

PosePilot AI

An AI photo-direction assistant that turns one photo of a location into a personal, ready-to-shoot plan.

Scene understanding · Grounded retrieval · Pose, composition and camera guidance in under 10 seconds

CONTENTS

Table of contents

01	Document control	3
02	Executive summary	4
03	Problem statement	5
04	Target users and jobs to be done	6
05	Market and competitive landscape	7
06	Goals, non-goals and success metrics	8
07	Solution overview and core user journey	9
08	Functional requirements	10
09	User stories	11
10	AI system design	12
11	Non-functional requirements	14
12	Responsible AI, privacy and safety	15
13	Data model	16
14	UX requirements and edge cases	17
15	Experimentation and launch plan	18
16	Roadmap and milestones	19
17	Monetization	20
18	Risks, dependencies and open questions	21
19	Appendix: glossary and assumptions	22

SECTION 01

Document control

This PRD defines the MVP of PosePilot AI and is ready for engineering, design and ML review.

Product	PosePilot AI (mobile-first web app; iOS and Android to follow)
Document type	Product Requirements Document, MVP (v1.0)
Status	Draft for review
Product owner	Harshal Borole, Product Manager
Reviewers	Engineering lead, ML lead, Design lead, Legal / Privacy
Related artifacts	PosePilot knowledge-base schema (PostgreSQL / Supabase), Figma prototype, evaluation rubric

PORTFOLIO NOTE

PosePilot AI is a concept product. User counts, research findings, pricing and targets in this document are planning assumptions and hypotheses, not results from a shipped launch, and are labeled where they appear. The photography knowledge base (8 tables, seeded and validated on PostgreSQL) is a built artifact.

Change log

VERSION	DATE	CHANGE
1.0	Oct 3, 2026	Full MVP PRD: requirements, AI design, metrics, launch plan
0.1	Earlier draft	Concept PRD and knowledge-base schema

Approvals

ROLE	NAME	DECISION	DATE
Product	Harshal Borole		
Engineering			
Machine learning			
Design			
Legal / Privacy			

SECTION 02

Executive summary

PosePilot AI turns one photo of a location into a ready-to-shoot plan in under 10 seconds: where to stand, how to pose, how to frame, and which camera settings to use.

<10 s

p90 time from photo to full shot plan

3 to 5

ranked, scene-specific shot ideas per plan

142

pose-to-scene links grounding every suggestion

35%

North Star target: plans shot within 7 days (hypothesis)

The problem

People arrive at a great location and still come home with stiff, poorly framed photos. Pose inspiration on Pinterest and Instagram is generic: it ignores the light, background and space the person is actually standing in.

The solution

The user snaps the scene. A vision model reads it; a retrieval layer matches it to a curated knowledge base; a language model writes a short, personal shot list with pose cards, framing overlays and device settings.

■ Why the hybrid design matters

The knowledge base holds 55 poses, 20 scene types, 6 styles, 20 composition rules and 25 camera-setting rules, linked by 142 pose-to-scene relationships. Grounding the LLM in this structure keeps advice specific and correct, prevents invented poses or impossible settings, and makes every suggestion traceable to a rule a photographer can audit.

■ MVP scope

In scope

- Scene analysis with confidence handling
- 3 to 5 ranked shot ideas with pose cards
- Composition overlay on the user's photo
- Phone and camera settings
- Shoot mode, save, share and feedback

Out of scope

- Live AR coaching through the camera
- Editing, filters or retouching
- Social feed or public profiles
- Professional and studio workflows

NORTH STAR METRIC

Planned shoots completed is the share of generated plans where the user marks at least one shot as taken within 7 days. Launch target: 35% (hypothesis, reviewed after four weeks of beta data).

SECTION 03

Problem statement

Casual photographers lack a fast way to turn a specific location into specific shot ideas, so they waste time, light and good locations.

USER PROBLEM

"When I am at a location I want to photograph, I don't know what poses, angles or settings will work *here*, so I take the same few shots, feel awkward, and leave without the photos I wanted."

