AI REIMAGINE PROJECT · FILE 02

AI Feature Spec & PRD

FitSense - an AI Fit-Advisor for the Myntra PDP

Prepared by Abhishek Katkar · June 2026

Continues from File 1. File 1 selected friction F1: even with a size chart and a purchase-history recommendation on the PDP, a shopper still can't trust which size will fit - the recommendation is garment-blind, unexplained, and absent at cold-start (first-ever, new brand, new category). This file specifies the fix: FitSense, an AI fit-advisor. The win is scoped honestly - it does not try to out-predict the history signal for repeat buyers; it covers the moments that signal can't reach and makes every call explainable and confidence-scored.

1. Feature Overview

Field	Detail
Feature name	FitSense - AI Fit-Advisor
One-liner	On the product page, FitSense recommends the size most likely to fit this specific garment, with a confidence level and a one-line reason - and when it doesn't know enough, it says so and offers a 10-second way to sharpen the call.
Where it appears	On the apparel PDP, in the size-selection area - upgrading the existing "We recommend L based on your purchase history" banner (see File 1, Fig 3.4).
Interaction model	Hybrid: a passive, always-on explained recommendation for every shopper, plus an optional "Refine my fit" expansion that collects 1-2 inputs to raise confidence.

Before → After

Before (today)	After (FitSense)
"We recommend L based on your purchase history." No reason, no confidence, and nothing useful on a first-ever or new-brand buy. Size chart is measurements only.	"We suggest M · High confidence. This Henley runs slim, and you prefer a relaxed fit - so M gives room." A confidence pill sits beside it. At cold-start: "Low confidence - tell us your usual fit (10s) to sharpen this," with a Refine my fit CTA.

2. Problem & Goals

Goal type	Statement
User problem	At the size-decision moment, shoppers can't tell which size will actually fit this garment - especially on a first-ever, new-brand, or new-category buy where purchase history gives no signal - so they guess, bracket (order two

Goal type	Statement
	sizes), or abandon.
User goal	Pick the right size on the first try, with enough reason to trust the call.
Business goal	Reduce size-driven returns (and the reverse-logistics cost behind them) and lift PDP conversion - concentrated on the cold-start / cross-brand cohort that the existing engine cannot serve.
Product goal	Raise the share of recommendations that are shown with a reason + confidence and accepted by the user, and cut size-related returns for the FitSense-exposed cohort.

3. AI Capability Used

Primary capability: Personalisation + Retrieval-Augmented Generation (RAG) over the garment's size chart and construction attributes, with an LLM reasoning over the user's fit profile to produce a size, a confidence level, and a natural-language rationale.

Why this capability fits the friction

- The task is to reconcile heterogeneous, partly-unstructured signals - per-garment cut/stretch (often only in the product description), the per-brand size chart, and the user's fit history/preference - and then explain the call. That is reasoning + retrieval, not a deterministic lookup.
- Classification alone (size = f(measurements)) can't absorb cross-brand or per-garment variance, and produces no reason. A rule table can't generalise across 3,500+ brands' inconsistent charts.
- Confidence estimation is part of the capability: the system must estimate how sure it is (from input completeness, consistency, and retrieval quality) and surface that honestly - the core differentiator from today's silent "we recommend L."

4. Input → Processing → Output

Input given to AI	AI processing / logic	Output shown to user
Garment context: brand, category, this SKU's size chart (chest, length), construction (fit type, stretch) - extracted from the product description where unstructured.	Retrieve the garment's chart + construction (RAG); normalise against a cross-brand size reference.	Recommended size (e.g., "M").
User fit profile: past-purchase fit outcomes (size + kept/returned), optional body measurements, optional fit preference (relaxed / true / snug).	LLM reasons over garment + profile to choose a size and draft a one-line rationale.	Confidence level - High / Medium / Low - shown as a pill.

Input given to AI	AI processing / logic	Output shown to user
Cold-start case: no history for this brand/category.	Confidence estimator scores the result; if below threshold, route to the low-confidence path.	One-line plain-language reason (“runs slim; you like a relaxed fit, so M gives room”).
Optional refine inputs: fit preference and/or one body measurement.	Recompute with the added inputs to raise confidence.	Low-confidence state: honest message + “Refine my fit (10s)” CTA.
Post-purchase signal: size accepted, then kept or returned for fit.	Log for evaluation and update the fit profile.	(Over time) sharper recommendations for that user.

