Myli — Product Requirements Document

Document status

Version 1.0 of a concept-stage PRD for Myli, an ongoing personal case study. This document defines what the Figma designs build; success metrics are benchmark-informed targets, not validated findings.

Area	Where it lives	Status
Onboarding (account, documents, context review, first work area)	Figma · High-fidelity UI	Designed
Real estate workflow: Market Intelligence canvas and Saved Work	Figma · High-fidelity UI	Designed
AI-adaptive components (adaptive box, generation feedback, sources, approval, saved-work cards)	Myli Base Design System	Designed
End-to-end approval screen	Figma · High-fidelity UI	In progress
Mobile experience (4 screens so far)	Figma · High-fidelity UI	In progress
Agent actions (scheduling, tasks, sending)	This PRD	Planned
Finance & Investment and Laboratory workspaces	This PRD	Planned
Usability testing	This PRD	Planned

At a glance

- **What:** a standalone AI workspace that builds editable canvases around a goal, inside a stable app shell.
- **Who:** Amanda Costa, a NYC luxury agent researching an expansion into the Catskills.
- **MVP:** one real estate workflow, from intent to a sourced market analysis to a saved canvas, on desktop and mobile.
- **Trust:** sources on every figure; approvals scaled to risk before anything leaves Myli.

- **Success:** at least 25% faster than a chat-only tool with no loss in accuracy, and every user knows what an approval will do.
- **Status:** core flow designed; approval screen and mobile in progress; agent actions and more industries planned.

Overview

Myli is a standalone AI-adaptive workspace for complex professional work. It turns a goal into a persistent, editable interface, without replacing predictable navigation with unpredictable generation.

Problem. Conventional AI assistants put the conversation at the center. Professional work also needs durable context, navigable information, visible evidence and control over how outputs evolve. Real estate agents already use AI, but mostly for writing; analysis work such as market research, comparisons and client briefs is still spread across portals, documents and chats.

Positioning. Not another chatbot or fixed dashboard: a structured environment where AI helps people think, analyze and act.

Hypothesis: stable structure, adaptive content. A predictable app shell paired with task-specific canvases lets AI assemble analysis, maps, charts and notes inside an information architecture that stays understandable and user-controlled.

Thesis. Myli chooses the right interaction for each task: a conversational answer for a simple question, an assembled canvas for multi-part analysis, and an explicit approval step for anything consequential.

Users and jobs to be done

The MVP serves one persona: Amanda Costa, a Manhattan and Brooklyn luxury residential agent ($1M–$10M) assessing an expansion into the Catskills, with an adjacent interest in finance and investing. She is a composite; her goals and frictions are hypotheses to validate in interviews.

Core job to be done: “Help me assess a new market and communicate evidence-backed opportunities without adding administrative work.”

Goals

- Evaluate Catskills demand, pricing and inventory before expanding.

- Match high-value properties with client criteria.
- Prepare credible market and showing briefs.
- Explore investment implications alongside real estate research.

Needs

- A workspace that keeps NYC and Catskills research context.
- Comparisons, maps, charts, documents and notes in one place.
- Source-linked findings she can check before using them with clients.
- Analysis she can reuse across Real Estate and Finance & Investment.

Hypothesized frictions

- Research scattered across portals, documents and conversations.
- Repeated setup when switching markets or professional domains.
- Unclear provenance or overconfident AI recommendations.

Market context

- **Agents work from their phones.** Realtors work a median 35 hours a week, and 96% use a smartphone for business daily or nearly every day (NAR 2026 Member Profile).
- **They adopt tools to save time.** 81% adopt new technology mainly to save time, and 23% use AI daily (NAR 2026 Technology Survey).
- **AI use is mostly writing.** The top AI tasks are listing descriptions (75%), social posts (56%) and emails (52%) (RISMedia). Analysis work is largely unserved.
- **Use is high, value is low.** 90% of real estate professionals use AI, but only 35% find it genuinely helpful; 47% call checking AI output their biggest AI time drain, and 1 in 5 runs AI outside their main workflow. Small sample (n = 89), directional only (HousingWire, Fyxer report).
- **Manual research takes time.** Most agents spend over 30 minutes on a single comparative market analysis; 31% spend more than 45 (RISMedia, homegenius/RRC 2022).

Product principles

Seven principles decide every trade-off in this document.