Current workarounds and their gaps

WORKAROUND	WHAT USERS DO	GAP
Pinterest / Instagram boards	Save pose screenshots before the trip	Generic; not matched to the real light, background or space
The person behind the camera	Rely on a friend's eye	Inconsistent; most friends are not photographers
Phone camera auto mode	Point and shoot	Handles exposure, not pose, framing or story
Paid photographer	Book a shoot	Costly (\$150 to \$400 per session at typical local rates; assumption) and needs planning
General AI chatbots	Describe the scene in text	Advice is vague, unstructured and sometimes wrong

Root causes

- 1. **Tacit knowledge.** Pose and composition skill takes photographers years to learn.
- 2. **Context dependence.** Good advice depends on light direction, background clutter and depth. A text search can't capture.
- 3. **Time pressure.** Users have seconds, not minutes: friends are waiting and the light is moving.

Evidence plan: 10 to 12 interviews and a 3-mockup concept test before build

ID	HYPOTHESIS	VALIDATION SIGNAL
H1	Users feel awkward posing without direction	≥7 of 12 interviewees describe it unprompted
H2	Users will take a photo of the scene before shooting	≥70% complete the capture step in concept test
H3	A plan in under 10 seconds is fast enough on site	Majority rate wait time acceptable at 10 s

SECTION 04

Target users and jobs to be done

The MVP targets travelers and social creators shooting on a phone; hobbyist camera owners are a secondary segment.

PERSONA	PROFILE	CORE NEED	PRIMARY PAIN	PRIORITY
Maya, the traveler	27, shoots on iPhone, 4 to 6 trips a year	Great travel photos of herself and friends	Freezes at landmarks; photos look like everyone else's	PRIMARY
Jordan, the creator	23, posts 4× a week, 8K followers	Fresh, on-trend content at new spots	Runs out of pose ideas; shoots take too long	PRIMARY
Priya, the hobbyist	34, mirrorless camera, learning manual mode	Learn why a shot works	Unsure of settings for tricky light	SECONDARY
Sam, behind the camera	Shoots for a partner or friend	Take photos the other person likes	Doesn't know how to direct or frame	SECONDARY

■ Jobs to be done

Functional

When I arrive at a location, help me decide what to shoot and how, so I get keeper photos in minutes.

Emotional

When I'm being photographed, help me feel confident and natural, so I don't feel awkward in public.

Social

When I share photos, help them look intentional and unique, so my feed stands out.

Learning (secondary)

When a shot works, show me why, so I can repeat it without the app.

ANTI-PERSONAS: NOT DESIGNED FOR

Professional photographers running paid shoots, and studio photography with controlled lighting.

SECTION 05

Market and competitive landscape

No mainstream product combines scene understanding, pose direction and camera settings in one on-site plan; that gap is PosePilot's wedge.

CATEGORY	EXAMPLES	STRENGTH	WHERE IT FALLS SHORT	POSEPILOT ADVANTAGE
Inspiration platforms	Pinterest, Instagram	Huge visual library	Not location-aware; browsing, not directing	Plan built from the user's own scene
Pose-library apps	Pose reference apps	Many pose images	Static catalogs; no scene or settings	Poses ranked for this scene and light
Native camera aids	Grid lines, portrait mode, scene detection	Built in, zero friction	Optimizes exposure, not people or story	Adds pose, framing and story layer
Editing apps	Lightroom, VSCO	Strong post-processing	Fixes after the fact; can't fix a bad pose	Gets it right in camera
General AI assistants	Multimodal chatbots	Flexible, can see images	Generic, unstructured, may invent settings	Grounded in an auditable rule base; mobile-first output

POSITIONING STATEMENT

For travelers and creators who want better photos on the spot, **PosePilot AI** is a photo-direction assistant that turns a picture of any location into a personal shot list. Unlike pose galleries or chatbots, it is grounded in a curated photography knowledge base, so every suggestion fits the real scene.

■ Defensibility over time

- **Knowledge base depth:** a growing, curated library of poses, scenes and rules.
- **Feedback data:** which poses actually get shot in which scenes, a proprietary ranking signal.
- **Personalization:** each user's saved style improves relevance and switching cost.

SECTION 06

Goals, non-goals and success metrics

Success means users act on the plan: they shoot it, keep the photos and come back for the next location.

