FlavorCharter: Development of a Framework for Quantifying and Visually Analyzing Flavor Perception

BY

FRANCESCO BOTTO
B.S, Politecnico di Torino, Turin, Italy, 2023

THESIS

Submitted as partial fulfillment of the requirements
for the degree of Master of Science in Computer Science
in the Graduate College of the
University of Illinois at Chicago, 2025

Chicago, Illinois

Defense Committee:

G. Elisabeta Marai, Chair and Advisor
Khairi Reda
Stefano Di Calro, Politecnico di Torino

ACKNOWLEDGMENTS

I want to thank my family for supporting me while studying far from home, and my friends in Chicago for making me feel at home.

This work is dedicated to Enrica and all the people who are fighting cancer.

FB

TABLE OF CONTENTS

TABLE OF CONTENTS (continued)

CHAPTER **PAGE**

TABLE OF CONTENTS (continued)

CHAPTER **PAGE**

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

HNC	Head and Neck Cancer
IRB	Institutional Review Board
UI	User Interface
API	Application Programming Interface
SDK	Software Development Kit
JSON	JavaScript Object Notation
REST	Representational State Transfer
SD	Standard Deviation

SUMMARY

Taste perception plays a fundamental role in human well-being, influencing appetite, nutrition, and overall quality of life. However, quantifying how individuals perceive flavor remains challenging, as taste is inherently subjective and influenced by a complex interaction of biological, psychological, and cultural factors. This challenge becomes particularly relevant in clinical contexts such as head and neck cancer, where taste loss (dysgeusia) is a frequent and persistent side effect of radiotherapy. Patients often struggle to describe the changes in their perception of food, and clinicians lack standardized data to interpret these symptoms and complaints effectively.

This thesis presents FlavorCharter, a data-driven research framework designed to quantify and visually analyze the multidimensional perception of flavor. The project establishes the technological and methodological foundations for collecting, organizing, and interpreting sensory data at scale, with the goal of supporting clinical studies on taste loss.

The framework focuses on data collection from healthy individuals to establish a normative baseline of flavor perception. It integrates two complementary components: a mobile application for large-scale, crowd-sourced sensory data collection, and an interactive visualization tool for exploratory analysis.

The mobile platform allows users to rate everyday foods across ten sensory dimensions. Alongside these ratings, the application collects key demographic and dietary information enabling stratified and population-level analysis. The collected data are stored in a structured cloud database and processed for visualization.

The visualization environment enables clinicians to explore aggregated patterns, compare demographic groups, detect outliers, and investigate variability in flavor perception across populations. The design focuses on interpretability and scalability, providing an effective link between data acquisition and visual reasoning.

By providing a standardized framework for collecting and interpreting sensory data, FlavorCharter contributes to the development of precise, data-driven methodologies for characterizing flavor perception and its alteration in disease. Ultimately, this work aims to support more effective communication between clinicians and patients and to establish a foundation for future clinical and nutritional studies focused on sensory recovery, personalized nutrition, and quality-of-life improvement.

CHAPTER 1

INTRODUCTION

Food enjoyment is a fundamental component of human quality of life, and most people take the ability to perceive food flavor for granted. However, not everyone is that fortunate. Taste loss, also known as dysgeusia, is a common and often long-term side effect among head and neck cancer (HNC) patients, particularly after radiotherapy. Clinical studies have reported that dysgeusia affects between 51% and 100% of patients at the end of treatment, and 23% to 50% even one to two years later with no clear end date [1].

This dysfunction arises from radiation-induced damage to soft tissues, taste buds, and nerves in the oral cavity [2–5], and is recognized as a significant toxicity impacting patients' quality of life [6–8]. Those affected may experience food as dull, distorted, or even unpleasant. Over time, these changes can lead to reduced appetite and altered eating behavior, ultimately contributing to malnutrition and sarcopenia [9,10], conditions that correlate with poorer treatment outcomes and increased mortality [11–13].

Beyond oncology, taste loss can also manifest as a symptom of infections (such as COVID-19 or sinusitis), neurodegenerative diseases (e.g., Alzheimer's, Parkinson's, multiple sclerosis), medication side effects, or simply aging. Yet, despite its broad clinical relevance, dysgeusia has been primarily studied in head and neck cancer [14,15], where its prevalence and consequences are most severe.

1.1 Problem Statement

Efforts to mitigate taste dysfunction typically involve multidisciplinary interventions between oncologists, nutritionists, and dietitians [16, 17]. However, these strategies are often impaired by two main limitations:

1. The absence of a standardized language that allows patients and clinicians to describe and compare sensory alterations in a precise and quantifiable way.
2. The lack of large-scale, structured data describing how flavor is perceived in healthy individuals, clinicians currently lack any quantitative reference to interpret patients' taste-related complaints.

Patients frequently report vague descriptions such as “things taste different” or “food tastes like cardboard,” which offer little actionable information to clinicians. This communication gap reflects a deeper scientific limitation: the difficulty of quantifying the inherently subjective experience of flavor.

1.2 Vision

The FlavorCharter project aims to build a comprehensive, data-driven foundation for the study of human flavor perception. The vision is:

1. to define a standardized language that enables both clinicians and patients to communicate taste experiences in a consistent and quantifiable way; and
2. to establish a normative baseline describing how healthy individuals perceive the flavor of common foods.

Developing a standardized language is a crucial step toward building this baseline. It ensures that sensory descriptions are consistent across participants and interpretable by clinicians and researchers alike. The resulting healthy participant baseline provides population-level context and demographic-level tends. It allows researchers and clinicians to identify deviations from typical perception patterns, supporting more informed interpretation of patient reports.

This thesis focuses on designing and implementing the digital infrastructure required to collect, organize, and analyze large-scale sensory data in healthy populations.

1.3 Approach

To develop this framework, we adopted an *activity-centered design* methodology, engaging directly with oncologists from the MD Anderson Cancer Center to understand their analytical needs and the challenges involved in interpreting sensory data. This iterative process helped define the system requirements and guided the design of both data collection and visualization components.

The core components of the framework are:

- a mobile application for large-scale sensory data acquisition; and
- an interactive visualization tool for exploratory data analysis.

The **mobile application** serves as the primary tool for large-scale, crowd-sourced sensory data collection. Through a clear and accessible interface, participants can evaluate familiar foods across a set of predefined sensory dimensions, while also providing demographic and dietary information. Each dimension is accompanied by a concise textual description within the app, ensuring consistent interpretation among users with different cultural or linguistic backgrounds. All data are securely stored in a structured cloud database, designed for modularity, scalability, and privacy. This architecture enables population-level aggregation and the creation of a robust, queryable baseline.

Complementing the mobile app, the **visualization tool** transforms the collected data into interpretable, multivariate visual representations. It allows researchers to explore aggregated patterns, compare demographic groups, identify outliers, and examine the variability of flavor perception across populations. The visualization environment was designed with a focus oninterpretability and scalability, enabling a seamless transition from raw data to analytical insight.

Together, these two components bring to life the methodological core of *FlavorCharter*: connecting quantitative sensory data collection with interactive visual analysis.

1.4 Research Goals

The overarching goal of this thesis is to establish the technological and methodological foundations of the *FlavorCharter* framework, enabling the quantitative study of human flavor perception through data collection and visualization. The work focuses on four main objectives:

- **Define a standardized language:** Develop a structured and consistent way to describe flavor perception, enabling both researchers and clinicians to communicate sensory experiences in a quantifiable and comparable form.
- **Develop a mobile data collection platform:** Design and implement a scalable mobile application that applies this standardized language to collect crowd-sourced sensory

data. The app allows participants to evaluate familiar foods across the defined dimensions while also providing demographic and dietary information, ensuring consistency and reproducibility in the collected data.

- **Design an interactive visualization environment:** Create a visual analytics tool that enables exploratory data analysis, supporting the identification of aggregated patterns, variability across demographic groups, and potential outliers.
- **Build a baseline for future clinical applications:** Collect and organize sensory data from healthy individuals to establish a normative dataset of flavor perception, which will serve as a reference point for future comparisons with patients experiencing taste dysfunction.

By achieving these objectives, *FlavorCharter* lays the groundwork for a new, data-driven methodology in sensory science, transforming subjective taste experiences into measurable information and providing the basis for future research in clinical and nutritional contexts.

CHAPTER 2

STATE OF THE ART

Flavor perception is a complex sensory process that integrates gustatory, olfactory, and textural cues. While this multidimensional nature is widely recognized across culinary, biomedical, and sensory-science domains, current research still lacks a unified and standardized way to describe how individuals perceive the flavor of everyday foods. As highlighted in recent interdisciplinary studies, most existing approaches either focus on isolated taste qualities or rely on context-specific testing protocols, making it difficult to compare results across populations, studies, or clinical settings [18, 19].

A major limitation emerging from both clinical and sensory literature is the scarcity of quantitative data describing how healthy individuals typically perceive flavor. Without such a baseline, clinicians and researchers have limited support for interpreting individual variability, demographic differences, or pathological taste alterations. Moreover, the methods traditionally used to study taste tend to capture only narrow aspects of flavor, are not scalable to large populations, and are rarely deployed in everyday environments.

This chapter reviews the main approaches currently used to measure taste and flavor perception, highlights the methodological constraints that limit their comparability and scalability, and identifies the lack of standardized, multidimensional baseline data as a critical gap in the literature. These limitations motivate the development of new frameworks capable of capturing flavor perception in a structured, interpretable, and population-wide manner.

2.1 Related Work

Research on taste and flavor perception spans biomedical assessment, sensory science, and culinary studies, yet these domains provide only partial and often incompatible views of how individuals experience flavor. Existing approaches differ substantially in their goals, measurement conventions, and scalability, which complicates cross-study comparison and limits the development of standardized, population-wide baselines. The following subsections review three major groups of methods that dominate the current literature and highlight the constraints that motivate new frameworks for multidimensional flavor quantification.

2.1.1 Clinical Assessments of Taste Function

Clinical assessment of taste loss, particularly in head and neck cancer (HNC), relies on a combination of patient-reported outcomes, psychophysical tests, and coarse sensory disturbance scales. These tools are widely used in oncology and supportive care but provide only partial insight into the multidimensional nature of flavor perception.

The M.D. Anderson Symptom Inventory Head and Neck Module (MDASI-HN) is the commonly used to monitor taste dysfunction in HNC patients [19]. It includes a single symptom item on "problems with tasting," scored from 0 to 10 based on the patient's experience over the previous 24 hours. While clinically practical and routinely administered throughout radiotherapy, this measure captures only the severity of taste impairment and does not distinguish which sensory qualities are affected. As noted in clinical reports, patients often struggle to describe their symptoms beyond general statements such as "food tastes different" or "it tastes bland," limiting the interpretability of MDASI-HN data.

Objective psychophysical assessments, such as taste strips, liquid tastants, and edible films, aim to measure detection thresholds for the five classical tastes (sweet, sour, salty, bitter, umami) under controlled clinical conditions [18]. These tests can help identify basic taste deficits, but they evaluate only one taste at a time and do not account for higher-order components of flavor such as aroma, texture, and other trigeminal components. Moreover, they require standardized preparation and in-clinic administration, making them unsuitable for large-scale or at-home assessment.

Additional tools, including Likert-based disturbance scales such as CiTAS, attempt to categorize gustatory dysfunction along dimensions such as intensity reduction, discomfort, phantogeusia, and parageusia [20]. Although these scales allow finer symptom characterization than single-item PROs, their categories remain broad and do not map cleanly onto recognized sensory dimensions. As a result, they offer limited resolution for identifying which aspects of flavor perception are affected or how these changes evolve over time.

Overall, clinical assessments tend to prioritize feasibility and symptom monitoring over detailed perceptual characterization. They provide valuable information about the presence and severity of taste dysfunction but do not capture the multidimensional sensory structure of flavor, nor do they scale to population-level baselines of healthy individuals.

2.1.2 Sensory Science and Psychophysics

Within sensory science, one of the few structured instruments relevant to flavor-related research is the Taste Liking Questionnaire (TasteLQ), a validated tool developed to measure liking for foods associated with specific taste and oral sensations [21]. The questionnaire includes items representing basic tastes (sweet, sour, salty, bitter, umami) as well as additional sensations such as fat sensation, pungency, and astringency, allowing researchers to capture preference patterns across multiple taste-related qualities.

However, despite its broader coverage of taste-associated sensations, TasteLQ measures liking rather than perception. It does not quantify how individuals actually experience the sensory attributes of foods, nor does it capture dimensions such as aroma and texture, or the trigeminal components that contribute to flavor. Moreover, the instrument is designed for use in controlled survey contexts and has been validated only within the Danish population, limiting its applicability for large, diverse, or internationally representative studies.

2.1.3 Culinary Literature and Qualitative Frameworks

Culinary literature offers intuitive and culturally grounded frameworks for describing flavor, though these remain qualitative and non-standardized. The Flavor Bible organizes ingredients by associated aromatic notes, complementary pairings, and sensory impressions, reflecting how chefs conceptualize flavor synergies and balance [22]. Similarly, Nosrat's Salt, Fat, Acid, Heat identifies four core levers that shape the sensory experience of food preparation, saltiness, richness, acidity, and heat/spiciness, serving as a practical guide for modifying flavor rather than measuring perception [23]. Online guides and visual heuristics, such as crowd-sourced "everyday taste profile" diagrams, further illustrate common intuitions about flavor categories but lack empirical grounding or demographic generalizability [24].

Although these sources provide rich descriptive vocabularies, they do not offer quantitative scales, standardized descriptors, or population-level data suitable for scientific or clinical interpretation.

2.1.4 Sensor-Based Taste Measurement

A separate line of work investigates taste through sensor-based systems, most notably electronic tongues. These devices analyze the chemical properties of foods, rather capturing the flavor perception, using sensor arrays and are typically applied to specific products such as beer or vacuum-packed minced beef [25,26]. This type of analysis is largely driven by food producers and manufacturers, who use electronic sensing to monitor product quality or predict sensory attributes.

While useful for chemical characterization, these systems measure objective physicochemical signals rather than human perception. As a result, sensor-based approaches do not capture the multidimensional sensory experience of flavor and cannot provide population-level insight into how individuals perceive everyday foods.

2.2 Limitations of Existing Approaches

Despite contributions from clinical assessment, sensory science, and culinary literature, the current body of research presents several structural limitations that hinder the development of a standardized and scalable understanding of flavor perception.

A first major limitation concerns the lack of a multidimensional and interpretable vocabulary for describing flavor. Clinical tools typically measure symptom severity or basic taste detection, focusing on broad constructs such as "problems with tasting" or unspecific disturbances. Sensory-science instruments such as TasteLQ capture liking for foods associated with specific taste qualities, but do not characterize how individuals perceive these sensory dimensions [19–21]. Culinary frameworks, while rich in descriptive terminology, rely on qualitative heuristics that are not standardized across populations. As a result, the literature lacks a unified descriptive language capable of capturing the full complexity of flavor perception [22, 23].

A second limitation is methodological: most existing protocols are designed for controlled laboratory or clinical environments and cannot be deployed at scale. Taste-function assessments such as taste strips, liquid stimuli, and edible films require standardized preparation and in-clinic administration, which limits both accessibility and ecological validity [18]. Consumer-based instruments, including the Taste Liking Questionnaire (TasteLQ), are easier to administer but measure liking or preference rather than the perceptual structure of flavor [21]. As a result, current methods do not support large, distributed, or real-world data collection capable of capturing how individuals perceive the flavor of everyday foods.

A third limitation concerns the absence of normative, population-level datasets describing flavor perception in healthy individuals. Clinical studies on dysgeusia highlight substantial variability in how patients articulate taste changes, but lack reference distributions against which individual deviations can be interpreted. Sensor-based approaches such as electronic tongue systems provide only chemical or quality-control measurements and offer little insight into human sensory experience; moreover, objective taste documentation remains sparse, as this work is typically driven by food producers rather than perceptual research [25,26]. The resulting landscape is fragmented, with no large-scale datasets integrating demographic, cultural, and dietary factors known to influence taste sensitivity.

2.3 Gap Analysis and Motivation for This Work

The review of existing literature reveals that current approaches provide only fragmented and domain-specific views of flavor perception. Clinical tools reduce taste to a coarse symptom score, consumer instruments capture preference rather than perception, culinary frameworks offer rich but qualitative descriptors, and sensor-based methods quantify chemical properties without addressing human experience. None of these approaches offers a unified, scalable, or interpretable way to describe how people perceive the flavor of everyday foods.

What emerges is a clear and persistent gap: the field lacks a comprehensive framework that bridges subjective sensory experience, standardized descriptors, and population-level variability. Despite the clinical relevance of taste dysfunction and the widespread interest in sensory evaluation across disciplines, no existing method provides the foundational data needed to establish normative patterns of flavor perception or to contextualize individual differences within a broader population.

This gap motivates the need for a new paradigm—one that treats flavor perception as a multidimensional phenomenon that can be measured consistently, collected at scale, and explored through interpretable visual representations. Such a framework must be able to operate outside laboratory environments, reach diverse populations, and produce data that support both scientific inquiry and future clinical translation.

The work presented in this thesis responds directly to this need. By addressing the absence of standardized descriptors, scalable data-collection methods, and population-level baselines, it lays the groundwork for a unified, data-driven approach to understanding human flavor perception. This approach seeks not only to fill a methodological void but also to enable future research on taste dysfunction, demographic variability, and the sensory determinants of food experience.

In the longer term, the framework introduced here has the potential to support both patients and clinicians in head and neck cancer care. A standardized perceptual language can empower patients to articulate their sensory changes more clearly and consistently, reducing the ambiguity that currently characterizes taste-related symptom reports. At the same time, structured reference data provide clinicians with an objective context through which these descriptions can be interpreted, facilitating more informed conversations and offering a quantitative foundation for monitoring taste alterations over the course of treatment. In this way, the system contributes not only to methodological innovation but also to improving communication and understanding within clinical practice.

CHAPTER 3

SYSTEM DESIGN

This chapter presents the overall design of the FlavorCharter framework, outlining the architecture, data model, and design principles that connect the mobile data collection application with the visualization tool. The goal of this phase is to translate the conceptual vision introduced in the previous chapter into a coherent system architecture capable of supporting large-scale data acquisition and interactive visual analysis.

3.1 Overview of the Framework

The FlavorCharter framework is composed of two main components: a mobile application and a web-based visualization tool. These modules operate in synergy to enable the collection, storage, and analysis of sensory perception data. While each component can function independently, their integration ensures a continuous flow of information from user input to visual interpretation.

The framework is designed to collect large-scale, structured data describing how individuals perceive the flavor of everyday foods and to make this information visually explorable by researchers and clinicians. FlavorCharter follows a distributed client–server architecture built around a cloud-based backend. The mobile application constitutes the primary data acquisition interface and is supported by a Backend–as–a–Service (BaaS) infrastructure based on Firebase

(Firestore and Auth). The visualization tool functions as an analytical front-end that retrieves and processes the stored data through the same backend interface.

At a high level, the system workflow proceeds as follows:

1. **Participant recruitment:** participants are invited to join the study through institutional or public channels and install the FlavorCharter mobile application.
2. **Data acquisition:** within the app, users provide demographic and dietary information and evaluate a predefined set of foods across several perceptual dimensions.
3. **Data transmission:** each evaluation is securely transmitted to the cloud database using the Firebase SDK. All communications are authenticated and encrypted.
4. **Data storage and organization:** collected data are stored in Firestore following a structured schema that separates user metadata and sensory ratings.
5. **Data visualization and analysis:** the visualization tool accesses the aggregated dataset to generate interactive visual representations that support filtering, comparison, and exploratory analysis.

This end-to-end workflow connects data generation and visual reasoning within a unified infrastructure. New submissions from the mobile client are synchronized in real time with the cloud and immediately available to the visualization environment.

3.2 Design Methodology

The design of the FlavorCharter framework followed an activity-centered design approach, grounded in the observation of real user needs and analytical workflows rather than abstract system requirements [27]. The objective was to develop a platform that could simultaneously support two distinct user groups: participants, who provide sensory evaluations through an intuitive mobile interface, and clinicians or researchers, who interpret the resulting data through a visual analytics environment. This approach ensured that the system was designed around actual usage contexts, bridging the gap between technical feasibility and practical applicability.

A key aspect of this process was the direct collaboration with clinicians and oncologists from the MD Anderson Cancer Center. These collaborations provided valuable insights into the clinical challenges associated with taste dysfunction, particularly in head and neck cancer patients, and highlighted the need for standardized, quantifiable descriptors of flavor perception. Through a series of iterative discussions and design reviews, the clinical partners contributed to defining the type of data to be collected, the desired level of granularity, and the essential criteria for interpretability in medical and nutritional contexts. Their feedback directly influenced several design decisions, including the definition of the perceptual dimensions, the organization of the food list, and the metadata structure used for stratified analysis.

The following sections describe the data model and system architecture that emerged from this collaborative design process.

3.3 Data Model

The data model of FlavorCharter defines the core entities captured by the system and their relationships. The model ensures that data collected through the mobile application can be easily queried, aggregated, and visualized to support exploratory analysis and future clinical comparisons.

The design of the data schema was informed by both clinical expertise and literature in sensory science and culinary theory, and includes the definition of perceptual dimensions, metadata categories, and the food list aimed at creating a model that could capture the complexity of human flavor perception.

3.3.1 Sensory Dimensions

A central design decision in FlavorCharter concerned the definition of the sensory dimensions used to characterize flavor perception. The goal was to create a perceptual model that was at once scientifically grounded, clinically meaningful, and intuitive enough to be understood by a general population of participants.

The definition process combined insights from multiple disciplines and sources. Initial reference was taken from biomedical literature addressing taste alteration and loss, as well as from the Taste Liking Questionnaire (TasteLQ) [21, 28]. Complementary inspiration was taken from culinary frameworks to ensure that the perceptual model remained both comprehensive and intuitive for non-specialist participants [22, 23]. To adapt these scientific taxonomies to a digital, crowd-sourced setting, the team engaged in a series of iterative design sessions with clinicians and oncologists from the MD Anderson Cancer Center. The discussions focused on identifying a minimal but comprehensive set of perceptual dimensions that could capture the most relevant aspects of flavor experience while remaining usable in a mobile interface.

Three main criteria guided the selection:

1. **Perceptual relevance:** each dimension should correspond to a distinct and widely recognizable sensory quality, supported by physiological or neural evidence.

2. **Clinical interpretability:** the dimensions must be relatable to known alterations in dysgeusia and usable as standardized descriptors in clinical communication.
3. **Cognitive simplicity:** participants with no sensory training must be able to understand and consistently apply each term with minimal instruction.