1. **Stable navigation, adaptive content.** AI writes into the information architecture; it never replaces it. Users always know where they are.

2. **Persistent context.** Work survives individual prompts and sessions.
3. **Inspectable evidence.** Sources, assumptions and freshness stay visible, and facts are distinguished from AI interpretation.
4. **Human control.** Consequential actions require clear approval; everything else can be undone.
5. **Conversation for intent, UI for precision.** Users can change any canvas by prompt or by direct manipulation.
6. **AI as the operating layer, not the brand.** No glowing buttons or AI feature lists; a quiet command bar invites requests.
7. **Same capabilities on every device.** Desktop and mobile differ in presentation, not in what users can do.

Scope

The MVP proves one end-to-end real estate workflow, intent to analysis to saved canvas, before the same building blocks extend to other industries.

MVP: Real Estate

- Account, secure sign-in and professional profile (markets, price range, goals).
- Optional document upload into the Knowledge library, with a review of what Myli extracted.
- Real Estate work areas: Market Intelligence, Property Analysis, Clients & Meetings, Deal Pipeline. Market Intelligence is the fully designed path.
- An adaptive canvas for the Catskills Market Outlook: market comparison, research notes, asking vs. sold chart, and a map with an AI-written brief.
- Editing by prompt or directly; sources and assumptions on every result.
- Save, resume and Saved Work.
- Approval before consequential actions.
- Desktop and mobile.

Planned

- Finance & Investment workspace, the persona's adjacent interest, reusing comparison and opportunity analysis.
- Laboratory operations workspace.
- More industries through Add Industry, as demand shows.

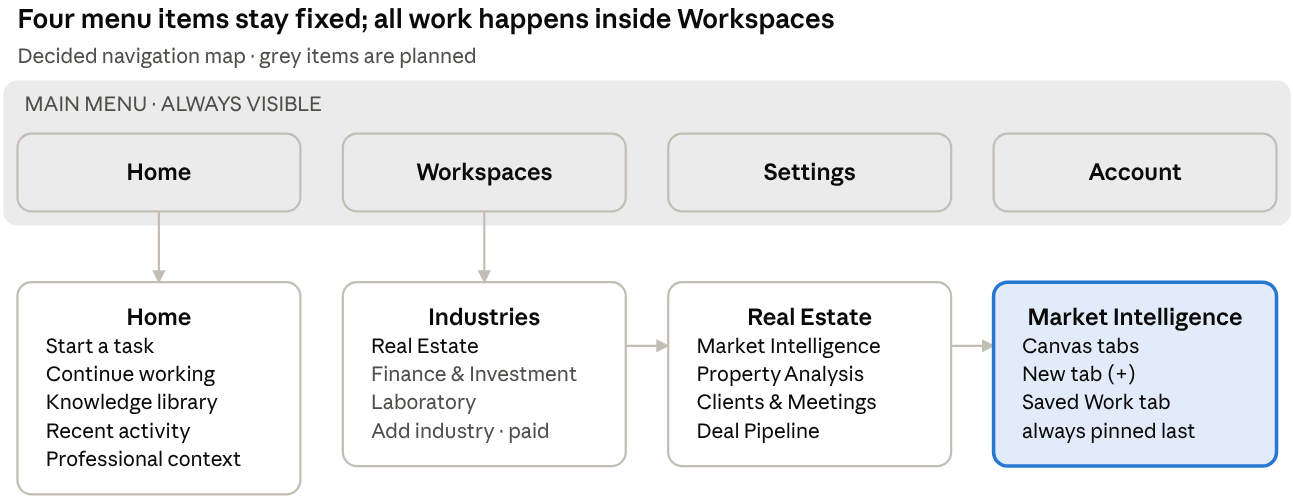
- Agent actions: scheduling showings, creating tasks, sending approved briefs and updating records.

Out of scope for the MVP

- Unrestricted UI generation; the canvas uses only approved components.
- Autonomous high-impact actions.
- Full feature parity across industries.
- A separate AI-assistant login (ChatGPT, Claude) as a requirement.
- Live MLS data. MVP figures are illustrative and labeled as such.

Information architecture and navigation

The main menu has four fixed items: Home, Workspaces, Settings and Account. Industries live inside Workspaces, never in the main menu, so adding one never changes the shell.



