



AIN SHAMS UNIVERSITY

Faculty of Engineering

Department of Computer and Systems Engineering

INTELLIROOM AI

Graduation Project

Submitted by:

Youssef Ossama Sayed	2101240
Mohamed Sameh Abdelrahman Ahmed	2100401
Marwan EL Nabawy Abd elraheem	2100564
Badr Sayed Abdelaziz	2100282
Ahmed Karam Abdel Hamid	2101767
Mohamed Hani Hamdi	2100915
Alharith Mujeeb Hasan	2101358

Supervised by:

Prof. Gamal A. Ebrahim

June 2026

Submitted in partial fulfillment of the requirements for the degree of
Bachelor of Science in Computer and Systems Engineering

Declaration

We hereby declare that this graduation project entitled “**IntelliRoom: AI-Powered Interior Design Platform with Spatial Intelligence and Egyptian E-Commerce Integration**” is our own original work and has not been submitted previously for any degree or diploma at any other university or institution.

All sources of information and references used in this project have been duly acknowledged and cited in accordance with academic standards using IEEE citation format.

We further declare that the work presented in this document represents our genuine effort and contribution to the field of AI-powered interior design, specifically tailored for the Egyptian market.

Date: June 2026

Abstract

Online furniture retail in Egypt suffers from a critical visualization gap: shoppers cannot see how products will look in their own spaces, leading to cart abandonment rates exceeding 90%. Meanwhile, professional interior design services cost 500–2,000 EGP per room—prohibitive for most households. IntelliRoom bridges this gap with an AI-powered interior design platform built specifically for the Egyptian e-commerce market.

IntelliRoom leverages a production-ready generative AI engine built on **Flux2Klein 4B**—a state-of-the-art Diffusion Transformer—combined with SAM2 for precision image segmentation, ControlNet for structural preservation, and Florence-2 for intelligent object detection, all orchestrated through ComfyUI workflows. Users upload a room photograph and receive high-fidelity redesign results while the system simultaneously identifies and retrieves visually matching furniture from Egyptian retailers.

The furniture retrieval engine implements a novel two-plane architecture: a local FastAPI orchestrator (ChromaDB vector database, semantic taxonomy snapper) connected via zrok tunneling to a remote GPU inference server running **RF-DETR** for furniture detection, **SigLIP2 so400m** for 1152-dimensional visual embeddings, and **Qwen3-VL 4B**—selected through an empirical ablation study over 2B/4B/8B variants—for taxonomy tagging. This pipeline enables accurate, semantically-constrained product retrieval from the Egyptian furniture catalog at commercial scale.

The platform delivers a complete full-stack solution: a React.js frontend, a Node.js/Express REST API with asymmetric JWT (RS256) authentication, MongoDB with 14+ schemas, a Fawaterk-integrated subscription system, a community gallery, and a plugin marketplace. The result is a commercially viable platform that democratizes interior design visualization for Egyptian homeowners and directly converts design inspiration into purchasable local furniture.

Keywords: Artificial Intelligence, Interior Design, Flux2Klein, Computer Vision, Image Segmentation, Generative AI, Egyptian E-Commerce, Furniture Retrieval, RF-DETR, SigLIP2, Qwen3-VL, ComfyUI, Vector Database, SAM2, ControlNet

Table of Contents

List of Figures

List of Tables

List of Abbreviations

Abbreviation	Definition
AI	Artificial Intelligence
API	Application Programming Interface
AR	Augmented Reality
ASU	Ain Shams University
CAGR	Compound Annual Growth Rate
CNN	Convolutional Neural Network
CLIP	Contrastive Language-Image Pre-Training
ComfyUI	Comfy User Interface (Node-based AI workflow tool)
ControlNet	Neural network architecture for controlling diffusion models
CSS	Cascading Style Sheets
CSRF	Cross-Site Request Forgery
DPI	Dots Per Inch
DiT	Diffusion Transformer
EGP	Egyptian Pound
ERD	Entity-Relationship Diagram
FK	Foreign Key
GPU	Graphics Processing Unit
HTML	HyperText Markup Language
HTTP	HyperText Transfer Protocol

IEEE	Institute of Electrical and Electronics Engineers
img2img	Image-to-Image (generation technique)
JSON	JavaScript Object Notation
JWT	JSON Web Token
ML	Machine Learning
NSFW	Not Safe For Work
ODM	Object-Document Mapper
PK	Primary Key
PDF	Portable Document Format
REST	Representational State Transfer
RF-DETR	Roboflow Detection Transformer
SAM	Serviceable Addressable Market
SAM2	Segment Anything Model 2
SD	Stable Diffusion
SigLIP	Sigmoid Language-Image Pre-training
SigLIP2	Sigmoid Language-Image Pre-training 2
SOM	Serviceable Obtainable Market
SQL	Structured Query Language
TAM	Total Addressable Market
UI	User Interface
URL	Uniform Resource Locator
UX	User Experience
VAE	Variational Autoencoder
VDB	Vector Database
VGG16	Visual Geometry Group 16-layer network
VLM	Vision-Language Model

VRAM	Video Random Access Memory
WebXR	Web Extended Reality
XSS	Cross-Site Scripting
2D	Two-Dimensional
3D	Three-Dimensional

Chapter 1

Introduction

The interior design industry stands at a pivotal transformation point, driven by advances in artificial intelligence and computer vision technologies. While AI-powered design tools have proliferated globally, a critical gap remains: online furniture shoppers cannot visualize products in their own spaces before purchase. This chapter introduces IntelliRoom, an AI-powered interior design platform specifically developed to address this visualization gap in the Egyptian e-commerce market through spatial intelligence and direct furniture marketplace integration.

1.1 Problem Statement

The global interior design technology market has witnessed remarkable growth, expanding from USD 18.32 billion in 2024 toward projected valuations exceeding USD 184 billion by 2032. Despite this growth, a critical practical barrier persists for Egyptian homeowners and online furniture shoppers: the inability to visualize furniture in their actual living spaces before purchase.

Online furniture retailers in Egypt report cart abandonment rates exceeding 90%—far above the e-commerce average of 70–75%—driven primarily by insufficient visualization tools. Meanwhile, professional interior design consultations cost 500–2,000 EGP per room, making them inaccessible to most households.

Existing AI interior design platforms such as COOHOM, Interior AI, RoomGPT, and Spacely AI offer generative capabilities but fail to address the core Egyptian market needs:

- No integration with Egyptian furniture retailers or local pricing in EGP
- No direct connection between AI-generated designs and purchasable products
- No support for Egyptian payment methods (Fawry, InstaPay)
- No e-commerce checkout or order management workflow

This visualization-to-purchase gap creates a substantial commercial opportunity for a platform that connects AI-powered spatial redesign directly to the Egyptian furniture marketplace.

1.2 Project Objectives

The IntelliRoom platform aims to bridge the gap between AI-powered spatial design and Egyptian e-commerce integration through the following objectives:

1.2.1 Primary Objectives

1. **Develop an AI-powered spatial redesign engine** powered by Flux2Klein 4B for high-fidelity room transformations, with SAM2 for precision segmentation, ControlNet for structure-preserving generation, and Florence-2 for room analysis, all orchestrated through ComfyUI workflows.
2. **Implement an advanced AI furniture retrieval pipeline** using RF-DETR for furniture detection, SigLIP2 for visual embeddings, Qwen3-VL for taxonomy tagging, and ChromaDB for vector similarity search.
3. **Create an Egyptian furniture retrieval framework** that connects users with local manufacturers and craftsmen, enabling seamless transition from AI-generated designs to purchasable products.
4. **Build a community-driven design platform** where users can share designs, discover inspiration, and collaborate on culturally-relevant interior solutions.
5. **Establish a sustainable business model** through credit-based access tiers, marketplace integration, and premium feature offerings.

1.2.2 Secondary Objectives

1. Develop comprehensive UI/UX designs that accommodate Arabic language support and right-to-left interfaces
2. Integrate Egyptian payment methods including InstaPay and Fawry
3. Create educational content and tutorials for users unfamiliar with AI-powered design tools
4. Implement gamification features to enhance user engagement and platform retention

1.3 Project Scope

The IntelliRoom project scope has been carefully defined to ensure deliverable completion within the graduation project timeline while maintaining ambitious technical and market objectives. Table 1.1 presents the scope boundaries.

Table 1.1: Project Scope Definition

Category	Items
In Scope	• Room photo upload and AI-powered preprocessing • Style transformation engine with Flux2Klein generative work-flows • Furniture detection and selective replacement • Egyptian furniture catalog integration • User authentication and profile management • Credit-based usage system • Community gallery for design sharing • 2D floor planner interface • Basic 3D visualization capabilities • Complete shopping cart, order registration, and subscription checkout with Fawaterk
Out of Scope	• Direct payment processing for physical furniture purchases (handled via retailer affiliate links) • Real-time video generation for walkthroughs • Mobile native applications (iOS/Android) • Physical furniture delivery logistics • Multi-language support beyond English and Arabic
Future Enhancements	• Augmented reality furniture placement preview • Video generation for virtual room walkthroughs • Multi-room consistency for whole-home redesigns • Reinforcement learning from user preferences • Voice-controlled AI design assistant

Note: Project scope is subject to refinement based on stakeholder feedback and technical feasibility assessment during the implementation phase.

1.4 Project Timeline

The IntelliRoom project follows a structured development methodology spanning two academic semesters, from August 2025 through July 2026. The project timeline is organized into six major phases as illustrated in Figure 1.1.




Figure 1.1: Project Timeline and Development Phases

1. **Project Analysis (August–September 2025):** Initial problem identification, market research, and feasibility assessment
2. **Planning (September–October 2025):** Requirements gathering, technical architecture design, and resource allocation
3. **Design (October–December 2025):** System design, database schema, UI/UX mockups, and workflow documentation (completed)
4. **Implementation (January–April 2026):** Backend API development, frontend implementation, AI pipeline integration (current phase)
5. **Testing and Quality Assurance (April–May 2026):** Unit testing, integration testing, user acceptance testing

6. **Deployment and Documentation (May–July 2026):** Production deployment, final documentation, project presentation

1.5 Methodology

The IntelliRoom project adopts an **Agile development methodology** with iterative sprint cycles, enabling continuous refinement based on emerging AI breakthroughs and technical discoveries. This approach is particularly suited for AI-intensive projects where model performance and user experience require iterative optimization, especially as new models and techniques are released by the research community.

Key methodological elements include:

- **Sprint-based development:** Two-week sprint cycles with defined deliverables
- **Continuous integration:** Automated testing and deployment pipelines
- **User-centered design:** Regular feedback collection from target Egyptian users
- **Documentation-driven development:** Comprehensive technical documentation maintained throughout the project lifecycle

1.6 Document Organization

This graduation project document is organized into five chapters, structured to present the complete design phase deliverables:

- **Chapter 1 (Introduction):** Presents the problem statement, objectives, scope, and project timeline
- **Chapter 2 (Literature Review):** Reviews market analysis, competitive positioning, and related academic work in AI-powered interior design
- **Chapter 3 (System Requirements):** Documents functional and non-functional requirements, use cases, and user stories
- **Chapter ?? (System Design and Architecture):** Presents technical architecture, database design, UI mockups, and system workflows
- **Chapter ?? (Conclusion and Future Work):** Summarizes achievements, expected outcomes, and implementation roadmap

Supporting materials including the complete market analysis, detailed use case diagrams, business model canvas, and additional UI mockups are provided in the appendices.

Chapter 2

Literature Review

This chapter presents a comprehensive review of the technological foundations, relevant academic research, and the research gap that informs the IntelliRoom platform development. The focus is on AI technologies for interior design, computer vision techniques, and related academic contributions. Complete market analysis is provided in Appendix ??.

2.1 Technology Review

The IntelliRoom platform leverages state-of-the-art AI technologies that have matured significantly in recent years. This section reviews the key technological foundations.

2.1.1 Generative Diffusion Models

Diffusion models represent the current state-of-the-art in image generation, operating through a process of iteratively denoising random noise into coherent images guided by conditioning signals [1]. Key developments include:

- **Stable Diffusion:** Open-source latent diffusion model enabling high-quality image generation with reasonable computational requirements [2]
- **ControlNet:** Architecture for adding spatial conditioning to diffusion models, enabling structure-preserving transformations essential for interior design [3]
- **Flux2Klein 4B:** A state-of-the-art Diffusion Transformer (DiT) architecture for high-fidelity image generation and inpainting, used as IntelliRoom’s primary generative engine under an Apache 2.0 license [4]

2.1.2 Computer Vision for Interior Design

Modern computer vision enables precise understanding and manipulation of room images:

- **SAM2 (Segment Anything Model 2):** Meta’s foundation model for zero-shot image segmentation, providing surgical precision in isolating furniture, walls, and architectural elements [5]

- **Florence-2:** Microsoft’s vision foundation model for object detection, captioning, and visual grounding, enabling automatic room analysis [6]
- **Depth Estimation:** Monocular depth prediction models enable 3D understanding from 2D photographs, essential for realistic furniture placement

2.1.3 ComfyUI Workflow Architecture

ComfyUI provides the foundational infrastructure for IntelliRoom’s AI pipeline, offering:

- Node-based visual programming for complex AI workflows
- Modular pipeline construction enabling detection → segmentation → generation → compositing sequences
- Custom node ecosystem for specialized operations
- API accessibility for backend integration

2.2 Related Academic Work

Several academic contributions inform the IntelliRoom approach to AI-powered interior design.