After different phases of refinement, the final model converged on ten core dimensions that jointly span the sensory space of flavor:

- **Sweet** — the sensation of sugars and sweet foods.
- **Salty** — the taste of salt and mineral-rich foods.
- **Sour** — the sharp, acidic sensation found in citrus or vinegar.
- **Bitter** — the taste often found in dark chocolate, coffee, or leafy greens.
- **Umami (Savory)** — the "meaty" or broth-like taste present in cheese, mushrooms, and cured meats.
- **Fatty** — the mouth-coating sensation linked to foods rich in oils or fats.
- **Astringent** — the dry, puckering feeling caused by foods like tea, wine, or unripe fruit.
- **Aromatic** — the smell and aroma of food sensed through the nose.
- **Texture** — how food feels in the mouth, such as creamy, crunchy, or soft.
- **Piquancy** — the spicy "heat" felt when eating chili peppers or pepper.

The final set of ten dimensions represents a synthesis of sensory, clinical, and cognitive considerations, forming the core vocabulary through which participants describe their flavor perception.

During the early design stages, a broader range of candidate descriptors was also examined, some attributes were progressively excluded after iterative evaluation and discussion with clinicians. Terms such as *metallic* and *mouth-coating* were judged to be redundant or strongly correlated with other existing categories, such as bitterness, astringency, or fattiness, and thus offered limited additional interpretive value. Conversely, *temperature* and *heat* were excluded because they primarily relate to the physical or culinary context of food preparation rather than to intrinsic gustatory or trigeminal stimuli, making them difficult to standardize and less relevant to taste dysfunction [23]. The X factor, term occasionally used in sensory analysis to refer to the holistic, emotional, or even spiritual dimension of eating, was omitted as it lacked operational clarity and cross-cultural consistency [22].

Each dimension is rated by participants on a five-point Likert scale (0 = not perceived, 5 = strongly perceived). During the design stage, special attention was given to linguistic clarity: each label was accompanied by a short description to ensure consistent interpretation across languages and cultural contexts.

The final set of ten dimensions provides a standardized yet flexible language for describing human flavor perception and serves as the core semantic model linking data collection and visual analysis within the FlavorCharter ecosystem.

3.3.2 Demographic and Dietary Metadata

To contextualize sensory data and support stratified analysis, each participant provides a set of demographic and dietary descriptors. These attributes were chosen based on known clinical and cultural factors that influence taste perception, as well as the need for population-level grouping in statistical and visual analyses. The collected metadata include:

- **Age group** — grouped to account for age-related decline in taste sensitivity.
- **Gender**
- **Nationality and Ethnicity/Race** — to investigate cross-cultural variability and genetic influences on taste.
- **Regional Cuisine** — a categorical variable capturing the participant's typical culinary exposure (e.g., Indian, Mexican, Chinese, Thai, other).
- **Diet type** — self-reported information on dietary habits (e.g., balanced, vegetarian, vegan, low-carb, gluten-free).
- **Typical foods consumed** — an open-ended field used to capture habitual food preferences.

3.3.3 Food List

The food list defines the items participants evaluate within the mobile app. It was curated through a structured selection process combining culinary reasoning, clinical usability, and data consistency requirements. The selection was informed by the Danish Taste Liking Questionnaire (TasteLQ) and refined through discussions with MD Anderson clinicians [1]. Four main principles guided the curation process:

- **Familiarity:** items must be widely recognized and easily available to the target population.
- **Clear taste profile:** each food should have a clear, dominant taste profile corresponding to one primary sensory dimension.
- **Ease of evaluation:** foods must be ready-to-eat, requiring minimal or no preparation.
- **Balanced representation:** each sensory dimension is represented by at least three food items.

Items with high variability (e.g., dependent on brand, ripeness, or preparation method) were excluded or standardized through explicit labeling (e.g., "plain unsweetened yogurt", "yellow ripe banana"). The final set includes 36 food items grouped by their dominant dimension (Table I), providing a consistent, easy-to-evaluate, and taste-specific foundation for reliable sensory assessment.

TABLE I: Foods grouped by dominant flavor dimension.

Sweet	Sweet Potato, Strawberry, Nutella, Brownie, Honey, Yellow Banana, Orange Juice, Milk Chocolate, Pear, Syrup
Salty	Soy Sauce, French Fries, Potato Chips, Parmesan
Sour	Yogurt, White Vinegar, Pickles, Lemon Slice, Orange Juice
Bitter	Dark Coffee, Dark Chocolate, Tonic Water
Umami	Soy Sauce, Parmesan, Salami, Garlic, Bacon
Fatty	Olive Oil, Avocado, Milk, Almonds, Eggs, Dark Chocolate, Parmesan, Bacon
Astringent	Red Wine, Black Tea, Walnuts, Dark Coffee
Aromatic	Strawberry, Garlic, Parmesan, Dark Coffee, Dark Chocolate, Black Tea, Yellow Onion
Piquancy	Black Pepper, Chili, Ginger, Yellow Onion

3.4 System Architecture

The FlavorCharter system follows a modular client–server architecture designed to enable scalable data collection, secure storage, and interactive analysis of sensory perception data. At a high level, FlavorCharter connects participants who contribute sensory evaluations through a mobile application with a centralized cloud backend and a web-based visualization tool. The system operates as a continuous pipeline: data collected from users are securely stored and synchronized in the backend, then accessed in read-only mode for visualization and exploratory analysis.

The two-layer structure ensures a clear separation of concerns:

- The **Data Collection System** focuses on data input, validation, authentication, and secure transmission.
- The **Visual Analysis System** focuses on data aggregation, interpretation, and visual reasoning.

This separation enhances scalability (each layer can evolve independently), security (write vs. read-only access), and modularity (the visualization environment can be extended or replaced without affecting data collection). Figure 1 illustrates the overall system architecture, showing the interaction between data collection and visualization layers.

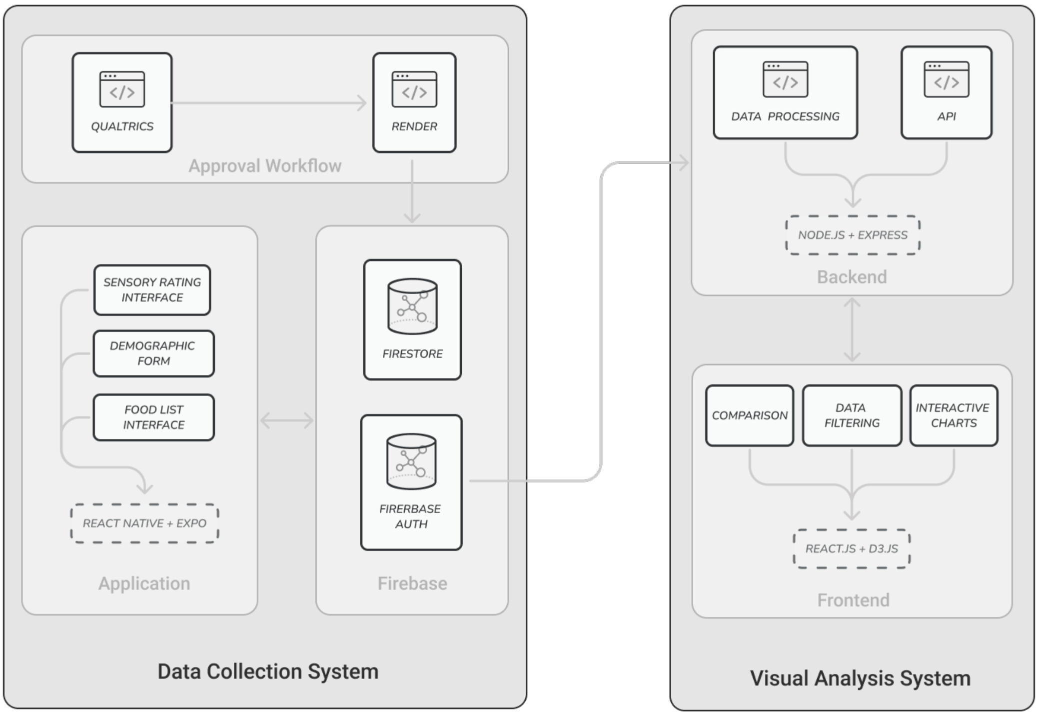

Figure 1: High-level overview of the FlavorCharter architecture

3.4.1 Data Collection System

The **Data Collection System** is responsible for acquiring, validating, and securely storing user-generated data. It comprises the mobile client, backend infrastructure, and user approval pipeline. While this layer is discussed in greater depth in *Data Collection System* chapter, this section provides a structural overview of its components and internal data flow.

- **Mobile Application (Expo/React Native)** Serves as the primary interface for participants to evaluate foods, submit sensory ratings, and provide demographic and dietary information. It includes client-side validation, offline caching, and secure data submission through Firebase SDKs.
- **Firebase Authentication** Handles user identity management and session control. Authentication tokens issued by Firebase Auth are used to validate and secure all write operations in Firestore.
- **Firebase Firestore** Acts as the structured database for storing all collected data, including users, foods, and sensory ratings. Its document-oriented schema enables real-time synchronization and scalability, allowing new submissions to be immediately available to the visualization tool.
- **Participant Approval Workflow** Participant registration and approval are managed through an automated workflow that integrates *Qualtrics* and a secure *Render* webhook. This process ensures that only verified and consented participants gain access to the FlavorCharter mobile application.

3.4.2 Visual Analysis System

The **Visual Analysis System** provides an analytical environment for researchers and clinicians to explore the aggregated data collected through FlavorCharter. This component is described in greater detail in **Visual Analysis System** chapter, but an architectural overview is provided below.

- **Visualization Tool (React.js)** A web-based application developed using React framework and D3.js for data-driven visualization. The tool enables users to filter, compare, and explore aggregated flavor perception data through interactive charts and coordinated visual views. It communicates with the backend through REST APIs, retrieving preprocessed analytical data while maintaining a responsive and user-friendly interface.
- **Backend (Node.js + Express)** Acts as an intermediary layer between the visualization client and Firestore. This service handles authenticated queries, performs lightweight data transformations such as aggregation, normalization, and statistical computation, and returns the processed data to the client. Centralizing analytical logic within the backend ensures consistency across users, optimizes query efficiency, and prevents direct client access to the database.
- **Data Source (Firestore)** Serves as the central data repository shared across both the collection and visualization layers. The backend retrieves data from Firestore in read-only mode, applying query filters that aggregate and anonymize records before sending them to the visualization tool. This approach preserves data integrity and privacy while maintaining real-time synchronization with new entries from the mobile application.

3.4.3 Cross-layer Data Flow

The two layers of FlavorCharter are linked through a continuous and secure data pipeline that connects participant input to analytical visualization. This architecture enables near real-time propagation of new data from the point of collection to the visualization environment, ensuring that analytical results always reflect the most recent submissions.

The end-to-end process can be summarized as follows:

1. **Participant Approval (Qualtrics and Webhook Integration)** Participant onboarding is managed through an approval workflow that connects a **Qualtrics** survey with a secure webhook service. Upon completing the survey, Qualtrics communicates the participant's information to the webhook, which grants the user permission to create an account within the FlavorCharter mobile application.

2. **Data Collection (Mobile App → Firebase)** Once approved, the participant signs in via Firebase Authentication and gains access to the mobile application. The app allows users to submit sensory evaluations and demographic metadata.

3. **Data Access and Processing (Firestore → Backend)** The visualization backend operates in read-only mode, retrieving aggregated and anonymized data from Firestore. It performs lightweight analytical processing—such as computing averages, standard deviations, and demographic distributions—and exposes the results through secure REST API endpoints.

4. **Exploration and Visualization (Backend → Client)** The visualization client requests preprocessed data from the backend and renders it using interactive D3-based components. Researchers can filter, compare, and explore the data across sensory dimensions or demographic groups. Because the pipeline is synchronized with Firestore, any new submission or update from the mobile app becomes instantly visible in the visualization environment.

3.5 Summary

The FlavorCharter system provides a modular and scalable framework for the structured collection, storage, and visualization of sensory perception data. By clearly separating data acquisition from analytical exploration, the architecture ensures maintainability, flexibility, and real-time consistency across components. Secure authentication, controlled data access, and a unified data model guarantee integrity throughout the entire workflow. The following chapters present the two core components of this architecture in detail: the **Data Collection System** chapter outlines the structure and operation of the mobile application, while the **Visual Analysis System** chapter examines the design principles and interactive functionalities of the visualization tool.

CHAPTER 4

DATA COLLECTION SYSTEM

The mobile application represents the core component of the FlavorCharter framework's Data Collection System. It provides the primary interface through which participants contribute sensory evaluations, demographic information, and dietary metadata. As the system's entry point, the app plays a crucial role in translating subjective flavor experiences into structured, analyzable data that feed the broader FlavorCharter ecosystem.

4.1 Overview and Objectives

The application was designed to support large-scale, crowd-sourced data collection while maintaining scientific rigor, usability, and security. It bridges human–computer interaction principles with technical infrastructure, enabling users from diverse backgrounds to participate in flavor perception research with minimal friction. Through a combination of an intuitive interface, real-time data synchronization, and robust authentication, the app ensures data quality and participant trust throughout the collection process.

In addition to technical and usability considerations, the development of the mobile application followed Institutional Review Board (IRB) guidelines to ensure ethical data collection and participant protection. All procedures, including informed consent, data anonymization, and secure data handling, were designed in accordance with approved research protocols.

In summary, the FlavorCharter mobile app serves as both a research instrument and a digital infrastructure component, designed to transform individual sensory experiences into standardized data suitable for large-scale analysis. The following sections describe its system design, technologies, and user experience in greater detail.

4.2 System Architecture and Technologies

The FlavorCharter mobile application was developed as a cross-platform client built with *React Native* and the *Expo SDK*, connected to a cloud-based backend powered by *Firebase*. This architecture enables scalable deployment across iOS and Android devices while maintaining a unified codebase and native performance. The overall design follows a client–cloud model where the app acts as the data acquisition interface, Firebase functions as a *Backend-as-a-Service* (BaaS), and additional middleware components support user approval and data flow integration with external research tools, Figure 2 shows the overall architecture of the Data Collection System.

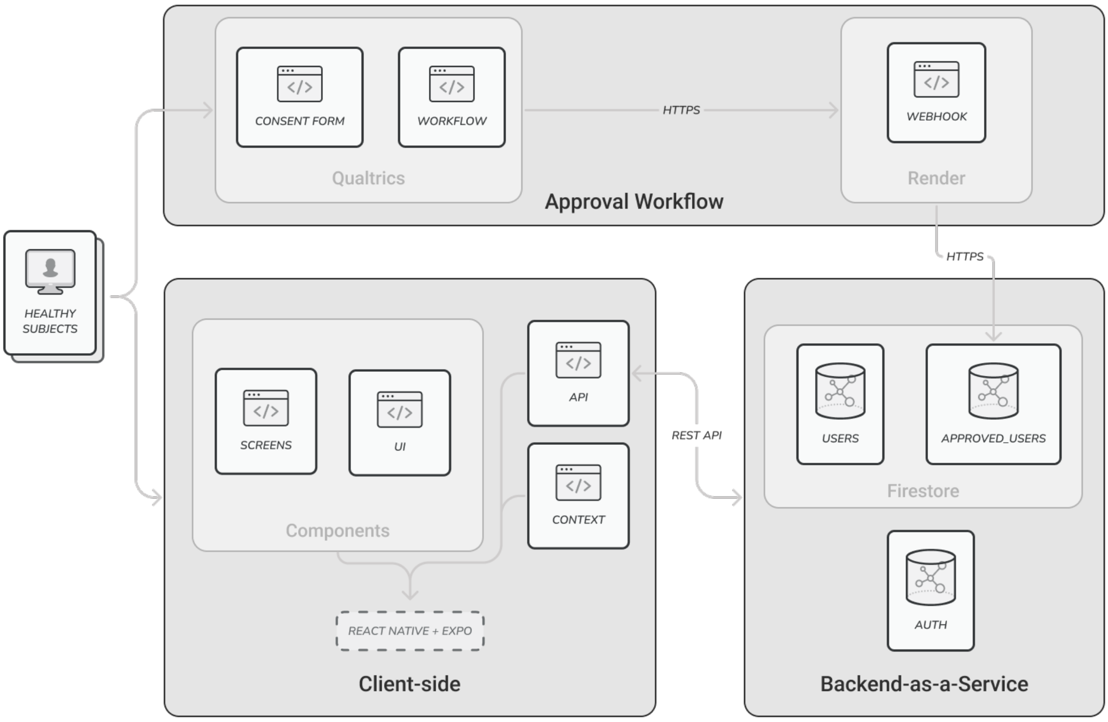

Figure 2: Data Collection System architecture.

4.2.1 Technology Stack

The FlavorCharter data collection ecosystem is composed of three primary layers that interact seamlessly to manage authentication, data submission, and real-time synchronization:

- **Client-side architecture:** The mobile app is built using *React Native* with the *Expo* SDK, which simplifies development, testing, and deployment through a single codebase. The app uses a modular, component-based structure, allowing each feature, such as authentication, sensory rating, or demographic forms, to be encapsulated within reusable UI and logic components. Application logic and persistent state are managed via the *Context API*, providing lightweight global state sharing without introducing the complexity of Redux.

- **Backend-as-a-Service (BaaS):** Firebase provides all essential backend functionalities, including Authentication and Firestore. *Firebase Authentication* manages secure user identity and session persistence. *Firestore*, a real-time NoSQL document database, handles structured storage for users, foods, and sensory ratings. By using Firebase as a BaaS, the system eliminates the need for a custom backend while maintaining scalability, low latency, and real-time data propagation between devices.
- **External services:** The app integrates with external services used for participant recruitment and access control. *Qualtrics* is employed as the initial survey external platform to gather participant consent and eligibility data. Once a participant completes the survey, a secure *webhook* hosted on *Render* transmits the approval information to Firestore, where it is stored in the list of approved users . When the user attempts to register in the app, this record is verified to determine whether access should be granted.

4.2.2 Component Architecture

The mobile app follows a modular file and component structure designed for scalability and code reusability. Each functional area is isolated into distinct directories to separate concerns between presentation, logic, and state management. The high-level organization of the project is as follows:

- **screens/** – Contains the main application screens that define navigation flow and user interaction logic (e.g., `LoginScreen.js`, `SignupScreen.js`, `FormPage.js`, `RadarPage.js`). Each screen orchestrates interactions between UI components and shared state.
- **components/** – Includes reusable interface elements grouped by functionality: *auth/* for authentication forms and consent screens, *form/* for demographic and dietary input fields, *radar/* for sensory rating widgets (radar charts and sliders), and *UI/* for general-purpose elements like buttons, modals, and progress indicators. This separation promotes reusability and visual consistency across the app.
- **store/context/** – Implements global state management using the *Context API*. Dedicated contexts manage user data, selected foods, radar chart values, and authentication state (`user-context.js`, `foods-context.js`, `radar-context.js`, `auth-context.js`). This structure provides centralized but lightweight state management, reducing prop drilling and improving maintainability.
- **store/data/** – Contains static data definitions and configuration files, such as the list of available foods (`foods-options.js`, `form-options.js`), form field options, and shared styles.
- **util/** – Contains helper functions for HTTP communication and utility logic (`http.js`), which facilitate modular integration with external APIs, including the approval webhook.

4.2.3 State Management and Navigation

State management in FlavorCharter relies on React's built-in *Context API* combined with hooks such as *useState* and *useContext*. This approach provides a simple yet effective mechanism to share data between components and preserve user input during navigation without introducing external dependencies. Each functional context (e.g., authentication, user demographics, rating values) maintains a synchronized copy of the current application state. When a user interacts with an interface element, for instance, modifying a demographic field or adjusting a flavor dimension, the change is immediately propagated both to the local component state and to the corresponding global context. This design improves responsiveness and ensures that data remain consistent across screens and can be accessed or restored at any point in the navigation flow.

The navigation flow is implemented using the *React Navigation* library, structured as a stack-based navigator that organizes screens in a linear hierarchy (Figure 3). This hierarchical design mirrors the logical progression of the study protocol, ensuring that key steps, such as onboarding, demographic entry, and sensory evaluations, are presented in a fixed and guided sequence. This prevents users from skipping mandatory tasks or navigating back into incomplete sections, thereby preserving the integrity of the data-collection process.

Once authentication is completed, the user gains access to the full application context. At this point, the bottom-tab navigator provides stable and persistent entry points to the main sections of the app (Home, Ratings, Profile), enabling repeated participation and seamless re-entry across sessions. Modal screens are employed for focused, high-priority tasks such as completing the demographic form or rating a food item, maintaining a clear separation between general navigation and task-specific workflows.

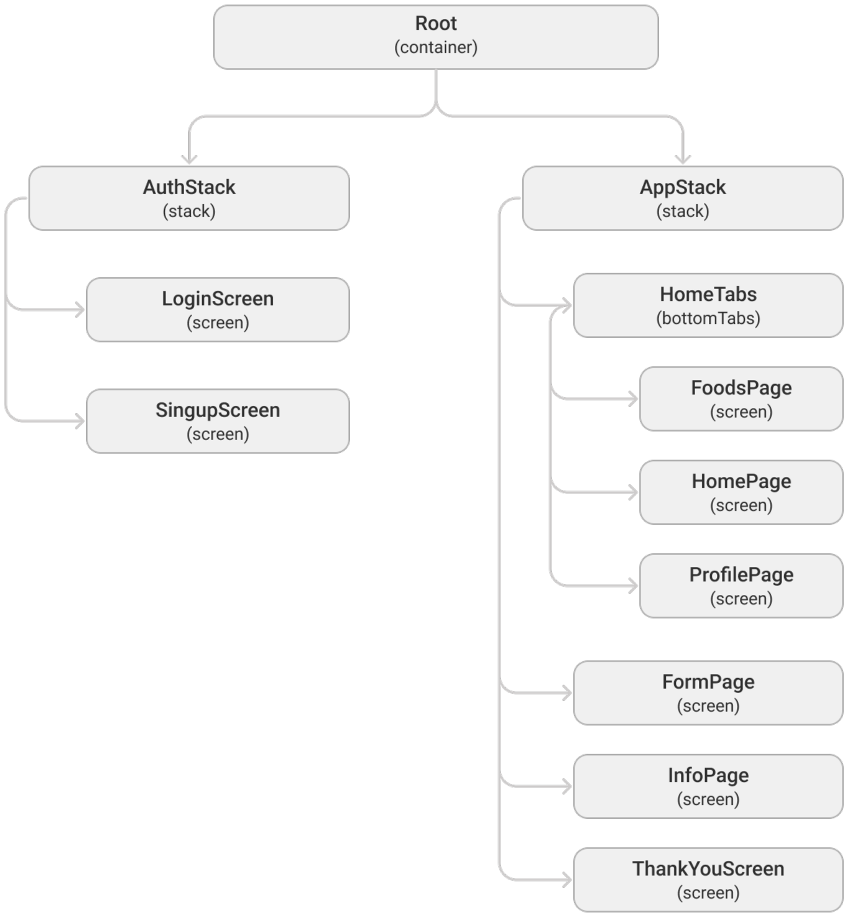

Figure 3: Page hierarchy schema.