Myli navigation map · 4 menu items, 3 levels inside Workspaces

A user drills down from Workspaces to an industry, then a work area, then a canvas; Market Intelligence is the fully designed path. A breadcrumb (“Real Estate / Market Intelligence”) always shows the current location, and the sidebar can be expanded or minimized without losing it. The Knowledge library and Recent activity live on Home because users check them from there rather than work in them.

Requirements

Nine P0 requirements define the MVP; eight are designed and one, approval, is in progress.

ID	Requirement	Priority	Done when	Status
R1	Account and professional profile	P0	Profile (role, markets, price range) persists across sessions	Designed
R2	Knowledge library upload	P1	Extracted context can be reviewed and edited; skipping never blocks setup	Designed
R3	Predictable navigation	P0	Market Intelligence in ⩽ 3 steps; breadcrumb always visible	Designed
R4	Start a task by intent	P0	Command bar and suggested topics on Home and every canvas	Designed
R5	Adaptive canvas assembly	P0	Built from approved components only; each box shows its source	Designed
R6	Edit by prompt or directly	P0	Reorder, resize, add or remove boxes; every change can be undone	Designed
R7	Sources and assumptions	P0	Every figure links to a dated source; unverified claims flagged	Designed
R8	Save and resume	P0	Saved Work tab pinned last; reopened canvases keep full state	Designed
R9	Approval for consequential actions	P0	Nothing high-risk happens without explicit confirmation	In progress
R10	Share a canvas	P1	Recipients and contents shown first; treated as high-risk	In progress

ID	Requirement	Priority	Done when	Status
R11	Mobile capability parity	P1	Every capability available, adapted for small screens	In progress
R12	Activity log	P1	Every AI change and approved action logged with sources	Planned
R13	Agent actions	P2	Showings, tasks, sending and record updates, all behind approval	Planned
R14	Additional industries	P2	Finance & Investment and Laboratory reuse the same components	Planned

Adaptive canvas and component system

Myli adapts the interface only when the task needs it, and only with components from its own design system. That keeps generated layouts predictable, accessible and consistent.

When Myli adapts

The user asks for…	Myli responds with…	Example
A simple fact or explanation	An answer in the assistant panel; the canvas stays as it is	“What does days on market mean?”
Multi-part analysis	A canvas assembled from components	“Compare Woodstock, Hudson and Saugerties.”
A change to an existing canvas	A preview of the change, with undo	“Replace the comparison with a price-trend chart.”
An action that leaves Myli	An approval step (see the approval framework)	“Send this brief to my client.”

Component contract

Every canvas box follows the same contract, whatever it contains:

- **Anatomy:** a title, a one-line scope (“Median asking price · 3 towns”), the content, and a footer with source and freshness.
- **Controls:** move left or right, expand, open sources, remove. On mobile: move up or down.
- **States:** generating, ready, unverified, source changed, error.
- **Color meaning:** each content type has its own palette family, so color means the same thing everywhere. Comparisons use Sea Urchin Lavender, maps Kelp Sage, charts Sea Glass, insights Coral Mist, research Shell Sand.

Design system

The Myli Base Design System extends the Obra shadcn/ui kit with Myli components: navigation (expanded and minimized), prompt bar states, canvas tabs, contextual canvas actions, scrollbar, and five AI-adaptive pattern sets:

1. Adaptive content box, with resize and rearrangement.
2. AI generation and feedback.
3. Sources and evidence: inline citation, source summary, freshness, unverified claim, source replaced.
4. AI approval and confirmation: proposed change, before/after compare, approve, reject, undo, needs permission.
5. Workspace and saved-work cards.

The design system evolves with the product; new patterns are added as the case study grows.

Approval framework

The more an action affects other people or can’t be undone, the more Myli asks before doing it. The building blocks are designed in the design system; the end-to-end approval screen is in progress and will be built to this spec.

Risk level	Example actions	What the user does	On mobile
Low: changes the user’s own canvas	Rearrange boxes, add a chart, refresh data	Nothing extra; undo is always available	Same

Risk level	Example actions	What the user does	On mobile
Medium: changes saved work	Replace an analysis, delete a box, overwrite a saved canvas	Previews the change, confirms, can undo afterward	Same, with before/after as a toggle
High: leaves Myli or affects others	Share a brief with a client, send an email, book a showing	Reviews exactly what goes out, to whom, and which claims are unverified; then confirms explicitly	Same, plus Face ID or passcode (hypothesis to test)

What every high-risk approval shows

- The action in plain words, and who it affects.
- Exactly what is included, with illustrative or unverified figures flagged.
- What can't be undone once confirmed.
- Three choices: approve, edit or cancel.