2.2.1 Furniture Set Retrieval

IntelliRoom implements a zero-shot furniture retrieval pipeline: RF-DETR detects and crops furniture from room images, SigLIP2 so400m generates 1152-dimensional visual embeddings, and ChromaDB performs cosine similarity search with metadata filtering against the Egyptian furniture catalog [7], [8].

2.2.2 Style Transfer in Interior Design

Research in neural style transfer has evolved from artistic applications to practical interior design scenarios [9]. Recent advances enable:

- Content-aware style application that respects spatial relationships
- Multi-style blending for hybrid aesthetic creation
- Semantic-guided transfer that distinguishes furniture, walls, and decorative elements

2.3 Research Gap and Contribution

The literature review reveals a clear research and market gap: while AI interior design technology has achieved remarkable capabilities, the application of these technologies to culturally-specific contexts remains largely unexplored. IntelliRoom addresses this gap through:

1. **Flux2Klein Integration:** Leveraging a state-of-the-art DiT-based generative model for high-fidelity spatial redesign, orchestrated via ComfyUI workflows
2. **Two-Plane Inference Architecture:** A novel local FastAPI orchestrator connected via zrok tunneling to a remote GPU inference server, enabling commercial-scale AI generation without cloud lock-in
3. **VDB Furniture Retrieval:** A zero-shot visual retrieval pipeline (RF-DETR + SigLIP2 + Qwen3-VL + ChromaDB) that connects AI-generated room designs to purchasable Egyptian furniture catalog items
4. **Integrated Egyptian E-Commerce:** Direct connections between AI-generated designs and purchasable local products with EGP pricing and Fawaterk payment integration

2.4 Market Context

The Egyptian interior design market provides a compelling context for IntelliRoom's development. Key market indicators include:

- Egyptian interior design market: USD 940M (2024) growing to USD 1.51B (2033) at 6% CAGR
- Population of 113.5 million with 75% internet penetration
- Furniture e-commerce growing at 27.5% CAGR—the fastest-growing sector
- 78% of furniture buyers require AR/3D visualization tools
- 90.75% cart abandonment rate indicating unmet visualization needs

Analysis of competitors including COOHOM, Interior AI, RoomGPT, and Spacely AI revealed universal absence of direct e-commerce integration with local Egyptian furniture retailers, presenting a clear market opportunity for IntelliRoom.

Note: For the comprehensive market research including detailed competitive analysis with feature comparisons, Egyptian market size and growth projections with visualizations, user segmentation and demographic profiles, business strategy and revenue model, and risk assessment with mitigation strategies, see Appendix ??.

Chapter 3

System Requirements

This chapter presents the comprehensive requirements specification for the IntelliRoom platform, documenting the functional and non-functional requirements, use cases, and user stories that guide the implementation phase. The requirements have been gathered through stakeholder interviews, competitive analysis, and user research targeting Egyptian market users. Implementation status is noted for each requirement.

3.1 Functional Requirements

The functional requirements define the specific behaviors and capabilities that the IntelliRoom system provides. Table 3.1 presents the categorized functional requirements with implementation status.

Table 3.1: Functional Requirements Specification

ID	Category	Requirement Description	Priority	Status
FR-01	Authentication	The system shall allow users to register accounts using email and password	High	Complete
FR-02	Authentication	The system shall support social login via Google and Facebook	Medium	Complete
FR-03	Authentication	The system shall implement secure password recovery via email	High	Complete
FR-04	Profile	The system shall allow users to update profile information including name, photo, and bio	Medium	Complete
FR-05	Profile	The system shall display user design history and saved projects	High	Complete
FR-06	Room Upload	The system shall accept room photographs in JPEG and PNG formats	High	Complete

ID	Category	Requirement Description	Priority	Status
FR-07	Room Upload	The system shall validate image quality and provide feedback on unsuitable uploads	High	Complete
FR-08	Room Upload	The system shall support direct camera capture from mobile devices	Medium	Planned
FR-09	Room Analysis	The system shall automatically detect room type (bedroom, living room, kitchen, etc.)	High	Complete
FR-10	Room Analysis	The system shall identify existing furniture items in uploaded photographs	High	Complete
FR-11	Room Analysis	The system shall analyze current room style and provide style classification	Medium	Complete
FR-12	Style Selection	The system shall offer predefined style presets including Arabic Traditional, Modern Egyptian, Mediterranean, and Contemporary	High	Complete
FR-13	Style Selection	The system shall allow mood-based design preferences (cozy, minimalist, luxurious)	Medium	Complete
FR-14	AI Generation	The system shall generate 3-5 design variations for each request	High	Complete
FR-15	AI Generation	The system shall preserve room structure (walls, windows, doors) during transformation	High	Complete
FR-16	AI Generation	The system shall enable selective furniture replacement while maintaining room context	High	Complete
FR-17	Visualization	The system shall provide before/after comparison with interactive slider	High	Complete

ID	Category	Requirement Description	Priority	Status
FR-18	Visualization	The system shall support zoom and detail inspection of generated designs	Medium	Complete
FR-19	2D Planning	The system shall provide a 2D floor planner for room layout design	Medium	Complete
FR-20	2D Planning	The system shall support drag-and-drop furniture placement in 2D view	Medium	Complete
FR-21	3D Visualization	The system shall convert 2D layouts to 3D visualization	Low	Complete
FR-22	Marketplace	The system shall display furniture product recommendations based on generated designs	High	Complete
FR-23	Marketplace	The system shall link to Egyptian furniture retailers and manufacturers	High	Complete
FR-24	Marketplace	The system shall display prices in Egyptian Pounds (EGP)	High	Complete
FR-25	Credits	The system shall implement a credit-based usage system	High	Complete
FR-26	Credits	The system shall display remaining credit balance in user dashboard	High	Complete
FR-27	Credits	The system shall support subscription purchase via Fawaterk (supporting Fawry, cards, and wallets)	Medium	Complete
FR-28	Community	The system shall allow users to share designs to a public gallery	Medium	Complete
FR-29	Community	The system shall support design likes and comments	Low	Complete
FR-30	Export	The system shall enable high-resolution download of generated designs	High	Complete

ID	Category	Requirement Description	Priority	Status
FR-31	Export	The system shall generate PDF reports with furniture specifications	Medium	Complete
FR-32	Tutorial	The system shall provide onboarding tutorials for new users	Medium	Planned
FR-33	E-commerce	The system shall provide shopping cart functionality	High	Complete
FR-34	E-commerce	The system shall support order management and tracking	High	Complete
FR-35	E-commerce	The system shall implement product wishlist functionality	Medium	Complete
FR-36	E-commerce	The system shall support product search and filtering by category	High	Complete
FR-37	Projects	The system shall allow users to save and manage multiple design projects	High	Complete
FR-38	Projects	The system shall store 2D/3D scene data for each project	High	Complete
FR-39	Plugin Market	The system shall provide a community plugin marketplace for custom styles and presets	High	Complete
FR-40	Plugin Market	The system shall allow plugin creators to publish plugins with pricing in credits	Medium	Complete
FR-41	Plugin Market	The system shall support plugin search, filtering, and category browsing	High	Complete
FR-42	Plugin Market	The system shall display plugin ratings, reviews, and download counts	Medium	Complete
FR-43	Contact	The system shall provide a contact form for user inquiries and support requests	Low	Complete

3.2 Non-Functional Requirements

Non-functional requirements define the quality attributes and constraints that the system must satisfy. Table 3.2 presents these requirements.

Table 3.2: Non-Functional Requirements Specification

ID	Category	Requirement Description
NFR-01	Performance	The system shall generate initial design variations within 60 seconds of request submission
NFR-02	Performance	The system shall support concurrent design requests from at least 100 simultaneous users
NFR-03	Performance	Page load times shall not exceed 3 seconds on standard broadband connections
NFR-04	Scalability	The system architecture shall support horizontal scaling to accommodate user growth
NFR-05	Availability	The system shall maintain 99% uptime during standard operating hours
NFR-06	Security	User passwords shall be stored using bcrypt hashing with salt
NFR-07	Security	All API communications shall use HTTPS encryption
NFR-08	Security	The system shall implement rate limiting to prevent abuse
NFR-09	Usability	The interface shall be responsive and functional on devices with screen widths from 320px to 2560px
NFR-10	Usability	The system shall support both English and Arabic languages
NFR-11	Usability	Color contrast ratios shall meet WCAG 2.1 AA accessibility standards
NFR-12	Reliability	The system shall implement automatic retry mechanisms for failed AI generation requests
NFR-13	Maintainability	The codebase shall follow documented coding standards and include thorough unit tests
NFR-14	Compatibility	The system shall function correctly on Chrome, Firefox, Safari, and Edge browsers (latest two versions)

3.3 User Engagement and Attention Design

Modern digital users face unprecedented challenges with attention and focus. Studies indicate that average attention spans have decreased significantly, with users expecting immediate gratification and frequent engagement signals. IntelliRoom addresses this reality by transforming potential friction points into engagement opportunities.

3.3.1 The Attention Economy Challenge

AI-powered design generation requires 30-60 seconds of processing time, a duration that risks user abandonment in an era of instant gratification. Rather than viewing this latency as a limitation, IntelliRoom leverages it as an engagement opportunity through carefully designed interactive experiences.

3.3.2 Gamification During Generation

During AI processing, users participate in meaningful micro-activities:

- **Style Preference Voting:** Users vote on design element pairs (warm vs. cool tones, minimalist vs. detailed), simultaneously personalizing their experience while contributing to community preference data
- **Design Trivia:** Quick interior design tips and cultural insights keep users engaged while subtly educating them about design principles
- **Progress Visualization:** Animated progress indicators with meaningful stage descriptions (“Analyzing room structure...”, “Applying style transformation...”) maintain user connection to the ongoing process
- **Preview Teasers:** Partial glimpses of the generating design create anticipation and dopamine-driven excitement

3.3.3 Dopamine-Driven Positive Engagement

IntelliRoom channels the neurological drive for reward signals toward constructive outcomes:

- **Achievement Unlocks:** Users earn badges for exploring different styles, completing rooms, and sharing designs
- **Streak Rewards:** Daily engagement streaks provide bonus credits, encouraging consistent platform usage
- **Community Recognition:** Featured designs and trending tags provide social validation
- **Progressive Reveals:** Design variations are revealed sequentially rather than simultaneously, creating multiple reward moments from a single generation

3.3.4 Reducing Cognitive Load

The platform minimizes decision fatigue through:

- **Smart Defaults:** AI-recommended starting points based on room type and detected preferences
- **One-Click Actions:** Common operations accessible without navigation complexity
- **Visual-First Interface:** Imagery over text reduces processing effort
- **Contextual Guidance:** Just-in-time tips appear exactly when relevant

This engagement-first design philosophy transforms IntelliRoom from a passive tool into an active, rewarding experience that respects modern attention patterns while guiding users toward successful interior design outcomes.

3.4 Use Case Diagrams

This section presents the primary use cases that define user interactions with the IntelliRoom system. Figure ?? provides a high-level overview of user roles and their access to system capabilities. The following subsections illustrate the five most critical use cases in detail; complete use case documentation is provided in Appendix ??.

3.4.1 Use Case 1: User Registration and Authentication

Figure ?? illustrates the authentication use case, encompassing user registration, login, and password recovery flows.

Use Case Description:

- **Actor:** Unregistered User, Registered User
- **Preconditions:** User has access to a web browser
- **Main Flow:** User navigates to registration page → Enters email and password → System validates input → Account created → Confirmation email sent
- **Postconditions:** User account exists in database with initial credit allocation

3.4.2 Use Case 2: Room Photo Upload and Analysis

Figure ?? presents the room upload use case, showing the flow from image capture to AI analysis.

Use Case Description:




Figure 3.1: System Use Case Overview: Guest and Registered User roles mapped to E-Commerce, Community, and AI Design capabilities.




Figure 3.2: Use Case Diagram: User Authentication System




Figure 3.3: Use Case Diagram: Room Photo Upload and Analysis

- **Actor:** Registered User, AI System
- **Preconditions:** User is logged in; user has available credits
- **Main Flow:** User selects upload option → Chooses image from device or captures with camera → System validates image quality → Florence-2 analyzes room type → SAM2 segments furniture items → Results displayed to user
- **Postconditions:** Room image stored; furniture inventory generated; room ready for design transformation

3.4.3 Use Case 3: AI Design Generation

Figure ?? illustrates the core design generation workflow.




Figure 3.4: Use Case Diagram: AI Design Generation Process

Use Case Description:

- **Actor:** Registered User, ComfyUI Engine

- **Preconditions:** Room photo uploaded and analyzed; style preferences selected; sufficient credits available
- **Main Flow:** User selects style preset (e.g., Arabic Traditional) → Optionally specifies mood → Confirms furniture replacement preferences → System queues generation request → ComfyUI processes multi-stage pipeline → 3-5 design variations generated → Results displayed with before/after comparison
- **Postconditions:** Design variations saved to user history; credits deducted; furniture recommendations generated

3.4.4 Use Case 4: Furniture Marketplace Interaction

Figure ?? shows the marketplace integration use case.




