, any found "use by" marking (being a safety ceiling) always dominates derived values, whereas "best before" relates to sensory quality and can remain tolerable where storage protocols were upheld, and the food exhibits no spoilage indicators [49]. Correspondingly, for groupings with acute perishability (e.g., uncooked meats and aquatic products), the mapping imposes markedly tight chilled windows (measured in single-digit days), and for thermized dairy it enforces roughly week-scale bounds. For shelf-stable preserved goods, it deploys extended quality horizons (months up to years) yet still defers to the printed stamp. These figures constitute cautious orientational brackets meant to trigger prompts, not nutritional safety rulings, and the application visibly elevates any producer-stamped marking it successfully reads [50].

To address OCR failures arising from motion smear, dim environments, or focus misalignment, the handset client incorporates a quality-gating capture loop that judges frame sufficiency in real time and discharges the shutter solely when basic conditions are satisfied. Via CameraX, the engine tallies three per-frame signals, namely luminance taken as the average Y on a reduced grid, sharpness taken as the Laplacian spread over Y, and motion taken as the exponentially weighted moving mean of gyroscopic magnitude. The overlay renders these as illumination, clarity, and steadiness readouts, and the snapshot subroutine withholds the shot until all three metrics beat their predefined cutoffs—specifically luminance no lower than 70, Laplacian spread no lower than 18, and gyroscopic magnitude no higher than 0.5 radians per second—or until a brief fallback timeout of roughly 2.5 s elapses. Immediately before capture, the client invokes centre autofocus together with autoexposure using a metering region and awaits focus lock, then retries should the stored JPEG byte count suggest a failed grab.

Upon upload, frames are standardized through orientation correction, moderate edge crispening, and tonal range equalization, then routed through a dual-pass detection head that couples a holistic scan with a tiled scan to improve retrieval for small subjects. OCR is attempted against each extracted patch initially, and where no valid date is successfully decoded, the system degrades gracefully to a cautious shelf-life estimate that takes account of product grouping and packaging type.

The final realization provides capture, remote classification, canonicalization, and immediate on-screen portrayal through a slender Android front-end paired with a provisioned server-free backend. The source structure is compartmentalized and test-harness-friendly, qualities that ease future augmentations such as refined priors, new classifiers, and altered notification configurations without requiring modifications to the mobile deliverable.

Application Testing and Empirical Evaluation

A practical validation exercise was mounted to benchmark two distinct faculties of the software: the correctness with which it names grocery articles, and the correctness with which it settles on a spoilage date. In any instance where the

wrapper bore no imprinted use-by marking or where such a marking resisted machine decoding, the platform invoked a stored lookup table—organized by food category—that maps taxonomic groupings onto typical keeping spans, thereby supplying a surrogate spoilage estimate. Each forecast emerging from the system mated a product identifier either with a date successfully wrested from the packaging via OCR or with a date derived through that same categorical fallback mechanism.

Experimental setup and dataset construction

The evaluation session was conducted on 24 September 2025, sited inside a working grocery outlet with a Samsung Galaxy A54 serving as the test handset. All captures were obtained under quotidian operating parameters—standard overhead fluorescent illumination, the device held in one hand. The image corpus was built from two scene archetypes: straightforward compositions where a lone product type filled the frame, and multi-unit compositions where packs or specimens of identical merchandise sat grouped side by side on aisles or nested within shopping carts. Packaging faithfully represented the Romanian consumer landscape, with textual content rendered overwhelmingly in Romanian intermixed with English. Measurement proceeded on a per-product-type, per-scene basis: a photograph containing numerous indistinguishable jars or pieces yielded exactly one observational record, and the yardstick applied was whether the pipeline successfully identified that product type and attached a spoilage date—harvested directly from the label through OCR, or, failing that, estimated via the taxonomy-driven fallback rule.

To faithfully reproduce the rhythm of genuine household shopping trips, we deliberately assembled a heterogeneous collection spanning raw vegetables and fruit, milk-derivative goods, muscle meats and aquatic catches, baked dough products, refreshments, impulse snacks, tinned and bottled preserves, deep-frozen fare, milled grains, and cooking lipids with condiments. The composition of this test basket was calibrated against published purchasing statistics reporting that approximately 82% of households regularly acquire dairy lines, roughly 80% buy fresh greens and fruit, around 76% purchase snack-type foods, some 71% select bread and viennoiserie, close to 70% choose frozen products, an estimated 65% opt for canned provisions, roughly 65% fill their carts with pasta and rice, about 62% take home butcher-counter meats, and approximately 38% buy fish and seafood; our test manifest was consequently proportioned to echo the anatomy of a typical market basket while deliberately incorporating items that pose tough recognition and OCR challenges [51].