4.3 User Experience and Interaction Flow

The user experience design of the FlavorCharter mobile application placed significant emphasis on visual clarity, intuitive navigation, and the accurate representation of sensory dimensions. Given the scientific nature of the study, the interface was designed to make complex perceptual evaluations accessible to a non-specialist audience while preserving experimental rigor. Each interaction was developed to reduce cognitive effort and guide participants through the process with clear feedback and minimal visual clutter.

4.3.1 User Journey

The user journey was structured as a guided, multi-session process designed to ensure proper consent, controlled access, and consistent data collection. Unlike a single-session survey, FlavorCharter allows participants to contribute multiple ratings over time, progressively enriching the dataset until the required quota is reached. The complete flow—from recruitment to submission—is illustrated in

1. **Recruitment and App Download:** Participants are first invited through institutional recruitment channels. They receive a link to download the FlavorCharter mobile application, which serves as the primary entry point to the study.

2. **Consent and Approval:** Before registration, participants are required to review and sign a digital informed consent form accessed through an external link. Upon submission, the consent information is processed and approval is granted, enabling the participant to proceed with account creation and authentication within the application.

3. **Authentication and Onboarding:** Following approval, participants create an account through Firebase Authentication, which establishes a secure and unique identity within the system. This step initiates the onboarding sequence, a guided process that requires participants to complete all preliminary steps before beginning the study. The onboarding flow ensures that users provide the necessary demographic and dietary information, review the study guidelines, and gain a clear understanding of how to navigate and use the application correctly.

4. **Demographic and Dietary Information:** The onboarding process begins with a structured form that collects demographic and dietary metadata, including age, gender, ethnicity, race, cuisine preferences, diet type, and eating habits. Completing this step is mandatory, and participants are prompted to provide all required information before proceeding to the next phase of the study.

5. **App Guide:** Before starting the first evaluation, participants are guided through an in-app tutorial that outlines the overall flow of participation and explains how to complete the study. The tutorial introduces each step of the process and provides clear instructions on how to evaluate items accurately.

6. **Food Selection and Sensory Rating Sessions:** After completing the onboarding and tutorial, participants can begin their first sensory evaluation session. In each session, they select a food item from a categorized list organized by dominant flavor dimension (e.g., Sweet, Salty, Bitter, Umami). Each selected food is evaluated across the ten perceptual dimensions using selector-based inputs.

7. **Iterative Participation and Data Update:** Participants can re-enter the application at any time to continue contributing new data or to refine existing entries. The app allows users to select and rate additional foods, update previously submitted ratings, and modify demographic or dietary information when necessary. This flexibility supports an iterative participation model in which users progressively complete their evaluations over multiple sessions. To ensure sufficient sampling across flavor categories, participants are required to rate at least three foods per dominant dimension before final submission is enabled.

8. **Submission and Completion:** Once the participant has completed the required number of ratings, the app prompts submission. A confirmation screen marks the end of the contribution, while additional sessions remain available for voluntary further input.

4.4 Core Functionalities

The core functionalities of the FlavorCharter mobile application were designed to support the complete cycle of participant interaction. Each functionality is tightly integrated into the system's architecture, ensuring consistency between user experience, data structure, and backend synchronization. Together, these components allow the application to serve as a robust instrument for collecting structured sensory data at scale.

At a high level, the mobile platform performs six primary functions:

1. **Approval and Registration:** Ensures that only eligible participants can access the study.

2. **Sensory Dimensions Rating:** Enables participants to evaluate foods across ten perceptual dimensions using an intuitive graphical interface based on radar and selector components.

3. **Demographic and Dietary Form:** Captures key participant attributes, including demographic and dietary factors, through a structured, validated form.

4. **Food List Management:** Manages a locally stored list of foods organized by dominant flavor dimension.

5. **Conclusion of the Study:** Handles completion logic, including validation of the minimum required ratings per category and final submission confirmation.

6. **Account Deletion:** Allows participants to request the permanent removal of their data.

The following modules describe the individual modules in detail, outlining their design rationale, technical implementation, and integration within the overall system architecture.

4.4.1 Approval and Registration

Before participants can register or access the application, they must complete an approval workflow designed to guarantee that only individuals who have provided informed consent can participate in the study, as required by IRB guidelines. This workflow integrates external and internal services into a secure pipeline that connects the Qualtrics survey platform, a webhook hosted on Render, and the Firebase backend.

From the registration page of the mobile application, users are first prompted to open an external link to the consent form. This step is mandatory: the app prevents opening of registration form until the consent form has been opened. The form, hosted on Qualtrics, collects the participant's email address, screening information and consent agreement, which serve as the identifiers for subsequent approval.

Upon submission of the consent form, Qualtrics triggers an automated *Workflow* that sends a secure `POST` request to a custom webhook hosted on Render. This webhook, implemented in `Node.js` and `Express`, verifies the request's authenticity through an API key and writes a new record to Firebase Firestore under the `approved_users` collection, marking the participant as `approved: true`. This approval step effectively bridges the external consent process and the internal authentication mechanism.

To mitigate delays caused by Render's cold-start latency (which may reach up to 50 seconds in the free tier), the application sends a preliminary `GET /ping` request to the webhook as soon as the consent link is opened. This "pre-warming" strategy ensures that the server is active and ready to process the subsequent approval request almost immediately after the participant submits the form.

Once the user starts the registration process, the mobile application enters a short polling cycle that periodically checks the participant's approval status in Firestore. This mechanism compensates for the brief synchronization delay between the webhook update and database propagation, ensuring a seamless transition to the registration phase.

After the approval record (`approved: true`) is detected, the participant is authorized to proceed with account creation using Firebase Authentication. At this stage, the user provides an email and password, which are securely validated through Firebase SDKs. The resulting account is linked to the previously approved record, creating a verified and traceable identity associated with all subsequent data stored in Firestore.

This dual-step design, approval followed by authentication, ensures that every participant has explicitly provided informed consent before any data collection occurs. It aligns the system with ethical research practices and IRB requirements, while maintaining data integrity and participant verification throughout the study lifecycle.

4.4.2 Sensory Dimensions Rating

The sensory rating interface constitutes the core of the FlavorCharter mobile application, enabling participants to evaluate foods across ten perceptual dimensions — Sweet, Sour, Salty, Bitter, Umami, Fatty, Astringent, Aromatic, Texture, and Piquancy. The design of this interface combines visual expressiveness and simplicity of interaction, ensuring that complex sensory judgments can be captured in a standardized and accessible way.

The evaluation screen integrates two primary components: a graphical radar chart providing a visual overview of the participant's ratings, and a series of interactive selectors used to input individual dimension values. Each selector presents the dimension name with its descriptive extremes (e.g., "Mild" to "Hot") and a row of discrete numeric options ranging from 0 ("NA") to 5. The selectors are designed for clarity and precision, employing large, high-contrast touch targets and consistent color codes drawn from the global style palette. Color consistency is maintained throughout the application by referencing a shared color dictionary, ensuring perceptual coherence between the radar visualization, selectors, and other interface elements.

The radar chart serves as a dynamic visual summary, instantly reflecting the participant's input across all ten dimensions. It uses a D3-based rendering pipeline to generate wedge-shaped arcs proportional to the intensity of each rating, with labels distributed evenly around the circumference. The chart is intentionally minimal, focusing on clarity and proportional representation rather than decorative graphics. A more detailed discussion of its visual encoding design is provided in the Visual Analysis System chapter.

The process for saving ratings follows a unified logic that handles both new entries and updates within the same workflow. When a participant submits an evaluation, the application checks whether the selected food already exists in the local context. If so, the record is revised locally and synchronized with the remote database; otherwise, a new entry is created and stored. The `RadarContext` functions as a local data layer that mirrors the user's Firestore records, enabling immediate access to rated foods and real-time updates to the radar chart without redundant network requests. By combining local caching with asynchronous synchronization to Firestore, this design ensures a smooth and reliable experience even under unstable connectivity, while maintaining consistency between local and remote data across sessions.

Figure 4 illustrates the sensory dimension rating interface for two foods.

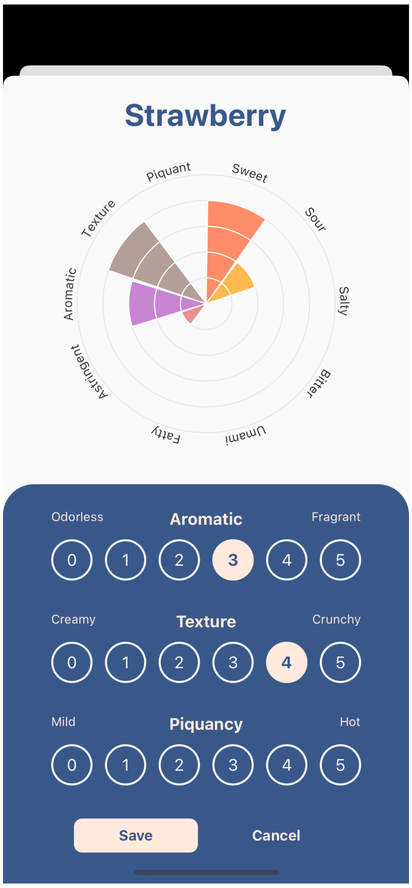

(a) Strawberry rating.

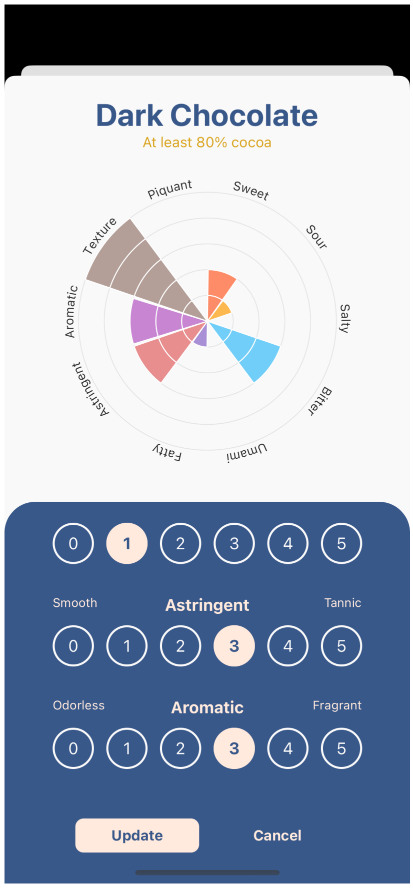

(b) Dark chocolate rating.

Figure 4: Two views of the food rating interface.

4.4.3 Demographic and Dietary Form

The demographic and dietary metadata form is designed to capture essential participant information before any sensory data submission. It provides a structured interface divided into three sequential sections, each focusing on a different aspect of the participant's background:

- **Step 1 — Personal Information:** collects basic personal details such as name and surname (collected for administrative completeness but not used in the analysis), along with age and gender.
- **Step 2 — Cultural and Geographic Background:** records nationality, race, and ethnicity to provide cultural and contextual information.
- **Step 3 — Dietary Habits:** gathers data on diet type, preferred cuisine, and includes an open-ended field for describing typical eating patterns.

To improve usability, a progress indicator at the top visually communicates advancement across the three steps, and each section must be completed and validated before the next becomes accessible.

The selectable options for categorical fields such as gender, nationality, race, ethnicity, diet, and cuisine are defined in a dedicated configuration module (`store/data/form-options.js`). This centralized approach ensures consistency across the form, facilitates future updates, and allows the addition or localization of new values without modifying the form logic itself. Each set of options is exported as a structured array of {`label, value`} objects, which are dynamically rendered by the corresponding dropdown components.

Validation occurs locally within the client application. At each step, the system checks both completeness and data type consistency—ensuring. profile Invalid entries are visually flagged in real time, and progression to the next section is blocked until all required inputs are valid. Once the final step is complete, the data are aggregated into a structured object and passed to the application's `UserContext`, which acts as a local store for user metadata. This local state is then asynchronously synchronized with Firebase Firestore through a secure API call, ensuring that updates are propagated to the cloud while maintaining a responsive user experience.

Participants can access and modify their demographic and dietary information at any time through the profile section of the application. Any changes are validated locally and synchronized with Firestore, ensuring that updates are consistently reflected across sessions and analyses.

This form is integrated into the onboarding process. When a user logs in or registers, the system verifies whether demographic data already exist in the local context or in Firestore. If none are found, the participant is automatically redirected to the demographic form before accessing the main study interface. This mechanism guarantees that no sensory data can be collected without a complete set of participant metadata, reinforcing both methodological rigor and IRB compliance.

Figure 5 illustrates the three steps of the demographic and dietary form interface.

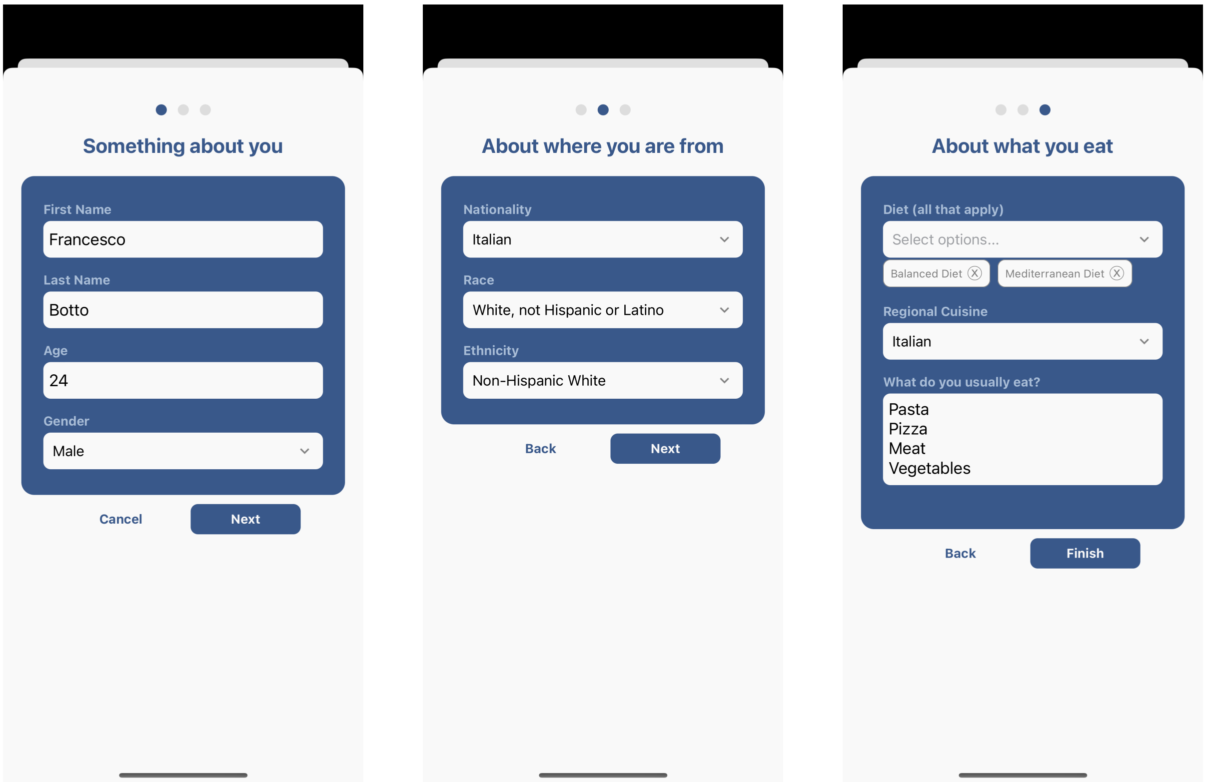

(a) First form page (b) Second form page (c) Third form page

Figure 5: Three steps of the demographic and dietary form

4.4.4 Food List Management

The food list constitutes as the reference dataset for sensory evaluations. Unlike other information stored dynamically in the backend, this list is implemented locally within the mobile application, ensuring immediate accessibility, uniform presentation, and consistent operation even under limited network connectivity.

Each food entry is defined in a dedicated configuration file (`store/data/foods-options.js`), structured as a collection of JavaScript objects containing key attributes such as:

- **name:** the label of the food item (e.g., *Strawberry*, *Soy Sauce*);
- **category:** a general classification describing the type of food (e.g., *Fruit*, *Condiment*, *Dairy*);
- **group:** an array of dominant flavor dimensions (e.g., `["Sweet", "Aromatic"]`), allowing a single item to belong to multiple sensory categories;
- **image:** a local visual asset ensuring that each item is displayed with a consistent and recognizable illustration;
- **comment (optional):** additional contextual information used to standardize tasting conditions (e.g., *"Unflavored, unsweetened, full fat"* for yogurt).

At runtime, the food list is dynamically filtered according to the sensory dimension selected by the participant. This allows focused evaluation within one taste category while also supporting multi-category foods that express more than one dominant flavor quality. The filtered results are rendered through reusable components that form a responsive grid layout. When a participant completes the evaluation of a food, the interface updates in real time to display a checkmark overlay on that item, clearly distinguishing completed evaluations from pending ones.

The category navigation bar enables quick switching between flavor dimensions, while the current progress for each dimension is computed in real time through the `FoodsContext`. The

`FoodsContext` returns, for each flavor category, the total number of foods already rated by the participant. This information is then used by the food list component to compute progress relative to the study's target of three foods per category. Based on this comparison, the interface dynamically displays a contextual message—such as *"You still need two foods"* or *"Category completed!"*—providing participants with immediate feedback on their advancement within each category. Although the detailed logic behind category completion is discussed in the following section, this real-time progress tracking encourages user engagement and ensures structured data acquisition across all sensory categories.

Figure 6 illustrates the foods list interface in two conditions: when a category is fully completed and when some foods are still missing.

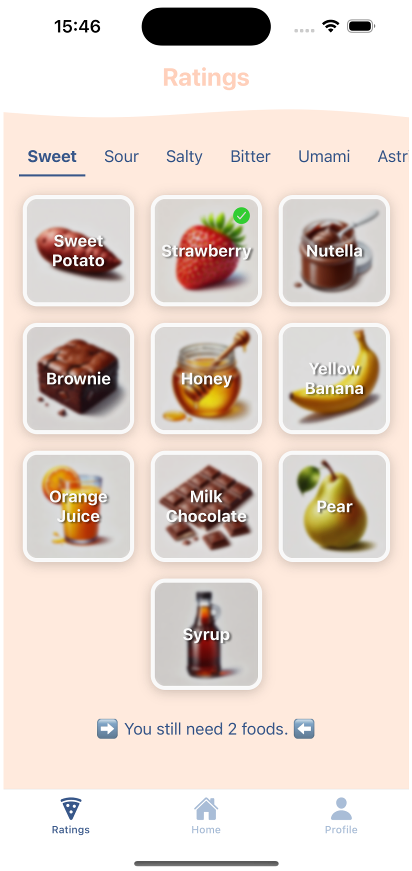

(a) Not completed category.

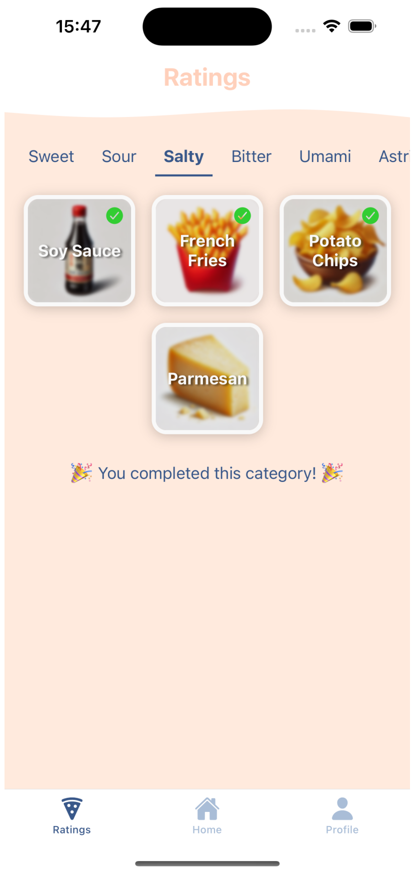

(b) Completed category.

Figure 6: View of the food list of two dimensions.

4.4.5 Conclusion of the Study

The conclusion of the study is determined through a structured completion logic that ensures each participant provides sufficient and balanced data across all sensory dimensions. The system defines a threshold of three foods per category as the minimum requirement for study completion. This number was selected to balance the participant's effort with statistical robustness, ensuring that the dataset captures intra-individual consistency while remaining feasible for untrained users. Foods associated with multiple categories (e.g., *Parmesan*, classified as Salty, Umami, Fatty, and Aromatic) contribute simultaneously to each corresponding category, thereby reflecting the multidimensional nature of real flavor experiences.

Progress tracking and completion validation are handled through the `FoodsContext` component, which aggregates rating data retrieved from the `RadarContext`. For every sensory dimension, the context computes the number of completed items by comparing the participant's rated foods with the static list of all available foods. Each category's completion status is updated in real time and stored in the local state as an array of objects containing the category name, the number of completed items, and the total available. Once all nine categories reach the required threshold, the system updates both its internal state and the participant's record in the cloud database to reflect study completion. This update adds the `all_categories_completed` flag to the user's data structure, indicating that all required evaluations have been performed and that the study has been successfully completed.

The participant's progress is visually represented in the application's home interface. Each flavor category is displayed as a tile within the `CategoryGrid`, which dynamically adjusts its label based on the number of rated foods per category. Once all categories are complete, the home screen overlays a semi-transparent confirmation view displaying a congratulatory message and a *"Save and Quit"* button, as shown in

When the participant confirms the end of the study, the application performs a final update to record the completion status. Instead of resubmitting data, this operation adds a completion flag to the participant's record in the cloud database, marking the study as finalized. At the same time, the user's approval is revoked through a secure backend process, preventing any subsequent access to the application. Once these updates are confirmed, the participant is redirected to the `ThankYouScreen`, which acknowledges the successful conclusion of the study and informs the user that their session is permanently closed.

Figure 7 illustrates the interfaces involved in the conclusion of the study process.

(a) Homepage. (b) Categories completed. (c) Thank you page.

Figure 7: Conclusion of the study interface

4.4.6 Account Deletion

The account deletion feature is implemented to comply with both platform policies and research ethics requirements. App stores such as Google Play and Apple App Store mandate that all published applications offer users a permanent account removal option. In addition, the Institutional Review Board (IRB) overseeing this study requires that participants be able to withdraw from the research at any time. This functionality therefore serves a dual purpose: ensuring regulatory compliance and safeguarding participant autonomy.

When a participant initiates the account deletion process, the application first displays a secure confirmation modal requiring password re-entry to verify the user's identity. This step is mandatory because Firebase enforces a re-authentication procedure before performing any sensitive operation, such as deleting an account. The participant's credentials are re-validated to ensure that the request originates from an authenticated session rather than from an expired or compromised token. Once the password is successfully confirmed, the system proceeds with two coordinated backend actions:

1. A deletion flag (`deleted = true`) is written to the user's document in the Firestore database, marking the account as withdrawn from the study.
2. The user's authentication record is permanently removed from Firebase Authentication, effectively disabling any future login.