Figure 3.5: Use Case Diagram: Furniture Marketplace Interaction

3.4.5 Use Case 5: Community Design Sharing

Figure ?? presents the community gallery use case.




Figure 3.6: Use Case Diagram: Community Design Sharing

Complete use case diagrams with detailed alternate flows and exception handling are provided in Appendix ??.

3.5 User Stories

User stories capture requirements from the end-user perspective, following the format: “As a [role], I want [feature], so that [benefit].” Table ?? presents prioritized user stories.

Table 3.3: User Stories Specification

ID	Role	I want to...	So that...	Priority
US-01	Guest	Try the platform without registering	I can evaluate before committing	High
US-02	User	Upload a photo of my room	I can receive AI-generated redesign suggestions	High
US-03	User	Select an Arabic Traditional style	My redesign respects my cultural preferences	High
US-04	User	See before/after comparisons	I can evaluate the transformation quality	High
US-05	User	Replace specific furniture items	I can keep pieces I like while changing others	High
US-06	User	Find Egyptian furniture matching my design	I can purchase items locally	High
US-07	User	See prices in Egyptian Pounds	I can budget appropriately	High
US-08	User	Pay with InstaPay or Fawry	I can use familiar payment methods	Medium
US-09	User	Share my design to the community	I can inspire others and receive feedback	Medium
US-10	User	Download high-resolution images	I can share or print my designs	High
US-11	User	Create a 2D floor plan	I can plan furniture arrangement before generation	Medium
US-12	User	Receive a furniture shopping list	I can efficiently purchase recommended items	Medium
US-13	User	Earn credits through engagement	I can access more generations without paying	Low
US-14	Designer	Showcase my portfolio	I can attract potential clients	Low
US-15	Admin	Monitor generation queue	I can ensure system performance	High

3.6 Requirements Traceability

Requirements traceability ensures that each user story maps to functional requirements and use cases. Table ?? provides this mapping for key requirements.

Table 3.4: Requirements Traceability Matrix (Partial)

User Story	Functional Requirements	Use Case
US-02	FR-06, FR-07, FR-08	UC-2
US-03	FR-12, FR-13	UC-3
US-04	FR-17, FR-18	UC-3
US-05	FR-16	UC-3
US-06	FR-22, FR-23	UC-4
US-07	FR-24	UC-4
US-09	FR-28, FR-29	UC-5
US-10	FR-30	UC-3

3.7 Implementation Status Summary

This section provides an overview of the current implementation progress against the defined requirements.

3.7.1 Requirements Completion Overview

Table ?? summarizes the implementation status across requirement categories.

3.7.2 Completed Components

The following major components have been fully implemented:

- **Backend API Infrastructure:** Complete REST API with Express.js, organized into modular routes, controllers, and services
- **User Authentication & Security:** Asymmetric JWT (RS256) authentication with separate public/private key pairs for auth and refresh tokens, httpOnly secure cookie sessions, refresh token rotation tracked in MongoDB, Zod input validation schemas, bcrypt password hashing, and automated monthly credit reset logic for free-tier users
- **Database Layer:** MongoDB schemas for all data entities including Users, Projects, Products, Cart, Orders, and Gallery Posts
- **E-commerce Module:** Full CRUD operations for products, categories, shopping cart, orders, and wishlists with search, pagination, and filtering

Table 3.5: Requirements Completion by Category

Category	Total	Complete	Partial	Planned	Progress
Authentication	3	3	0	0	100%
Profile & Projects	4	4	0	0	100%
Room Upload	3	2	0	1	67%
Room Analysis	3	3	0	0	100%
Style Selection	2	2	0	0	100%
AI Generation	3	3	0	0	100%
Visualization	2	2	0	0	100%
2D/3D Planning	3	3	0	0	100%
Marketplace	3	3	0	0	100%
Credits	3	3	0	0	100%
Community	2	2	0	0	100%
Export	2	2	0	0	100%
Tutorial	1	0	0	1	0%
E-commerce	4	4	0	0	100%
Plugin Market	4	4	0	0	100%
Contact	1	1	0	0	100%
Overall	43	41	0	2	95%

- **Design Project Management:** Project creation, retrieval, update, and deletion with flexible sceneData storage for 2D/3D scenes
- **ComfyUI Integration:** REST API-based service for AI workflow execution, image upload, and result retrieval
- **AI Vector Recommendation Engine:** Two-plane architecture with a local FastAPI orchestrator (ChromaDB vector store, semantic taxonomy snapper) and a remote GPU inference server (RF-DETR for detection, SigLIP2 for embeddings, Qwen3-VL 4B for taxonomy tagging) connected via zrok tunneling
- **Community Features:** Gallery post schema with likes, comments, and tags support
- **Credit System:** User credits, transaction tracking, and subscription plan support (free, monthly, yearly) with automated monthly reset
- **Visualization Module:** Interactive before/after comparison tools and multi-level zoom inspection features for high-resolution AI-generated design variations
- **Export Engine:** High-resolution image downloader and PDF report generator for furniture specification details
- **Docker Configuration:** Development containers for frontend, backend, and MongoDB services

3.7.3 In-Progress Components

The following components are actively being developed:

- **Frontend UI:** React.js structure established; component implementation ongoing across multiple screens

3.7.4 Planned Components

The following components are planned for future implementation:

- Mobile camera capture
- Onboarding tutorials

Chapter 4

System Design and Architecture

This chapter presents the comprehensive system design for the IntelliRoom platform, documenting the technical architecture, database schema, API design, and implementation details. The designs reflect the current implementation state with clear indication of completed and planned components.

4.1 System Architecture Overview

The IntelliRoom platform follows a modern three-tier architecture with containerized deployment, separating concerns between frontend presentation, backend API services, AI processing pipelines, and data storage layers. Figure ?? illustrates the high-level system architecture.




Figure 4.1: High-Level System Architecture Diagram

4.1.1 Architecture Layers

The implemented architecture consists of the following layers:

1. **Presentation Layer:** React.js-based web application with component-based architecture. The frontend implements responsive design principles and is structured with dedicated folders for components, screens, services, context, hooks, and utilities.
2. **API Layer:** Node.js with Express framework serving RESTful endpoints on port 5000. The backend implements:
 - **CORS middleware** for cross-origin request handling
 - **Static file serving** for uploaded images and ComfyUI outputs
 - **Modular route organization** by domain (ecommerce, design2D, upload)
3. **Service Layer:** Domain-specific services organized by functionality:
 - **E-commerce Service:** Product catalog, shopping cart, orders, categories, wishlist management
 - **Design 2D Service:** Project management, floor planning, 3D asset handling
 - **Upload Service:** Image upload with ComfyUI processing integration
 - **Authentication Service:** Complete JWT-based authentication with asymmetric RS256 signing, refresh token rotation, and cookie-based session management
 - **Plugin Marketplace Service:** Community plugins with CRUD operations, ratings, and reviews
 - **Contact Service:** Contact form submission handling
4. **AI Processing Layer:** Logic for managing generative AI workflows and model inference:
 - **ComfyUI Engine:** Node-based engine executing complex generative workflows including Stable Diffusion, ControlNet, and SAM2.
 - **Direct Integration:** The Node.js backend communicates directly with the ComfyUI API to trigger workflows and retrieve results.
 - **AI Recommendation Service:** A Python-based FastAPI orchestrator running locally, connected to a remote GPU inference server via zrok tunneling. Uses RF-DETR [7] for detection, SigLIP2 [8] for embeddings, Qwen3-VL [10] for taxonomy tagging, and ChromaDB for vector storage and retrieval.
5. **Data Layer:** MongoDB 7 document database with Mongoose ODM:
 - **MongoDB:** Primary database for all structured data
 - **File Storage:** Local storage for uploaded images and generated outputs

4.2 Containerization and Deployment

The IntelliRoom platform utilizes Docker for containerized deployment, enabling consistent development environments and simplified deployment.

4.2.1 Docker Compose Configuration

The system is orchestrated using Docker Compose with the following services:

The Compose configuration defines a development container that runs the React frontend and Node.js backend with live source mounting and port exposure for both services. It depends on a MongoDB 7 database container that is configured with authentication, persistent storage, and a dedicated bridge network. A named volume is used to ensure database persistence across restarts.

4.2.2 AI Recommendation Service

A separate Docker configuration handles the AI recommendation service:

The AI container runs a Python 3.12 runtime with FastAPI and Uvicorn serving the local orchestrator on port 7860. Dependencies are managed via `pyproject.toml` with `uv` as the package manager. The service provides three endpoints: `/health` for status monitoring, `/search` for room-photo-to-recommendation queries, and `/ingest/single` for incremental catalog ingestion. ChromaDB data is persisted via a local volume mount. The remote GPU inference plane is accessed through a zrok public tunnel, configured via the `REMOTE_GPU_BASE_URL` environment variable.

4.3 Backend API Architecture

The backend follows a modular MVC-like architecture with clear separation between routes, controllers, models, and services.

4.3.1 Directory Structure

The backend source code is organized as follows:

The backend is organized into configuration, controllers, middleware, models, routes, services, and a server entry point. Controllers are grouped by domain (Design 2D, E-commerce, Upload, Authentication, Plugins, and Contact), while models are split into design, commerce, user/community, credit/object, plugin, and contact collections. Service utilities handle integrations such as the ComfyUI connector.

4.3.2 API Endpoints

Table ?? documents the implemented REST API endpoints.

Table 4.1: Implemented API Endpoints

Method	Endpoint	Description
E-commerce Routes (/api/ecommerce)		
GET	/products	Get products with search, pagination, sorting
GET	/products/:id	Get single product by ID
POST	/products	Create new product
PUT	/products/:id	Update product
DELETE	/products/:id	Delete product
GET	/categories	Get all categories
POST	/categories	Create category
GET	/cart	Get user's shopping cart
POST	/cart	Create new cart
PUT	/cart/:id	Update cart items
DELETE	/cart/:id	Delete cart
GET	/wishlist	Get user's wishlist
POST	/wishlist	Add item to wishlist
DELETE	/wishlist/:id	Remove from wishlist
GET	/order	Get user's orders
POST	/order	Create new order
Design 2D/3D Workspace Routes (/api/design2D3D)		
GET	/	Get all user projects
GET	/:id	Get project by ID
POST	/	Create new project
PUT	/:id	Update project
DELETE	/:id	Delete project
GET	/assets	Get available 3D assets by category
GET	/assets/categories	Get available 3D asset categories
GET	/textures	Get available planner textures
GET	/textures/categories	Get available texture categories
Upload Routes (/api/uploadImage)		
POST	/	Upload image for AI processing
Authentication Routes (/api/auth)		

Method	Endpoint	Description
GET	/me	Verify active user authentication session
POST	/login	User login with JWT (RS256) cookie session
POST	/signup	User registration with bcrypt password hashing
POST	/refreshToken	Refresh expired auth token using refresh cookie
POST	/logout	Invalidate session and clear auth cookies
Plugin Marketplace Routes (/api/plugins)		
GET	/	Get all plugins with author info
GET	/:id	Get plugin details by ID
POST	/	Create new plugin
PUT	/:id	Update plugin (author-only)
DELETE	/:id	Delete plugin (author-only)
Subscription & Billing Routes (/api/subscription)		
GET	/me	Get current active subscription and usage limits
POST	/subscribe	Subscribe to a new pricing plan
POST	/change-plan	Upgrade or downgrade active plan
POST	/fawaterk-webhook	Secure webhook endpoint for Fawaterk integration
Administrative Base Routes (/api/admin)		
POST	/signin	Admin login with dedicated cookie session
POST	/refresh	Refresh admin session token
GET	/me	Get current admin profile details
POST	/logout	Invalidate admin session and clear cookies
POST	/addAdmin	Add new admin account (superadmin only)
GET	/all	List all admin accounts (superadmin only)
PATCH	/unlock/:id	Unlock suspended admin account (superadmin only)

Method	Endpoint	Description
GET	/logs	View system audit and authentication logs
Administrative Dashboard Routes (/api/admin/dashboard)		
GET	/users	List all registered users with subscription tiers
PATCH	/users/:userId/suspend	Suspend active user subscription
PATCH	/users/:userId/change-plan	Manually edit user subscription plan
Pricing Plan Configuration Routes (/api/plans)		
GET	/	List all available pricing plans
POST	/	Create a new pricing plan
PUT	/:id	Edit pricing plan configuration
DELETE	/:id	Delete a pricing plan
User Profile & Community (/api/updateProfile, /api/community)		
PUT	/updateProfile	Update user profile information
GET	/community/posts	Fetch public gallery posts with pagination
POST	/contact	Submit a new contact/support inquiry
Contact Routes (/api/contact)		
POST	/	Submit contact form message

4.3.3 Server Configuration

The main server entry point configures Express with middleware and route mounting:

The server initializes environment variables, connects to MongoDB, enables JSON parsing and CORS, and exposes static directories for uploaded images and generated outputs at `/api/uploads` and `/api/comfyOutputs`. It mounts domain routes for e-commerce, design, and uploads, then starts listening on a configurable port. The ComfyUI integration is instantiated during startup using a configurable host address.