The benchmark suite ultimately comprised 59 discrete product entities fanned across a dozen classificatory divisions: leafy and root vegetables, citrus fruits, other fresh produce, aromatic and potted herbs, meats alongside fish, dairy-derivative goods, shell eggs, potable liquids, tinned and glass-jarred preserves, bottled oils and culinary sauces, bread and morning goods, and sugar-based treats with chocolate articles. The atomic measurement unit for any given frame was the product type observable within that frame. Shots enclosing a solitary type contributed a single record. Shots capturing multiple replicates of that same type—a shelf-edge full of identical cartons, to take one example—also contributed merely a single record, given that the investigative aim was to test whether the pipeline correctly resolved the product category and output a meaningful spoilage date for it; extra identical cartons added no new evidentiary weight.

For every product in the set, we methodically noted down the single top-ranked predicted designation together with the confidence value attached to that prediction. A further categorical note recorded whether the physical packaging actually carried a use-by or best-before imprint or whether no such imprint existed. When an imprint was present, we categorized its readability into one of three states: plainly legible in its entirety, only partly legible, or entirely illegible. The result of the OCR pass was captured and assigned to one of three bins: fully successful extraction, extraction with errors, or complete failure to extract. Where a date string could be decoded, we stored that value. Finally, we flagged the pathway that produced the definitive date ultimately shown to the user—either direct OCR decoding or algorithmic estimation driven by the category-to-shelf-life rule base.

For scoring purposes, when a printed label was physically present on the package, that printed date constituted the reference truth. When no printed label was there, the reference truth defaulted to a mutually agreed category-linked shelf-

life assignment. Naming accuracy was simply judged as right or wrong. Date accuracy was likewise judged as right or wrong when compared against whichever reference truth applied. For those dates that the system channelled through the estimation rule, we granted a leniency band of plus or minus a single calendar day, purely to neutralize edge effects from day-boundary rounding. It bears reiterating that whenever the system cannot find a readable on-pack date, the expiry estimate it furnishes is rooted exclusively in general, category-flavoured storage guidance meant only to provide an orientational nudge. These figures have no standing as safety evaluations. The actual interval during which a food remains safe and palatable turns on real domestic storage temperatures, the integrity of its wrapping, and formulation details specific to each manufacturer. In all cases, the detection of a manufacturer-affixed “use by” stamp automatically overrides and discards any computed estimate in favor of that regulatory date. Consistently, throughout the analyses presented below, every estimated date is treated strictly as a provisional placeholder, useful for driving on-screen countdowns and reminder prompts, and remains open to user amendment based on how the item was actually stored and on any visible spoilage signals.

Representative test cases

While conducting the practical evaluation, we documented a gallery of characteristic scenarios that together sketch the system's operational profile under retail-floor conditions: (i) unqualified successes where both product identity and OCR-harvested date hit the mark, (ii) situations absent any machine-readable label necessitating activation of the categorical shelf-life fallback, (iii) puzzling cases where OCR extracted a perfectly valid date but the product was given the wrong name, (iv) perceptual collisions among articles that share a close visual kinship, and (v) instances where the classifier landed on the correct identity despite shockingly low internal confidence. Each vignette reproduced below bundles the source photograph, its ground-truth category, the model's best guess with confidence, the method that yielded the data, the specific date string, and the overall pass/fail verdict.

(a) Case A—Correct product and correct OCR date (factory-sealed merchandise)

Here, the system pinned the right name on a pre-packaged article and simultaneously pulled the imprinted spoilage date off the label without error. The final pairing attached the correct identity to an OCR-derived timestamp that agreed with the reference date, as visible in [Figure 5](#).




Figure 5. Packaged items with the correct product ID and correct OCR date (Source: Authors' supermarket testing, 24 September 2025).

These snapshots capture the textbook “no-surprises” trajectory: the pipeline both called out the correct name and succeeded in reading a real spoilage imprint from one capture. Across these examples, the top prediction matched the ground-truth label and the OCR-decoded date matched what was stamped on the package, yielding a wholly correct product–date combination.

(b) Case B—Correct product, no readable date, category-based fallback (loose, label-free fresh items)

Confronted with a loose fresh item having no printed or scannable date whatsoever, the pipeline correctly supplied the item's name and then generated a spoilage estimate derived from the categorical shelf-life table, results that stayed within the expected ranges for that category, as documented in **Figure 6**.




Figure 6. Fresh produce with no readable printed date, expiration estimated via category-based shelf-life rule. Source: Authors' supermarket testing, 24 September 2025.

The bulk of fruits and raw vegetables were named accurately, typically with high confidence values attached, and since products in this class seldom bear date stamps of any kind, the application uniformly resorted to category-driven shelf-life projections consistent with the produce lookup table.