In accordance with IRB policy, data that have been anonymized and cannot be traced back to individual participants may be retained for aggregate analysis. Conversely, any personally identifiable information (PII) linked to the deleted account is removed or irreversibly dissociated from the dataset during this process.

The account deletion interface is accessible directly from the *App Guide* page, where participants can confirm their decision through an in-app modal dialogue. This ensures transparency, ease of access, and adherence to both ethical and technical standards governing human-subject research in digital environments.

4.5 Database architecture and Data Handling

The database architecture of *FlavorCharter* provides the technological backbone that enables secure, scalable, and synchronized data management across the system's components. This layer connects the mobile front-end, where data are generated, with the backend infrastructure responsible for storage, authentication, and access control.

At the core of this architecture lies **Firebase**, a cloud-based *Backend-as-a-Service* (BaaS) developed by Google. Unlike traditional server-based backends, a BaaS abstracts the complexity of managing servers, APIs, and databases by providing pre-built cloud services that handle these tasks automatically. This approach allows developers to focus on data logic and user experience rather than infrastructure maintenance.

For this implementation, the system relies on Firebase's free-tier plan, which provides sufficient performance for the means of this study.

Within the *FlavorCharter* framework, Firebase underpins two essential cloud services:

- **Firestore**, a NoSQL document-oriented database that stores user metadata, sensory ratings, and study completion states;
- **Firebase Authentication**, which securely handles user registration, login, and session management;

The mobile application connects to Firebase through the official client SDKs. This direct integration allows the front-end to read and write data securely without the need for a custom server, while all access is governed by authentication tokens and Firestore Security Rules. Each interaction is transmitted through authenticated API calls that operate under the user's session identity, ensuring data integrity and privacy compliance.

4.5.1 Firestore Schema and Collections

The *FlavorCharter* framework relies on a document-oriented data model implemented through **Firebase Firestore**. Firestore's design aligns with the system's requirements for scalability, modularity, and secure user-level data isolation, making it particularly suited for a multi-user research platform such as *FlavorCharter*.

At the root level, the Firestore database is organized into two primary collections that separate participant authorization from study data storage (Figure 8):

- **`approved_users`:** This collection contains one document per participant who has completed the consent procedure. Each document is indexed by the participant's email and includes a Boolean field `approved`, which determines whether the user is authorized to access the mobile application. The collection is automatically updated through the Qualtrics–Render webhook pipeline, ensuring that only participants who have signed the informed consent form can proceed with registration and data submission.
- **`users`:** This collection contains one document per registered participant, identified by their unique Firebase Authentication ID (`uid`). It stores all user-specific data and study progress indicators, serving as the central repository for both sensory evaluations and participant metadata.

Within each `users` document, data are organized into three main logical components:

- `radarData` (subcollection): Contains one document per food evaluated by the participant. Each entry stores the food identifier, its sensory ratings across the ten perceptual dimensions, and optional user comments. The subcollection design allows the system to efficiently scale to multiple evaluations per participant while maintaining clean separation between users.
- `demographics` (subcollection): Holds static personal, cultural, and dietary metadata collected during onboarding. These values are grouped under a nested object named `profile` (document), which contains fields such as `age`, `gender`, `ethnicity`, `dietType`, and `preferredCuisine`. Since these attributes are collected once and rarely modified, they are stored as a single structured document rather than as a subcollection.
- **System-level metadata fields:** Several attributes are stored directly within the root of the user document to track study and account states:
  - `email` — used to link authentication credentials with the database record.
  - `hasSeenInfo` — marks completion of the initial app guide during onboarding.
  - `all_categories_completed` — indicates that the participant has rated at least three foods in each flavor category.
  - `save_and_quit` — signals that the participant has explicitly exited the study after completing all evaluations.
  - `deleted` — marks withdrawn or permanently deleted accounts while preserving data integrity.

This hierarchical schema provides several advantages:

- **Scalability and Performance:** Each user has an independent document, minimizing conflicts and ensuring efficient read and write operations even under concurrent access.
- **Efficient Queries and Cost Optimization:** The separation between user-level metadata and food ratings enables selective data retrieval, avoiding unnecessary network or billing overhead.
- **Separation of Concerns:** Each user's data are logically and physically isolated, simplifying maintenance, debugging, and compliance management.

The resulting structure balances flexibility and rigor, supporting both large-scale crowdsourced data collection and controlled clinical research environments.

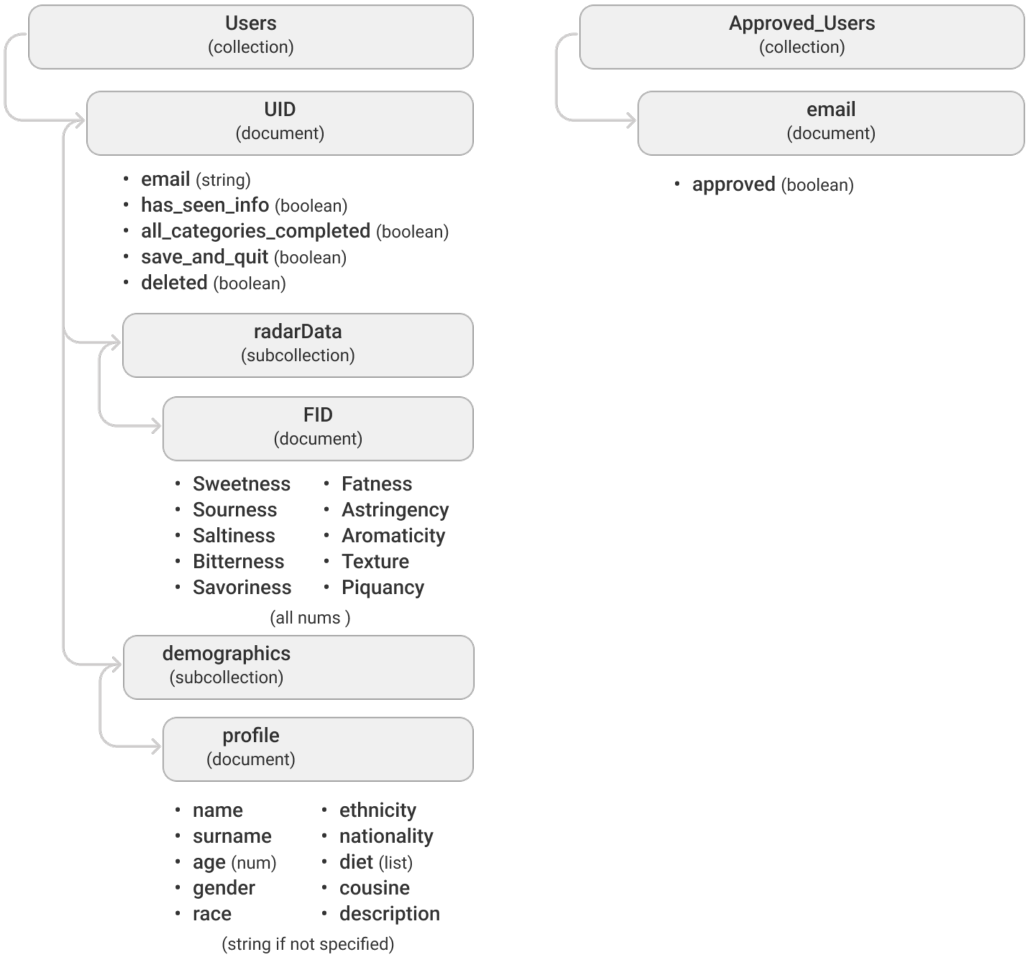

Figure 8: Firebase Firestore structure.

4.5.2 Firebase Authentication

User authentication within the FlavorCharter system is managed through **Firebase Authentication**. This module provides a secure and scalable mechanism to manage user identities, handle session persistence, and enforce access control across the mobile client and the Firestore database.

FlavorCharter employs a standard *email–password* authentication flow. When a participant registers through the mobile app, their credentials are securely transmitted to Firebase Auth, which issues a unique user identifier (`uid`). This `uid` acts as the primary key linking the user's authentication record to all related data in Firestore, ensuring a one-to-one correspondence between identity and stored information.

Access to the database is mediated entirely through Firebase's SDKs, which enforce authentication tokens for all read and write operations. Only authenticated sessions are allowed to access user-specific documents under the `users/{uid}` path, and all backend rules are designed to prevent cross-user data exposure. This security model eliminates the need for a custom authentication server, while maintaining compliance with institutional data protection requirements.

Finally, Firebase Authentication provides automatic token refresh and session persistence across app restarts, ensuring a seamless user experience while preserving data integrity.

4.5.3 API Specifications

The following list summarizes the core API functions implemented in the `util/http.js` module. Each endpoint or function interacts directly with Firebase Firestore or Authentication through the SDK. All operations are asynchronous, authenticated, and scoped by user identity (`uid`).

- **postRadarData(radarData)** Creates a new sensory evaluation document under the authenticated user's `radarData` subcollection.
  - *Request Body:* { `Food, Sweetness, Sourness, Saltiness, Bitterness, Umami, Fatness, Astringency, Aromaticity, Texture, Piquancy` }
  - *Response:* { `id` }
  - *Access:* Authenticated user only.
- **updateRadarData(id, radarData)** Updates an existing sensory record under `users/{uid}/radarData`.
  - *Request Body:* { `Food, Sweetness, Sourness, Saltiness, Bitterness, Umami, Fatness, Astringency, Aromaticity, Texture, Piquancy` }
  - *Response:* { `status` }
- **getRadarData()** Retrieves all sensory evaluations for the authenticated participant.
  - *Response:* `[` { `Food, Sweetness, Sourness, Saltiness, ...`}, { `Food, Sweetness, Sourness, Saltiness, ...`} `]`
- **postDemographics(demographicsData)** Saves or updates the user's demographic and dietary metadata under `users/{uid}/demographics/profile`.

- *Request Body:* { `name, surname, age, gender, nationality, ethnicity, race, diet, cuisine, description` }
- *Response:* { `status` }

- **markAllCategoriesCompleted()** Sets the flag `all_categories_completed = true` in the user document.
  - *Response:* { `status` }

- **markSaveAndQuit()** Marks the study as concluded by setting `save_and_quit = true`.
  - *Response:* { `status` }

- **checkApprovedUser(email)** Verifies whether an email is authorized in the `approved_users` collection.
  - *Response:* { `approved` }

- **revokeUserApproval(email)** Revokes a participant's approval after completion or withdrawal by setting `approved = false`.
  - *Response:* { `status` }

- **deleteAccount(password)** Deletes the Firebase Authentication account and flags the Firestore record with `deleted = true`.
  - *Request Body:* { `password` }
  - *Response:* { `status` }

  - *Access:* Requires re-authentication.

- **createUserWithEmailAndPassword(email, password)** Registers a new participant once their consent has been approved.
  - *Request Body:* { `email, password` }
  - *Response:* { `uid` }
- **signInWithEmailAndPassword(email, password)** Authenticates an approved participant and retrieves a valid session token.
  - *Request Body:* { `email, password` }
  - *Response:* { `uid, accessToken` }
- **sendPasswordResetEmail(email)** Sends a password reset link to the participant's registered email.
  - *Request Body:* { `email` }
  - *Response:* { `status` }
- **POST /webhook** Handles participant approval upon completion of the Qualtrics consent form.
  - *Request Body:* { `email` }
  - *Response:* { `status, code` }
  - *Access:* Restricted to trusted Qualtrics workflows authenticated via a secret key. The endpoint validates the `x-api-key` header against the server's environment variable

`SECRET_KEY`. If the key and email are valid, the server writes a record to Firestore under `approved_users/{email}` with the following structure: { `email, approved` }.

- **GET /ping** Keeps the Render webhook service active by sending a lightweight request at the beginning of the Qualtrics survey.
  - *Response:* { `status` }
  - *Access:* Public endpoint triggered automatically by the Qualtrics workflow (*Fire and Forget* mode), without requiring authentication or response handling.
- **POST /webhook** Called automatically by Qualtrics after a participant completes the consent form.
  - *Request Headers:* { `API-KEY`}
  - *Request Body:* { `email` }
  - *Response:* { `status, code`}
  - *Access:* Restricted to trusted Qualtrics workflows authenticated via a secure `api-key`.

4.5.4 Data Flow

The FlavorCharter data flow connects external services, Firebase backend, and the mobile client into a secure end-to-end system for user approval, authentication, data collection, and withdrawal. All interactions occur through asynchronous, authenticated API calls—either via the Firebase SDK on the client side or HTTPS requests handled by the Render webhook middleware.

- **Approval Workflow (Qualtrics → Render → Firebase)** The data flow begins when a participant completes the consent form hosted on Qualtrics. A `POST /webhook` request is automatically sent to the Render server, which validates the request through a secret key and writes a record to Firestore under `approved_users/{email}`. Before this, a lightweight `GET /ping` request is triggered to pre-warm the Render instance and prevent latency due to cold starts. On the client side, the application periodically calls `checkApprovedUser(email)` to confirm that the user has been authorized to proceed.
- **Registration and Authentication (Firebase Auth)** Before account creation or login, the system verifies that the participant's email address appears in the Firestore collection `approved_users`, confirming that informed consent has been recorded and approved through the external workflow. During registration, the application calls the funcion `checkApprovedUser(email)` before invoking `createUserWithEmailAndPassword(email, password)`. If the approval record is not yet present, a polling mechanism automatically rechecks the collection at regular intervals (every five seconds, up to 30 seconds total) to account for latency between Qualtrics submission and backend update. Only once approval is confirmed can the participant proceed to account creation, at which point a Firebase Authentication record is generated and a unique user ID (`uid`) is created as the root document in Firestore (`users/{uid}`).

  Subsequent logins follow the same logic: the app first validates approval status in the `approved_users` collection through `checkApprovedUser(email)`, then authenticates the user via the function `signInWithEmailAndPassword(email, password)`. This doublelayer validation ensures that only users with valid consent and active approval can access the system. If credentials are lost, participants can recover access using the function `sendPasswordResetEmail(email)`, which triggers Firebase's built-in password recovery flow.

- **Onboarding and Demographic Data Collection** After registration, the user is automatically directed to the onboarding phase, which ensures that all required information is provided before beginning the study. At login, the app verifies whether demographic data exist under `users/{uid}/demographics/profile` using the `getDemographics()` API. If no record is found, a mandatory modal appears on the Home Page, prompting the participant to complete the demographic and dietary form. This modal cannot be dismissed until valid data are submitted, ensuring that every participant provides the required personal information before accessing the study interface. The demographic form itself is submitted through `postDemographics(demographicsData)`, which saves structured metadata to Firestore and links the participant's email at the root of their user document. Once submitted, the data are stored locally in context and synchronized with Firestore, allowing future edits through the profile section.

  In addition to demographic data, the onboarding sequence includes an informational step that verifies whether the user has already viewed the study instructions. This is managed through the `hasSeenInfo` flag stored in the user's Firestore document. At the first access, the app calls `checkUserHasSeenInfo()`, and if the flag is absent or false, a guide modal appears prompting the user to open the instruction page. Once the participant reads the guide, the app records this event using `markHasSeenInfo()`, ensuring that the informational step is not shown again in subsequent sessions. Together, these mechanisms guarantee that every participant has both completed their demographic data and reviewed the study guidelines before providing sensory evaluations.

- **Rating Collection and Synchronization** During active participation, sensory evaluations are stored and updated using `postRadarData(radarData)` and `updateRadarData(id, radarData)`. Each record is saved within the subcollection `users/{uid}/radarData`. All submissions are first added to the local context (for real-time reactivity) and then synchronized asynchronously with Firestore. The data model ensures that local and remote copies remain consistent, even in cases of temporary offline access.
- **Study Completion** The completion status of each participant is monitored locally and in the database. When all categories reach the minimum number of rated foods, the flag `all_categories_completed = true` is written via `markAllCategoriesCompleted()`. Upon the participant's decision to finish, the study is closed by invoking `markSaveAndQuit()`, which adds the field `save_and_quit = true` to the user document. This marks the dataset as finalized and disables further data entry.
- **Approval Revocation and Account Deletion** After study completion or withdrawal, the app calls `revokeUserApproval(email)`, setting `approved = false` in `approved_users`. This revocation prevents future logins through Firebase Authentication. If the participant requests full account removal, `deleteAccount(password)` is executed. Firebase requires password re-entry for reauthentication before deletion; once confirmed, the account is re-

moved from Authentication and the Firestore document is flagged with `deleted = true`. This ensures that the identity is anonymized, while de-identified sensory data remain preserved for research integrity under IRB compliance.

4.5.5 Data Validation and Security Rules

Data validation and access control in FlavorCharter are implemented through a combination of client-side checks and backend constraints enforced by Firestore Security Rules. On the client side, every input field is validated locally before any database operation is triggered, preventing incomplete or malformed submissions from being uploaded. On the backend, Firestore applies strict authentication-based policies that restrict all reads and writes to the authenticated user's own documents, ensuring that no participant can access or modify another user's data.

In the client-side, all input fields undergo local validation before any write or authentication operation is permitted. During demographic submission, each step of the form checks completeness, type correctness, and formatting constraints. Age must be numeric; categorical fields such as gender, race, ethnicity, and cuisine must be non-empty strings; and multi-select fields such as diet must contain at least one element. Invalid inputs are visually highlighted and the user is prevented from continuing until the errors are resolved.

A similar validation strategy is applied during authentication. Before issuing any Firebase Auth request, the app verifies that the email contains a valid structure, that passwords meet minimum length requirements, and, during registration, that confirmation fields match their originals. If these conditions are not met, the authentication call is not triggered and the user is shown context-aware error messages. This ensures that malformed credentials never reach the backend and prevents unnecessary failed authentication attempts.

For sensory evaluations, data validation is handled primarily at the UI level. Users can only input values through constrained interface components (sliders or radar controls), guaranteeing that all submitted ratings fall within the expected ranges and maintain the correct structure. Since the food list is fully controlled and locally defined within the app, the interface only allows selection of predefined items, preventing the creation of inconsistent or undefined entries. Both updates and new records follow the same internal checks, ensuring consistent handling of every submission.

Server-side protection is implemented through Firestore Rules, a strict ruleset that governs all reads and writes in the database. Access to the `approved_users` collection is public for reads, enabling the app to verify participant eligibility even before authentication. Write permissions for this collection are restricted to authenticated requests, such as those issued by the webhook. The `users/{userId}` namespace is fully protected: both read and write operations require a valid Firebase Authentication session, and access is permitted only if the authenticated user's UID matches the document path. These constraints automatically cascade to all subcollections, including `radarData` and `demographics`, preventing horizontal access across accounts.

Certain fields are intentionally controlled only through system-level API calls rather than through client writes. These include `all_categories_completed`, `save_and_quit`, `deleted`, and `hasSeenInfo`. The app never writes these values directly during standard user interaction. Instead, they are modified exclusively through dedicated, authenticated API functions such as `markAllCategoriesCompleted()`, `markSaveAndQuit()`, and `deleteAccount()`, ensuring that participants cannot manipulate study-completion status or privacy flags.

Because every Firestore operation is tied to Firebase Authentication, no anonymous or cross-user access is possible. Even if a malicious user acquires an internal document path, Firestore Rules deny all reads and writes unless the UID matches. All communication with Firebase occurs over encrypted HTTPS channels, further protecting data during transmission.

4.5.6 IRB Compliance and Privacy Protocols

The design and implementation of the FlavorCharter mobile application strictly adhere to the requirements established by the Institutional Review Board (IRB) for the approved study (StudyID: STUDY2025-0477). All procedures for consent, data collection, storage, access control, and participant rights follow the ethical safeguards mandated by the protocol and by the UIC Office for the Protection of Research Subjects.

Participation in the study is entirely voluntary, and no user may access the application before providing informed consent. Consent is obtained externally through Qualtrics, where participants review the IRB-approved document and electronically sign the consent form. Only after this step is completed does the approval workflow authorize account creation within the mobile application, ensuring that no data are collected from individuals who have not formally agreed to participate, in full compliance with the protocol's requirements.

The application collects only the fields explicitly approved by the IRB: name, surname, age, gender, nationality, race, ethnicity, diet, regional cuisine, eating habits, and self-reported flavor ratings. No medical records, health identifiers, or diagnostic information are stored, and the interface does not allow submission of any data beyond what is listed in the protocol.

All data are stored in Firebase, which provides encryption both in transit and at rest. Access to identifiable data is restricted to authorized study personnel; no external developers, automated scripts, or third-party processors have access to participant information. In accordance with the IRB's data management plan, exported datasets are de-identified prior to analysis: identifiers such as name and email remain only in Firebase and are excluded from any research dataset shared outside the secure backend.

Participants retain the right to withdraw at any time, consistent with IRB guidelines and Stores constraints. If a subject requests withdrawal, their account can be deleted directly from the app interface. This operation removes the Firebase Authentication record and updates the corresponding Firestore document with a `deleted = true` flag, preventing further login attempts and marking the account as withdrawn from the study. In accordance with IRB policy, any data already exported and de-identified for aggregate analysis cannot be retroactively withdrawn.

The study poses minimal risk to participants, limited primarily to the potential loss of confidentiality. To mitigate this risk, the system enforces strict authentication, Firestore security rules, controlled data flow, encrypted storage, and restricted researcher access.

4.6 Testing and Feasibility Study

The testing and feasibility assessment of the FlavorCharter system was conducted through a two–stage evaluation strategy designed to validate usability, technical robustness, and the scientific plausibility of the sensory data collected. The goal of this evaluation was to ensure that both the mobile application and the visualization workflow could support a real study before broader deployment.

Two complementary forms of evaluation were employed:

Alpha Testing (At–Home Evaluation). This phase was conducted internally by three members of the research team, who used the mobile application autonomously from their homes, rating foods they already had available in their everyday environment.

Hallway Testing (Controlled Environment). A second evaluation was conducted with twelve participants in a supervised setting, where testers rated standardized food samples under direct observation.

Together, these two evaluation modes provide a comprehensive view of both technical performance and scientific feasibility, demonstrating that the system is usable, reliable, and capable of generating interpretable sensory data in both naturalistic and controlled conditions.

4.6.1 Alpha Testing (At–Home Evaluation)

The at–home evaluation was conducted internally by members of the research team, who used the mobile application autonomously in their everyday environment. This phase provided the first complete technical validation of the end–to–end data collection workflow, as well as early insights into the usability and robustness of the interface under naturalistic conditions.

The goal of the alpha testing phase was to validate the stability of the early application, verify that the onboarding and rating flow worked correctly across devices, and identify usability issues that emerge during unsupervised, at–home usage. Three members of the research team tested the app on their personal devices, providing rapid qualitative feedback that guided the next iteration.

The main issues identified were:

- authentication improvements (password reset, clearer error messages, cold–start delays);
- cross–platform UI inconsistencies (blur fallback, button rendering, press feedback);
- terminology and form clarity (labels, clearer food descriptors);
- navigation and interaction improvements (category progress clarity, post–rating prompts);
- rating interface refinements (N/A option, clearer dimension visibility);
- general usability enhancements (guide page, improved text readability, conclusion of the study).