4.3.4 Authentication Architecture

The IntelliRoom backend implements an asymmetric JWT authentication system using the RS256 algorithm. Separate RSA key pairs are used for authentication tokens (15-minute TTL) and refresh tokens (7-day TTL). Tokens are delivered via `httpOnly`, `secure`, `sameSite=none` cookies to prevent XSS and CSRF attacks. Figure ?? illustrates the complete authentication lifecycle including token refresh and rotation.

Key security features of the authentication implementation include:




Figure 4.2: JWT Authentication Flow: RS256 token generation, cookie-based session management, and refresh token rotation sequence.

- **Asymmetric Signing:** Private keys sign tokens; public keys verify them, enabling microservice-compatible verification without sharing secrets.
- **Refresh Token Rotation:** Each refresh invalidates the previous token and issues a new pair, limiting the window of token compromise.
- **Input Validation:** All authentication inputs are validated with Zod schemas before processing.
- **Credit Management:** The login handler integrates a monthly credit reset system that automatically refreshes free-tier user balances every 30 days.

4.4 Database Design

The IntelliRoom platform employs MongoDB as its primary database, selected for its flexible document model, horizontal scalability, and native JSON support that aligns naturally with the JavaScript/Node.js backend stack. The database architecture is organized into two logical domains reflecting the platform’s modular structure: Core & Commerce operations and Design & Community features.

4.4.1 Database Architecture Rationale

Why MongoDB?

Several factors influenced the selection of MongoDB over relational databases:

1. **Schema Flexibility:** The `sceneData` field in Project documents stores complex 2D/3D scene hierarchies with varying structures. MongoDB’s document model accommodates this variability without rigid table schemas.
2. **Developer Productivity:** Mongoose ODM provides TypeScript-like schema validation while maintaining MongoDB’s flexibility, enabling rapid iteration during development.
3. **Scalability:** Built-in sharding and replica sets support horizontal scaling as user base grows beyond initial Egyptian market.
4. **JSON-Native:** Direct JSON storage eliminates ORM impedance mismatch, particularly beneficial for REST APIs returning JSON responses.

Document vs. Relational Design Decisions

The schema employs a hybrid approach combining embedded documents and references:

- **Embedded Documents:** `CartItem` and `OrderItem` are embedded within their parent documents (`Cart`, `Order`) since they have no independent lifecycle and are always accessed with their parent.
- **References:** `User`, `Product`, and `Category` maintain separate collections with `ObjectId` references, enabling independent queries and updates without duplication.
- **Denormalization:** Product name and price are duplicated in `OrderItem` to preserve historical order data even if product details change later.

4.4.2 Entity-Relationship Overview

The database schema is organized into two primary domains: (1) Core User and E-commerce entities, and (2) Design and Community entities. This separation reflects the modular architecture of the platform while maintaining the `User` entity as the central linking point across both domains.

Core & Commerce Schema

Figure ?? presents the core user management and e-commerce entities. The schema centers on the `User` entity, which serves as the authentication and authorization anchor for the entire system.

Key Relationships in Core & Commerce Schema:

- **User → CreditTransaction (1:N):** Each user accumulates a history of credit transactions (spent/earned), enabling audit trails and usage analytics. The separate collection supports efficient querying of transaction history without loading full user documents.
- **User → Cart (1:1):** Each user maintains a single persistent shopping cart. The cart stores a reference to the user and contains an embedded array of `CartItem` subdocuments.
- **Cart → CartItem (1:N, Embedded):** `CartItems` are embedded within the `Cart` document. Each item references a `Product` and stores the quantity. The price is dynamic and fetched from the referenced `Product` at retrieval and checkout times, meaning price changes in the catalog are reflected in the active cart. The price is frozen only upon checkout when items are copied to the `OrderItem` document.
- **User → Order (1:N):** Users can place multiple orders over time. Orders are separate documents to support independent querying, status updates, and historical analysis.
- **Order → OrderItem (1:N, Embedded):** `OrderItems` are embedded within `Order` documents, capturing product name, price, quantity, and image at purchase time. This denormalization ensures order history remains accurate even if products are modified or deleted.




Figure 4.3: Core & Commerce Schema: User, Credits, Cart, Order, Product, Category

- **CartItem/OrderItem → Product (M:1):** Both cart and order items reference Product documents via `ObjectId`. This maintains data integrity while allowing product updates without affecting existing carts/orders (due to denormalized fields).
- **Product → Category (N:1):** Products belong to a single category (e.g., "Living Room Furniture", "Bedroom Furniture"). Categories enable faceted search and filtering in the marketplace.

Design Decisions:

- The `plan` field in User stores subscription tier (free, monthly, yearly) as an enum, enabling feature gating and usage limit enforcement.
- The `credits` field tracks remaining balance, updated atomically via MongoDB `$inc` operator to prevent race conditions.
- Cart persistence across sessions reduces friction in the purchasing funnel, addressing the high cart abandonment rates (90%+) documented in the Egyptian e-commerce market.

Design & Community Schema

Figure ?? illustrates the design workflow and community features. This schema centers on the creative pipeline from Project creation through AI generation to community sharing.

Key Relationships in Design & Community Schema:

- **User → Project (1:N):** Users create multiple design projects. A single unified `Project` model is used where the entire layout state (2D floor plan, 3D furniture coordinates, active materials, and camera perspectives) is stored within a flexible `data` field (Mongoose Mixed type), avoiding separate collections for 2D designs and 3D scenes.
- **Project → GalleryPost (1:N):** Users share designs publicly by posting them to the gallery. A separate `GalleryPost` collection is used, storing a reference to the project ID, user ID, tags, and caption, decoupling the active creative workspace from the community feed.
- **GalleryPost → Comment (Embedded):** Comments are embedded directly as subdocuments inside the parent `GalleryPost` document. This eliminates the need for separate query roundtrips when fetching a post and its comments.
- **User → Reaction (1:N):** Users react to posts (e.g., likes), tracked via the `Reaction` entity, which stores references to the `user_id` and the target `post_id`.
- **Style → Project (1:N):** Styles represent reusable design templates. The style ID is referenced inside projects when applying cultural presets.




Figure 4.4: Design & Community Schema: Project, 2D/3D Design, Style, Assets, Gallery

- **Project → Object (1:N):** Projects reference multiple 3D object assets representing furniture and decor items, referencing their specific mesh URL and manufacturer catalog information.

Design Decisions:

- The `sceneData` field in `Project` uses MongoDB's flexible schema to store varying scene structures (2D floor plans, 3D scenes, mixed layouts) without requiring schema migrations as features evolve.
- The `Style` entity separates design aesthetics from project data, enabling users to apply community-created styles to their designs. The `config_data` field stores AI parameters in JSON format.
- `FloorPlanObject` and `Assets` entities maintain catalogs of reusable 3D assets with metadata (dimensions, categories, constraints) for drag-and-drop functionality in the 2D/3D planner.
- Gallery posts decouple social features from design storage, allowing users to share designs without exposing full project data. The `tags` array enables content discovery and trending topic tracking.

4.4.3 Schema Definitions

User Schema

The User schema stores account information, subscription plans, and credit balance:

Key fields include a unique user identifier, username, email, hashed password, profile data, credit balance, and subscription plan. Timestamps are recorded for account creation and updates. A first-time flag supports onboarding logic, and the wishlist stores product references.

Project Schema (2D/3D Design)

The Project schema stores design projects with flexible `sceneData` for 2D/3D scenes:

Each project references its owner, includes a title and version, and stores full 2D/3D scene data as a flexible structure. Optional cover and thumbnail images support previews, while timestamps capture creation and update times.

Product Schema (E-commerce)

The Product schema supports the furniture marketplace:

Product records include name, description, price, category reference, inventory stock, and a list of image URLs with optional alt text. Review data is captured via rating and review count, while a featured flag supports merchandising. Timestamps are enabled for auditing.

Gallery Post Schema (Community)

The GalleryPost schema enables community sharing:

Community posts store a unique post identifier, user reference, optional design reference, post text, like counts, and tags. Comments are embedded with user identifiers, text, and timestamps. Each post maintains creation and update timestamps for moderation workflows.

Credit Transaction Schema

The CreditTransaction schema tracks credit usage:

Each credit transaction records a unique identifier, amount, and transaction type (spent or earned). A timestamp and descriptive note are stored for each transaction. The collection is stored separately and does not use automatic timestamps.

3D Object Schema

The Object schema stores 3D furniture assets for the floor planner:

3D asset records include a unique identifier, object name and type, preview image URL, model file URL, default color, configuration metadata, and the associated e-commerce provider. Timestamping is enabled for catalog tracking.

Plugin Schema (Community Marketplace)

The Plugin schema enables the community plugin marketplace:

Each plugin record stores a unique identifier, plugin name, description, author reference (linked to User), rating, reviews array, included features list, download count, and pricing in credits. Timestamps track creation and updates, enabling plugin discovery by recency.

Contact Schema

The Contact schema stores contact form submissions:

Contact records capture the sender's name, email, subject line, and message body. Automatic timestamps enable follow-up tracking and support queue management.

Subscription and Billing Schemas

To support premium subscription tiers and usage tracking, four MongoDB schemas are implemented:

- **Plan Schema:** Stores pricing and feature tier configurations. Key fields include the tier name (free, pro, or enterprise), price in EGP, currency, Fawaterk product mappings, and specific tier limits (renderLimit, model3DLimit) along with available features.

- **Subscription Schema:** Manages active customer subscriptions. It references the user and plan, tracking subscription status (active, canceled, expired, trial, or past_due), billing cycle (monthly or yearly), validity dates, and Fawaterk tokens for recurring billing.
- **Payment Schema:** Tracks all invoice transactions. It records the associated user, subscription, payment gateway provider (Fawaterk), gateway order identifiers, transaction status (pending, paid, failed, or refunded), amount, and payment timestamp.
- **Usage Schema:** Implements runtime throttling. It tracks consumed and remaining renders for the active billing cycle.

4.5 ComfyUI Integration Service

The ComfyUI integration represents a key technical component enabling AI-powered image processing. The service implements REST API polling for workflow status monitoring.

4.5.1 Service Architecture

The ComfyUIService class provides the following capabilities:

1. **REST API Polling:** HTTP-based status polling to ComfyUI server for workflow monitoring
2. **Image Upload:** Upload images to ComfyUI input directory for processing
3. **Workflow Execution:** Submit and track node-based AI workflows
4. **Completion Polling:** Periodic status checks until workflow history is available
5. **Result Retrieval:** Download processed images from ComfyUI output
6. **Zrok Support:** Compatibility with zrok tunneling for remote ComfyUI access

4.5.2 ComfyUI Service Implementation

The service uses HTTP polling to track workflow execution status:

The service polls the ComfyUI server's history endpoint at regular intervals to check workflow completion status. It resolves a job once the workflow history contains outputs. A five-minute timeout protects the system from stalled executions. For remote access via zrok, the client adds a custom header to bypass the interstitial screen.

4.5.3 Image Upload and Workflow Execution

The service provides methods for uploading images and running workflows:

Image uploads are sent to the ComfyUI upload endpoint with optional subfolder placement. Workflow execution is requested through the prompt endpoint, after which the service polls history until outputs are available. A five-minute timeout protects the system from stalled executions.

4.5.4 Sample Workflow Definition

The repository includes ComfyUI workflow definitions that are imported and executed by the ComfyUI engine. These workflow files serve as the reference implementations for generation tasks and can be executed directly in ComfyUI.

4.6 AI Vector Recommendation Engine

The AI recommendation engine enables users to upload a room photograph and receive visually similar furniture recommendations from a curated catalog. The system employs a two-plane architecture that separates resource-intensive GPU inference from lightweight control logic.

4.6.1 Two-Plane Architecture

The recommendation system is split into two cooperating services, as illustrated in Figure ??.

1. **Local Control Plane** (`intelliroom_local_orchestrator.py`): A FastAPI application running on the host machine that orchestrates the end-to-end pipeline. It manages the ChromaDB vector database, applies taxonomy normalization via a semantic snapper, builds metadata query filters, and exposes the public API endpoints (`/health`, `/search`, `/ingest/single`).
2. **Remote Inference Plane** (GPU Server): A FastAPI server running on a dedicated GPU-accelerated instance, exposed to the local orchestrator through a **zrok** public tunnel. This plane hosts three inference models:
 - **RF-DETR**: Roboflow Detection Transformer for furniture bounding box extraction and crop generation.
 - **SigLIP2 so400m**: Vision encoder producing normalized 1152-dimensional embedding vectors for visual similarity search.
 - **Qwen3-VL 4B Instruct**: Vision-language model that generates structured taxonomy metadata (color, material, style, room fit, complementary styles) and natural-language descriptions from furniture crop images.




Figure 4.5: Two-Plane Architecture: Local control plane (FastAPI orchestrator, ChromaDB, Semantic Snapper) communicates with the remote GPU inference plane (RF-DETR, SigLIP2, Qwen3-VL) via a zrok public tunnel.