(c) Case C—Correct OCR date, incorrect product label (packaged items with deceptive visual signatures)

The pipeline successfully read the spoilage imprint from the wrapping in these cases, but it attached the wrong name to the article, an error whose most plausible root lies in distracting on-pack visual elements—among them, large-print ingredient mentions such as “milk,” prominent graphic illustrations, or imagery cues that bridge categories—that inadvertently pulled the classifier toward other goods. The outcome is a date that is correct mated to a product label that is not, as laid out in **Figure 7**.




Figure 7. Correct OCR date but misclassified product (packaged goods with misleading cues) (Source: Authors' supermarket testing, 24 September 2025).

Despite grappling with shiny, highly reflective wrapper surfaces and with date imprints that were both tiny and low in contrast, the system nonetheless managed to lock onto and correctly interpret the spoilage date through OCR. This illustrates a meaningful degree of pipeline resilience even under adverse visual conditions of the sort that routinely hamper legibility for human readers and machines alike.

(d) Case D—Product confused within a visually similar cluster

In frames featuring fresh produce items that closely resemble one another, the classifier occasionally returned a label that, while not exactly right, belonged to a near neighbour within the same broad food family. Since these specimens carried no date markings, the resulting date output was driven entirely by the category-level shelf-life rule, as shown in [Figure 8](#).




Figure 8. Product confusion within visually similar families of fresh produce. Source: Authors' supermarket testing, 24 September 2025.

These instances typify a widely acknowledged ceiling of fine-grained recognition applied to raw produce: the predicted label remains within the correct botanical or culinary grouping but may wander at the cultivar or variety level. The date output remains sensible because all members of that grouping share the same fallback shelf-life window.

(e) Case E—Correct product identity at extremely depressed confidence

The classifier produced the right identity label in two separate captures despite confidence scores hovering in the neighbourhood of 1%, as evidenced in **Figure 9**. A credible explanation for this extraordinarily faint conviction is the presence of heavy scene clutter, non-optimal framing, and a profusion of surrounding distractor objects unrelated to the target item.




Figure 9. Correct product predictions at extremely low confidence levels (Source: Authors' supermarket testing, 24 September 2025).

These observations show that output correctness and internal model confidence are not always tightly coupled in practice. Actionable mitigation strategies include deploying calibrated decision thresholds and, perhaps most usefully, surfacing a lightweight “confirm item?” dialog whenever the system’s self-assessed certainty falls below a defined floor.

Surveyed as a whole, this set of representative test vignettes illuminates both what the platform handles well and where friction remains when operating inside an actual grocery environment. The software proved itself capable of dependably

linking product identities with legitimate spoilage dates under routine conditions while also showing a degree of hardness in demanding contexts such as gloss-coated packaging, faint date imprints, and visually busy frames. At the same time, the documented instances of product naming errors and of predictions accompanied by ultra-low confidence highlight specific territories—fine-grained discrimination among look-alike foods and confidence estimation quality—that call for additional attention in subsequent development rounds. On balance, the field test confirms that the system can serve as a reliable aid for everyday pantry management, while simultaneously charting concrete directions for future refinement in both recognition robustness and the design of user-facing interactions.

Evaluation summary

The present segment consolidates the pooled findings laid out in [Table 2](#) and distills the platform's behaviour across the entirety of the 59-product, 12-category collection. We condense overall hit rates for (i) article naming and (ii) spoilage-date reading via OCR, then trace connections back to the paradigmatic vignettes discussed earlier.

Table 2. Consolidated evaluation results by product category.

Category	No. of products	Products correctly identified	Date present on packaging	Expiration date not visible	Received estimated date	Date correctly read (OCR)	Product identification success rate	Expiration date identification success rate *
Fruits	10	7	0	0	10	-	70%	-
Citrus	4	2	0	4	4	-	50%	-
Vegetables (incl. herbs)	13	9	0	13	13	-	69.2%	-
Meat and fish	2	2	2	0	0	2	100%	100%
Dairy and eggs	6	6	5	1	1	5	100%	100%
Beverages	3	3	3	0	0	3	100%	100%
Canned and preserved	4	2	3	1	1	2	50%	75%
Condiments and oils	4	4	4	4	0	3	100%	75%
Cereals and bakery	6	5	5	1	1	4	83.3%	80%
Sweets and confectionery	7	3	7	0	0	5	42.9%	71.4%
Total	59	43	30	29	29	25	72.88%	83.33%

Source: Authors' supermarket testing, 24 September 2025. We computed the expiration-date identification success rate using the columns "Date present on packaging" (denominator) and "Date correctly read (OCR)" (numerator). We computed the product identification success rate using "No. of products" (denominator) and "Products correctly identified" (numerator).