This early internal evaluation ensured that major usability and workflow issues were resolved before proceeding to hallway testing.

Beyond usability testing, the at–home phase also served as a preliminary feasibility study. Since participants evaluated foods they already had available, only a small subset of items overlapped across the three responders. The items shared by the entire group were *Yogurt*, *Strawberry*, *Dark Chocolate*, and *Tonic Water*. Despite the limited overlap, these common foods allowed for an initial assessment of whether the mobile collection pipeline could capture coherent perceptual patterns under real domestic conditions.

Figure 9 shows the flavor–profile ratings for *Strawberry*, the food that exhibited the clearest cross–participant consistency. All three respondents identified sweetness as the dominant dimension, confirming a stable perceptual baseline. Some variability emerged in secondary attributes: one participant did not report any aromatic note, while another indicated the presence of a mild umami component.

Even with these natural variations, which may or not be due to variability in the food, the overall multivariate profile of *Strawberry* remained highly consistent across individuals, supporting the feasibility of collecting reliable sensory data outside of controlled environments.

A similar pattern was observed for *Tonic Water*, where all participants identified bitterness as the dominant attribute, with mild disagreement on secondary notes such as astringency or umami. These results further suggest that the FlavorCharter mobile workflow is capable of capturing stable, interpretable flavor–perception patterns even in small samples and uncontrolled at–home settings.

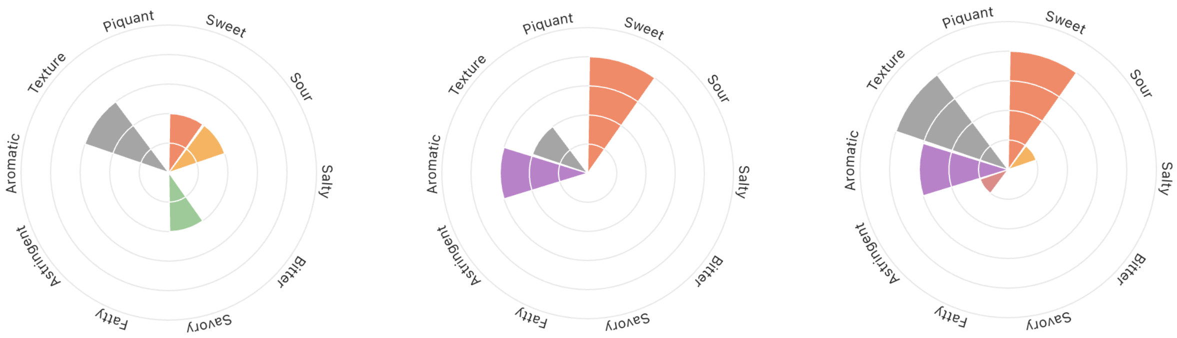

(a) Nationality: Chinese

(b) Nationality: Spanish

(c) Nationality: American

Figure 9: At–home ratings for Strawberry

4.6.2 Hallway Testing (Controlled Environment)

After the at–home phase, a second evaluation was conducted in a controlled environment ("hallway testing") with twelve participants. Unlike the autonomous at–home setting, this phase allowed direct observation of user interactions, enabling rapid identification of usability issues and more consistent comparison across participants, as all individuals tasted the same set of foods under identical conditions.

During the hallway testing sessions, participants were asked to register, navigate the onboarding flow, and complete at least one rating session while thinking aloud. This setup enabled real–time observation of user behavior and revealed issues that do not typically emerge during remote or unsupervised evaluation.

The main observations included:

- Participants frequently attempted to create an account before opening the consent link, suggesting that its placement and role were not immediately clear.
- Several users did not open the consent document at all, assuming it was optional or a non-interactive element of the consent form.
- Some users clicked form buttons multiple times, expecting faster visual feedback or assuming the first interaction had not registered.
- Users often forgot the meaning of certain sensory dimensions during rating.

In addition to usability observations, the controlled environment sessions provided an opportunity to conduct a second feasibility assessment using foods that were served directly during the session. All participants evaluated the same set of items, enabling a more controlled comparison of their flavor profiles and revealing whether the system could capture consistent perceptual patterns across individuals.

Three foods were overlapping between all twelve paricipants: *Dark Chocolate*, *Tonic Water*, and *Strawberry.* Despite expected individual variability, all three items produced stable and interpretable profiles across participants.

For *Dark Chocolate*, showed at Figure 10 and Table II, all respondents consistently identified high bitterness, minor secondary notes such as fatness and aroma were detected with moderate variability.

Almond, see Figure 11 and Table III, showed the strongest agreement on texture: all participants identified the characteristic crunchy component. Fatness was also rated with moderately high intensity, though with noticeably greater variability across respondents.

For *Strawberry*, see Figure 12 and Table IV, all participants identified sweetness as the dominant attribute, while variations appeared in aromaticity and texture, but the overall profile remained coherent across users, mirroring patterns observed in the at–home phase.

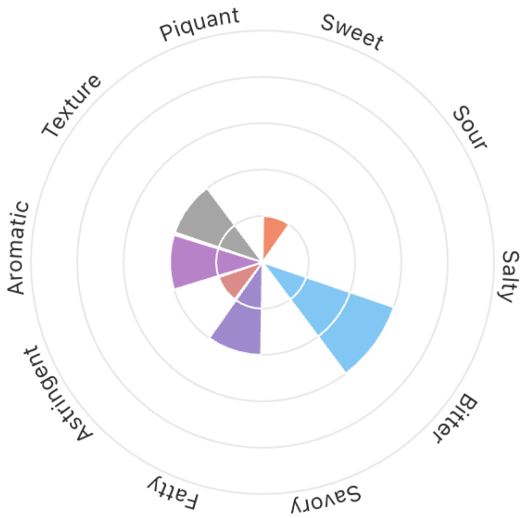

Figure 10: Dark Chocolate profile.

Dimension	Mean ± Stdev
Sweet	1.42 ± 1.00
Sour	0.50 ± 1.17
Salty	0.33 ± 0.65
Bitter	3.50 ± 1.09
Umami	0.08 ± 0.29
Fatty	1.58 ± 1.16
Astringent	1.00 ± 1.65
Aromatic	1.67 ± 0.98
Texture	2.17 ± 1.40
Piquant	0.00 ± 0.00

TABLE II: Dark Chocolate metrics.

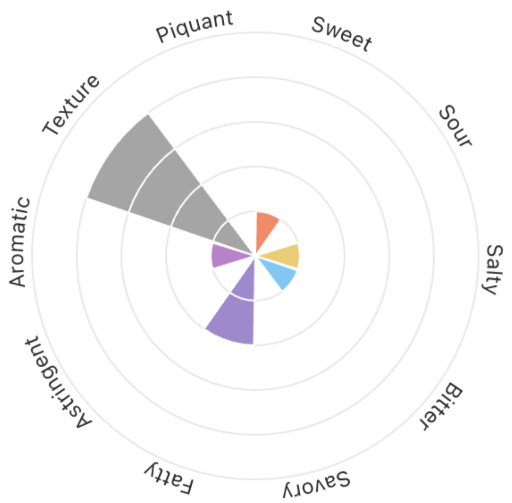

Figure 11: Almond profile.

Dimension	Mean ± Stdev
Sweet	0.75 ± 1.22
Sour	0.00 ± 0.00
Salty	0.83 ± 1.11
Bitter	0.83 ± 1.59
Umami	0.17 ± 0.39
Fatty	2.17 ± 1.47
Astringent	0.42 ± 0.90
Aromatic	0.50 ± 0.67
Texture	4.25 ± 1.48
Piquant	0.00 ± 0.00

TABLE III: Almond metrics.

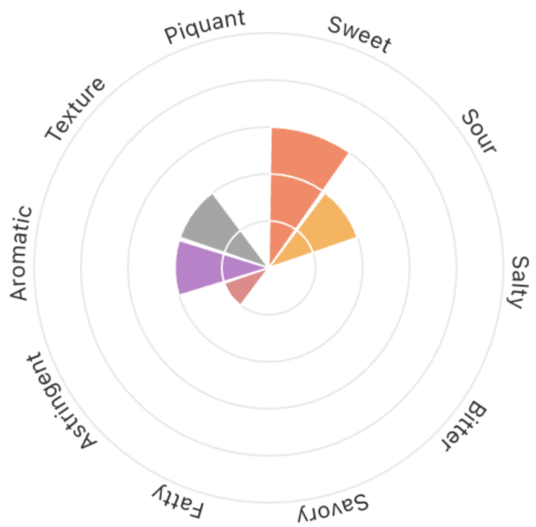

Figure 12: Strawberry profile.

Dimension	Mean ± Stdev
Sweet	3.00 ± 0.82
Sour	2.38 ± 1.19
Salty	0.23 ± 0.44
Bitter	0.38 ± 1.12
Umami	0.08 ± 0.28
Fatty	0.00 ± 0.00
Astringent	0.69 ± 0.75
Aromatic	2.23 ± 1.42
Texture	1.85 ± 1.28
Piquant	0.00 ± 0.00

TABLE IV: Strawberry metrics.

4.7 Summary

The Data Collection System establishes the technical and methodological foundation through which FlavorCharter acquires structured sensory data at scale. This chapter detailed how the mobile application integrates cross-platform development technologies (React Native + Expo), Firebase services for authentication and data storage, and an externally managed approval workflow to enforce IRB-compliant access control.

Through a guided onboarding process, including informed consent verification, demographic and dietary metadata collection, and an instructional walkthrough, the application ensures that all participants enter the study with consistent knowledge and complete prerequisite information. The core rating interface combines an intuitive selector–based input system with a real-time radar visualization, enabling participants to evaluate foods across the ten perceptual dimensions in a standardized and repeatable manner. Local context storage and real-time synchronization with Firestore guarantee reliability even under intermittent connectivity, while strict validation and security rules preserve data integrity and confidentiality across the entire pipeline.

The chapter also described the structured Firestore schema, the separation of approval and user data, the role of Firebase Authentication in enforcing identity-bound access, and the API abstractions connecting the mobile client, backend workflows, and external services. Together, these components form a secure and scalable infrastructure that supports continuous data acquisition while remaining aligned with IRB protocols and platform requirements.

Finally, the chapter introduced the validation process conducted through internal alpha testing and a two–phase feasibility study, demonstrating that the system is both usable in real-world conditions and capable of producing consistent, interpretable sensory data across participants.

With the Data Collection System fully implemented and validated, the FlavorCharter ecosystem is equipped to collect high-quality, standardized sensory data and to support the visualization and analytical methods introduced in the following chapter.

CHAPTER 5

VISUAL ANALYSIS SYSTEM

The Visual Analysis System is the analytical component of the FlavorCharter framework. While the mobile application is responsible for acquiring structured sensory data, this system provides the tools needed to interpret those data through interactive, multivariate visual representations. Its purpose is to transform raw ratings and demographic metadata into interpretable patterns, comparisons, and summaries that would not be accessible through numerical inspection alone.

This chapter describes the requirements guiding the design of the visualization tool, the abstractions applied to the underlying dataset, and the visual encodings selected to represent multidimensional flavor information. Beyond its technical role, the visualization layer also establishes the interpretive foundation upon which future clinical analyses can be built, linking the population baseline collected in this study to the sense-making processes clinicians and researchers rely on when evaluating flavor perception data.

5.1 Design Process

The design process for the visualization system was guided by a structured approach aimed at defining how flavor-perception data should be explored and interpreted. Before selecting specific visual encodings, the development focused on clarifying the analytical tasks, the constraints of the dataset, and the needs that emerged in collaboration with clinicians and researchers. The following sections outline the methods and principles that shaped this design workflow.

5.1.1 Design Approach

The project followed an *activity–centered design* methodology, grounded in the concrete analytical actions clinicians must perform when interpreting sensory data [27]. These activities were not inferred abstractly but emerged through direct discussions and iterative review sessions with clinicians and researchers at MD Anderson Cancer Center and the University of Iowa. From these sessions, a set of core interpretive tasks was identified, such as comparing demographic groups, evaluating deviations from expected patterns, exploring relationships among sensory dimensions, and interpreting variability at both individual and population levels. These real analytical activities shaped the structure of the system, determining which views were necessary, how they should be coordinated, and which interactions were essential for clinical sense–making.

5.1.2 Visualization Design Principles

The selection of visual encodings followed Munzner's *Why–What–How* framework [29]. High–level analytical goals (*Why*) were translated into task abstractions (*What*) including comparison, summarization, trend identification, and anomaly detection. These abstractions guided the choice of visualization and interaction strategies (*How*), ensuring that each design choice supported interpretability and analytical clarity rather than aesthetic preference.

5.1.3 Prototyping Workflow

The development process employed the *5–Sheets Design Process*, used to iterate rapidly through conceptual alternatives before implementation [30]. The five sheetsì, *brainstorm*, *initial* *designs*, and *realization*, served as a structured scaffold for exploring multiple design paths, validating assumptions with clinicians, and refining the interface before building high–fidelity prototypes.

This combined approach grounded the visualization system in real clinical activities, formal design principles, and iterative prototyping. The following sections illustrate how these methods informed the structure and interaction patterns of the FlavorCharter visualization tool.

5.2 Requirements Analysis

The design of the Data Visualization System is guided by the analytical needs of clinicians and researchers who must interpret multivariate sensory data collected through the FlavorCharter framework. The following requirements define the capabilities that the visualization tool must support.

5.2.1 Users

The visualization tool is primarily intended for clinicians working in oncology, especially those treating Head and Neck Cancer (HNC) patients. These clinicians must interpret perception profiles, explore foods and compare demographic-specific data against population-level baselines. The system is therefore designed to support quick comparisons, low cognitive load, and intuitive exploration.

5.2.2 Tasks

The system must enable the following analytical tasks:

- Inspect the complete flavor profile of a food, displaying mean, standard deviation and dimension-level distributions.

- Evaluate the reliability and accuracy of ratings for a selected food.
- Identify the dominant sensory dimension(s).
- Sort foods by number of ratings and accuracy.
- Group foods based on the dominant dimension(s) and combine it with sorting.
- Identify outliers at the level of individual food and sensory dimension.
- Compare the flavor profile of a food between a demographic group and the baseline population.
- Analyze how a dimension is perceived across different demographic groups.
- Determine which foods score highest on a specific dimension with lowest variability.
- Examine the demographic composition of the dataset and understand its influence on aggregated results.

5.2.3 Data Requirements

To support the tasks described above, the visualization layer operates on the baseline dataset collected from healthy individuals and stored in Firebase Firestore. Each participant record includes demographic metadata and a set of food evaluations, each containing intensity ratings for all sensory dimensions on a numerical scale. From this raw data, the system derives the abstractions required for visual analysis:

- **food-level summaries** such as mean ratings, variability measures, and dimension distributions;

- **rating volume and agreement indicators** used to assess reliability and interpretability;
- **demographic group structures** enabling stratified comparisons across subpopulations;
- **dimension-level outlier flags** to support data quality control.

Together, these elements allow the visualization layer to transform individual sensory ratings into structured, interpretable representations aligned with the analytical needs of clinicians and researchers.

5.2.4 Flow

The exploration process supported by the visualization tool follows a workflow centered on food selection rather than dimension-first navigation. A typical clinician workflow begins by browsing the list of available foods and selecting an item of interest. Once a food is chosen, the system presents its complete multivariate flavor profile, including mean intensities, variability, and rating counts.

From this view, the clinician may explore the distribution of ratings for each sensory dimension, inspect possible outliers, and identify the dominant perceptual attributes of the food. The workflow then extends to comparison tasks: users can contrast the selected food's profile across demographic groups or examine the overall trends of a chosen subgroup.

The flow is intentionally non-linear. Clinicians may begin from any food, switch to another, apply demographic filters, or broaden the analysis to examine patterns across the full set of foods. This flexibility supports open-ended inquiry and aligns with the exploratory nature of clinical interpretation.

5.2.5 Nonfunctional Requirements

Beyond functional capabilities, the visualization environment must satisfy several nonfunctional constraints essential for clinical adoption. Visual encodings must support users with varying levels of visualization literacy. Regarding performance requirements: filters, comparisons, and view updates must occur within milliseconds to maintain fluid interaction and avoid disrupting analytical reasoning. The system must also remain scalable as additional participants are added to the baseline dataset.

5.3 Data and Task Abstraction

This section formalizes the structure of the dataset and the analytical tasks the visualization layer must support. The organization follows the framework proposed by Munzner, which defines data abstraction and task abstraction as two complementary steps for mapping raw information to appropriate visual encodings [29]. By characterizing the dataset through structured taxonomies and identifying the operations users must perform on it, we establish the conceptual foundation that guides both the visual and interactive design of the system.

5.3.1 Data Abstraction

The dataset collected through the FlavorCharter application consists of multivariate sensory evaluations provided by healthy individuals. Following Munzner's framework, the data can be characterized through three taxonomic perspectives that define its type, structure, and analytical implications.

Taxonomy 1: Data and Dataset Types. The dataset consists of individual rating events, where each user evaluates a specific food across all sensory dimensions defined in the

FlavorCharter framework. Each of these events constitutes a single item in the dataset, resulting in a multidimensional table in which rows represent rating instances and columns correspond to the attributes associated with the user and the evaluated food. The attribute space is heterogeneous and combines perceptual intensity values with demographic descriptors. In detail, the dataset includes the following attribute types:

- ordinal numerical variables for the intensity ratings of each flavor dimension,
- categorical variables for demographic metadata such as gender, race, ethnicity, diet, and cuisine background,
- a sequential numerical variable representing the participant's age.

Table V shows all the attribute types and data characteristics.

The dataset is static, as all records come from a baseline collection of healthy individuals. There are no temporal components, and items do not maintain explicit relational links to one another. Each rating event remains independent, resulting in a flat, multidimensional structure consistent with the abstraction described for multivariate tabular datasets in the workbook.

Attribute	Attribute Type	Attribute Nature	Range
Flavor Dimension	Ordinal (Ordered)	Sequential	[0, 5]
Age	Quantitative (Ordered)	Sequential	[18, 99]
Gender	Categorical	Nominal	
Ethnicity	Categorical	Nominal	
Race	Categorical	Nominal	
Nationality	Categorical	Nominal	
Diet	Categorical	Nominal	
Regional Cuisine	Categorical	Nominal	

TABLE V: Attribute types and data characteristics.

Taxonomy 2: Spatial vs. Abstract. The dataset is entirely abstract, with no inherent spatial or geographic embedding. Sensory dimensions represent perceptual attributes rather than physical coordinates, and foods or users do not possess spatial relationships. As a consequence, visualization techniques must operate on abstract attribute spaces rather than spatial layouts.

Taxonomy 3: Dimensionality and Attribute Count. The dataset is multivariate and heterogeneous, combining ten sensory dimensions with multiple demographic attributes. All foods share the same sensory schema, forming an ensemble of comparable multivariate profiles.

The data contain no temporal component, and the attribute structure requires visualizations capable of handling high-dimensional, non-spatial information. In summary, the dataset can be defined as an *abstract, multivariate, multidimensional table of items with sequential attribute values.*

5.3.2 Task Abstraction

The analytical operations supported by the visualization tool can be described through a task abstraction process that identifies the types of questions users must answer and the actions they must perform on the data. Following Munzner's framework, tasks are characterized by their intent (such as querying, comparing, or summarizing) and by the target attributes or items they operate on. These abstractions clarify how the visualization layer should structure information and which views or interactions are required to support effective analysis.

Query. Query tasks focus on identifying specific properties or extracting direct summaries from the data. They operate primarily on individual attributes or subsets.

- Flavor profile: Inspect the complete multivariate profile of a food *(Action: Query → Summarize) (Target: Attributes → Distribution, All Data → Trends)*
- Dominant dimension: Identify the strongest perceptual attribute(s) of a food *(Action: Query → Identify) (Target: Attributes → Extremes)*
- Outliers: Detect anomalous ratings at food or dimension level *(Action: Query → Identify) (Target: All Data → Outliers)*
- Compare dimension: Compare the intensity of a single dimension across groups *(Action: Query → Compare) (Target: Attributes–One)*

Search. Search tasks involve locating foods or subsets that satisfy predefined criteria. They rely on sorting, filtering and browsing operations.

- Sort: Order foods by number of ratings or accuracy *(Action: Search → Lookup) (Target: Attributes → Many)*
- Grouping: Organize foods by dominant dimension(s) *(Action: Search → Browse) (Target: Attributes → Many → Similarity)*
- Sort dimension: Rank foods by intensity and variability within a dimension *(Action: Search → Locate) (Target: Attributes → Extremes)*

Analyze. Analyze tasks involve discovering patterns, relationships, or differences across groups. They operate on aggregated or multi-attribute structures.

- Accuracy: Assess the reliability of ratings for a food *(Action: Analyze → Derive) (Target: Attributes → Many)*
- Compare profile: Compare a food's multivariate profile between groups and baseline *(Action: Analyze → Discover) (Target: Attributes → Many)*
- Population: Examine demographic composition and its effect on aggregated patterns *(Action: Analyze → Discover) (Target: All Data → Trends, Attributes → Distributions)*

These task types map naturally to the hierarchical structure of analytical activities involved in exploring the dataset. The process includes:

- **Exploring the overall flavor landscape:** inspecting global distributions, understanding population composition, and gaining an overview of how flavors are perceived. This includes activities such as examining aggregated multivariate profiles, exploring how individual dimensions vary across the population, and understanding large-scale demographic trends.
- **Organizing and structuring food profiles:** sorting and grouping foods based on perceptual attributes, dominant dimensions, or consistency metrics. This covers operations such as ordering foods by statistical properties, categorizing them by their prevailing sensory characteristics, and identifying groups of items with similar perceptual patterns.
- **Analyzing individual food profiles:** examining multivariate ratings, variability, distribution shapes, and potential outliers for a specific item. This involves evaluating the reliability of ratings, inspecting dimension-level distributions, and detecting anomalous or inconsistent evaluations.
- **Comparing demographic groups:** evaluating shifts in perception across subsets and determining how subgroup characteristics influence flavor evaluations. This includes comparing multivariate profiles between groups and baseline, or assessing how a single flavor dimension differs across demographic categories.

5.4 Visual Encoding

The visualization layer of the FlavorCharter system relies on a coordinated set of visual encodings designed to represent multivariate sensory data, population-level patterns, and group-specific differences in a clear and interpretable manner. Since the dataset is abstract, heterogeneous, and intrinsically multivariate, the visual design must support both high-level exploration and detailed comparison across foods, dimensions, and demographic groups.

The encodings adopted in this system follow a mapping between data attributes and visual channels that aligns with the analytical tasks identified in the previous section. Each view is tailored to a specific subset of these tasks, such as summarizing a food's flavor profile, inspecting variability, detecting outliers, or comparing demographic groups.

The following modules describe the visual encodings used in the system, outlining the rationale behind each representation and clarifying how it supports the analytical goals of clinicians and researchers.