Goals (MVP)

1. Deliver a scene-specific shot plan in under 10 seconds (p90).
2. Make suggestions feel personal and correct, with $\geq 80\%$ of plans rated helpful.
3. Prove repeat value: users plan shoots at more than one location.

Non-goals (MVP)

- Real-time AR pose overlay through the live camera
- Editing, filters or retouching
- A social feed, followers or public profiles
- Professional or studio workflows

■ Metrics framework

All targets are launch hypotheses, reviewed after four weeks of beta data.


LAYER	METRIC	DEFINITION	TARGET
North Star	Planned shoots completed	Plans with ≥ 1 shot marked "taken" within 7 days ÷ plans generated	35%
Activation	First plan generated	New users who generate a plan in their first session	70%
Quality	Plan helpfulness	Thumbs-up ÷ (thumbs-up + thumbs-down) per plan	$\geq 80\%$
Quality	Suggestion relevance	Share of shot ideas rated relevant by an expert panel (offline eval)	$\geq 85\%$
Engagement	Shot ideas used per plan	Average shots marked taken per completed plan	≥ 2.0
Retention	W4 retention	Users who generate a plan in week 4 after sign-up	20%
Growth	Share rate	Plans or pose cards shared ÷ plans generated	10%
Guardrail	Latency	p90 time from upload to full plan	<10 s
Guardrail	Unsafe / invalid suggestions	Plans flagged for unsafe locations or impossible settings	<0.5%
Guardrail	Cost per plan	Model + infrastructure cost per generated plan	<\$0.04
Guardrail	Bias check	Helpfulness gap across body types, genders and skin tones in the eval set	<5 pts

SECTION 07

Solution overview and core user journey

The core loop is capture → plan → shoot → feedback, completed on site in under two minutes.

■ Core user journey (happy path)



Style options at step 2 are casual, editorial, romantic, adventurous, minimal and playful; people options are solo, couple and group; device options are phone and camera.

■ Key product decisions

DECISION	CHOICE	RATIONALE
Input	One still photo, not live video	Fast, cheap and works on weak signal; live AR is a v2 bet
Output length	3 to 5 shot ideas	Enough choice without decision fatigue on site
Grounding	Knowledge base + LLM, not LLM alone	Specific, auditable advice; fewer hallucinations
Personalization	Lightweight context chips + saved style	Personal results without a long onboarding
Explanations	One "why this works" line per shot	Builds trust and serves the learning job

SECTION 08

Functional requirements

Fourteen requirements define the MVP; the nine Must-haves gate launch.

ID	REQUIREMENT	PRIORITY	ACCEPTANCE CRITERIA
FR-01	Capture or upload a scene photo	MUST	Camera and gallery both work; JPEG / PNG / HEIC up to 15 MB accepted; image downscaled to ≤1600 px before upload
FR-02	Context chips (people, style, device)	MUST	Each chip is one tap; defaults pre-filled from profile; plan generates if user skips all
FR-03	Scene analysis	MUST	Returns scene type, lighting, light direction, depth and key elements; maps to one of 20 scene types with a confidence score
FR-04	Low-confidence handling	MUST	If scene confidence <0.6, user picks from the top 3 scene options before the plan generates
FR-05	Ranked shot plan	MUST	3 to 5 shot ideas, each with pose name, description, position cue and "why it works" line; no duplicate poses in one plan
FR-06	Composition overlay	MUST	Overlay (rule of thirds, leading lines, framing, etc.) drawn on the user's own photo for each idea
FR-07	Camera and phone settings	MUST	Phone: mode, lens, exposure tip. Camera: aperture, shutter, ISO range. All values come from the 25 setting rules
FR-08	Shoot mode	MUST	Full-screen pose card, ≤2 coaching cues, "Mark as taken" control; usable one-handed
FR-09	Plan feedback	MUST	Thumbs up / down per plan and per shot; optional reason tags ("not my style", "not possible here", "wrong settings")
FR-10	Save plans to trips	SHOULD	Plans saved to named trips; viewable offline after first load
FR-11	Share pose card	SHOULD	Share as image or link to messages and social apps; no personal data in the link
FR-12	Regenerate / "more like this"	SHOULD	Regenerate a plan or request variants of one shot; max 3 regenerations per plan
FR-13	Style profile	COULD	User picks 3 sample photos they like; style weights stored and applied to ranking
FR-14	Golden-hour timing tip	COULD	If location permission is granted, show best time of day for the scene's light

WON'T HAVE IN THIS RELEASE

Live AR overlay, auto-editing, public feed, multi-photo scene stitching.