Spanning the complete collection of test articles, the software assigned the right name to 43 out of 59 entries (72.9%). Among those articles whose wrappers actually bore a printed date, the OCR subsystem decoded it correctly in 25 out of 30 instances (83.3%). For articles lacking any visible imprinted date (29 out of 59), the category-driven backup heuristic generated an estimated date in every single case (29 out of 29, representing total coverage). The upshot is that every observational record in the dataset was paired with a functional spoilage date, whether extracted directly or algorithmically imputed.

Peak performance manifested within groupings characterized by straightforward, uncluttered packaging and sharply rendered date imprints—most notably Meat and Fish, Dairy and Eggs, and Beverages—where both article recognition and OCR extraction either attained or came extremely close to the 100% mark. Condiments and Oils delivered flawless article naming alongside OCR success at 75%, a pattern aligning with sporadic low-contrast date stamps in that bracket. Cereals and Bakery posted 83.3% for article identification and 80% for OCR decoding. Canned and Preserved along with Sweets and Confectionery turned in noticeably weaker article recognition at 50% and 42.9% respectively, a drop consistent with the mislabelling dynamics exemplified in Case C. Categories dominated by loose fresh produce seldom featured imprinted date markings at all; article naming landed in a moderate band spanning roughly 50 to 70 percent, and the spoilage dates furnished to users rested entirely on the taxonomy-driven shelf-life projections detailed in Cases B and D.

Instances of mistaken identity clustered within visually congested compositions, within packages sporting trans-category visual signals—salient ingredient names standing out as a prime culprit—and within taxonomically tight-knit groupings such as citrus varieties. OCR shortfalls traced back predominantly to reflective or bowed wrapper surfaces and to stamps combining diminutive size with weak contrast, though notably the pipeline frequently triumphed even under those very conditions, as Case C attests. Confidence readings dipping to the neighbourhood of 1 percent occasionally sat alongside correct labels in densely populated frames, a phenomenon illustrated in Case E. These patterns collectively argue for deploying calibrated confidence guardrails, lightweight verification nudges when certainty sags, and strategies such as multi-shot consolidation or focused data enrichment to dampen confusion rates.

Under the routine conditions of a functioning grocery aisle, the platform dependably delivers a workable spoilage date—pulled straight from the package when possible, approximated when not—and hits strong article-recognition marks on factory-sealed goods. The outcomes simultaneously flag well-defined avenues for refinement: broadening training coverage for fine-grained horticultural subclasses, hardening OCR against wrapper gloss and surface curvature, adopting multi-angle (Scan 360°) frame merging for temporal stability, and engineering user-interface workflows that adapt to the system's self-assessed certainty.

The degree to which these findings travel beyond the present study is bounded by the modest specimen count (59 items) and the constrained single-venue, single-handset methodology. Accordingly, the accuracy figures reported here ought to be read as proof-of-viability signals rather than as population-level estimators. Expanding the sampling footprint across diverse retail chains, wrapper substrates, illumination regimes, linguistic contexts, and camera modules would inject the variance necessary to properly probe generalizability and to dilute site- or hardware-specific skew. We envisage that scaling up the corpus and dispersing data-collection locations would generate performance readings that more faithfully mirror what the platform can achieve in everyday deployment. A pre-planned multi-venue, multi-device replication employing stratified sampling by category is therefore slated as a follow-on undertaking.

Conclusion

The present inquiry put forward a workable and extensible blueprint for residential food oversight by marrying article recognition to dependable extraction or imputation of spoilage dates. The pipeline advanced here illustrates that a lean mobile front-end mated to a serverless cloud substrate can furnish perceptually immediate responses while holding operator burden to a minimum and preserving user confidence.

During a live grocery-outlet trial, the platform correctly named 43 of 59 discrete articles and faithfully decoded a stamped spoilage date for 25 of the 30 items that carried one. Where no imprinted date could be spotted, the regulation-based fallback generated a functional estimate for each and every member of that subset. These figures attest to dependable reach across heterogeneous categories and capture settings and bear out the proposition that a commonplace smartphone can deliver both fidelity and briskness for the designated workload.

These findings speak squarely to the first trio of investigative questions: (RQ1) the categorizer upholds sturdy performance inside naturalistic, operator-shot compositions, (RQ2) the OCR module deciphers printed dates with reliability across assorted wrapper types and lighting conditions, and (RQ3) the taxonomy-driven fallback furnishes blanket coverage whenever textual information is missing or indecipherable, with its absolute precision earmarked for calibration against independent benchmarks.