5.4.1 Radar Chart

The radar chart (Figure 13) provides a compact representation of a food's flavor profile by arranging sensory dimensions along fixed angular axes and encoding perceived intensity through radial distance. Its purpose is to offer an immediate, shape-based overview of how a food is perceived across all taste attributes, enabling users to recognize dominant notes, weak dimensions, and the overall balance of the profile at a glance.

The encoding relies on a circular layout in which each sensory dimension is assigned a constant angular position (mark: wedge), while intensity is mapped to radial extent (channel: position). Color is used only as a categorical differentiator between dimensions, and no area is computed or interpreted, avoiding misleading associations between surface size and magnitude. Unlike traditional spider or polygonal radar charts, where connected lines can imply correlations between adjacent attributes, this design uses independent wedges to emphasize that the flavor dimensions in this dataset are not inherently ordered and bear no topological relationship to one another.

The data displayed in this encoding consist of the aggregated mean intensity ratings for each of the ten sensory dimensions. Each wedge corresponds to the average value computed across all participants who evaluated the food, rounded up to match the simplified representation used in the application interface.

Analytically, this visualization supports early-stage exploration and inspection tasks, including Query $\rightarrow$ Identify for recognizing dominant attributes and Search $\rightarrow$ Locate for quickly scanning the relative strengths of multiple taste dimensions. Its limited quantitative precision is compensated by its high perceptual salience, making it effective for rapid comparisons across many foods.

The choice of this encoding is motivated by the need for an intuitive, compact, and multivariate representation that remains visually interpretable at small sizes. Alternative visualizations such as bar charts or line-based spider plots either require more space or risk implying false correlations among dimensions. The independent wedge design provides a clear overview of intensity distributions without suggesting structural relationships not present in the data.

Within the system, this radar chart is used exclusively in the food overview panel, where each food is represented by a small icon summarizing its perceptual signature. To maximize recognizability at this reduced scale, the encoding is simplified: axis labels and gridlines are omitted, preserving only the wedge geometry and color coding. This minimalist variant ensures that flavor profiles remain legible even when many food items are displayed simultaneously.

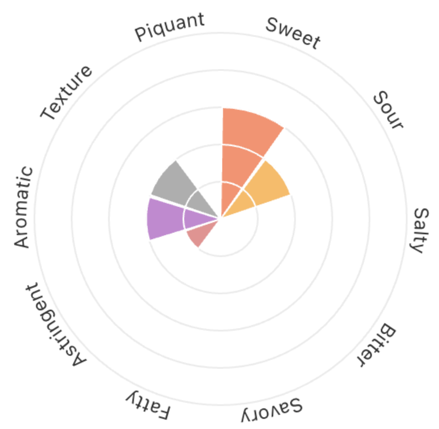

Figure 13: Radar chart encoding.

5.4.2 Distribution Radar Chart

The distribution radar chart (Figure 14) extends the basic radial representation to convey not only the average flavor profile of a food but also its underlying variability and the presence of outliers. Its purpose is to provide a more complete characterization of how participants perceive each sensory dimension, enabling clinicians and researchers to understand the stability, consensus, and dispersion of ratings beyond the aggregated mean.

The visual encoding combines multiple channels: angular position encodes the sensory dimension, radial position encodes rating magnitude, and shaded bands represent the distribution of observed values across participants. The outer line represents the aggregated mean profile. Individual outliers, ratings that deviate substantially from the majority, are drawn as distinct marks, emphasizing extreme or inconsistent evaluations that may reflect unusual perceptual responses.

The data feeding this encoding include, for each sensory dimension, the mean rating, the full distribution of observed values, and the set of outlier ratings, defined as values lying more than two standard deviations from the mean. This threshold ensures that only substantial deviations are marked visually in the chart.

From an analytical perspective, the distribution radar chart supports tasks such as Query → Summarize by presenting the multivariate pattern of intensity and variability, and Query → Identify by highlighting anomalous or extreme values. It also aids Analyze → Discover when users interpret variability patterns to assess the reliability of perception across the population.

The choice of this encoding stems from the need to represent multivariate distributions in a compact, holistic view. Alternatives such as per-dimension histograms or small multiples would require substantially more space and reduce the ability to perceive cross-dimensional patterns. The radial layout maintains comparability across dimensions while enabling the integration of variability and outlier information in a single visual form.

Within the system, the distribution radar chart is used in the food detail panel to provide a comprehensive representation of how a selected food is perceived. It serves as the detailed counterpart to the simplified radar glyphs shown in the food list, offering a richer depiction. The inclusion of outliers makes this view particularly relevant when assessing the consistency of ratings and identifying potential measurement anomalies.

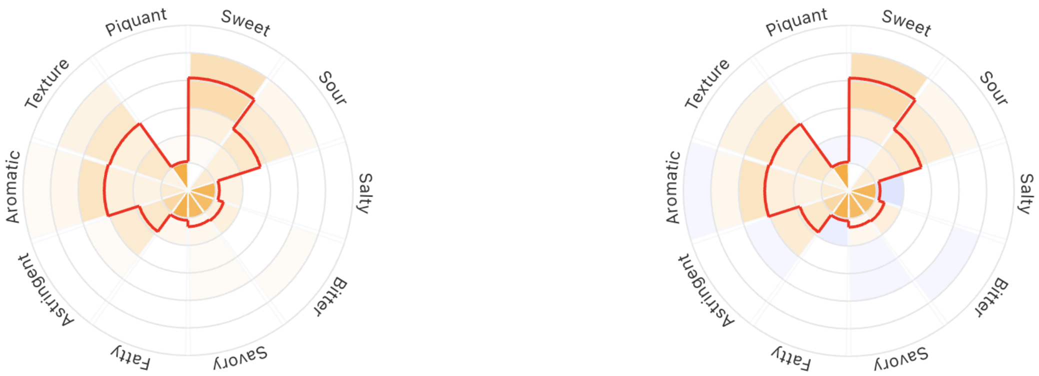

(a) No outlier detection.   (b) Outlier detection activated.

Figure 14: The two views of the distribution radar chart.

5.4.3 Z-Score Radar Chart

The Z-Score Radar Chart (Figure 15) provides a comparative view of how a demographic group's mean perceptions differ from the baseline population across all sensory dimensions [31]. Its purpose is to reveal systematic upward or downward shifts in flavor intensity, enabling clinicians to identify whether a specific group tends to over-perceive or under-perceive certain tastes relative to the reference population.

The encoding places each sensory dimension on a fixed angular axis, as in a standard radar chart, but the radial position does not represent raw intensity. Instead, it encodes the standardized deviation from the baseline. For each dimension, a point is positioned along the axis according to how the group's mean rating compares to the baseline mean: points lie on a central baseline ring when there is no difference, move outward for positive deviations, and shift inward for negative deviations. Color serves as an additional channel, with red indicating dimensions where the group scores above the baseline and blue indicating lower perception. The points are connected with a smooth closed curve, and the regions where the curve lies consistently above or below the baseline are filled in red or blue respectively, highlighting contiguous ranges of deviation.

The data required for this encoding include, for each sensory dimension, the mean rating of the baseline population, the mean rating of the selected group, and the baseline standard deviation. Deviations are computed as standardized differences:

$$Z_d = \frac{\text{compareMean}[d] - \text{baselineMean}[d]}{\text{baselineStDev}[d]},$$

which are then normalized into discrete radial levels used to build the glyph. This produces a compact representation of relative differences rather than absolute intensities.

From the perspective of task abstraction, this visualization supports Query $\rightarrow$ Compare and Analyze $\rightarrow$ Discover, enabling users to inspect group-to-baseline deviations across multiple dimensions simultaneously and to identify patterns of over-perception or under-perception.

The glyph makes directional differences salient and facilitates reasoning about demographic contrasts at the multivariate level.

The choice of this encoding is motivated by the need to represent deviations in a way that is both holistic and directionally interpretable. Traditional radar charts capture absolute profiles but do not convey whether a group differs meaningfully from a reference. An alternative approach would be to overlay two radar charts—one for the baseline and one for the selected group—but such a design quickly becomes visually cluttered, makes overlapping regions difficult to interpret, and shifts the focus toward absolute values rather than differences. Since the analytical interest lies in how the group deviates from the baseline, not in the exact magnitude of each raw intensity, the Z-Score Radar Chart provides a more effective and concise representation. It preserves the compactness of the radial layout while repurposing it to encode relative rather than absolute values, making deviations immediately visible without requiring users to compare two overlaid shapes.

Within the system, this visualization is used in two complementary contexts. At the food level, it appears in the group comparison panel, where it provides an immediate and interpretable depiction of how a selected demographic group deviates from the baseline profile for the currently selected item. At the dimension level, the same encoding is used to summarize how the group differs from the baseline across all foods. In this case, the standardized deviation is first computed separately for each food evaluated by the group, using the food-level formula already introduced,

$$Z_d^{(food)} = \frac{\text{compareMean}_d^{(food)} - \text{baselineMean}_d^{(food)}}{\text{baselineStDev}_d^{(food)}},$$

and these food-specific deviations are then aggregated across all foods in which the group has provided ratings. The global deviation for each dimension is obtained by averaging these values,

$$Z_d^{(global)} = \frac{1}{K_d} \sum_{f=1}^{K_d} Z_d^{(food_f)},$$

yielding a population-wide deviation profile that reflects how the selected demographic group differs from the baseline in its overall perception of each sensory attribute. In both settings, the chart complements the heat map representation by providing a shape-based summary of deviations and by making contiguous regions of over- or under-perception visually salient.

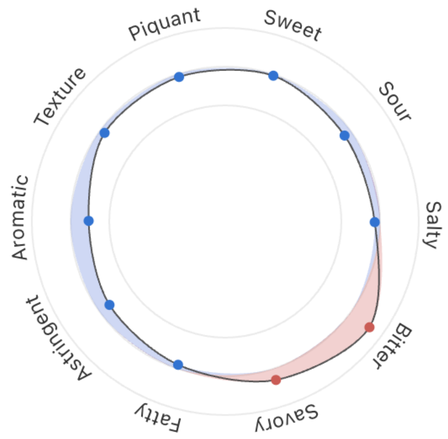

Figure 15: Z-score radar chart encoding.

5.4.4 Heat Map

The heat map (Figure 16) encodes deviations between a selected demographic group and the baseline population across sensory dimensions using a tabular, color-based representation. Its purpose is to provide a compact and immediately interpretable view of how perceptions differ across taste attributes, enabling users to identify dimension-level patterns and contrasts between demographic groups.

The encoding uses a matrix layout in which rows represent demographic groups and columns represent sensory dimensions. Each cell is filled with a color that reflects the direction and magnitude of deviation: warmer hues indicate that the group perceives the dimension more intensely than the baseline, while cooler hues indicate lower perception. Color hue and saturation therefore serve as the primary visual channels, while the spatial organization of rows and columns supports structured, cross-dimensional comparison.

The data used in this encoding are derived from the same standardized deviation logic employed in the Z-Score Radar Chart. At the food level, deviations reflect how a group's perception of a specific item differs from the baseline across each dimension. At the dimension level, the same per-food deviations are aggregated across all foods evaluated by the group, producing a population-wide summary of perceptual shifts. Using a consistent deviation metric across both contexts ensures interpretive continuity between the two views.

From a task abstraction perspective, the heat map supports Query $\rightarrow$ Compare and Analyze $\rightarrow$ Discover by enabling users to examine deviation patterns across multiple dimensions simultaneously and to identify systematic perceptual differences across demographic groups. Its structured grid layout facilitates efficient scanning along both dimensions and group categories.

The choice of this encoding is motivated primarily by its ability to provide an immediate overview of deviation patterns across all sensory dimensions, allowing users to quickly identify where perceptual differences between demographic groups are concentrated. This broad, dimension-by-dimension summary acts as an entry point for more detailed inspection, which is then supported by the Z-Score Radar Chart. While the radial encoding offers a holistic, shape-based representation of deviations, it is less effective for precise per-dimension analysis. The heat map complements it by presenting a dense yet readable grid of localized deviations, reducing cognitive load when scanning for specific attributes. Alternative solutions such as small multiples or parallel coordinates were considered, but proved less suitable for the fine-grained comparisons required by clinical users.

Within the system, the heat map appears in two locations. At the food level, it is used in the group comparison panel to show how the selected group differs from the baseline for the current item. At the dimension level, it aggregates deviations across all foods, offering a broader view of how the group's overall perception deviates from the baseline across the full set of sensory dimensions. In both contexts, the heat map operates in tandem with the Z-Score Radar Chart to support both detailed and holistic comparison tasks.

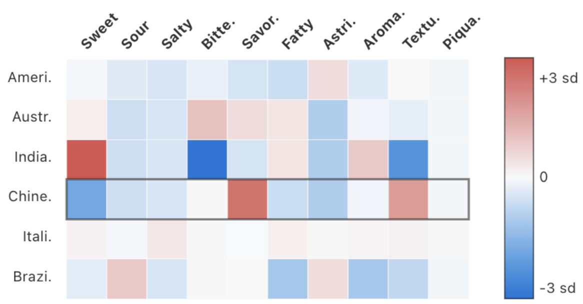

Figure 16: Heatmap encoding.

5.4.5 Scatter Plot

The scatter plot (Figure 17, Figure 18) addresses the need to evaluate pairwise relationships between quantitative attributes that characterize rating quality and dimensional stability. Its purpose in the FlavorCharter system is to reveal general trends in rating consistency and to identify foods that best represent a specific sensory dimension by combining high average intensity with low variability.

The encoding uses a set of points positioned along two orthogonal axes, mapping each item to a pair of quantitative attributes. No additional channels are introduced, allowing users to focus purely on spatial relationships such as clustering, separation, and distance from the origin. Depending on the analytical context, the scatter plot displays either agreement versus number of ratings or mean versus standard deviation for a selected sensory dimension.

The underlying data consist of aggregated food-level statistics. For agreement analysis, each point represents a food characterised by its number of collected ratings and its agreement score, which reflects how consistently participants evaluated it. For dimension-level analysis, each point encodes the mean intensity and standard deviation of a specific dimension across foods, enabling assessment of both central tendency and variability.

From the perspective of task abstraction, the scatter plot supports Query $\rightarrow$ Identify by making extreme or distinctive items immediately visible, and Analyze $\rightarrow$ Discover by enabling users to detect overall trends in consistency or stability. Its two-axis layout facilitates recognition of foods that deviate substantially from expected patterns or fall into desirable regions of high intensity and low variance.

The choice of this visualization is motivated by its effectiveness in capturing second-order relationships that cannot be inferred from univariate summaries. Alternative encodings such as bar charts or aggregated glyphs would obscure the joint behaviour of the attributes and limit the ability to reason about generalization trends or dimensional representativeness. The scatter plot, by contrast, preserves the structure of the joint distribution while remaining lightweight and easy to interpret.

Within the system, a single scatter plot component adapts its displayed attributes depending on whether a sensory dimension is selected. In the default overview, it visualizes the relationship between agreement and number of ratings to indicate whether additional sampling improves consensus or increases dispersion. When a dimension is selected, the plot automatically switches to mapping mean intensity against standard deviation, enabling users to identify foods that best represent that attribute through high intensity and low variability.

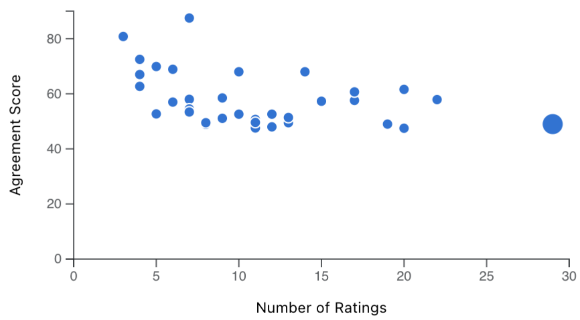

Figure 17: Scatter plot agreement/number of ratings.

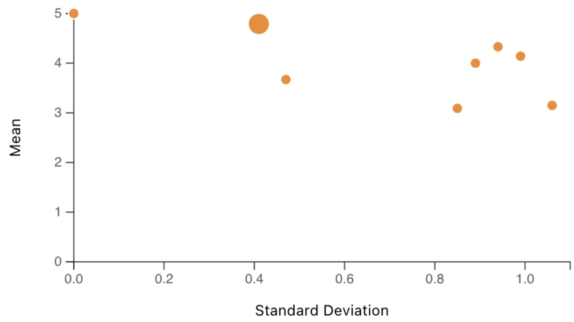

Figure 18: Scatter plot mean/standard deviation.

5.4.6 Pie Ring

The pie ring (Figure 19) provides a compact summary of demographic composition, helping clinicians interpret flavor–perception patterns in relation to how strongly each subgroup is represented in the dataset. Its purpose is not to analyze sensory ratings directly but to contextualize comparisons by showing whether observed deviations arise from well–sampled or sparsely represented groups.

The encoding uses a circular ring divided into colored angular sectors. Each arc corresponds to a demographic subgroup (mark), and its angle encodes the subgroup's percentage within the selected population (channel). Color encodes subgroup identity, while the hollow center improves readability at small sizes. Hover interactions reveal absolute counts and percentages.

The input data consist exclusively of aggregated demographic frequencies, either:

- globally across all participants, or

- restricted to users who rated the selected food.

From a task–abstraction perspective, the Pie Ring supports *Query* → *Summarize* by revealing distribution proportions, and *Analyze* → *Discover* by exposing demographic imbalance that may influence perceptual differences.

The choice of this encoding is motivated by its low cognitive load, strong part–to–whole semantics, and compact spatial footprint. Alternative representations such as bar charts offer higher precision but require more space and visual attention. For quick demographic awareness, the ring form provides intuitive interpretability without burdening the clinician.

Within the system, Pie Rings appear in two locations: the food–specific Pie Ring, showing the composition of users who rated the selected food, and the population–level Pie Ring, which contextualizes global Z–Score and heatmap analyses.

(a) Total population.

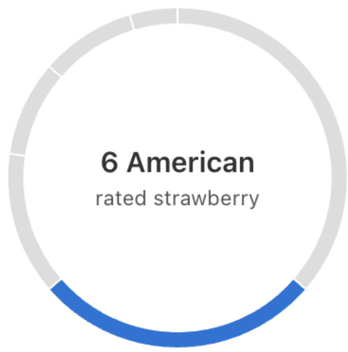

(b) Food–specific population.

Figure 19: The two views of the pie ring encoding.

5.4.7 Agreement Ring

The Agreement Ring (Figure 20) provides a compact visual summary of how consistently a food is perceived across participants. Its purpose is to indicate, at a glance, whether the underlying sensory evaluations are reliable enough to support meaningful comparison or whether substantial variability is present in the data.

The encoding consists of a circular ring divided into two angular segments: one segment represents the proportion of agreement and is colored according to its reliability level, while the remaining segment is rendered in a neutral grey to indicate the distance from full consensus.

The visualization uses angle to encode the agreement percentage, color to convey qualitative thresholds, and central text to report both the agreement value and the number of ratings available for the food. Color serves as a strong perceptual cue: green marks high reliability (75% or higher), orange indicates moderate agreement(50% or higher), and red signals low consistency among participants.

The agreement score displayed in the ring is derived directly from the variability of the ratings. For each sensory dimension d, agreement is computed from its standard deviation σ_d using:

$$A_d = \frac{100}{1 + \sigma_d}.$$

The overall food-level agreement is then obtained by averaging the ten dimension-specific scores:

$$A_{\text{food}} = \frac{1}{10} \sum_{d=1}^{10} \frac{100}{1 + \sigma_d}.$$

This final value, shown in the ring alongside the total number of ratings, provides a concise indicator of the stability and reliability of the food's multivariate sensory profile.

From an analytical perspective, the Agreement Ring supports rapid summarization and helps clinicians determine whether a food's profile is sufficiently stable before examining more detailed.

The design was chosen for its compactness, perceptual clarity, and visual harmony with the radial encodings used elsewhere in the interface.

Within the system, the Agreement Ring is located in the upper area of the right panel, next to the distribution radar chart and the scatter plot. It serves as an immediate reliability indicator that contextualizes the food's sensory profile and informs interpretation of subsequent analytical views.

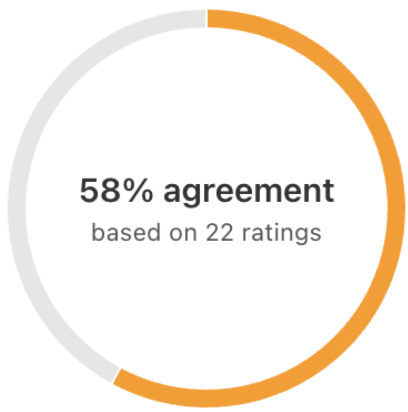

Figure 20: Agreement ring encoding.

5.5 Interaction Design

The interaction design of the visualization interface enables fluid exploration, comparison, and interpretation of multidimensional flavor–perception data. The system is organized as a coordinated collection of linked visual views, where user actions performed on one component propagate across the entire layout, ensuring a coherent analytical workflow.

5.5.1 Sorting and Grouping

The food list supports structured navigation through:

- **Sorting options**, allowing items to be ranked by *agreement score* or *number of ratings*;
- a **dominant–dimension filter**, restricting the list to foods whose strongest perceptual attribute corresponds to a selected flavor dimension (e.g., Sweet, Salty, Bitter).

5.5.2 Demographic Filters

Demographic filtering operates through two distinct controls that affect the interface in different ways.

1. Grouping variable. The first selector determines how the population is partitioned (e.g., by gender, age, nationality, diet). Changing this variable restructures all population–level views: the heatmap rows are regenerated according to the chosen demographic attribute, and the population distribution charts are re-computed using the corresponding groups.

2. Group value. Once a grouping variable is selected, the second selector chooses a specific subgroup (e.g., *male* within *gender*, or *american* within *nationality*). This value determines which subgroup is highlighted in the distribution charts and is used to compute subgroup-specific z-scores and mean for the comparative views (Z-Score Radar Charts and Heat Maps).

5.5.3 Comparison Views

The interface presents a coordinated set of side-by-side comparison views designed to contrast both food-specific and population-level trends. For each selected food, the system displays three complementary components:

- **Heat Map:** compares the overall population averages with the mean profile of each demographic group, allowing users to assess how a specific subgroup deviates from the normative baseline.
- **Z–Score View:** highlights normalized differences for the selected demographic subgroup, emphasizing which sensory dimensions are most atypical relative to the global mean.
- **Population Distribution:** shows how the selected subgroup is represented within the entire participant pool, grounding the comparison in demographic context.

5.5.4 Details on Demand

Hover-based tooltips provide contextual, fine-grained information without adding visual clutter to the interface. Each visualization exposes different details depending on its analytical role:

- **Distribution Radar Chart:** tooltips display the full empirical distribution for each flavor dimension, including counts per rating level and the computed mean.
- **Heat Maps (food–specific and dimension–level):** hovering reveals the exact z–score associated with each cell, enabling immediate interpretation of group–level deviations.

- **Population Pie Charts:** tooltips show both the absolute number of users in the selected demographic group and their percentage relative to the population.
- **Mean–Stdev Scatter Plot:** hovering a point displays the food name along with its precise mean and standard deviation values.