SECTION 09

User stories

Stories are grouped by epic and traced to requirements; the highest-risk flows carry Given / When / Then acceptance tests.

■ Epic 1: Plan a shoot

US-01 As a traveler, I want to photograph a location and get shot ideas, so I don't waste time guessing. FR-01, FR-03, FR-05

Given I have taken a scene photo, **when** I tap "Plan this spot," **then** I see 3 to 5 ranked shot ideas within 10 seconds.

US-02 As a couple, we want ideas that fit two people, so neither of us looks awkward. FR-02, FR-05

Given I selected "Couple," **when** the plan loads, **then** every pose shown is a two-person pose.

US-03 As a user in a scene the app can't read, I want to confirm what it is, so the ideas still fit. FR-04

Given scene confidence is below 0.6, **when** analysis finishes, **then** I pick from 3 scene options before seeing the plan.

■ Epic 2: Shoot the plan

US-04 As the person behind the camera, I want to see where to stand and how to frame, so I can direct confidently. FR-06

US-05 As a hobbyist, I want the right settings for this light, so the shot is sharp and well exposed. FR-07

Given I selected "Camera," **when** I open a shot, **then** I see aperture, shutter speed and ISO ranges with one line on why.

US-06 As someone posing, I want one or two simple cues, so I can follow them without stopping the shoot. FR-08

■ Epic 3: Learn and return

US-07 As a user, I want to rate a plan, so future plans get better. FR-09

US-08 As a traveler, I want to save plans to a trip, so I can use them later even without signal. FR-10

US-09 As a creator, I want to share a pose card with my photographer, so we stay in sync. FR-11

US-10 As a returning user, I want results to match my style, so I don't re-explain it every time. FR-13

SECTION 10

AI system design

PosePilot uses a four-stage pipeline (perceive, retrieve, rank, generate), so the LLM writes the plan but the knowledge base decides what is in it.

The knowledge base decides what goes in the plan; the LLM only writes it




Figure 1. PosePilot AI pipeline: four stages and one validation gate. A plan that fails validation loops back to generation; user feedback flows back into ranking.

■ Stage 1: Scene understanding (vision model)

A multimodal model receives the downscaled photo and returns strict JSON: scene type (one of 20), lighting (golden hour, overcast, harsh midday, backlit, low light, indoor mixed), light direction, depth (shallow, medium, deep), key elements (leading lines, frames, symmetry, water, stairs) and clutter level, each with a confidence score.

■ Stage 2: Retrieval (knowledge base)

Scene type and context chips query PostgreSQL. The pose to scene relationship table returns candidate poses with a fit score; composition and camera-setting rules are filtered by scene, lighting and device.

■ Stage 3: Ranking (deterministic scorer)

$$\text{Score} = 0.40 \times \text{SceneFit} + 0.25 \times \text{StyleMatch} + 0.15 \times \text{GroupFit} + 0.10 \times \text{LightFit} + 0.10 \times \text{Feedback}$$

Weights are starting values, tuned with offline evaluation and later with feedback data. A diversity rule keeps at most two poses from the same pose category in one plan.

■ Stage 4: Plan generation (LLM)

The LLM receives only the top candidates, the matched rules and the scene JSON. It writes the user-facing plan: position cue, coaching cues and “why it works” lines. It may not introduce poses or settings outside the provided candidates; a validator rejects and regenerates any plan that does.