5. AI Workflow

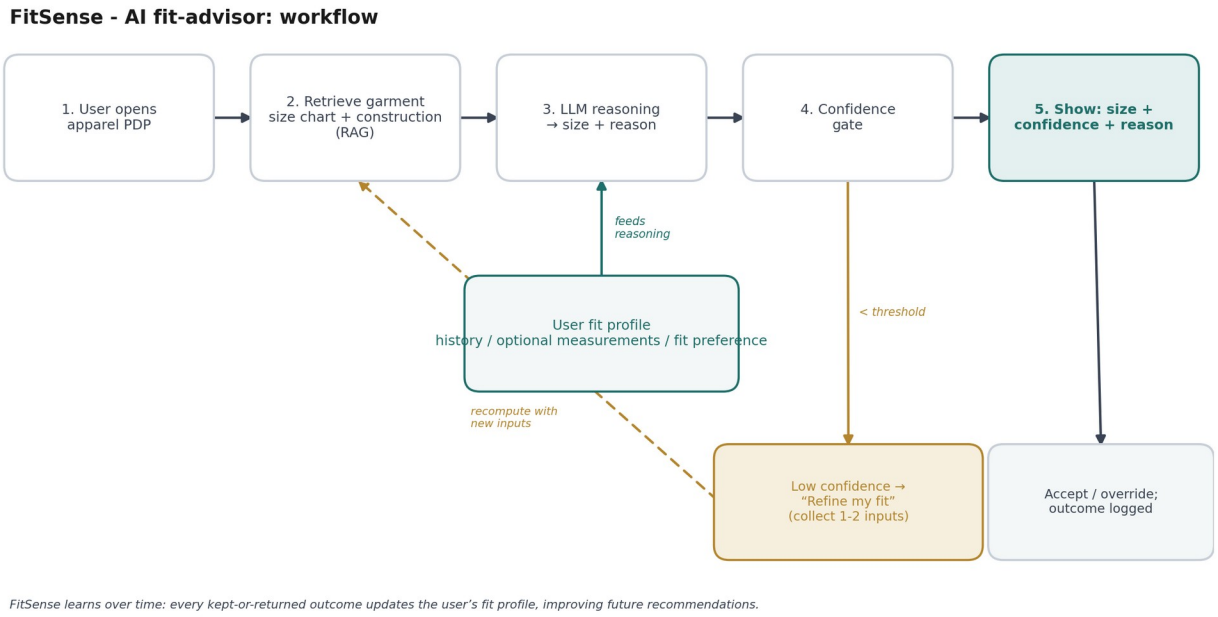

Figure 5.1 - FitSense workflow. Teal = the shown recommendation; amber = the low-confidence “Refine my fit” path that recomputes; the fit profile feeds reasoning and learns from outcomes.

1. User opens an apparel PDP; FitSense reads the garment ID.
2. System retrieves the garment’s size chart and construction attributes (RAG; extracts fit type/stretch from the description if unstructured) and loads the user’s fit profile.
3. The LLM reasons over garment + profile and drafts a candidate size with a one-line reason.
4. A confidence estimator scores the result from input completeness, consistency, and retrieval quality.
5. Confidence gate: at or above threshold → show size + confidence + reason; below threshold → show the honest low-confidence state with “Refine my fit.”

6. (Optional) the user refines with 1-2 inputs; the system recomputes at higher confidence.
7. The user accepts or overrides; the post-delivery kept-or-returned outcome is logged, feeding evaluation and updating the fit profile.

6. Data Needs

- **Garment catalogue:** per-SKU size chart (chest, length…), brand, category, and construction attributes (fit type, stretch) - partly structured, partly extracted from descriptions.
- **Cross-brand size reference:** a normalisation table so a brand's "M" can be compared to another's.
- **User fit profile:** past-purchase fit outcomes (size ordered, kept/returned/exchanged), optional self-reported body measurements, optional fit preference.
- **Knowledge base for RAG:** the size charts and construction text the model retrieves and grounds on.
- **Feedback logs:** recommendation shown, accepted/overridden, and post-delivery kept/returned - for evaluation and profile updates.

> **Sensitive data.** Body measurements and purchase history are personal data. Handling: measurements are strictly opt-in; data is encrypted in transit and at rest; used only for fit, never shared with sellers; the user can view and delete it; retention is minimised. This is specified in full in File 3.

7. User Stories

Written to INVEST - each is independent, negotiable, valuable, estimable, small, and testable.

US1 As a **first-time buyer of a brand**, I want to **get a size recommendation with a clear reason**, so that I can pick a size confidently even though I have no purchase history here.

US2 As a **returning shopper**, I want to **see why a size is recommended and how confident the system is**, so that I can trust it, or know when to refine it.

US3 As a **shopper unsure of fit**, I want to **add my fit preference in a couple of taps**, so that I can get a recommendation that matches how I like clothes to feel.

US4 As a **value-conscious shopper**, I want to **be told when the system isn't sure**, so that I can not blindly order a size and get stuck returning it.