5.5.5 Multiple Coordinated Views

The visualization interface operates as a system of *multiple coordinated views*, in which all components share the same global application state. Any interaction performed by the user, selecting a food, changing a demographic grouping, choosing a specific subgroup, or activating a flavor dimension—immediately propagates across all visual modules.

Food selection. Choosing a food from the list updates all food–specific views: the distribution radar shows its sensory profile, the mean-stdev plot highlights it, the heatmap and -Score Radar Charts recompute group differences for that item, and the population chart shows how the selected demographic group is distributed among those who rated that food.

Demographic grouping variable. When the user selects a demographic category (e.g., gender, age, nationality), all comparison views reorganize according to that grouping: heatmaps reorder their rows, -Score Radar Charts show group–level deviations, and population distributions update to reflect the selected variable.

Demographic group value. Selecting a specific subgroup (e.g., “Female”, “25–34”, “Italian”) controls which group is used for Z-Score comparison. The interface highlights this subgroup in the population chart and updates both Heat Maps and Z-Score Radar Charts to display differences between the global baseline and the chosen demographic segment.

Flavor dimension switching. Activating a flavor dimension changes the analytical context of several views. The scatter plot transitions from the global *agreement vs. number of ratings* overview to a *mean vs. standard deviation* layout, showing only foods associated with that dimension. Simultaneously, radar charts, heatmaps, and z–glyphs update to reflect the chosen attribute, allowing focused exploration of dimension–specific perceptual patterns.

5.6 System Architecture

The visualization system adopts a lightweight client–server architecture that enables real–time exploration of the multidimensional flavor–perception dataset gathered through the mobile application (Figure 21). Its design separates analytical computation from interactive rendering, ensuring both performance and clarity. The **backend** is implemented as a Node.js server using Express and the Firebase Admin SDK. The **frontend** is a React single–page application that uses D3.js to render coordinated visualizations. Communication between the two components occurs through a set of lightweight HTTP API endpoints, each providing the aggregated metrics and demographic summaries required by the visualization system.

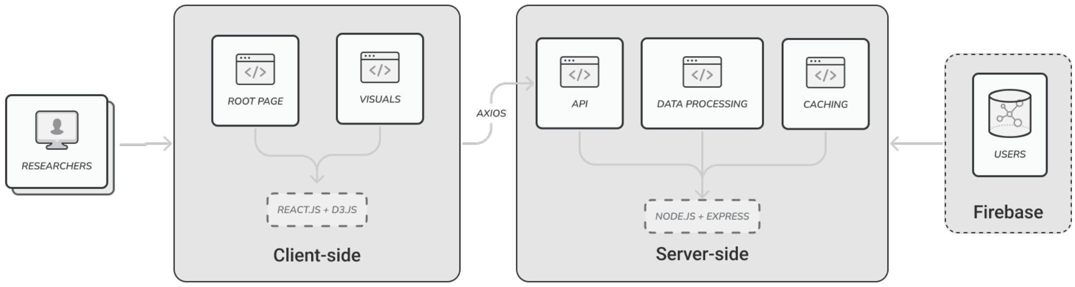

Figure 21: Visual Analysis System architecture.

5.6.1 Technology Stack

The visualization system is implemented as a two–layer architecture composed of a React frontend and a Node.js backend.

The client interface is built using:

- **React** for component–based UI design and state management.
- **D3.js** for all custom visualizations, including radar charts, heatmaps, Z–Score, pie charts, and scatter plots.
- **Axios** for HTTP communication with the backend API.
- **React Bootstrap** for lightweight layout utilities and responsive grid structure.

The analytical backend is implemented using:

- **Node.js + Express** to expose a lightweight REST API consumed by the frontend.

- **Firebase Admin SDK** for secure server–side access to Firestore.
- **In–memory caching** of all user data to avoid repeated Firestore reads and ensure fast computation of group statistics.

5.6.2 Data Source and Memory Caching

All backend computations rely on the data stored in Firebase Firestore, which serves as the authoritative source for both user demographics and sensory ratings. Access to Firestore occurs through the Firebase Admin SDK, authenticated via a service account key stored securely in the backend environment.

To ensure consistency across requests, the backend performs an initial synchronization step on startup: it retrieves all documents from the `users` collection, loads each user's demographic profile, and aggregates all rating records found in the corresponding `radarData` subcollection. As part of this process, personally identifiable fields such as `name`, `surname`, and `email` are explicitly removed before caching, ensuring that the backend operates exclusively on de–identified data in accordance with IRB privacy requirements.

The resulting dataset is stored in memory as a unified structure mapping each user ID to its demographics and complete rating history. Once this initial synchronization is complete, the backend does not continuously query Firestore; instead, all analytical endpoints operate directly on the cached dataset.

This in–memory caching strategy provides two key advantages:

- it eliminates redundant Firestore reads during interactive use, significantly reducing latency and database load;

- it ensures that all statistical summaries, comparisons, and group–level computations are derived from a stable and coherent snapshot of the dataset.

5.6.3 Data Processing

All analytical computations required by the visualization layer are performed server–side by the Node.js backend. Rather than storing pre–computed summaries, the system derives all statistical measures on demand from the in–memory dataset populated during the synchronization stage. This approach ensures that every request reflects a coherent snapshot of the data while remaining computationally efficient.

The backend exposes a set of processing routines that implement the core analytical logic:

- `computeAveragesAndStdevs` Computes the mean and population standard deviation for each flavor dimension. Values are aggregated over all ratings in the input set, and the output includes one pair `(mean, stdev)` per dimension.
- `computeDistribution` Computes the frequency of each rating value (0–5) for every flavor dimension and derives percentage distributions and an agreement score based on the dimension's variability. The output includes counts, percentages, and agreement metrics for each dimension.
- `detectOutliers` Identifies outlier ratings using a fixed–threshold rule based on standard deviation. A rating is marked as an outlier for a dimension whenever

$$|v - \mu| > k \cdot \sigma$$

where $k = 2$ in the deployed system. Output is organized per dimension and lists all values exceeding the threshold.

- `computeCharacterization` Determines which dimensions most characterize a food by evaluating both intensity and contrast. For each dimension d, the function computes:

$$\text{score}_d = \frac{\mu_d}{1 + \sigma_d}$$

and the contrast with the mean of the other dimensions. The final score is:

$$\text{FinalScore}_d = \text{score}_d \times (\mu_d - \overline{\mu}_{\neg d})$$

A dimension is marked as *characterizing* if $\text{FinalScore}_d > 2.0$. The function returns both per–dimension scores and the ranked list of characterizing dimensions.

These computations collectively enable all visual encodings used in the frontend

5.6.4 API Specifications

The visualization frontend communicates with the backend through a set of lightweight REST endpoints. Each endpoint performs a specific analytical operation over the in–memory dataset and returns a JSON response optimized for visual rendering.

- **/api/summary/:foodName** Returns all computed statistics for a specific food: means, standard deviations, distributions, agreement scores, and outliers.
  - *Query Parameters:* `filterKey`, `filterValue` (optional demographic filters)

  - *Response:* { `food, count, averages, stdevs, distribution, agreementScoreAvg, outliers` }
  - *Access:* Internal.

- **/api/stats** Returns global statistics about the dataset.
  - *Response:* { `totalUsers, totalRatings, distinctFoods` }
  - *Access:* Internal.

- **/api/group-distribution** Returns demographic distributions globally or for a specific food.
  - *Query Parameters:* `foodName` (optional)
  - *Response:* { `demographicKey: { group: {count, percent} }` }
  - *Access:* Internal.

- **/api/foodList** Returns the complete list of foods with aggregated statistics.
  - *Response:* { `food, count, averages, agreementScoreAvg, dominantDimensions` }
  - *Access:* Internal.

- **/api/group-means/:foodName** Computes per–group mean values for a selected demographic variable.
  - *Query Parameters:* `filterKey` (e.g., gender, nationality, diet, age)

  - *Response:* { `groupValue: {dimension: mean}` }
  - *Access:* Internal.

- **/api/group-dimension-zscores** Computes demographic Z–Scores by first computing food-specific standardized deviations relative to each food's baseline, and then aggregating these per-food deviations into a global dimension-level deviation profile for the selected demographic group.
  - *Query Parameters:* `filterKey`
  - *Response:* { `groupValue: {dimension: zScore}` }
  - *Access:* Internal.

5.6.5 Frontend Layout and Components

The frontend (Figure 22) is organized into a structured three–column layout that distributes the analytical workspace into four functional regions: a navigation panel on the left and three stacked analytical areas on the right (upper, central, and lower). Each region hosts a coordinated set of React components whose purpose is to support navigation, food–level inspection, demographic comparison, and population–level analysis. Each of the visual components receives updated data via props derived from the global application state maintained in `FlavorAnalyzer.jsx`.

Left Panel — Food Navigation. This region provides access to the full food catalogue and contains all controls for sorting, grouping, and searching. At the top, the interface includes a sorting selector (agreement or number of ratings), a search bar, and a dominant–dimension filter. Below these controls, the `FoodList.jsx` component displays foods as a responsive grid of tiles; each tile updates its appearance when selected and drives the content of all views in the right panel. This panel acts as the primary navigation mechanism of the interface.

Right Panel (Upper Area) — Food Profile. This region shows the population–level sensory profile of the selected food. It combines three components: `OutlierRadarChart.jsx` showing means, distributions, and optional outliers enabled by a checkbox; `PlotChart.jsx` that positions the food relative to others using either agreement–based or dimension–specific metrics depending on the current context; and `AgreementRing.jsx` summarizing overall agreement and rating count. Together, these views establish a detailed overview of the food prior to demographic comparison.

Right Panel (Central Area) — Group Comparison. Below the profile area, a dedicated control bar allows users to choose a demographic grouping variable (e.g. gender, age, nationality) and a corresponding subgroup. The selected values drive a coordinated set of comparison views displayed side by side. This block includes: `ZGlyph.jsx` showing subgroup–specific standardized differences; food–specific `HeatMap.jsx` comparing group means with population averages; and `PieChart.jsx` illustrating how the selected subgroup is represented among participants who rated the chosen food. This region supports food–level demographic inspection.

Right Panel (Lower Area) — Population Trends. The bottom portion extends the analysis from the selected food to global trends across all foods. This block includes `ZGlyph2.jsx` visualizing dimension–level Z–Scores for the chosen demographic subgroup, `HeatMap2.jsx` showing group means across all foods for the current grouping variable, and the second `PieChart.jsx` representing how the demographic groups are distributed in the overall population. A distinction exists between the components `ZGlyph` and `ZGlyph2`, and similarly between `HeatMap` and `HeatMap2`. The food–specific components (`ZGlyph`, `HeatMap`) receive raw population means and subgroup means directly from the backend and compute Z–Scores internally, since the baseline depends on the selected food. In contrast, the population–level components (`ZGlyph2`, `HeatMap2`) operate on Z–Scores that are precomputed by the backend across all foods. Because these Z–Scores rely on aggregated statistics derived from the entire dataset, they cannot be recomputed locally within the components; instead, they must be provided explicitly as input. This separation ensures that food–specific and population–level analyses remain consistent with their respective statistical baselines.

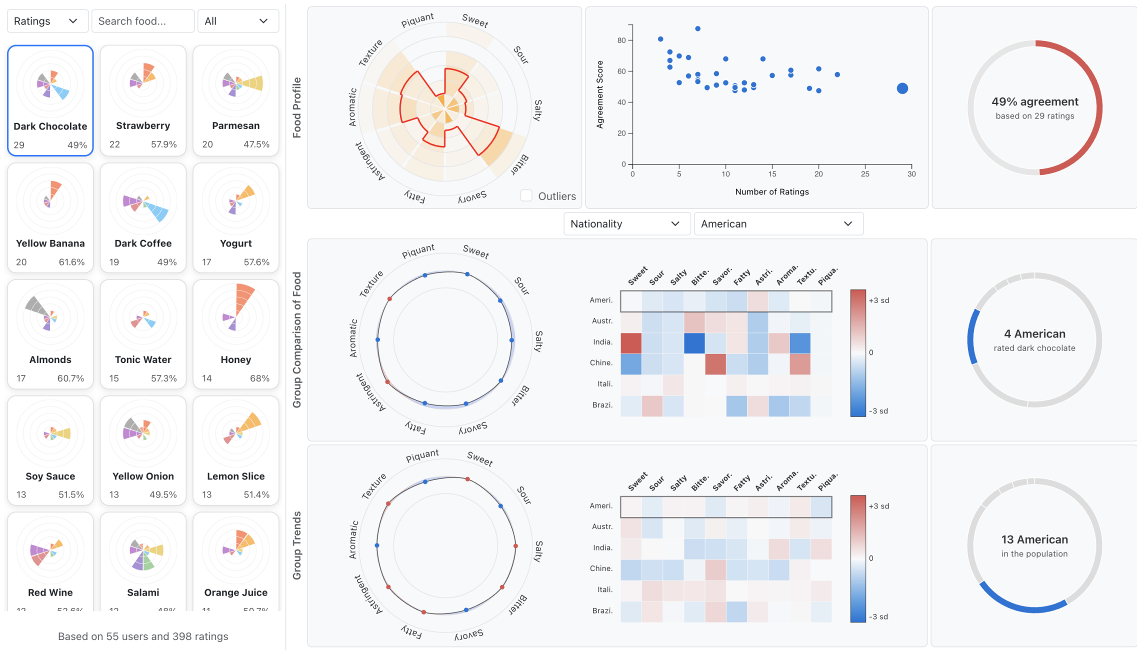

Figure 22: Visual Analysis System interface.

5.6.6 Data Flow

The data flow of the visualization interface follows a unidirectional, state–driven architecture in which all information originates from the analytical backend and is retrieved through a fixed set of REST API endpoints. The frontend coordinates when each endpoint is called based on changes in the global application state (`selectedFood`, `groupType`, `groupValue`, `dominantDimension`).

On startup, the client loads three global datasets through the following API calls:

- `GET /api/foodList/` retrieves the full food list with agreement scores and rating counts; these data populate the `FoodList` component and the `PlotChart` overview.
- `GET /api/group-distribution/` retrieves the population–level demographic distribution (across all foods); used by the population–level `PieChart`.
- `GET /api/stats/` retrieves aggregate statistics such as total users and total ratings; displayed in the footer of the left panel.

These datasets remain constant during the entire session and support sorting, filtering, and high–level demographic analysis.

Food selection. When the user selects a food, two food–specific endpoints are triggered:

- `GET /api/summary/{foodName}/` returns the complete statistical summary of the food, including means, standard deviations, distribution profiles, and outliers; this dataset feeds `OutlierRadarChart`, `PlotChart`, and the `AgreementRing`.
- `GET /api/group-distribution/?foodName={foodName}` returns the demographic breakdown of users who rated that food; used by the food–specific `PieChart`.

Changing the demographic grouping variable (`groupType`). Selecting a new demographic axis (e.g. gender, nationality, age) triggers two additional backend requests:

- `GET /api/group-means/{foodName}/?filterKey={groupType}` returns the mean values of the selected food for *every subgroup* belonging to the chosen demographic axis; passed to the food–specific `HeatMap` and `ZGlyph`.

- `GET /api/group-dimension-zscores/?filterKey={groupType}` returns dimension–wise Z–Scores computed *across all foods* for each subgroup within the chosen axis; used by the population–level `ZGlyph2` and the population `HeatMap2`.

These endpoints reorganize comparison views, update the food–specific heatmap, and regenerate the dimension–level heatmap and Z–Glyphs.

Changing the demographic subgroup (`groupValue`). Selecting a specific subgroup (e.g. "female", "25–34", "italian") does *not* trigger new backend calls. Instead, the frontend extracts the relevant subgroup values from the previously retrieved datasets:

- the subgroup row in the food–specific group means (for `HeatMap` and `ZGlyph`),
- the subgroup row in the Z–Scores returned by `group-dimension-zscores` (for population `ZGlyph2` and `HeatMap2`),
- the subgroup slice in the population distribution data (for `PieChart`).

Flavor dimension switching. Changing the active flavor dimension affects only the local state. It reorganizes the `PlotChart` into its mean–stdev mode and highlights relevant foods, but does *not* trigger any backend request. Radar charts, heatmaps and Z–Glyphs also update locally to emphasize the selected dimension.

React's `useEffect` hooks ensure that each API is called strictly when its dependencies change. All D3 visualizations receive updated information exclusively through props, maintaining a stable unidirectional flow:

User Interaction $\rightarrow$ State Update $\rightarrow$ API Request $\rightarrow$ New Data $\rightarrow$ Visual Update.

5.7 Usage Scenarios

The visualization system is designed to support clinicians and researchers in interpreting deviations from normative flavor perception. Because the mobile application collects detailed sensory ratings from healthy individuals, the interface allows users to contrast a patient's complaint with population–level baselines, demographic subgroups, and dimension–specific trends. The following usage scenarios illustrate how the system can be applied in clinical and research contexts, highlighting its relevance for head and neck cancer care and sensory–function monitoring.

5.7.1 Scenario 1

Evaluating altered perception of a specific food after treatment. A head and neck cancer clinician meets with a patient undergoing radiotherapy who reports that banana tastes significantly less sweet than before treatment. To determine whether this change is clinically meaningful, the clinician opens the system (Figure 23) and selects *Strawberry* from the food list.

The interface immediately displays:

- the healthy–population profile via the distribution radar chart,
- the mean–standard deviation position of banana among all foods,
- the group comparison views showing sweetness deviations across demographic groups.

The clinician filters the data by the patient's demographic characteristics, setting nationality to american. The heatmap and Z–Glyph show that healthy individuals in matching demographics consistently rate banana as above–average in sweetness.

Since the patient now reports abnormally low intensity relative to this normative pattern, the clinician concludes that the alteration is likely treatment–related rather than a typical perceptual variation. This supports clinical decision–making regarding taste–function monitoring and rehabilitation planning.

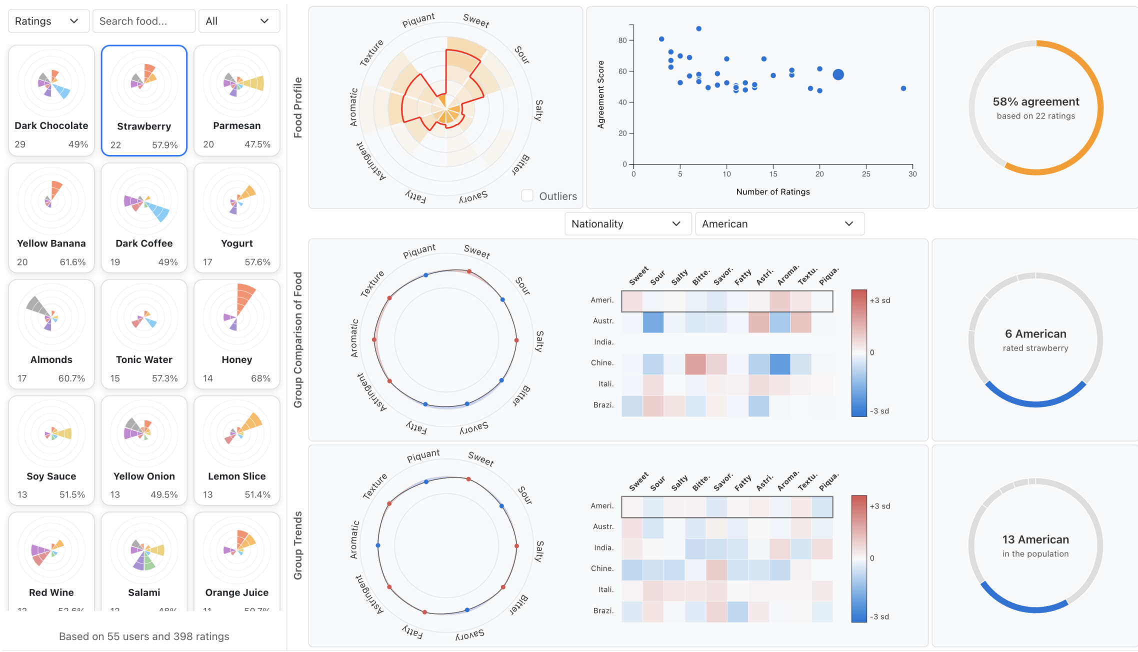

Figure 23: Interface of Scenario 1

5.7.2 Scenario 2

Assessing whether a patient's sensory complaint reflects a demographic trend. A nutrition specialist collaborating with an oncology rehabilitation team monitors a patient who states: "Sweet foods taste flat; I barely perceive any sweetness anymore." To determine whether this reduced sensitivity is common within the patient's dietary habits group, the clinician uses the demographic filters in the interface (Figure 24).

She selects:

- Grouping variable: race,

- Group value: hispanic or latino.

The system recomputes:

- population–level Z–scores across all foods for the selected subgroup,
- subgroup–specific mean profiles for each food,
- demographic distributions contextualizing representation.

The sweetness Z–scores demonstrate that this demographic group does not exhibit lower sweetness perception than the global baseline. Thus, the clinician interprets the patient's reduced perception as an individual impairment rather than a demographic trend. This distinction guides further assessment of taste dysfunction and supports appropriate clinical follow–up.

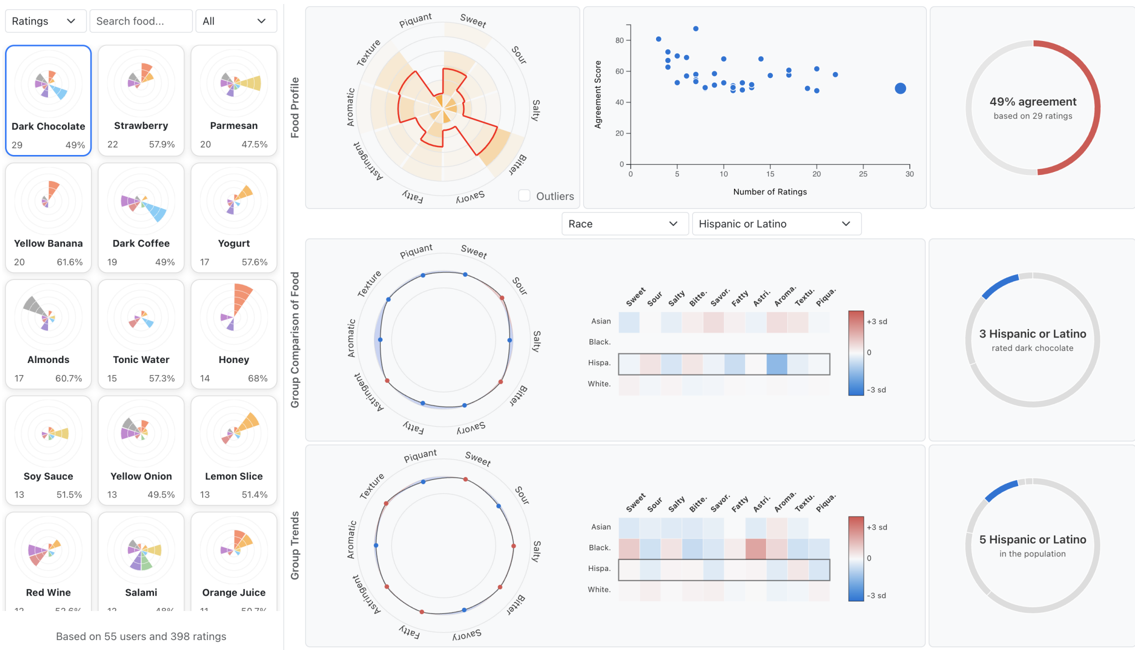

Figure 24: Interface of Scenario 2

5.7.3 Scenario 3

Selecting reference foods for sweetness assessment in clinical evaluations. A clinician designing a simplified sensory test aims to select foods that strongly and consistently express *Sweetness* for use in patient follow–up sessions.