■ Model choices and trade-offs

COMPONENT	MVP CHOICE	ALTERNATIVE CONSIDERED	RATIONALE
Scene understanding	Hosted multimodal LLM with JSON schema	Custom CNN classifier	No labeled data yet; faster to ship; revisit once 50K labeled scenes exist
Retrieval	SQL over curated tables	Vector search over pose images	Structured rules are auditable; vector search added in v2 for visual pose similarity
Ranking	Weighted rules	Learning-to-rank model	Needs feedback volume first
Generation	Small, fast LLM	Largest available model	Latency and cost guardrails; quality held by grounding

■ Evaluation plan

LAYER	METHOD	BAR
Golden set	300 labeled scene photos across all 20 scene types, balanced across lighting conditions	Regression-tested every release
Offline metrics	Scene-type accuracy; expert-rated shot relevance; settings validity; hallucination rate	$\geq 85\%$ · $\geq 85\%$ · 100% in range · 0 poses outside KB
Human review	Two photographers score 50 plans per release on relevance, clarity and feasibility (1 to 5)	Release blocked if average <4.0
Online	Helpfulness rating, shots marked taken, regeneration rate as a dissatisfaction signal	Tracked against metrics framework

FALLBACKS

If the vision call fails or times out (>6 s), show the low-confidence scene picker. If generation fails, render the ranked candidates as plain cards straight from the knowledge base with no LLM text.

SECTION 11

Non-functional requirements

Speed and cost are launch gates because the product is used on site, on mobile data, with people waiting.

AREA	REQUIREMENT
Performance	p50 plan time <6 s, p90 <10 s on 4G; first screen interactive <2 s
Availability	99.5% monthly uptime for plan generation
Scalability	50 plans / minute at beta, 500 / minute at public launch without redesign
Cost	<\$0.04 model + infrastructure cost per plan; alert at \$0.05
Offline	Saved plans and pose cards viewable offline; capture queues upload until signal returns
Security	TLS 1.2+ in transit, AES-256 at rest; row-level security on user data in Supabase
Accessibility	WCAG 2.1 AA; pose cues readable by screen readers; 44 × 44 pt minimum touch targets; one-handed use
Device support	Last 3 major iOS and Android versions; mobile web on Safari and Chrome
Observability	Latency per pipeline stage, model errors, validator rejections and cost per plan logged and dashboarded
Localization	English at launch; all strings externalized for later languages

SECTION 12

Responsible AI, privacy and safety

PosePilot analyzes places, not people: no face recognition, no body scoring, and scene photos are deleted after 24 hours by default.

Privacy

- Scene photos deleted within 24 hours unless the plan is saved.
- No facial recognition, identity inference or body measurement.
- Photos never used for training without separate, explicit opt-in.
- EXIF location stripped on upload; location used only for golden-hour tips with permission.
- Zero-data-retention terms with model providers.
- CCPA / CPRA and GDPR rights: in-app export and account deletion.

Fairness and inclusion

- Seated, standing and adaptive pose variants where possible.
- Illustrations show diverse body types, ages, skin tones and mobility aids.
- Neutral, encouraging language, with no “slimming” or body-judging cues.
- Bias guardrail: helpfulness gap across demographic slices <5 points.

Physical safety

- Risky poses (edges, railings, water, traffic) are tagged and never suggested when the scene shows height, cliffs, roads or tracks.
- Safety reminder on any plan with water or elevation.
- “Report unsafe” on every suggestion; reviewed within 48 hours.

Transparency

- Every shot shows “why it works,” tied to a composition rule.
- Plans are labeled as AI-generated suggestions.
- Feedback controls on every plan and every shot.

SECTION 13

Data model

The knowledge base is an 8-table PostgreSQL schema on Supabase, seeded and validated against a live instance; app data sits in separate user-scoped tables.