Upcoming screen (v1.1). Myli proposes sending the Catskills Market Outlook to a client. Amanda sees the recipient, the included boxes and the flagged illustrative figures. She approves, edits or cancels; the sent brief is then logged under Recent activity.

Acceptance criteria

- Before approving, the user can say in their own words what will happen.
- No high-risk action happens without explicit confirmation.
- Every approved action appears in Recent activity with its sources.

Desktop and mobile

Same capabilities, adapted presentation: everything Myli does is available on both, and only the interaction changes with screen size. Mobile is in progress, with four screens designed so far: workspace home, the Catskills canvas, expanded box controls and the first onboarding step.

Mobile matters for this persona because agents work from their phones between showings and meetings. Two rules keep the promise honest:

- **Never block.** Where a view is clearly better on a large screen, a gentle "Easier on a larger screen" hint appears, and the user can still continue.

- **Continue anywhere.** Workspaces persist, so a task started at a showing can be finished at a desk.

Capability	Desktop	Mobile treatment	Mobile status
Ask, review and resume saved work	Full canvas	Boxes stacked in one column	In progress
Rearrange boxes	Drag, width controls, move left or right	Move up or down, or ask Myli	Designed
Onboarding and document upload	Drag and drop	File picker, camera scan or upload later	In progress
Approve AI actions	Inline approval card	Full-screen sheet with the consequence spelled out	Planned
Sources and evidence	Side panel	Bottom sheet opened from the citation	Planned
Wide comparison tables	Side by side	First column frozen, horizontal scroll	Planned
Map and analysis together	Side by side	Full-screen map with a draggable analysis sheet	Planned
Before/after compare	Two panes	Toggle between current and proposed	Planned

Pricing model

The first industry workspace is free; each additional industry is a paid add-on through Add Industry. One account holds every workspace, so adding an industry never means a second login or rebuilding context.

This is a hypothesis. Agents already pay for tools out of pocket: in 2025, 34% spent $50–$250 a month on technology, 20% spent $251–$500 and 24% spent over $500 (NAR 2025 Technology Survey). Price points, trial length and team plans are open questions.

Success metrics

Myli succeeds if agents finish a sourced market comparison faster than with a chat-only tool, with no loss in accuracy, and know what will happen before every approval. These are benchmark-informed targets, to be validated in usability testing against a chat-only baseline.

Metric	Target	Benchmark it is based on
Completion of the core flow (build a sourced comparison, save, resume)	⩾ 85% of participants	Average task completion is 78% (MeasuringU)
System Usability Scale (0–100)	⩾ 75	Average SUS is 68 (MeasuringU)
Ease rating per task (1–7)	⩾ 5.5 core tasks; ⩾ 6 approval	Average is about 5.5 (MeasuringU)
Time to a sourced market comparison	⩾ 25% faster than chat-only; under 15 minutes	AI made consultants 25% faster on suitable tasks (HBS/BCG study); manual analysis takes most agents 30+ minutes
Correct conclusions vs. chat-only	No decrease	Outside AI's strengths, consultants were 19% less likely to be correct (HBS/BCG study)
Spots the flagged, unverified figure	⩾ 80% of participants	Only 1% of visits included a click on a source cited in AI summaries (Pew via Slashdot)
Opens a source before using a key figure	⩾ 50% of participants	66% rely on AI output without checking accuracy (KPMG / University of Melbourne)
Can explain why a box appeared	⩾ 80% of participants	Principle 3, inspectable evidence
States the outcome before approving a high-risk action	100%, zero unintended actions	High-stakes tasks require 100% completion (MeasuringU)

Metric	Target	Benchmark it is based on
Mobile task completion	Within 10 points of desktop	Principle 7, same capabilities everywhere

Risks and constraints

The biggest risk is an AI that is fast and confident but wrong; most mitigations make evidence and control visible.