The local orchestrator communicates with the remote plane via HTTP through the zrok tunnel. A custom header (`skip_zrok_interstitial`) bypasses the zrok landing page for programmatic access.

4.6.2 VLM Model Selection: Qwen3-VL Ablation Study

To select the optimal vision-language model for taxonomy tagging, an ablation study was conducted across three Qwen3-VL variants on representative furniture images. The evaluation criteria were: (1) JSON format compliance (the model must output valid JSON without markdown fences or extraneous text), (2) semantic accuracy of predicted attributes, and (3) inference latency on the GPU inference server.

Table 4.2: Qwen3-VL Model Comparison for Furniture Taxonomy Tagging

Model	Latency	JSON Valid	Notes
Qwen3-VL 2B	9s	× Failed	Prepends <code>json</code> markdown fence; parse fails at char 0
Qwen3-VL 4B	18s	✓ Pass	Clean JSON output; accurate color/material/style; good complementary tags
Qwen3-VL 8B	23.5s	✓ Pass	Slightly richer complementary tags; 31% slower than 4B for marginal gain

The 2B variant consistently failed JSON parsing due to prepending a `json` markdown fence before the JSON body, violating the structured output constraint. The 8B variant produced marginally richer complementary search tags (e.g., “crystal chandelier” vs. “gold floor lamp”) but at a 31% latency penalty over the 4B model. The **4B Instruct** variant was selected as the production model, offering the optimal balance of format compliance, semantic accuracy, and inference speed.

4.6.3 Offline Ingestion Pipeline

The ingestion pipeline processes catalog images into searchable vector entries. Figure ?? illustrates the complete ingestion flow.

1. **Health Check:** The orchestrator verifies the remote GPU backend is reachable via `/health`.
2. **Catalog Scan:** All JPEG/PNG images in the `furniture_assets/` directory are enumerated.
3. **Detection:** Each image is sent to the remote `/detect` endpoint. RF-DETR extracts bounding boxes for detected furniture items with confidence scores.




Figure 4.6: Offline Ingestion Pipeline: Catalog images are scanned, furniture crops are detected, parallel embedding and tagging are performed, metadata is normalized via the semantic snapper, and results are upserted into ChromaDB.

4. **Crop Processing:** For each detected crop, two parallel remote calls are made:
 - `/embed_b64`: SigLIP2 computes a normalized 1152-dimensional embedding vector.
 - `/tag_b64`: Qwen3-VL generates taxonomy metadata (color, material, style, room_fit, complementary_styles, description).
5. **Taxonomy Snapping:** Raw tags are normalized to the allowed taxonomy using a two-tier strategy:
 - **Tier 1 — Color Alias Map:** Common color synonyms (e.g., “Cream” → “Beige”, “Charcoal” → “Gray”) are resolved via a deterministic alias lookup.
 - **Tier 2 — Character N-Gram Cosine Similarity:** Remaining unmatched tags are compared against the allowed taxonomy using character n-gram vectorization and cosine similarity, selecting the closest canonical label.
6. **Metadata Flattening:** List-valued metadata fields (`room_fit`, `complementary_styles`) are serialized as comma-separated strings to comply with ChromaDB’s scalar metadata constraints.
7. **Upsert:** The embedding vector and flattened metadata are upserted into the ChromaDB `furniture` collection (cosine similarity space).

4.6.4 Online Search Pipeline

When a user uploads a room photograph, the search pipeline executes the sequence illustrated in Figure ??.

1. **Detection:** The room photo is sent to `/detect`. RF-DETR returns all detected furniture crops with labels, confidence scores, and bounding boxes.
2. **Crop Selection:** The orchestrator selects the highest-confidence crop for recommendation matching.
3. **Parallel Inference:** The selected crop is sent simultaneously to:
 - `/embed_b64` for embedding extraction (SigLIP2).
 - `/tag_b64` for taxonomy tagging (Qwen3-VL).
4. **Taxonomy Snapping:** Raw tags are snapped to canonical values using the same two-tier snapper.
5. **Filtered Vector Query:** A ChromaDB query is constructed with:




Figure 4.7: Online Search Pipeline: User uploads a room photo, the system detects furniture, performs parallel embedding and tagging, queries ChromaDB with metadata filters, and falls back to pure vector search if filtered results are empty.

- The embedding vector as the similarity anchor.
 - A `where` filter on `color`, `material`, and `style` metadata fields.
6. **Fallback:** If the filtered query returns no results (due to narrow metadata constraints), the system automatically retries with a pure vector search (no metadata filter), ensuring the user always receives recommendations.
 7. **Scoring:** Results are ranked by similarity score computed as $\text{score} = 1 - \text{cosine_distance}$ and returned with full metadata.

4.6.5 ChromaDB Data Model

Each ingested furniture item is stored in the ChromaDB `furniture` collection with the following schema:

- **id:** Unique crop identifier (format: `source_image_cropN`)
- **embedding:** 1152-dimensional normalized float vector (SigLIP2)
- **Scalar metadata:** `color`, `material`, `style`, `label`, `bbox` (JSON string), `source_image`, `description`
- **CSV metadata:** `room_fit` and `complementary_styles` stored as comma-separated strings

The `room_fit` and `complementary_styles` fields are persisted and returned in search results but are not yet used for retrieval ranking, providing a foundation for future multi-signal ranking.

4.6.6 AI Service Dependencies

The local orchestrator dependencies are managed via `pyproject.toml` with `uv` as the package manager:

- **API Layer:** FastAPI, Uvicorn, python-multipart
- **Vector Database:** ChromaDB (cosine similarity space)
- **HTTP Client:** HTTPX (async communication with remote GPU backend)
- **Data Processing:** NumPy, Pillow

The remote GPU backend loads PyTorch, HuggingFace Transformers, and model-specific processors (RF-DETR, SigLIP2, Qwen3-VL) at server startup. Models are cached in persistent storage to avoid re-downloading across sessions.

4.6.7 Single-View 3D Object Reconstruction Pipeline

To bridge the gap between 2D inspiration and 3D spatial planning, IntelliRoom incorporates a dedicated pipeline for Single-View 3D Reconstruction. This module allows users to upload a single photograph of a furniture piece (e.g., a heritage chair or a specific marketplace find) and converts it into a textured 3D mesh (.GLB format) suitable for AR visualization or CAD integration.

The pipeline utilizes the Hunyuan3D-2.0 architecture, orchestrated via ComfyUI. Unlike traditional photogrammetry which requires dozens of overlapping images, this generative approach infers 3D geometry from a single viewpoint using a two-stage process.




Figure 4.8: 3D Reconstruction Pipeline Part 1: Initial image loading, segmentation (InSpyReNet), and geometry generation via Hunyuan3D DiT.

Stage 1: Preprocessing and Segmentation

The input image first undergoes rigorous preprocessing to ensure clean geometry generation:

- **Background Removal:** The system employs InSpyReNet, a high-fidelity segmentation model, to remove the background and isolate the subject. This prevents background noise (e.g., floor tiles, shadows) from being interpreted as physical geometry.
- **Normalization:** The segmented image is resized (typically to 512×512) and centered to align with the training distribution of the diffusion model.

Stage 2: Geometry Generation (DiT)

The core 3D shape is generated using a Diffusion Transformer (DiT) model (hunyuan3d-dit-v2-0).

- **Latent Generation:** The model predicts a 3D implicit representation (Signed Distance Field - SDF) in latent space based on the 2D visual conditioning.

- **Marching Cubes Decoding:** A VAE Decoder converts the latent representation into a raw mesh using the Marching Cubes algorithm.
- **Mesh Cleaning:** A post-processing node removes “floaters” (disconnected geometry artifacts) and degenerate faces to simplify the mesh for web rendering.




Figure 4.9: 3D Reconstruction Pipeline Part 2: Texture generation, delighting, and multi-view synthesis.

Stage 3: Texture Synthesis and Delighting

Once the geometry is established, the pipeline applies high-fidelity textures:

- **Delighting:** A Hy3D Delight module analyzes the original image to remove baked-in lighting information (shadows and highlights). This results in an “Albedo” map, allowing the 3D object to react naturally to new lighting environments in the 3D viewer.
- **Multi-View Painting:** The system uses a specialized paint model (hunyuand-paint-v2) to hallucinate the “unseen” sides of the object (e.g., the back of the chair) based on the style of the front view.
- **UV Mapping & Inpainting:** The generated views are projected onto the mesh UVs. A CV2-based inpainting step fills any remaining texture gaps or seams.

Stage 4: Upscaling and Export

To ensure visual fidelity, the texture maps are upscaled (to 2048×2048) before being baked onto the mesh. The final object is exported as a .GLB (glTF binary) file, which is optimized for web-based rendering in the IntelliRoom frontend.




Figure 4.10: 3D Reconstruction Pipeline Part 3: UV mapping, inpainting, upscaling, and final .GLB export.

Summary: From Pixel to Polygon

- **The Challenge:** Users want to see their specific furniture in the room, not generic 3D models.
- **The Solution:** Single-Image-to-3D Pipeline using Hunyuan3D 2.0.
- **The Process:**
 - **Isolate:** AI removes background (InSpyReNet).
 - **Sculpt:** Diffusion Transformer creates the 3D shape.
 - **Paint:** AI hallucinates the “back” of the object and removes static shadows (De-lighting).
- **The Result:** A downloadable 3D .GLB file generated in under 45 seconds.

4.7 ComfyUI Workflow Catalog

This section documents the primary ComfyUI workflows that operationalize IntelliRoom’s AI capabilities. Each workflow is designed as a modular pipeline that can be called independently or routed through the Universal Master Workflow described at the end of this section.

4.7.1 Empty Room Furnishing Workflow

The Empty Room workflow addresses the requirement of furnishing an unfurnished space while preserving architectural structure. It operates as an image-to-image (img2img) pipeline guided by a CR Multi-ControlNet stack to maintain walls, windows, and perspective fidelity.




Figure 4.11: Empty Room Furnishing Workflow: Multi-ControlNet guided img2img pipeline preserving architectural structure while generating furniture.

Workflow Architecture

Input and Conditioning

- **Checkpoint:** RealisticVisionV6 (photorealistic interiors).
- **Input:** User-provided empty room image (Load Image node).
- **Prompting:** Simple positive prompts (e.g., “living room”) and strict negative prompts (e.g., “text”) to reduce artifacts.

Multi-ControlNet Stacking

- **Canny Edge (Strength 1.00):** Locks high-frequency architectural details (windows, fixtures).
- **M-LSD (Strength 0.40):** Preserves straight lines and perspective of walls and ceilings.
- **Depth (Off):** Disabled in the referenced snapshot when edge and line signals are sufficient.

Latent Generation

- **KSampler:** dpmpp_sde with karras scheduler, 12 steps, CFG 2.0.

- **Denoise:** 0.86 to allow furniture hallucination while retaining structure.

Output

- VAE decode followed by Image Comparer for before/after validation.

Integration with Requirements: Supports FR-15 (structural preservation during style transformations).

4.7.2 Object Replacement Workflow

This workflow enables selective regeneration of furniture items without altering the rest of the room, using Florence-2 for detection and SAM2 for precision masks.




Figure 4.12: Object Replacement Workflow: Florence-2 detection and SAM2 segmentation enable precision inpainting of specific furniture items.

Workflow Architecture

Intelligent Detection (Florence-2)

- Text-based detection (e.g., “carpet”) produces bounding boxes for target objects.

Precision Segmentation (SAM2)

- Generates pixel-level masks from Florence-2 boxes to isolate objects cleanly from background.

Inpainting and Generation

- **Inpaint Crop:** Crops to masked region to maximize resolution.
- **ControlNet:** Inpaint preprocessor + Canny edge to preserve surrounding geometry.
- **KSampler:** Denoise 1.00 for complete replacement of target pixels.

4.7.3 Object Replacement with Style Transfer

This workflow extends selective replacement by injecting a reference product image, enabling realistic “Shop This Look” previews.




Figure 4.13: Object Replacement with Style Transfer: IPAdapter injects reference product textures into the inpainting process for “Shop This Look” visualization.

Workflow Architecture

Reference Image Loading

- Secondary reference image (e.g., a specific rug) is provided to guide texture and pattern synthesis.

IPAdapter Injection

- IPAdapter Advanced encodes reference features and injects them into diffusion.
- Weight is set to 1.00 (linear) to closely mimic target patterns.

Structure-Aware Synthesis

- Florence-2 + SAM2 provide masks; ControlNet ensures perspective consistency.

- KSampler denoise at 0.95 balances texture transfer and spatial fidelity.

Integration with Requirements: Supports FR-16 (selective replacement) and FR-22 (marketplace visualization).

4.7.4 Sketch-to-Render Workflow

This workflow converts sketches or CAD wireframes into photorealistic renders, accelerating professional concept visualization.




Figure 4.14: Sketch-to-Render Workflow: Tri-layer ControlNet stack (M-LSD, Linear, DepthAnything) converts CAD wireframes into photorealistic renders.

Workflow Architecture

Input and Resolution Management

- Sketch input via Load Image node.
- SDXL Resolution Calculator determines optimal sampling dimensions.