The present manuscript is best understood as a viability demonstration and technical build study, not as a behavioural quantification of food-waste endpoints. Our offering operates at the architectural tier: we establish that an integrated mobile pipeline—stitching together image categorization, optical character reading, and heuristic date estimation—performs under realistic photographic conditions. We neither designed nor powered the investigation to detect shifts in domestic conduct, purchasing patterns, or measured discard volumes, nor did we execute longitudinal or randomized-intervention protocols. Generalizability beyond the study context is hemmed in by a single-location, single-device data pool and by locale-sensitive prior assumptions; the accuracy rates reported thus stand as feasibility signposts rather than population-wide projections. Subsequent inquiry will necessitate multi-location rollouts engaging varied user cohorts and controlled experimental architectures to put numbers on engagement depth, protocol adherence, and net waste-diversion impact.

The software delivers value to end consumers through swifter product logging, fewer overlooked goods, and transparent day-count readouts that steer everyday choices. The operator snaps a single frame and immediately receives an article label alongside either a harvested spoilage date or a defensible surrogate, a flow that cuts down manual keystroking and strengthens responsiveness to reminder prompts. Shop-floor personnel can harness the identical workflow to perform shelf triage after minimal onboarding, since the screen layout spotlights items edging toward expiry and buckets them by urgency tier. Warehouse crews can surveil batches carrying near-term risk and mobilize upon consolidated alert streams. Across each of these user roles, the marriage of on-the-spot recognition, uniform data canonicalization, and live refresh cycles bolsters trust in the displayed intelligence and compresses the interval separating observation from intervention.

The platform knits together three faculties that commonly sit walled off from one another in incumbent tools. It names the article, pulls an expiry timestamp when one is legible, imputes a credible surrogate when text is gone or garbled, and then surfaces a single canonicalized entry complete with computed residual days and a condition flag. The architecture executes on affordable, near-ubiquitous hardware and leans on provisioned cloud utilities, a combination that drives down total cost of ownership for households and institutional adopters alike. Precision lands strongly on factory-wrapped merchandise sporting clean imprints, and day-to-day usability stays elevated because the program demands nothing more than a brief photographic gesture followed by unattended backend processing. The workflow upholds one unified data schema regardless of intake mode—camera or manual—thereby flattening the complexity of sort, filter, and notification logic on both the device side and the server side.

The assessment brings to light foreseeable ceilings that will steer the next phase of development. Finely sliced subcategories nested within visually homogeneous families occasionally trigger mix-ups. Faint date stamps printed across glossy or arched wrapper surfaces can still defeat the OCR stage under ambient illumination. Confidence values hovering near the decision boundary sometimes walk hand in hand with correct labels, a pattern that troubles any simple threshold-cutting tactic. Loose fresh produce rarely wears a printed spoilage date at all, meaning the platform leans on categorical priors that are purposefully cautious. These limitations did not stop the system from putting a functional spoilage date beside every single data point in the test suite, but they do galvanize focused upgrades in both image acquisition and inferencing.

Concerning the unanswered investigative questions, we stopped short of instrumenting end-to-end latency, CPU and memory footprint, or battery drain on the handset, and we likewise refrained from mounting a longitudinal head-to-head comparison capable of quantifying user engagement or actual household waste contraction; these dimensions have been reserved for a dedicated successor study (RQ4–RQ5).

Subject to explicit operator authorization and robust privacy safeguards, we intend to preserve captured frames for the purpose of assembling a supervised machine-learning corpus destined to feed a convolutional neural network (CNN) that judges product condition along a fresh-to-spoiled continuum. The image collection will be annotated under a formally documented labelling protocol, partitioned into training, validation, and held-out test splits, and employed to train a compact CNN suited for on-device execution (for instance, a MobileNet or EfficientNet variant distilled to TensorFlow Lite). We will merge the CNN's output stream with OCR-harvested dates and categorical lifespan priors through a multimodal fusion architecture (whether early or late integration) and apply probability calibration (via isotonic regression or Platt scaling) to emit a consumption-window projection accompanied by a quantified uncertainty envelope. The evaluation suite for this phase will report AUROC, F1 score, expected calibration error (ECE), and ablation studies isolating the contribution of each signal pathway (CNN, OCR, category priors). A subsequent multi-site field deployment will then put generalizability and real-world operational footprint to the test, capturing instrumentation data on latency, CPU and memory load, battery impact, task completion times, confirmation rates for low-certainty outputs, and proxy indicators tied to waste reduction. This research agenda squarely tackles RQ4 and RQ5 and stands to generate the evidentiary basis required for rational tuning of decision thresholds, user prompts, and rollout strategies.