■ Photography knowledge base (seeded)

TABLE	SEED ROWS	HOLDS	USED IN
poses	55	Pose name, category, people count, difficulty, cues, safety tags	Retrieval, shoot mode
scenes	20	Scene type, typical elements, lighting traits	Scene mapping
styles	6	Casual, editorial, romantic, adventurous, minimal, playful	Context chips, ranking
composition_rules	20	Rule, description, overlay type	Overlay, “why it works”
camera_setting_rules	25	Scene + light + device → setting ranges	Settings output
pose_scene_relationships	142	Pose ↔ scene fit score	Candidate retrieval
+ 2 supporting tables	Varies	Lookup and linking tables per schema file	Joins

■ App data (new for MVP)

TABLE	KEY FIELDS	RETENTION
users	id, auth id, default style, default device	Until account deletion
plans	id, user id, scene JSON, ranked shot ids, created at	Saved: until deleted · Unsaved: 30 days
plan_feedback	plan id, shot id, rating, reason tag, taken flag	Anonymized after 12 months
trips	id, user id, name	Until deleted

STORAGE NOTE

Scene photos live in object storage, not the database, under the 24-hour deletion rule.

SECTION 14

UX requirements and edge cases

The interface is built for one hand, bright sunlight and ten-second attention spans.

Two taps to a plan

Capture, then “Plan this spot.” Everything else is optional.

Show, don't tell

Pose illustrations and overlays on the user's own photo beat text.

Glanceable in sunlight

High-contrast mode, large type in shoot mode, no low-contrast grey text.

Explain briefly

One “why it works” line, expandable for learners.

Key screens: Home (capture) · Context chips · Analyzing state · Plan (shot list) · Shot detail with overlay · Shoot mode · Trips · Profile and style.

■ **Edge cases**

CASE	EXPECTED BEHAVIOR
Photo is dark, blurry or mostly sky / ground	Prompt to retake with a tip (“step back to include the background”)
Photo contains people's faces	Process normally; faces are not analyzed or stored; no face-based output
Scene not among the 20 types	Low-confidence picker (FR-04); fall back to the closest general scene
Unsafe scene (cliff edge, road, rail track)	Exclude risky poses; show safety banner
Group of 6+ people	Use group poses; cap guidance at large-group layouts
No network at capture	Queue the photo; notify when the plan is ready
Indoor or night scene	Lighting-specific settings; flash-off and stabilization tips
Inappropriate or non-location image	Decline politely after a moderation check; no plan generated
User regenerates 3 times	Suggest “try a different angle of the spot” instead of a 4th regeneration

SECTION 15

Experimentation and launch plan

Launch is staged (internal, closed beta, open beta, public), and each stage advances only when its exit criteria are met.

■ Launch stages

STAGE	AUDIENCE	EXIT CRITERIA TO ADVANCE
Internal dogfood	Team and friends (~25)	No P0 bugs; p90 latency <10 s; offline eval targets met
Closed beta	300 invited travelers and creators	Helpfulness ≥75%; unsafe-suggestion rate <0.5%; cost / plan <\$0.05
Open beta	Waitlist (~3,000)	North Star ≥30%; W2 retention ≥25%
Public launch	App stores + web	All guardrails green for 2 consecutive weeks

■ Planned experiments

#	HYPOTHESIS	VARIANT VS CONTROL	PRIMARY METRIC	GUARDRAIL
E1	3 ideas convert better than 5 on site	3 vs 5 shot ideas	Planned shoots completed	Helpfulness
E2	“Why it works” lines build trust	With vs without explanation	Helpfulness rating	Time on plan
E3	Overlays on the user's photo beat stock pose images	Own-photo overlay vs stock image	Shots marked taken	Latency
E4	Asking for style up front personalizes better	Style chip at start vs inferred from feedback	W4 retention	Activation

Each test runs to 80% power at $\alpha = 0.05$, with a minimum detectable effect of 5 percentage points on the primary metric.

GO-TO-MARKET

Creator partnerships (10 travel micro-creators sharing pose cards), a Product Hunt launch, and SEO pages for top photo spots (“best poses at [landmark]”).

SECTION 16

Roadmap and milestones

The MVP reaches public launch in about nine months, and each phase advances only when its gate is met.

MVP ships in June 2027 after three quality gates




Figure 2. Delivery roadmap: four phases separated by three quality gates. Dates are planning targets.