Tri-Layer ControlNet Stack

- **M-LSD (0.85):** Preserves perspective and straight lines.
- **Linear (1.00):** Enforces drawing strokes and designer intent.
- **DepthAnythingV2 (1.00):** Provides volumetric depth cues.

High-Fidelity Generation

- KSampler denoise 1.00 with 40 steps (dpmpp_sde + karras) for full synthesis.

Integration with Professional Requirements: Supports rapid prototyping for the Professional Designer segment.

4.7.5 Ultimate Upscale Workflow

The upscale workflow generates high-resolution outputs for premium exports (FR-30), using tiled diffusion to avoid VRAM constraints.




Figure 4.15: Ultimate Upscale Workflow: Tiled diffusion with 4x_foolhardy_Remacri upscaler and Tile ControlNet for high-resolution exports.

Workflow Architecture

- **Upscaler:** 4x_foolhardy_Remacri.pth with scale factor 2.0.
- **Tile ControlNet:** SDXL union controlnet at strength 1.00 (SD 1.5 uses tile_resample).
- **Denoise:** 0.37 to add detail without altering composition.
- **Seam Fix:** Internal tiling seam correction (denoise 1.00, width 64).

4.7.6 Universal Master Workflow

The Universal Workflow consolidates all ComfyUI capabilities into a single logic-driven super-graph, enabling dynamic routing between modes without reloading separate JSON files. This “One-Graph-to-Rule-Them-All” approach allows the backend to switch between Empty Room,




Figure 4.16: Universal Master Workflow: Complete node graph showing the Rail-Switch architecture that unifies all generation modes into a single configurable pipeline.

Object Replacement, Style Transfer, and Upscale modes instantly by toggling boolean control signals.

The following subsections detail each functional block of the Universal Workflow, with zoomed views for clarity.

Control Center and Input Logic

The Control Center, located on the far left of the workflow, acts as the pipeline’s “BIOS.” A vertical stack of boolean switches controls the active state of the entire pipeline:




Figure 4.17: Control Center: Boolean switches (CTR1–CTR5) configure pipeline behavior without reloading the workflow.

- **[CTR1] Resize Source:** Toggles input image resizing logic.
- **[CTR2] Masking Mode:** Activates the specific style transfer logic path.

- **[CTR3] IP-Adapter:** Toggles reference image injection for “Shop This Look” functionality.
- **[CTR4] Controller Layer (In-Off):** Enables or disables specific ControlNet layers.
- **[CTR5] Invert Mask:** Toggles inpainting logic (preserve object vs. preserve background).

The Rail-Switch technique routes image data along a “rail” with boolean signals running in parallel. When a Control Center switch is set to `False`, the corresponding processing block is muted or bypassed entirely using “Fast Muter” and “Any Switch” nodes.

Dynamic Segmentation Block

When Florence-2 detection is enabled, the image is routed into the segmentation cluster for text-based object detection and mask generation.




Figure 4.18: Dynamic Segmentation Block: Florence-2 performs text-based detection (e.g., “carpet”) and SAM2 generates pixel-level masks.

- **Florence-2:** Receives text input and produces bounding box coordinates for target objects.
- **SAM2:** Takes bounding boxes and generates high-precision pixel masks.
- **Conditional Bypass:** If the segmentation switch is inactive (e.g., for Empty Room mode), this entire computational block is skipped, saving processing time.

Conditioning Stack

The central “Rail” manages the ControlNet stack, dynamically selecting which preprocessors and control signals are active based on the current mode.




Figure 4.19: Conditioning Stack: Logic gates dynamically swap between Canny (structure), Depth (volume), or M-LSD (lines) based on CTR settings.

Unlike the static Empty Room workflow with fixed ControlNets, the Universal Workflow uses logic gates to:

- Select appropriate preprocessors (Canny Edge, M-LSD, DepthAnything) based on mode.
- Adjust ControlNet strengths dynamically.
- Route structural guidance signals to the generation phase.

Style Injection Block

This block houses the IPAdapter logic for reference-based style transfer, enabling users to visualize specific marketplace products in their rooms.

- **Style Switch:** When the user uploads a reference image (e.g., a specific rug pattern), the switch closes and injects the style into the latent space.
- **No Reference Fallback:** If no reference is provided, the switch opens and the model relies solely on the text prompt.
- **Weight Configuration:** IPAdapter weight is typically set to 1.00 (linear) for close pattern matching.




Figure 4.20: Style Injection Block: IPAdapter encodes reference product features and injects them into the diffusion latent space.

Unified Generation and Output

Regardless of the path taken (segmentation, conditioning, or style injection), all data converges at a unified KSampler for final generation.

- **KSampler:** Central generation node receiving all conditioning signals.
- **Output Routing:** A final switch determines whether the image is saved immediately or sent to the Ultimate SD Upscale loop for high-resolution export.
- **VAE Decode:** Converts latent representation to final RGB image.

Strategic Value

The Universal Workflow satisfies Non-Functional Requirement NFR-01 (Performance) by eliminating workflow reload latency, and NFR-13 (Maintainability) by centralizing model updates. When the team upgrades the underlying checkpoint (e.g., swapping RealisticVisionV6 for a newer model), the improvement instantly propagates to Empty Room, Replacement, Sketch, and Upscale modes simultaneously.

4.7.7 Flux2Klein 4B Generation Workflow

In the second semester, the team introduced Flux2Klein 4B as a new generative backbone for high-quality image synthesis. This model was specifically selected for its **Apache 2.0 license**, which permits unrestricted commercial use, redistribution, and modification — a critical requirement for IntelliRoom’s planned commercial deployment.




Figure 4.21: Unified Generation and Output: Single KSampler path with optional routing to Ultimate SD Upscale for 4K refinement.

Architectural Advancement over Stable Diffusion

The earlier ComfyUI workflows (Sections ??–??) rely on Stable Diffusion 1.5/SDXL, which uses a **UNet** architecture with iterative denoising. Flux2Klein 4B represents a generational leap by adopting a **Transformer-based flow matching** architecture (derived from the Flux model family). Key architectural differences include:

- **Flow Matching vs. Denoising:** Instead of predicting noise at each timestep, Flux2Klein learns a direct velocity field that transports samples from noise to data along optimal paths. This reduces the number of sampling steps required for high-fidelity output.
- **DiT Backbone:** Replaces the UNet’s convolutional encoder-decoder with a Diffusion Transformer (DiT), enabling better global coherence across large image regions — critical for maintaining architectural consistency in room-scale scenes.
- **Image Edit Mode:** Unlike standard text-to-image generation, the Flux2Klein Image-edit variant accepts a **source image** alongside a text prompt, enabling precise structural preservation during restyling, object replacement, and atmospheric transformations.

Model Selection Rationale

The Flux2Klein family is available in two variants: a 4B-parameter model (Apache 2.0) and a 9B-parameter model (restricted license). While the 9B variant offers marginally higher visual

fidelity, the licensing constraints make it unsuitable for a commercial product. The 4B variant provides:

- **Licensing Freedom:** Apache 2.0 permits commercial use without royalty obligations or usage restrictions, enabling IntelliRoom to deploy the model in production and distribute generated outputs to end users.
- **Reduced VRAM Footprint:** The 4B model requires approximately 12 GB of VRAM in float16 precision, compared to ~ 22 GB for the 9B variant, enabling inference on consumer-grade GPUs and cost-effective cloud instances.
- **Competitive Quality:** Despite the smaller parameter count, the 4B model produces photorealistic interior design outputs that meet the platform’s quality standards.

Workflow Architecture




Figure 4.22: Flux2Klein 4B Image Edit Workflow: ComfyUI node graph detailing the generation pipeline, including text conditioning via the Qwen-3 4B text encoder, the DiT-based KSampler, and VAE decoding for photorealistic interior synthesis.

The Flux2Klein 4B workflow (Figure ??) is integrated into the ComfyUI pipeline and can be invoked through the same backend service infrastructure as the existing workflows. The pipeline accepts a source image and a text prompt, producing an edited output that preserves the structural layout of the input while applying the requested transformations.

Generation Results

To demonstrate the versatility of the Flux2Klein 4B generative backbone, the following figures showcase its capabilities across multiple design tasks. All results below were generated exclusively using the Flux2Klein 4B Image-edit workflow.

Room Staging. The most common use case: transforming an existing or empty room into a fully styled interior design.



(a) Original Room



(b) AI Generated Design

Figure 4.23: Flux2Klein 4B — Room Staging (1): Transforming a basic room into a styled interior.



(a) Original Room



(b) AI Generated Design

Figure 4.24: Flux2Klein 4B — Room Staging (2): Applying stylistic changes and furniture population.

Sketch-to-Render. The model converts hand-drawn architectural sketches into photorealistic renders, providing designers with rapid concept visualization without requiring the tri-layer ControlNet stack used by the older Stable Diffusion workflow.

Reference-Guided Object Replacement. A specific product image is provided as a reference, and the model replaces the target furniture item in the room while preserving the surrounding scene geometry, lighting, and shadows.



(a) Original Room (Empty)



(b) AI Generated Design

Figure 4.25: Flux2Klein 4B — Room Staging (3): Complete staging of an empty architectural space.



(a) Architectural Sketch Input



(b) Flux2Klein Photorealistic Render

Figure 4.26: Flux2Klein 4B — Sketch-to-Render: Converting a hand-drawn architectural sketch into a photorealistic interior render without ControlNet.



(a) Original Room



(b) Reference Product



(c) Generated Result

Figure 4.27: Flux2Klein 4B — Object Replacement: A reference bed product is seamlessly injected into the target room, preserving lighting and perspective.

Monochrome-to-Photorealistic Restyling. The model receives a black-and-white structural blockout and produces a fully detailed, colored, and lit photorealistic interior — demonstrating its ability to “understand” spatial layout from minimal input.



(a) Black and White Structural Blockout



(b) Flux2Klein Colorized Output

Figure 4.28: Flux2Klein 4B — Restyling: Transforming a monochrome structural blockout into a fully colored and detailed photorealistic interior.

Scene Decluttering. The model removes clutter and disorganized items from a room photograph, producing a clean and neatly arranged version of the same space — a high-value capability for real estate staging applications.

Lighting and Atmosphere Control. The model modifies the lighting conditions and atmosphere of a room while maintaining the physical layout of furniture and architecture — a capability that requires understanding of light physics, shadow casting, and material reflectance.



(a) Cluttered Room




(b) Decluttered Output

Figure 4.29: Flux2Klein 4B — Scene Decluttering: Removing clutter and disorganized items from a room photograph while preserving the underlying architecture and furnishings.



(a) Daytime Lighting



(b) Nighttime with Warm Artificial Lighting

Figure 4.30: Flux2Klein 4B — Lighting Control: Transforming a room from natural daylight to a warm nighttime atmosphere with artificial lamp lighting, while preserving furniture placement and room geometry.

4.7.8 SeedVR2 Video-Consistent Upscaling

While the Ultimate SD Upscale node provides adequate high-resolution exports for static imagery, the team integrated **SeedVR2** as an advanced alternative upscaling pipeline. SeedVR2 offers superior performance in cleaning up AI-generated artifacts and maintaining texture consistency.

Technical Rationale

Unlike standard tiled upscalers that can introduce “seam” artifacts or hallucinate unwanted details in large, uniform areas (such as blank walls or carpets), SeedVR2 employs advanced temporal and spatial consistency mechanisms. This makes it highly resilient against the high-frequency noise typically introduced during the latent diffusion process.

Upscaling Result

Figure ?? compares a standard low-resolution AI generation against its SeedVR2-upscaled counterpart, demonstrating significant enhancements in material texture clarity (e.g., fabric weave, wood grain) and edge definition without compromising the original composition.



(a) Low-Resolution Generation



(b) SeedVR2 4K Upscale

Figure 4.31: SeedVR2 Upscaling Result: The upscaler drastically improves texture fidelity, edge sharpness, and artifact reduction compared to the base generation.

4.8 Frontend Architecture

The frontend application uses React.js with a component-based architecture and React Router for navigation, functioning as a complete Single Page Application (SPA). It implements distinct layouts, role-based access control, and dynamic state management across an extensive suite of user-facing and administrative interfaces.

4.8.1 Directory Structure

The frontend codebase is structurally segregated into reusable components, context providers, and custom hooks. It features modularized page-level screens, separated layouts (e.g., Dashboard Layout vs. Marketplace Layout vs. Admin Layout), API service modules, shared aesthetic tokens, and utility helpers. The application entry point initializes the React application and mounts the comprehensive Routes component.

4.8.2 Implemented Screens

The application provides a robust suite of production-ready screens:

- **Authentication & Onboarding:** Dynamic Login and SignUp modals with JWT integration, a high-conversion Landing Page, and an informative About Page.
- **Core Design Tools:** Centralized User Dashboard, Profile Update interface, Image Upload Workspace for generative AI tasks, interactive 2D/3D Planner, and the Community Gallery.
- **E-Commerce & Billing:** Full marketplace discovery (Special Offers, Product Details, Category Listing, Shop-By-Room), cart and wishlist management, tiered Pricing Plans, and secure Checkout with Payment Success/Failure handlers.
- **Administrative Portal:** A protected layout offering granular control over users, logs, marketplace products, categories, pricing plans, room metadata, and active orders.
- **Plugin Marketplace:** Community marketplace featuring search and filter functionality, category browsing, plugin cards (displaying ratings and pricing), featured/trending sections, and a responsive grid layout.