Upgrades to both the underlying models and the training corpora offer another readily accessible route toward boosting dependability under ordinary usage scenarios. The shelf-life lookup matrix can be deepened to capture variety-level specificity within produce and baked goods, segments where keeping duration diverges considerably across subclasses. The pool of training examples stands to benefit from systematic enrichment with metallized reflective wrappers and tiny debossed date impressions that routinely trip up OCR engines, while the photographic intake module can be fitted with dynamic exposure adjustment and specular highlight suppression to contend with arced or glossy encasements. Text-line detection routines purpose-built for date-like motifs hold the potential to further prune false-positive readings that otherwise bubble up from ingredient enumerations and marketing copy.

Observations gathered on the shop floor bear out the fact that spoilage imprints are seldom positioned across the face of a package. They surface far more regularly along a lateral flank or the rear panel, and occasionally atop a screw cap or within a crimped seal. This positional tendency conspires with surface sheen and curvature to erode OCR accuracy whenever the system has only a single head-on snapshot to work with. A methodical 360-degree sweep modality stands to counteract these difficulties. The software can walk the operator through rotating the article in uniform angular increments, grab rapid burst sequences at each orientation, and then merge detections stitched across the full rotational sweep. Multi-perspective OCR swells the odds that no fewer than one frame will catch the date stamp under benign illumination and favorable surface geometry.

Aggregating classifications across several viewpoints firms up the product label via a temporal consensus mechanism and dulls the impact of visual noise confined to any lone frame. An actionable rule of thumb involves soliciting eight vantages encircling the object, spacing captures with a minimum inter-frame gap sufficient to suppress motion smearing, and terminating the loop early, the moment the pipeline attains a pre-specified certainty bar for both article identity and date reading. This augmentation slots cleanly into the existing system design, seeing that the front-end already directs the capture sequence and the cloud tier already coalesces canonicalized attributes into a unified item record.

Several natural extensions branch off from the empirical findings and from the expressed requirements of prospective users. The first thrust pertains to depot and retail-floor warning systems, which materialize as role-filtered consoles that bring forward items teetering on the edge of expiry, organized by physical zone, batch identifier, and vendor source. Notifications can be consolidated into morning and midday digests so as to curb alarm weariness while safeguarding the capacity to act without delay. The identical intelligence feeds enable shelf-facing rearrangement and surgical back-stock

retrievals that hold display aisles in lockstep with first-expiring-first-out discipline, a practice that tightens procedural adherence and pares back waste by intervening sooner.

A sister track homes in on self-acting restock triggers that couple near-expiry sensor streams with traditional replenishment thresholds for high-velocity stock-keeping units. Whenever the residual shelf horizon and the count of units physically on site jointly breach a mutually determined hazard boundary, the platform can float a resupply move with favored alternative items already prioritized in ranked order. This mechanism carries particular weight for merchandise that stays commercially viable under a secondary grading scheme once the primary expiry window has lapsed, and it forges a conduit linking the output of the recognition engine to tangible inventory operations.

Tariff governance for quality-stage progression emerges as a natural downstream application of these signals and buttresses retail strategies that apply graduated price reductions as goods march toward their consumption deadline or as they reclassify into a secondary quality bucket. The application is positioned to emit a daily candidate manifest bundling suggested markdown fractions pegged to residual shelf days and archival movement velocity, thereby equipping store supervisors with a grounded departure point. A self-tuning feedback mechanism can subsequently absorb per-category promotional elasticity and iteratively hone its suggestions absent any hand-cranking, which steadies outcomes across outlets exhibiting divergent footfall profiles.

The manner in which the system engages the human operator can likewise evolve to become sensitive to its own certainty, such that manual intervention is solicited only at junctures where it demonstrably adds value. A lightweight verification dialogue can surface whenever the classifier's self-assurance dips beneath a safe floor or whenever the OCR stage disgorges multiple candidate date strings, a design choice that preserves the briskness of the standard interaction path while still nailing down precision in the trickier edge cases catalogued throughout the evaluation. Should the inaugural lone-view snapshot fall short of netting an unambiguous date, the client can autonomously escalate to multi-view acquisition and synthesize the incremental evidentiary payload without forcing the operator to recommence the entire routine.