Risk	Mitigation
Faster, but to the wrong answer	Dated sources on every figure; accuracy measured against chat-only
Figures used without checking	One-tap sources; unverified figures listed before anything is sent
Fair Housing steering	Area recommendations use market data only, never protected traits or proxies (the Act still applies)
Listing-data licensing	MVP uses illustrative, labeled data; live MLS needs licensing
A changing interface adds cognitive load	Stable shell; user-controlled layout; undo everywhere
Learning a new app (cost of going standalone)	Onboarding useful even when skipped; MCP open as a later way in
Sensitive client documents	Permission-aware retrieval; nothing shared across industries without approval
Small screens limit complex views	Mobile treatments per capability; never block

Constraints. Controlled UI generation only, from approved components; permission-aware data access; grounded AI responses; accessible patterns (keyboard, focus and contrast states are defined in the style guide); auditable agent behavior; progressive loading.

Decision log

Seven decisions shaped the design. Each shows what was chosen, why, and when to revisit it.

1. A standalone app, not an AI-assistant plug-in

Why: The adaptive canvas, Saved Work and approvals need full control of the interface, and workspaces and permissions need one home. Agents shouldn't need a separate AI subscription.

Trade-off: Users learn a new app.

Revisit if: Agents want to start tasks from AI tools they already use. MCP stays open as an optional way in.

2. Real estate first

Why: It offered the clearest, most observable daily work. Watching high-volume luxury agents (e.g., Netflix's *Owning Manhattan*) gave concrete workflows: market research, client briefs, showings, deal tracking. The persona is a composite.

Revisit if: Interviews show the core job is weaker than assumed.

3. A Catskills expansion scenario

Why: An NYC agent learning a new market fast is the hardest version of the problem. It forces research, comparison and evidence checks without losing existing client context.

4. A four-item main menu: Home, Workspaces, Settings, Account

Why: Fewer top-level items keep the shell stable. Library and Activity are checked from Home, not worked in, and industries sit inside Workspaces.

Revisit if: Users can't find the Knowledge library or Recent activity in testing.

5. Canvas tabs and a pinned Saved Work tab, not a "booklet"

Why: Tabs are familiar for parallel work, and a permanent Saved Work tab makes resuming obvious without a new metaphor to learn.

Revisit if: Users struggle to manage many open canvases.

6. A light, frosted look, not a dark sidebar

Why: The first prototype's dark menu read as generic "AI SaaS." The ocean-and-reef identity gives Myli a calm, editorial character.

7. Full capability on mobile

Why: Agents work from their phones, and persistent workspaces let work move between devices. Extra confirmation for high-risk mobile approvals is a hypothesis.

Revisit if: Mobile approvals feel unsafe, or the extra step adds friction without value.

Open questions and next steps

The next milestone is the approval screen (v1.1), then a usability test of the core flow.

Next steps

- [ ] Design the end-to-end approval screen to the approval framework spec.
- [ ] Run usability tests of the core flow with 5–8 agents against a chat-only baseline, measuring the success metrics above.
- [ ] Interview agents on which tasks they start on mobile and which on desktop.
- [ ] Complete the mobile treatments marked Planned.
- [ ] Scope the Finance & Investment workspace and its data boundaries.

Open questions

- Does Face ID for high-risk mobile approvals build trust or just add friction?
- What price point and trial length fit an agent's existing tool budget?
- How many open canvases does an agent work with at once, and do tabs scale?
- Which Catskills data sources are credible enough for client-facing briefs?

Changelog and sources

Version	Date	Change
1.1	Upcoming	Approval screen designed to spec; findings from design review added

Version	Date	Change
1.0	November 2026	First portfolio PRD: standalone architecture, real estate MVP, approval framework, mobile principle, benchmark-informed metrics, decision log

Design files

- Myli — Product Design & UX Case Study
- Myli — Base Design System

Sources

- NAR 2026 Member Profile (HousingWire)
- NAR: How Realtors are using technology (2026)
- NAR 2026 Technology Survey highlights (RISMedia)
- NAR 2025 Technology Survey (Seattle Agent Magazine)
- AI use in real estate falls short, Fyxer report (HousingWire)
- homegenius/RRC CMA survey (RISMedia)
- MeasuringU: 10 benchmarks
- MeasuringU: four UX metrics
- HBS/BCG “jagged frontier” study (The Harvard Crimson)
- Pew study on AI summaries (Slashdot)
- KPMG / University of Melbourne: Trust in AI 2025 (Mondaq)
- HUD fair housing guidance withdrawals (Navigate Housing)