4.8.3 Split-Domain Routing Architecture

To ensure strict tenant and context segregation, the frontend implements a hybrid split-domain routing architecture. The application dynamically inspects the current hostname (`window.location.host`) and URI pathing to isolate administrative views:

- **Administrative Subdomain:** Subdomains starting with `admin` (e.g., `admin.localhost` or `admin.intelliroom.com`) mount the core administrative dashboard. This dashboard is responsible for user accounts, global audit logs, system configurations, and subscription plan management.
- **E-commerce Path-Based Admin:** Product and store management administrative panels (managing product listings, categories, rooms, and order tracking) are integrated on a path-based layout (`/ecomm/admin`) under the consumer root domain, allowing retailer-administrators to seamlessly navigate between store views and product catalogs.

Access to both portals is strictly enforced client-side by the `ProtectedAdminRoute` wrapper, and validated server-side via role-based token checks. To prevent cross-context session contamination, the backend separates session states into distinct cookie namespaces: `Authentication/Refresh` for consumers and `Admin-Authentication/Admin-Refresh` for administrators.

4.8.4 Component Architecture

The frontend implements a modular component architecture:

- **Common Components:** Reusable UI elements including Header and Footer
- **Plugin Marketplace Components:** `PluginCard`, `SearchBar`, `FilterDropdown`
- **Services:** API service modules for backend communication (e.g., `marketplaceService` for plugin data fetching)

4.8.5 State Management

The application implements React Context for global state management, handling:

- User authentication state
- Shopping cart contents
- Design project state
- Theme preferences

4.9 Technology Stack Summary

Table ?? summarizes the implemented and planned technology stack for IntelliRoom.

Table 4.3: Technology Stack Summary

Layer	Technology	Status	Purpose
Frontend	React.js	Done	Component-based UI framework
Backend	Node.js, Express	Done	REST API server
Authentication	JWT (RS256)	Done	Asymmetric cookie-based auth with refresh rotation
Document Database	MongoDB 7	Done	Primary structured data store
Vector Database	ChromaDB	Done	Furniture embedding storage and cosine retrieval
ODM	Mongoose	Done	Schema validation and queries
AI Engine	ComfyUI	Done	Node-based AI workflow execution
AI Orchestrator	Python, FastAPI	Done	Local recommendation service on port 7860
ML Framework	PyTorch	Done	Deep learning framework (remote GPU backend)
Transformers	HuggingFace	Done	Pre-trained model integration
Containerization	Docker Compose	Done	Development environment
Object Detection	RF-DETR	Done	Real-time furniture crop extraction
Object Detection	Florence-2	Done	Room and furniture analysis (ComfyUI workflows)
Visual Embeddings	SigLIP2 so400m	Done	1152-d normalized vectors for similarity search
VLM Tagging	Qwen3-VL 4B	Done	Structured taxonomy metadata generation
Segmentation	SAM2	Done	Precision image segmentation
Generation	ControlNet	Done	Style-preserving transformation
Generation	Flux2Klein 4B	Done	Apache 2.0 licensed generative model
Inference Tunnel	zrok	Done	Public tunnel for remote GPU access

4.10 User Interface Design

The IntelliRoom user interface implements a modern, responsive design system focused on accessibility and ease of use. This section presents the core screens that facilitate the primary user journey from discovery to design generation.

4.10.1 Landing Page

The landing page (Figure ??) serves as the entry point to the platform, introducing users to IntelliRoom’s capabilities through high-impact visuals and clear calls-to-action. Key features include immediate access to design tools via the “Start Designing Free” button, exploration of the community marketplace, and prominent display of the value proposition.




Figure 4.32: IntelliRoom Landing Page: Featuring hero section with direct access to design generation and community marketplace exploration.

The landing page design emphasizes:

- Clear value proposition: “Create Stunning Designs in Minutes”
- Dual call-to-action: authenticated start vs. guest trial
- Community showcase highlighting marketplace categories
- Footer navigation providing access to products, resources, and company information

4.10.2 User Dashboard

Upon authentication, users are directed to the Dashboard (Figure ??), which serves as the central hub for managing design activities. The dashboard provides an overview of recent projects, credit balance, and quick access to design tools and marketplace resources.




Figure 4.33: User Dashboard: Centralized interface for project management, credit tracking, and style discovery.

The dashboard implements:

- Project status tracking with visual indicators (Draft, Rendering, Completed)
- Credit balance display with direct link to pricing plans
- Curated style library showcasing design aesthetics (Minimalist Zen, Industrial Loft, etc.)
- Quick navigation to marketplace for community-created assets
- Sidebar navigation providing access to plugins, community, and learning resources

4.10.3 Design Workspace

The Design Workspace (Figure ??), referred to as the AI Studio or Room Enhancer, is the core functional interface where users interact with the generative AI pipeline. It implements a split-pane layout to guide users from spatial upload to generation.

Key workspace components and controls include:





Figure 4.34: Design Workspace (AI Studio): Room Enhancer split-screen interface featuring dual-image upload inputs, enhancement prompt field, and dynamic visualization output canvas.

- **Room Image Input:** An upload dropzone for the user’s current spatial photograph (source image).
- **Reference Style Input (Optional):** An optional upload dropzone for inspiration images (style reference images) to guide the visual output.
- **Enhancement Prompt:** A text area for entering descriptive natural language prompts (e.g., detailing textures, lighting preferences, or specific layout styles).
- **Generation Action:** A prominent “Generate Image” button that packages the inputs and triggers the asynchronous ComfyUI backend workflow via the zrok tunnel.
- **Output Canvas:** A dedicated right-hand visualization panel that displays an interactive loading state (“Awaiting your creation”) and renders the finalized high-fidelity room redesign upon completion.


This interface satisfies the user capabilities outlined in FR-06 (upload interface), FR-12 (style selection), and FR-14 (generation trigger).

4.10.4 E-Commerce Marketplace

The E-Commerce Marketplace (Figure ??) enables users to search, filter, and purchase furniture products that match their AI-generated room designs. By integrating spatial redesign with local retail discovery, the marketplace closes the loop between design visualization and product acquisition.




(a) Special Offers Section




(b) Featured Products

Figure 4.35: E-Commerce Marketplace Home: Curated promotional landing sections showcasing special offers and featured items from Egyptian furniture retail partners.



(a) Shop by Aesthetic



(b) Shop by Category

Figure 4.36: Marketplace Navigation and Browsing: Layouts allowing users to discover furniture grouped by design aesthetic themes or product categories.




Figure 4.37: Marketplace Search and Filter Results: Product search results view with a sidebar control panel for sorting, category drill-down, color, material, pricing, and stock availability.

Key marketplace features include:

- **Curated Banners:** Promotes the platform’s core workflow: “Design first. Shop second.”
- **Aesthetic Navigation:** Clusters furniture under design themes (Modern, Minimalist, Handcrafted, Elegant, Contemporary, and Cozy) matching the AI Studio’s generation outputs.
- **Category Drilldown:** Enables browsing by product types (Beds & Bedroom, Seating, Storage & Cabinetry, Lighting, Tables).
- **Advanced Filtering Sidebar:** Provides precise refinement tools based on category, sort order, color, material, price bounds, and stock status.
- **Product Cards:** Displays pricing (denominated in local currencies or USD), ratings, discounts, and visual thumbnails.

This interface satisfies requirements FR-20 (marketplace browsing) and NFR-05 (community engagement).

4.11 Floor Planner Architecture

The 2D/3D Floor Planner is a core feature of IntelliRoom (FR-19), enabling users to design room layouts from scratch or modify existing spaces before applying AI transformations. This section details the technical architecture and user interface of the integrated planning system.

4.11.1 Technical Stack and Architecture

The floor planner implements a multi-layered architecture combining 2D blueprint editing with real-time 3D visualization (Figure ??). The system is built on a modern JavaScript stack designed for performance and maintainability.




Figure 4.38: Floor Planner Technical Stack: Multi-layered architecture with Three.js for 3D rendering, React for UI components, and Redux for state management. The unified app flow connects Core Logic, State Management, 2D Visuals, and 3D Visuals layers.

The architecture consists of four primary layers:

- **Core Logic (ES6+/Webpack):** Business logic layer implementing room constraints, dimension calculations, and collision detection. ES6 modules provide clean separation of concerns while Webpack bundles optimize load performance.
- **State Management (Redux + Immutable.js):** Centralized state container managing scene data, user actions, and undo/redo history. Immutable.js ensures state immutability, preventing unintended mutations and enabling efficient change detection.
- **2D Visuals (react-svg-pan-zoom):** Blueprint editor rendering architectural floor plans as SVG graphics. Provides pan and zoom controls with precise measurement display. SVG format ensures resolution-independent scaling suitable for technical drawings.

- **3D Visuals (Three.js/WebGL):** Real-time 3D renderer providing photorealistic previews. Three.js abstracts WebGL complexity while maintaining high performance. Supports orbit controls for camera manipulation and PBR materials for realistic lighting.

This layered architecture enables the seamless synchronization between 2D blueprint editing and 3D preview that defines the user experience.

4.11.2 PBR Material System

The 3D visualization layer implements Physically-Based Rendering (PBR) to achieve photorealistic material representation (Figure ??). This approach simulates real-world light interaction with surfaces, crucial for accurate design preview.




Figure 4.39: PBR Rendering Pipeline: Brick wall demonstration showing diffuse and normal map application, GLSL shader calculations for shadow and lighting, and OrbitControls for interactive camera manipulation.

Key rendering features include:

- **Material Maps:** Diffuse (color) and normal (surface detail) maps provide realistic texture representation. Normal maps simulate fine geometric detail without additional polygons, improving performance.

- **GLSL Shaders:** Custom WebGL shaders compute lighting and shadow interactions in real-time. The shader pipeline processes vertex transformations and fragment coloring to achieve photorealistic results.
- **Interactive Camera:** OrbitControls enable intuitive 3D navigation through mouse/touch gestures. Users can rotate, pan, and zoom to inspect designs from any angle.
- **Dynamic Lighting:** Shadow calculations respond to scene geometry changes, providing immediate visual feedback when walls or furniture are repositioned.

This rendering system satisfies NFR-04 (visual quality) by providing professional-grade visualization suitable for client presentations and design validation.

4.11.3 2D Floor Planning Interface

The 2D planning interface provides architectural-grade drawing tools for precise room layout design (Figure ??). The system outputs dimensioned floor plans suitable for construction documentation.




Figure 4.40: 2D Floor Plan Output: Architectural top-down view with precise millimeter-level dimensions supporting professional room layout design (FR-19). Grid system ensures accurate placement and measurement.

The 2D editor supports:

- **Dimensional Accuracy:** All measurements displayed in millimeters with real-time updates as elements are resized or repositioned.

- **Grid Snapping:** Configurable grid system ensures alignment and maintains standard architectural dimensions (e.g., 3600mm module).
- **Furniture Placement:** Drag-and-drop furniture symbols with accurate footprints enable space planning before 3D rendering.
- **Export Capability:** Floor plans can be exported as vector graphics for use in external CAD systems or documentation.

4.11.4 3D Wireframe Visualization

The 3D wireframe view (Figure ??) provides spatial understanding of the complete interior structure, bridging the gap between 2D plans and photorealistic renders.




Figure 4.41: 3D Isometric Visualization: Wireframe rendering of complete interior structure with furniture placement. The transparent wireframe enables verification of spatial relationships and room proportions.

Wireframe benefits include:

- **Structural Clarity:** Transparent walls reveal internal layout and furniture positioning without visual occlusion.
- **Performance:** Lightweight rendering allows real-time updates even on lower-end hardware, supporting the free-tier user segment.
- **Technical Validation:** Designers can verify structural accuracy before committing to photorealistic rendering.

4.11.5 Interactive Wall Drawing Tools

The wall drawing interface (Figure ??) provides parametric controls for room construction with real-time dimension feedback.




Figure 4.42: 2D Floor Planner Interface: Wall drawing tools with real-time measurement display and property editing. Left sidebar provides drawing modes (Straight Wall, Arc Wall) and structural elements (Doors & Windows).

Tool features include:

- **Drawing Modes:** Straight wall tool for rectangular rooms, arc wall tool for curved architectural features.
- **Doors & Windows:** Pre-configured openings with standard dimensions (e.g., 900mm doors) can be inserted into walls.
- **Structural Elements:** Walls, floors, and ceilings are organized in collapsible panels for efficient navigation.
- **Property Inspector:** Selected elements display editable properties including coordinates, dimensions, thickness, and materials.

The interface design follows NFR-11 (WCAG accessibility) with keyboard navigation support and clear visual contrast.