Using the dominant–dimension filter in the tool (Figure 25), she restricts the food list to items where sourness is the strongest perceptual attribute. The scatter plot transitions to a mean–vs.–standard–deviation view specific to this dimension. From this plot, the clinician can immediately identify which foods exhibit the highest mean sourness while maintaining a low standard deviation, ensuring that the selected reference items are both intense and consistent across participants.

Foods such as *Honey* and *Syrup* emerge with consistently high sourness ratings and minimal inter–participant variance. These become suitable reference stimuli for clinical sensory evaluations because they offer:

- clear sensory expectations,
- low variability in the healthy population,
- stable comparison points for assessing recovery.

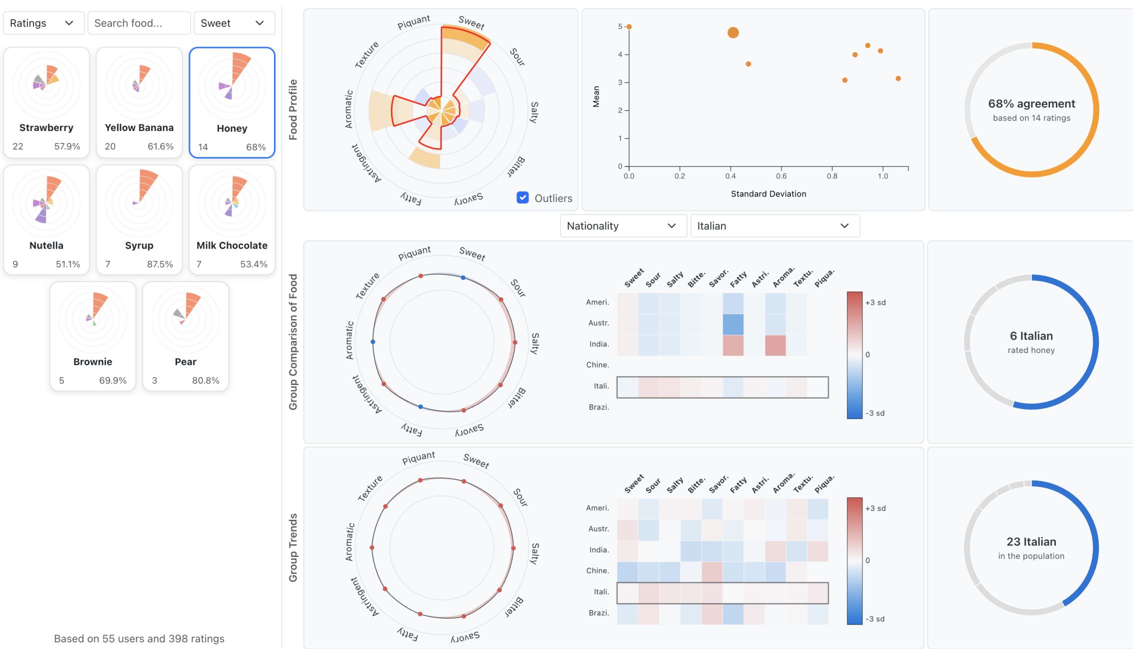

Figure 25: Interface of Scenario 3

5.7.4 Clinical Relevance

These usage scenarios illustrate how the visualization system functions as a decision–support layer. By grounding patient reports in quantitative population data, the interface helps clinicians distinguish between individual sensory changes and patterns that are typical for specific demographic groups.

The system is particularly valuable in three recurring clinical situations:

- **Evaluating perceived sensory loss:** clinicians can compare a patient's self–reported alteration with the expected intensity ranges observed in healthy individuals.
- **Understanding demographic context:** demographic filters reveal whether a sensory attribute is normally weaker or stronger within the patient's group, preventing misinterpretation of naturally occurring variability.
- **Selecting controlled stimuli:** foods with high intensity and low variability in a specific dimension can be identified as reliable reference items for monitoring sensory recovery.

5.8 Summary

The Visual Analysis System provides the interpretive layer of the FlavorCharter framework, transforming raw sensory ratings and demographic metadata into clinically meaningful multivariate patterns. This chapter presented the conceptual foundations of the system, beginning with the activity–centered design approach and the analytical requirements that emerged through collaboration with clinicians and researchers [27]. Through the application of Munzner's data and task abstraction framework, the visualization design was structured to support essential clinical activities such as comparing demographic subgroups, detecting deviations from population baselines, and evaluating the stability of flavor perception across sensory dimensions [29].

A coordinated set of visual encodings was introduced to represent the heterogeneous, multidimensional dataset in a structured and interpretable way. These encodings, combined with the interaction mechanisms enable users to fluidly transition between food–level inspection, dimension–specific exploration, demographic comparison, and population–level analysis.

The system architecture further supports these analytical workflows by separating computation from rendering. A Node.js backend performs all statistical processing on a synchronized, de–identified, in–memory dataset, while a React/D3.js frontend provides responsive, coordinated views driven by a unidirectional state model. Usage scenarios demonstrated how the system supports real clinical reasoning, from evaluating potential treatment–related sensory impairments to selecting reference foods for rehabilitation protocols.

The first functional prototype of the visualization system was reviewed by the clinical and research team involved in the FlavorCharter project, including collaborators from the MD Anderson Cancer Center and the University of Iowa. During these review sessions, clinicians provided qualitative feedback that guided both the refinement of the interface and the prioritization of analytical features, ensuring that the final design aligns with real clinical reasoning patterns and practical decision–support needs.

CHAPTER 6

DISCUSSION AND FUTURE WORK

The work presented in this thesis establishes the foundations of FlavorCharter, a unified framework for quantifying and interpreting human flavor perception through a combination of standardized descriptors, structured food-based stimuli, a mobile data-collection protocol, and a visual analysis system. Across its components, the project addresses a long-standing gap in both clinical and sensory research: the absence of a common, scalable, and data-driven language and baseline data for describing taste and flavor.

First, this work defines a ten-dimensional flavor schema grounded in culinary literature, biomedical research, and the clinical insights provided by head and neck cancer specialists. The resulting structure captures not only the classical taste qualities but also aroma, texture, and mouthfeel attributes that patients frequently report yet are rarely measured. Following on the same clinician-informed perspective, the schema is paired with a carefully curated food list designed to provide clear and controllable exemplars of each sensory dimension. Together with the collection of demographic and dietary data, these components enable the study of how factors such as age, nationality, and habitual diet shape flavor perception, supporting population-level comparisons and the analysis of group-specific patterns.

Building on this foundation, the FlavorCharter mobile application operationalizes the protocol into a reproducible and population-scale data-collection pipeline. Through this system, we collected preliminary real-world data from 55 users, yielding over 400 structured flavor ratings enriched with demographic and dietary information. The preliminary results demonstrate that individuals can reliably provide multidimensional flavor descriptions when guided by a shared vocabulary and consistent evaluation process.

Finally, the development of the visual analysis system illustrates how these data can be analyzed, compared, and contextualized through interactive multivariate representations, enabling both population-level summaries and the identification of individual variability.

6.1 Clinical and Scientific Implications

The FlavorCharter framework presented in this thesis has implications that extend beyond the immediate scope of data collection in healthy individuals. By defining a structured sensory language, establishing the foundations for population-level baselines, and introducing tools for visual analysis, this work lays the groundwork for future clinical, nutritional, and research applications. The main implications can be summarized as follows:

- **Improving Clinical Communication Through a Standardized Language.** The ten-dimensional schema provides a consistent vocabulary for describing flavor, addressing the long-standing issue of vague and subjective patient descriptions. Even though currently tested only in healthy individuals, this standardized language has the potential to support clearer, more actionable communication between patients and clinicians.
- **Providing the Foundations for Baselines That Empower Clinicians to Interpret Complaints.** The preliminary dataset from healthy users establishes early population-level baselines. These baselines, while not yet clinical references, demonstrate how structured sensory ratings may serve as quantitative anchors for interpreting deviations in

future patient-reported symptoms. This shift from subjective complaints to dimension-specific comparisons lays the foundation for more objective evaluations of taste dysfunction.

- **Supporting Dietary and Sensory Interventions Through a Visual Interface.** The visual analysis interface offers practical and clinically relevant value by enabling the exploration of how different foods and sensory dimensions behave across the healthy population. It allows researchers to identify which foods are most representative of specific flavor attributes, to observe global trends in sensory ratings, and to examine how these patterns vary across demographic groups. Although currently focused on healthy individuals, the visualization system establishes the analytical foundation required for later clinical comparisons.

- **Laying the Groundwork for Remote Monitoring and Longitudinal Assessment.** While the current implementation of the application focuses on healthy users, the underlying protocol, sensory schema, and visual tools provide a natural foundation for remote and repeated assessments over time. Such an extension could allow researchers and clinicians to monitor individual sensory changes, compare them against baselines, and support telehealth-oriented approaches to studying taste dysfunction and recovery.

6.2 Limitations

While the FlavorCharter framework demonstrates strong potential, the current stage of development comes with several limitations that should be acknowledged. These limitations highlight the challenges inherent to large-scale sensory data collection and outline important directions for future refinement.

- **Sample Size and Population Bias.** The preliminary dataset includes 55 healthy participants, which limits the representativeness and generalizability of the current findings. The planned expansion to approximately 2,000 users will be essential for constructing robust population-level baselines and for capturing greater demographic and sensory diversity.
- **Lack of Controlled Tasting Conditions.** Because the study relies on participants evaluating foods in their everyday environments, the tasting conditions cannot be standardized. Variability in factors such as ingredient brands, ripeness, temperature, portion size, or context of consumption is inherent to the crowd-sourced nature of the protocol and introduces noise into the data.
- **No Validation Against Existing Clinical Tools.** The current framework has not yet been compared against established clinical instruments such as the MDASI-HN, taste strips, or objective gustatory assessments. Future validation studies will be necessary to assess the concordance between FlavorCharter ratings and existing clinical or psychophysical benchmarks.
- **No Longitudinal Component.** The present deployment captures only single-time-point evaluations from healthy individuals. The system does not yet track changes over time, limiting its ability to model sensory trajectories or identify temporal patterns in flavor perception.

6.3 Future Work

Building on the foundations established by this work, several directions can further expand the scope and impact of the FlavorCharter framework. These avenues aim to strengthen the robustness of the baseline, extend the framework into clinical settings, and enable actionable, patient-centered applications.

- **Expand Baseline.** Increasing the size and demographic diversity of the healthy population dataset is essential for constructing robust, reliable baselines of flavor perception. A larger cohort will improve statistical power, allow subgroup analyses, and better represent natural variability across age, dietary patterns, and cultural backgrounds.
- **Clinical Comparison and Longitudinal Tracking in HNC Patients.** A natural next step is to apply the framework to HNC patients, enabling the comparison of their sensory profiles against the healthy baseline. Repeated assessments during and after radiotherapy would allow researchers and clinicians to characterize the temporal trajectories of taste dysfunction, quantify recovery patterns, and study how specific sensory dimensions are affected over the course of treatment.
- **Enhance the Visualization System.** The visual analysis interface can be extended to support clinical use by enabling the visualization of individual patient profiles and their evolution over time. A redesigned system would allow clinicians to track changes in flavor perception longitudinally, highlight deviations from healthy baselines, and identify emerging trajectories of taste dysfunction throughout and after treatment.

- **Develop Nutritional Guidelines and a Decision Support Tool.** Using insights from population-level data and longitudinal patient tracking, future work may develop data-driven nutritional guidelines and a clinical decision support tool. Such a system could help clinicians identify appropriate foods for sensory evaluations, guide dietary strategies, and support personalized interventions aimed at improving quality of life for patients experiencing taste dysfunction.

CHAPTER 7

CONCLUSION

Taste and flavor play a central role in quality of life, yet remain difficult to measure, describe, and interpret in both clinical and everyday contexts. Current tools provide limited granularity, and patients experiencing taste dysfunction often struggle to articulate their symptoms in ways that clinicians can interpret and act upon. This thesis addressed these challenges by introducing a unified, data-driven framework for characterizing flavor perception through a structured vocabulary, a reproducible tasting protocol, and a visual analysis system designed to interpret population-level patterns.

This work makes several contributions. First, it defines a ten-dimensional schema of flavor grounded in culinary and biomedical literature and informed by clinical experience in head and neck cancer care. Second, it presents a smartphone-based protocol that enables scalable, crowd-sourced data collection. Third, it reports the acquisition of real-world data from 55 healthy users and more than 400 flavor ratings, demonstrating the feasibility and consistency of this approach. Finally, it introduces a visual analysis interface that organizes these data into interpretable patterns and supports the exploration of trends, variability, and demographic influences within the baseline population.

Taken together, these components establish the foundations of FlavorCharter as a standardized framework for describing flavor perception and for building quantitative baselines of taste and flavor in healthy individuals. These baselines are a necessary prerequisite for studying deviation patterns in clinical populations and for interpreting patient complaints with greater precision. The framework also opens opportunities for scientific inquiry into demographic effects, dietary influences, and population-level flavor trends.

Several limitations remain. The dataset is preliminary and limited in size, tasting conditions are not controlled, and clinical validation against established instruments has not yet been performed. The current deployment also lacks longitudinal tracking, preventing an analysis of temporal changes in taste perception.

Looking ahead, future work will expand the baseline to a larger and more diverse population, compare healthy patterns with those of head and neck cancer patients, and track flavor perception longitudinally during radiotherapy to characterize trajectories of taste dysfunction. The visual analysis system will be extended to support patient-focused use cases, including individual trend visualization and comparisons against healthy baselines. Finally, the framework can evolve into a nutritional decision support tool that assists clinicians in selecting evaluation foods, designing sensory interventions, and guiding dietary strategies to improve patient well-being.

Taken together, these contributions establish *FlavorCharter* as the foundation for a standardized, data-driven understanding of flavor perception, one that can ultimately support improved communication, better nutritional strategies, and clinically meaningful monitoring of taste dysfunction in *Head and Neck Cancer* patients.

CITED LITERATURE

1. Gunn, L. et al.: Taste dysfunction following radiotherapy to the head and neck: A systematic review. Radiotherapy and Oncology, 157:130–140, 2021.

2. Wentzel, A., Hanula, P., et al.: Cohort-based T-SSIM visual computing for radiation therapy prediction and exploration. Trans. Vis. Comp. Graph., 26(1), 2019.

3. Wentzel, A., Hanula, P., van Dijk, L. V., et al.: Precision toxicity correlates of tumor spatial proximity to organs at risk in cancer patients receiving intensity-modulated radiotherapy. Radiotherapy and Onc., 148, 2020.

4. Wentzel, A., Luciani, T., van Dijk, L. V., et al.: Precision association of lymphatic disease spread with radiation-associated toxicity in oropharyngeal squamous carcinomas. Radiotherapy & Onco., 161, 2021.

5. Wentzel, A., Mohamed, A. S. R., Naser, M. A., et al.: Multi-organ spatial stratification of 3-d dose distributions improves risk prediction of long-term self-reported severe symptoms in oropharyngeal cancer patients receiving radiotherapy: development of a pre-treatment decision support tool. Front. in onc., 13, 2023.

6. Wentzel, A., Floricel, C., Canahuate, G., et al.: DASS Good: Explainable Data Mining of Spatial Cohort Data. Comp. Graph. For., 24(3), 2023.

7. Floricel, C., Wentzel, A., Mohamed, A., Fuller, C., et al.: Roses have thorns: Understanding the downside of oncological care delivery through visual analytics and sequential rule mining. IEEE Trans. Vis. Comp. Graph., 30(1), 2024.

8. Floricel, C., Nipu, N., et al.: THALIS: Human-Machine Analysis of Longitudinal Symptoms in Cancer Therapy. Trans. Vis. Comp. Graph., 28(1):151–161, September 2021.

9. Narici, M. V. and Maffulli, N.: Sarcopenia: characteristics, mechanisms and functional significance. British Medical Bulletin, 95(1):139–159, 03 2010.

10. Santilli, V., Bernetti, A., Mangone, M., and Paoloni, M.: Clinical definition of sarcopenia. Clinical Cases in Mineral and Bone Metabolism, 11(3):177–180, Sep 2014.

CITED LITERATURE (continued)

11. Beaudart, C., Zaaria, M., et al.: Health Outcomes of Sarcopenia: A Systematic Review and Meta-Analysis. PLoS One, 12(1):e0169548, January 2017.

12. Grossberg, A. J., Chamchod, S., et al.: Association of Body Composition With Survival and Locoregional Control of Radiotherapy-Treated Head and Neck Squamous Cell Carcinoma. JAMA Oncology, 2(6):782–789, jun 2016.

13. de Bree, R., van Beers, M. A., and Schaeffers, A. W. M. A.: Sarcopenia and its impact in head and neck cancer treatment. Current Opinion in Otolaryngology & Head and Neck Surgery, 30(2):87–93, apr 2022.

14. Wentzel, A., Attia, S., Zhang, X., et al.: Ditto: A visual digital twin for interventions and temporal treatment outcomes in head and neck cancer. Trans. Vis. Comp. Graph., 2024.

15. Sheu, T., Vock, D., Mohamed, A., et al.: Conditional survival analysis of patients with locally advanced laryngeal cancer: construction of a dynamic risk model and clinical nomogram. Scientific reports, 7(1), 2017.

16. Christopherson, K. M., Ghosh, A., Mohamed, A. S. R., Kamal, M., et al.: Chronic radiation-associated dysphagia in oropharyngeal cancer survivors: Towards age-adjusted dose constraints for deglutitive muscles. Clinical and translational radiation oncology, 18:16–22, 2019.

17. Elhalawani, H., Mohamed, A., Kanwar, A., Dursteler, A., et al.: Serial Parotid Gland Radiomic-based Model Predicts Post-Radiation Xerostomia in Oropharyngeal Cancer. Radiotherapy and Oncology, 127(Suppl. 1):S1167–S1168, 2018.

18. Epstein, J. B., Villines, D., Epstein, G. L., and Smutzer, G.: Oral examination findings, taste and smell testing during and following head and neck cancer therapy. Supportive Care in Cancer, 28(9):4305–4311, 2020.

19. Cleeland, C. S., Mendoza, T. R., Wang, X. S., Chou, C., et al.: Assessing symptom distress in cancer patients: The M.D. Anderson Symptom Inventory. Cancer, 89:1634—1646, 11 2000.

20. Silvio, S. et al.: The CiTAS scale for evaluating taste alteration induced by chemotherapy: state of the art on its clinical use. Acta Bio Medica: Atenei Parmensis, 90(Suppl 6):17–25, 2019.

CITED LITERATURE (continued)

21. Haydar, S. et al.: The Taste Liking Questionnaire (TasteLQ)–A tool for assessment of taste liking in the Danish population. Food Quality and Preference, 117, 2024.

22. Dornenburg, A. and Page, K.: The Flavor Bible: The Essential Guide to Culinary Creativity, Based on the Wisdom of America's Most Imaginative Chefs. Voracious, 2008.

23. Nosrat, S.: Salt, fat, acid, heat: Mastering the elements of good cooking. Simon and Schuster, 2017.

24. Reddit: A Cool Guide to Everyday Taste Profiles. https://www.reddit.com/r/coolguides/comments/1dxybbv/a_cool_guide_to_everyday_taste_profiles_oc/, 2024. Accessed: August 29, 2025.

25. Wilson, M. D., Hunt, I. G., Sawyer, S. M., Kocharunchitt, C., Li, Z., and Stanley, R. A.: Electronic tongue measurements as a predictor for sensory properties of vacuum-packed minced beef–A preliminary study. Meat Science, 220:109705, 2025.

26. Zhang, H., Zou, G., Liu, W., and Zhou, Z.: Beer Taste Detection Based on Electronic Tongue. Sensors & Materials, 32, 2020.

27. Marai, G. E.: Activity-centered domain characterization for problem-driven scientific visualization. Trans. Vis. Comp. Graph., 24(1), 2017.

28. Risso, D., Drayna, D., and Morini, G.: Alteration, reduction and taste loss: main causes and potential implications on dietary habits. Nutrients, 12(11):3284, 2020.

29. Munzner, T.: Visualization analysis and design. In Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Courses, pages 1–2, 2025.

30. Roberts, J. C., Headleand, C., and Ritsos, P. D.: Sketching designs using the five design-sheet methodology. IEEE transactions on visualization and computer graphics, 22(1):419–428, 2015.

31. Cao, N., Lin, Y.-R., Gotz, D., and Du, F.: Z-glyph: Visualizing outliers in multivariate data. Information Visualization, 17(1):22–40, 2018.

VITA

NAME	Francesco Botto
EDUCATION	
	Master of Science in Computer Science, University of Illinois Chicago, December 2025, USA
	Master of Science in Computer Engineering, Politecnico di Torino, December 2025, Italy
	Bachelor of Science in Computer Engineering, Politecnico di Torino, October 2023, Italy
WORK EXPERIENCE	
Jan - Dec 2025	Research Assistant - University of Illinois Chicago
	Designed and developed a crowd–sourced, data-driven React Native app (Expo) to quantify flavor perception in healthy patients. In collaboration with MD Anderson Cancer Center.
Jan - May 2023	Data Management and Workflow Automation - Denso
	Worked for IT department, mastering management software, collaborating on workflow creation, enhancing user interfaces, and utilizing SQL for data management and reporting.
PUBBLICATIONS	
2025	Development of a Smartphone App for Quantifying Food Flavor, Francesco Botto et al, pp. 1-8, Workshop on Computational Gastronomy: Data Science for Food and Cooking (CoGamy) at the International Conference on Data Mining (ICDM) 2025
SCHOLARSHIPS	
2025	RA position tuition waiver
2024	TOP-UIC and EXTRA-UE scholarship
2022	EXTRA-UE scholarship

VITA (continued)

TECHNICAL SKILLS	
Technical skills	Full-stack development, web development, neural networks, computer architectures, 3D design, computer vision, algorithms and data structures, and data visualization
Programming Languages	JavaScript, Java, C, C++, Python, SQL, HTML, CSS, ARM, 8086
Frameworks	React Native, Expo, React.js, Node.js, Express.js, Pytorch, Firebase, D3.js
LANGUAGE SKILLS	
Italian	Native speaker
English	Full working proficiency
	2024 - IELTS Band 8