Rounding out the picture, home-use engagement mechanics hold the capacity to transmute faithful recognition into lasting behavioural transformation with barely any added friction. Collective targets, consecutive achievement runs, and waste-averting insignia can mirror the count of goods finished ahead of their spoilage horizon and render forward strides visible across every occupant sharing a dwelling. Timely nudges can arrive bundled with straightforward recipe nudges conjured from provisions drifting toward the tail of their usable window, a pairing that lifts follow-through while demanding zero supplementary typing. In ensemble, these refinements safeguard the platform's approachable, economical character while simultaneously ratcheting up exactness and earned trust within both domestic contexts and operational deployments.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 12 Aug 2021

Revised: 18 Sep 2021

Accepted: 05 Nov 2021

Published online: 15 March 2022

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

Eurostat. Food waste and food waste prevention—estimates. Statist Explain [Internet]. 2024 [cited 2025 Jul 24]; Available from: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Food_waste_and_food_waste_prevention_-_estimates

Too Good To Go. Too Good To Go [Internet]. Copenhagen, Denmark: Too Good To Go; [cited 2025 Jul 24]. Available from: <https://toogoodtogo.com/en>

Munch. Bonapp [Internet]. [cited 2025 Jul 24]. Available from: <https://bonapp.eco/>

Olio. OlioApp [Internet]. [cited 2025 Jul 24]. Available from: <https://olioapp.com/>

NoWaste. NoWaste [Internet]. [cited 2025 Jul 24]. Available from: <https://www.nowasteapp.com/>

Al-Obadi M, Ayad H, Pokharel S, Ayari MA. Perspectives on food waste management: prevention and social innovations. *Sustain Prod Consum*. 2022;31:190-208.

Suprpto A. Laws and regulations regarding food waste management as a function of environmental protection in a developing nation. *Int J Crim Justice Sci*. 2022;17:223-37.

Batool F, Kurniawan TA, Mohyuddin A, Othman MHD, Aziz F, Al-Hazmi HE, et al. Environmental impacts of food waste management technologies: a critical review of life cycle assessment (LCA) studies. *Trends Food Sci Technol*. 2024;143:104287.

Johansson N. Why is biogas production and not food donation the Swedish political priority for food waste management? *Environ Sci Policy*. 2021;126:60-4.

Chia D, Yap CC, Wu SL, Berezina E, Aroua MK, Gew LT. A systematic review of country-specific drivers and barriers to household food waste reduction and prevention. *Waste Manag Res*. 2023;42(6):459-75.

Bunditsakulchai P, Liu C. Integrated strategies for household food waste reduction in Bangkok. *Sustainability*. 2021;13(14):7651.

Parsa A, Van De Wiel M, Schmutz U, Taylor I, Fried J. Balancing people, planet, and profit in urban food waste management. *Sustain Prod Consum*. 2024;45:203-15.

Gong L, Thota M, Yu M, Duan W, Swainson M, Ye X, et al. A novel unified deep neural networks methodology for use by date recognition in retail food package image. *Signal Image Video Process*. 2020;15(3):449-57.

Newsome R, Balestrini CG, Baum MD, Corby J, Fisher W, Goodburn K, et al. Applications and perceptions of date labeling of food. *Compr Rev Food Sci Food Saf*. 2014;13(4):745-69.

Florea V, Rebedea T. Expiry date recognition using deep neural networks. *Int J User-Syst Interact*. 2020;13:1-17.

UNEP DTU Partnership, United Nations Environment Programme. Reducing consumer food waste using green and digital technologies. New York, NY: UNEP DTU Partnership; 2022.

Too Good To Go. Impact report. Copenhagen, Denmark: Too Good To Go; 2023.

Olio. Our impact [Internet]. OlioApp; [cited 2025 Sep 12]. Available from: <https://olioapp.com/en/our-impact/>

BonApp. Întrebări generale [Internet]. BonApp Eco; [cited 2025 Sep 13]. Available from: <https://bonapp.eco/intrebari-generale/>

Lidl Plus [Internet]. Lidl România; [cited 2025 Sep 15]. Available from: <https://www.lidl.ro/c/lidl-plus/s10019652>

Aplicația Kaufland [Internet]. Kaufland România; [cited 2025 Sep 15]. Available from: <https://www.kaufland.ro/utile/aplicatia-kaufland.html>

Act for Good [Internet]. Carrefour România; [cited 2025 Sep 15]. Available from: <https://carrefour.ro/corporate/actforgood>

Aplicația Penny [Internet]. Penny România; [cited 2025 Sep 15]. Available from: <https://www.penny.ro/descarca-aplicatia-penny>

Carrefour Group. Tackling food waste: Carrefour joins forces with Too Good To Go. Bucharest, Romania: Carrefour Group; 2023.

Penny România. PENNY extinde parteneriatul cu Bonapp pentru reducerea risipei alimentare. Ștefăneștii de Jos, Romania: Penny România; 2024.

Mathisen T, Johansen F. The impact of smartphone apps designed to reduce food waste on improving healthy eating, financial expenses, and personal food waste: crossover pilot intervention trial studying students' user experiences. *JMIR Form Res*. 2022;6(10):e38520.

Peng W, Xu J, Wang H, Wang Y. An efficient on-device OCR system for expiry date detection using FCOS and CRNN. *Algorithms*. 2025;18(5):286.