■ Milestone detail

PHASE	WINDOW	KEY DELIVERABLES	EXIT GATE
1. Discover	Oct to Nov 2026	12 user interviews, 3-mock concept test, 300-photo golden set	H1 to H3 validated
2. Build MVP	Dec 2026 to Feb 2027	FR-01 to FR-09, evaluation pipeline and validator, internal dogfood	Offline eval targets met
3. Beta	Mar to May 2027	Closed beta (300), open beta (~3,000), experiments E1 to E4	All guardrails green
4. Launch v1.0	Jun 2027	App stores and web, Pro tier, creator partnerships	Launch complete

A missed gate extends the current phase rather than cutting scope from the Must-have list.

SECTION 17

Monetization

Freemium: free plans build the habit; PosePilot Pro charges for volume, personalization and trip features. Pricing is a hypothesis to test in open beta.

TIER	PRICE (HYPOTHESIS)	INCLUDES
Free	\$0	5 plans per month, 3 shot ideas per plan, save up to 2 trips
Pro	\$4.99 / mo or \$29.99 / yr	Unlimited plans, 5 ideas + variants, style profile, offline trips, golden-hour tips
Trip Pass	\$2.99 one-time, 7 days	Pro features for a single trip

■ Unit economics check (assumptions)

LINE	AT 15% STORE FEE	AT 30% STORE FEE
Pro monthly price	\$4.99	\$4.99
Net after store fee	\$4.24	\$3.49
Model + infra cost (30 plans × \$0.04)	\$1.20	\$1.20
Gross margin	~72%	~46%

Margin depends heavily on the store fee, so annual plans and web billing matter.

LATER OPTIONS OUTSIDE MVP

B2B licensing to tourism boards and hotels ("PosePilot spots"), and affiliate links for gear recommendations.

SECTION 18

Risks, dependencies and open questions

The biggest risk is plan quality: if the first plan feels generic, users won't return.

RISK	LIKELIHOOD	IMPACT	MITIGATION
Plans feel generic	MEDIUM	HIGH	Overlay on user's own photo; scene-specific cues; expert eval gate before each release
Vision model misreads scene	MEDIUM	HIGH	Confidence threshold + scene picker; golden-set regression tests
LLM invents poses or settings	LOW	HIGH	Constrained candidates + output validator; regenerate on violation
Latency over 10 s on mobile data	MEDIUM	MEDIUM	Image downscaling, small generation model, parallel retrieval, stream the first card
Model cost grows with usage	MEDIUM	MEDIUM	Cost alerts, caching common scenes, free-tier caps
Unsafe pose suggestion	LOW	HIGH	Safety tags, scene hazard exclusion, report flow
Knowledge base too small (55 poses)	HIGH	MEDIUM	Expand to 150 poses by v1.1, prioritized by feedback and scene coverage gaps
Platforms copy the feature	MEDIUM	MEDIUM	Move fast on personalization and feedback-driven ranking data

■ Dependencies

- Multimodal model API with JSON mode and zero-retention terms
- Supabase (PostgreSQL, auth, storage)
- Pose illustration set (design)
- Two photographer advisors for evaluation
- App-store review

■ Open questions

- ☐ Should the free tier allow 5 plans a month or 3 plans a week?
- ☐ Ship native apps at launch, or mobile web first?
- ☐ License or commission the pose illustrations?
- ☐ What share of users will grant location for golden-hour tips?
- ☐ Is 24-hour photo deletion the right default, or should unsaved photos be deleted right after the plan?

SECTION 19

Appendix

■ Glossary

TERM	MEANING
Shot plan	The 3 to 5 ranked shot ideas generated for one scene photo
Shot idea / pose card	One suggestion: pose, framing overlay, position cue, settings, "why it works"
Scene JSON	Structured output of the vision model describing the location
Grounding	Restricting the LLM to facts and options supplied by the knowledge base
Golden set	Labeled test images used to measure model quality before each release
Planned shoot completed	A plan where at least one shot is marked taken within 7 days

■ Key assumptions

1. Users will take a photo of the location before shooting (validated in interviews, H2).
2. A hosted multimodal model can classify the 20 scene types at $\geq 85\%$ accuracy without fine-tuning.
3. 55 curated poses cover most casual travel and creator shoots for the MVP.
4. Per-plan model cost stays under \$0.04 at current API pricing.
5. Pricing, targets and user counts in this PRD are hypotheses for a concept product, not measured results.

END OF DOCUMENT