US5 As a **repeat user**, I want to **have my kept/returned outcomes improve future recommendations**, so that I can get a better fit call over time without extra effort.

8. Acceptance Criteria

Given-When-Then, for the highest-value stories.

AC1 (US1/US4) **Given** a user has no purchase history for this brand or category, **When** they open the PDP, **Then** FitSense shows a recommendation with an explicit Low or Medium confidence label and a “Refine my fit” option, with the recommendation rendered within P95 ≤ 1.5s of PDP load.

AC2 (US3) **Given** a user taps “Refine my fit” and enters a fit preference, **When** they submit, **Then** the recommendation and confidence update within 1s and the one-line reason reflects the new input.

AC3 (US4) **Given** the system’s confidence is below the threshold (e.g., 0.6), **When** a recommendation would be shown, **Then** FitSense shows the honest low-confidence state and does not present any single size as certain.

AC4 (US2) **Given** any recommendation is shown, **When** the user views it, **Then** a one-line plain-language reason and a confidence indicator are always present - never a bare size.

9. Functional Requirements

- Show a garment-specific size recommendation on every apparel PDP.
- Always display a confidence level and a one-line reason with the recommendation - never a bare size.
- Provide an optional “Refine my fit” flow collecting a fit preference and (optionally) one body measurement.
- Enforce a confidence threshold; below it, show the low-confidence fallback instead of a definitive size.
- Recompute the recommendation in near-real-time when new inputs arrive.
- Log every recommendation, the user’s action, and the post-delivery kept/returned outcome.
- Let the user view and delete their stored fit data from within the app.
- Degrade gracefully: if the model/service is unavailable, fall back to the existing static size chart so the PDP never blocks.

10. Non-Functional Requirements

Quality attribute	Requirement / target
Latency	P95 ≤ 1.5s to render the recommendation on PDP load; ≤ 1s on “Refine.” Never block PDP render.
Reliability	99.9% uptime for the recommendation service; on failure, degrade to the static size chart rather than erroring.
Privacy	Measurements and history are personal data: opt-in, encrypted in transit and at rest, never shared with sellers, user-deletable, minimal retention.
Accessibility	Confidence and reason conveyed via text, not colour alone (WCAG 2.1 AA); screen-reader labels on the pill and CTA.

Quality attribute	Requirement / target
Security	Authenticated access to the fit profile; no unnecessary PII in logs; the recommendation endpoint is rate-limited.
Cost	Track cost-per-recommendation (token economics); cache garment-level retrieval so repeat PDP views don't re-pay for extraction.

11. Assumptions, Dependencies & Risks

Assumptions

- Size charts exist per SKU (verified - File 1, Fig 3.5), and some construction detail is extractable from product descriptions.
- Users will spend 1-2 taps to refine fit on higher-stakes or uncertain buys; purchase history exists for most repeat buyers.

Dependencies

- An LLM API (e.g., Claude / GPT / Gemini); the garment catalogue and size-chart data; a cross-brand normalisation dataset; a fit-profile store; design and engineering; and legal sign-off on data handling.

Risks

- **Noisy cold-start inputs.** 67-73% of shoppers don't measure accurately (File 1, Fig 4.4), so a confident-wrong recommendation is the central risk - mitigated by the confidence-honesty rule and the low-confidence fallback.
- **Sparse construction data.** Fit/stretch isn't structured for every brand, lowering coverage where extraction fails - degrade to confidence-scored chart guidance.
- **Latency & cost at PDP scale.** Mitigated by caching garment-level retrieval and a strict P95 budget.
- **Early-trust risk.** If the first few recommendations are wrong, users disengage - mitigated by honest confidence and a phased rollout (File 5).

12. Agile Backlog

Epic: FitSense - AI fit-advisor on the Myntra PDP. Priorities P0 (must-have for MVP) → P2; estimates in story points (Fibonacci).

User story	Priority	Story points	Notes
Garment-aware size recommendation + one-line reason on the PDP	P0	8	Core value; depends on catalogue + size-chart data.
Confidence estimation + threshold gate + low-confidence fallback	P0	8	The differentiator and the key responsible-AI control.

User story	Priority	Story points	Notes
“Refine my fit” input flow (fit preference, optional measurement)	P1	5	Lifts confidence for cold-start; opt-in only.
Fit-profile store + post-delivery kept/returned logging	P1	5	Enables learning loop and evaluation (File 5).
Cross-brand normalisation + construction extraction (RAG)	P1	8	Improves garment-awareness and cross-brand accuracy.
Privacy controls (view/delete fit data) + graceful degradation	P2	3	Required for launch; ties to File 3.

Carries into File 3 (Responsible AI) and File 4 (prototype). The confidence gate and low-confidence fallback specified here are the spine of the safety plan and the build.