Wang X-F, He Z-H, Wang K, Wang Y-F, Zou L, Wu Z-Z. A survey of text detection and recognition algorithms based on deep learning. *Neurocomputing*. 2023;556:126702.

Takeuchi Y, Suzuki K. Expiration date recognition system using spatial transformer network for visually impaired. In: Miesenberger K, Kouroupetroglou G, editors. *Computers helping people with special needs. ICCHP 2024. Lecture notes in computer science*. Berlin, Germany: Springer; 2024.

Gupta M, Jacobson N, Garcia E. OCR binarization and image pre-processing for searching historical documents. *Pattern Recognit*. 2017;40(1):389-97.

Maliński K, Okarma K. Analysis of image preprocessing and binarization methods for OCR-based detection and classification of electronic integrated circuit labeling. *Electronics*. 2023;12(11):2449.

Fraccascia L, Nastasi A. Mobile apps against food waste: are consumers willing to use them? A survey research on Italian consumers. *Resour Conserv Recycl Adv*. 2023;18:200150.

Lubura J, Pezo L, Sandu MA, Voronova V, Donsi F, Žlabur JŠ, et al. Food recognition and food waste estimation using convolutional neural network. *Electronics*. 2022;11(22):3746.

Liu D, Zuo E, Wang D, He L, Dong L, Lu X. Deep learning in food image recognition: a comprehensive review. *Appl Sci*. 2025;15(14):7626.

Dai X. Robust deep-learning based refrigerator food recognition. *Front Artif Intell*. 2024;7:1442948.

Google LLC. Firebase [Internet]. Mountain View, CA: Google LLC; [cited 2025 Aug 7]. Available from: <https://firebase.google.com/>

Google LLC. Use a TensorFlow Lite model for inference with ML Kit on Android [Internet]. Mountain View, CA: Google LLC; [cited 2025 Sep 15]. Available from: <https://firebase.google.com/docs/ml/android/use-custom-models>

Clark Q, Kanavikar D, Clark J, Donnelly P. Exploring the potential of AI-driven food waste management strategies used in the hospitality industry for application in household settings. *Front Artif Intell*. 2025;7:1429477.

Google LLC. Text recognition [Internet]. Mountain View, CA: Google LLC; [cited 2025 Aug 6]. Available from: <https://firebase.google.com/docs/ml-kit/recognize-text>

Google LLC. Image labeling [Internet]. Mountain View, CA: Google LLC; [cited 2025 Aug 7]. Available from: <https://firebase.google.com/docs/ml-kit/label-images>

Google LLC. Image labeling [Internet]. Mountain View, CA: Google LLC; [cited 2025 Sep 15]. Available from: <https://firebase.google.com/docs/ml/label-images>

Google LLC. Barcode scanning [Internet]. Mountain View, CA: Google LLC; [cited 2025 Aug 6]. Available from: <https://developers.google.com/ml-kit/vision/barcode-scanning>

Mastorakis G, Kopanakis I, Makridis J, Chroni C, Synani K, Lasaridi K, et al. Managing household food waste with the FoodSaveShare mobile application. *Sustainability*. 2024;16(7):2800.

Lim V, Funk M, Marcenaro L, Regazzoni C, Rauterberg M. Designing for action: an evaluation of Social Recipes in reducing food waste. *Int J Hum-Comput Stud*. 2017;100:18-32.

Grigorcea G-L, Lefter B-I. Expiry app interface prototype [Internet]. Figma; [cited 2025 Oct 1]. Available from: <https://www.figma.com/design/Ysm1cMH2C3IfnDfVrPnX3z/Expiry>

Firebase. Cloud Functions for Firebase [Internet]. Mountain View, CA: Google LLC; [cited 2025 Sep 15]. Available from: <https://firebase.google.com/docs/functions>

Firebase. Firebase Authentication [Internet]. Mountain View, CA: Google LLC; [cited 2025 Sep 15]. Available from: <https://firebase.google.com/docs/auth>

European Food Safety Authority (EFSA). *Listeria monocytogenes* [Internet]. Parma, Italy: EFSA; [cited 2025 Sep 10]. Available from: <https://www.efsa.europa.eu/en/topics/topic/listeria>

European Commission. Date marking [Internet]. Brussels, Belgium: Directorate-General for Health and Food Safety, European Commission; [cited 2025 Sep 10]. Available from: https://food.ec.europa.eu/food-safety/food-waste/eu-actions-against-food-waste/date-marking-and-food-waste-prevention_en

European Commission. Commission notice on the application of Regulation (EU) No 1169/2011 on the provision of food information to consumers. *Off J Eur Union*. 2017.

Drive Research. The state of grocery shopping. Syracuse, NY: Drive Research; 2024.