4.11.6 3D Camera Navigation Modes

The 3D planner implements two distinct camera modes to support both global structural inspections and immersive design walkthroughs:

- **Orbit Mode (Global Inspection):** Utilizes standard Three.js `OrbitControls` bound to mouse or touch gestures. This allows the user to rotate, pan, and zoom the camera around a central focal point to examine the overall layout, spatial proportions, and material distributions from isometric or perspective angles.
- **Walkthrough Mode (First-Person Perspective):** Integrates a custom first-person camera controller (`Viewer3DFirstPerson`) that uses the browser's Pointer Lock API to lock the mouse cursor and translate mouse drag movement directly to camera orientation (yaw and pitch). Movement is controlled via keyboard input (WASD or Arrow keys) with a velocity-based physics calculation.

To maintain physical realism and prevent disorientation in First-Person Mode, the custom controller applies spatial and kinematic constraints:

- **Height Clamping:** The camera's vertical position (Y -coordinate) is dynamically bound to a constant offset relative to the local floor level (simulating an average human eye-level of 1700 mm). This prevents the camera from falling through the floor geometry or flying through the ceiling.
- **Look Limits:** Pitch rotation is clamped between -85° and $+85^\circ$ to prevent the gimbal lock effect and ensure natural head-movement constraints.
- **Velocity Dampening:** Keyboard controls apply acceleration and deceleration curves to simulate inertia and prevent jarring movements.

4.11.7 3D Snapping System

To address the inherent difficulty of positioning objects in a 3D canvas via 2D mouse inputs, the planner incorporates a custom multi-mode 3D Snapping System (`snap-3d.js`). This utility executes three coordinate adjustment algorithms during object dragging:

1. **Wall Snapping (SNAP_3D_WALL):** Computes the perpendicular distance from the cursor's 3D intersection point on the ground plane to the nearest segment representing a wall center-line. If the distance falls below the snapping threshold (150 mm), the item snaps flush against the wall, applying a spatial offset equal to half the wall thickness plus half

the item's depth. It also automatically aligns the item's Y-rotation perpendicular to the wall's normal vector.

2. **Grid Snapping (SNAP_3D_GRID):** Rounds the object coordinates to the nearest point on a virtual 100 mm module grid, allowing for quick alignment of structural elements.
3. **Item-to-Item Snapping (SNAP_3D_ITEM):** Calculates the 2D bounding boxes (foot-prints) of all items based on their current dimensions and Y-rotations. When a dragged item approaches another item's edge, it snaps to a flush-contact position.

To prevent overlaps and stability problems, the snapping engine includes:

- **Collision Resolution (Pushback):** An iterative overlap solver detects bounding box intersections. If a collision occurs, it applies a pushback vector along the axis of minimum intersection depth. The system resolves overlapping geometry using up to 8 projection iterations.
- **Snapping Hysteresis:** To avoid rapid snapping and release oscillation (jitter), the system uses a dual-threshold hysteresis. The distance required to break a snap connection (250 mm) is defined as significantly larger than the initial activation threshold (150 mm).

4.11.8 Integrated 2D/3D Workflow

The synchronized dual-view interface (Figure ??) represents the core innovation of the floor planner, enabling simultaneous blueprint editing and 3D preview.




Figure 4.43: Integrated 2D/3D Workflow: Synchronized editing with simultaneous 2D blueprint and 3D preview. Changes in either view update the opposite view in real-time. Right panel provides parametric property editing with instant visual feedback.

Workflow capabilities:

- **Real-Time Synchronization:** Redux state management propagates changes bidirectionally. Modifying a wall in 2D view instantly updates the 3D geometry, and vice versa.

- **View Toggle:** Users can switch between 2D and 3D modes or display both simultaneously for comprehensive spatial understanding.
- **Properties Panel:** Right sidebar displays context-sensitive properties for selected elements. Wall properties include position (X1, Y1, X2, Y2), dimensions (Length, Height, Thickness), and material textures.
- **Material Assignment:** Texture selection (e.g., “bricks” in the example) applies PBR materials visible in 3D view. This enables style experimentation before AI generation.

4.11.9 Integration with AI Workflows

The floor planner output integrates with the ComfyUI workflows described in Section ??:

- **Scene Export:** The planner exports 2D floor plans and 3D renders as input images for the Empty Room and Universal workflows.
- **Structural Preservation:** ControlNet workflows leverage the precise wall dimensions and furniture placement from the planner to maintain architectural accuracy during style transformation.
- **Object Replacement:** The planner’s furniture catalog links to marketplace products, enabling the Object Replacement workflow to suggest compatible real-world items.
- **Viewport-Based Contextual Prompting:** In the first-person walkthrough view, the system utilizes WebGL frustum culling and bounding-box raycasting (`getVisibleFrameItemNames`) to dynamically identify all 3D furniture models visible in the user’s immediate field of view. When a generation is triggered, the system extracts the names of these visible objects and automatically constructs a descriptive text prompt (e.g., “a living room containing a sofa and a coffee table”). This contextually-enriched prompt is passed to the ComfyUI workflow to guide style diffusion, while the identified object categories are used to query the MongoDB e-commerce database, returning visually matched real-world Egyptian products for purchase.

This integration satisfies FR-19 (2D floor planner), US-11 (user story: create floor plan), and contributes to the professional design workflow required by the Designer Segment.

4.12 Implementation Status

Table ?? provides a summary of the current implementation status for each major component.

Table 4.4: Implementation Status Summary

Component	Status	Notes
Backend API Structure	Complete	Express routes, controllers, models
MongoDB Schemas	Complete	12+ collections defined
E-commerce CRUD	Complete	Products, cart, orders, categories
Design Project CRUD	Complete	Projects with sceneData storage
ComfyUI Integration	Complete	REST API polling, upload, workflow execution
User Authentication	Complete	JWT RS256, refresh rotation, cookie sessions
Plugin Marketplace	Complete	Full CRUD with ratings, reviews, pricing
Contact Form	Complete	Contact submission handling
Frontend UI	Complete	Fully responsive React SPA with dashboard layout, user profiles, marketplace shop, checkout, and admin dashboard
AI Recommendation Engine	Complete	Two-plane architecture with ChromaDB
RF-DETR Detection	Complete	Furniture crop extraction (remote GPU backend)
SigLIP2 Embeddings	Complete	1152-d visual similarity vectors
Qwen3-VL Tagging	Complete	Taxonomy metadata generation
SAM2/Florence-2	Complete	Integrated in ComfyUI workflows
ControlNet	Complete	Multiple workflow variants
Flux2Klein 4B	Complete	Apache 2.0 licensed generation workflows
Payment Integration	Complete	Integrated with Fawaterk API v2 payment gateway

Chapter 5

Conclusion and Future Work

This chapter summarizes the achievements of the IntelliRoom project and documents the current implementation status as of May 2026. It also outlines expected outcomes upon project completion and presents the roadmap for remaining development phases.

5.1 Project Summary

The IntelliRoom graduation project has made significant progress toward an AI-powered interior design platform tailored to the Egyptian e-commerce market. The project has completed its design phase and achieved substantial implementation milestones.

5.1.1 Problem Identification and Market Analysis

The project team conducted extensive market research, identifying a significant gap in culturally aware AI interior design solutions. Analysis of competitors including COOHOM, Interior AI, RoomGPT, and Spacely AI revealed universal absence of Middle Eastern aesthetic support, presenting a clear market opportunity for IntelliRoom.

Key market findings:

- Egyptian interior design market: USD 940M growing to USD 1.51B at 6% CAGR
- 113.5 million population with 75% internet penetration
- No existing platforms offering direct AI-to-e-commerce integration with Egyptian local furniture retailers
- Professional interior design costs prohibitive for most Egyptian homeowners

5.1.2 Technical Architecture Implementation

A robust three-tier architecture has been designed and implemented:

- **Frontend:** React.js-based web application with component architecture (multiple screens complete)
- **Backend:** Node.js/Express REST API with modular organization (complete)

- **Database:** MongoDB 7 with Mongoose ODM (complete)
- **Authentication:** Asymmetric JWT (RS256) with cookie sessions and refresh token rotation (complete)
- **AI Engine:** ComfyUI integration with REST API polling (complete)
- **AI Recommendation Service:** Python FastAPI orchestrator with ChromaDB, RF-DETR, SigLIP2, and Qwen3-VL 4B via remote GPU backend (complete)
- **Containerization:** Docker Compose development and production environments (complete)

5.1.3 Implementation Achievements

The following components have been fully implemented:

1. Complete Backend API:

- E-commerce routes: Products, categories, cart, orders, wishlist
- Design 2D routes: Projects, floor plans, 3D assets
- Upload routes: Image processing with ComfyUI integration
- Authentication routes: Full JWT (RS256) login, register, refresh, and logout with httpOnly cookie sessions
- Plugin Marketplace routes: Community plugins with CRUD, ratings, reviews
- Contact routes: Contact form submission handling

2. Database Schemas (12+ collections):

- User management with credits and subscription plans
- Design projects with flexible sceneData storage
- E-commerce models (Product, Cart, Order, Category)
- Community features (GalleryPost with comments/likes)
- Plugin marketplace (Plugin with author, ratings, pricing)
- Contact form submissions
- Credit transaction tracking
- 3D object catalog

3. ComfyUI Integration Service:

- REST API polling for workflow status monitoring

- Image upload to ComfyUI input directory
- Workflow execution with progress tracking
- Result retrieval and local storage
- Support for zrok tunneling

4. Frontend Implementation:

- React.js application with React Router navigation
- Plugin Marketplace screen with search, filter, and category browsing
- Reusable components (Header, Footer, PluginCard, SearchBar, FilterDropdown)
- API service integration for backend communication
- Loading and error state handling

5. AI Vector Recommendation Engine:

- Two-plane architecture: local FastAPI orchestrator + remote GPU inference server
- RF-DETR for furniture detection and crop extraction
- SigLIP2 so400m for 1152-dimensional visual embeddings
- Qwen3-VL 4B Instruct for taxonomy tagging (selected via ablation study over 2B/4B/8B variants)
- ChromaDB vector database with cosine similarity search and metadata filtering
- Semantic taxonomy snapper with character n-gram cosine similarity and color alias mapping
- Offline catalog ingestion and online room-photo search pipelines

6. Docker Development Environment:

- Dev container for frontend/backend (ports 3000, 5000)
- MongoDB 7 container (port 27017/27018)
- AI recommendation service container (port 7860)
- Persistent volume for database

5.2 Current Implementation Status

Table ?? summarizes the current implementation status by component.

Table 5.1: Implementation Status by Component

Component	Status	Details
Backend API Structure	100%	Express routes, controllers, middleware
MongoDB Schemas	100%	All data models defined and tested
E-commerce CRUD	100%	Full functionality with search/pagination
Design Project CRUD	100%	Create, read, update, delete operations
ComfyUI Integration	100%	REST API, upload, workflow, download
Docker Configuration	100%	Dev and production containers ready
User Authentication	100%	JWT RS256 with refresh rotation, cookie sessions
Plugin Marketplace	100%	Full CRUD with ratings, reviews, pricing
Contact Form	100%	Contact submission handling
Frontend Structure	70%	React app, routing, multiple screens
AI Recommendation Engine	100%	Two-plane architecture with ChromaDB
RF-DETR Detection	100%	Furniture crop extraction (remote GPU backend)
SigLIP2 Embeddings	100%	1152-d visual similarity vectors
Qwen3-VL Tagging	100%	Taxonomy metadata generation
SAM2 Segmentation	100%	Complete (ComfyUI workflow)
Florence-2 Detection	100%	Complete (ComfyUI workflow)
ControlNet Generation	100%	Complete (multiple workflows)
Flux2Klein 4B	100%	Apache 2.0 licensed generation workflows
Overall Progress	95%	Based on requirements completion

5.2.1 Code Metrics

The current codebase includes:

- **Backend:** 30+ JavaScript files across controllers, models, routes, and services
- **API Endpoints:** 35+ REST endpoints implemented
- **Database Collections:** 14+ MongoDB schemas
- **ComfyUI Service:** 350+ lines handling REST API communication and polling
- **AI Orchestrator:** 1500+ lines of Python (local orchestrator + remote GPU server)
- **Frontend:** React.js application with 15+ components and screens
- **Docker:** 2 compose files, 3 Dockerfiles (frontend, backend, AI service)

5.3 Expected Outcomes

Upon successful completion of the remaining implementation and testing phases, IntelliRoom will deliver:

5.3.1 Functional Platform

A fully operational web-based interior design platform enabling Egyptian users to:

1. Upload room photographs and receive AI-powered analysis
2. Generate high-fidelity redesign variations using Flux2Klein with style preset prompting
3. Discover and link to Egyptian furniture retailers matching generated designs
4. Share designs with a community of users for inspiration and feedback
5. Save and manage multiple design projects
6. Purchase recommended furniture through integrated e-commerce

5.3.2 Technical Innovation

The project demonstrates:

- Successful integration of state-of-the-art AI models (SAM2, Florence-2, RF-DETR, SigLIP2, Qwen3-VL, ControlNet, Flux2Klein 4B) for interior design applications
- Real-time AI workflow execution via ComfyUI REST API integration
- A novel two-plane architecture for GPU inference via zrok tunneling to a dedicated GPU server
- Vector database-powered furniture recommendation using visual similarity and taxonomy-constrained metadata filtering
- Empirical model selection via ablation study (Qwen3-VL 2B/4B/8B comparison)
- Novel Egyptian furniture retrieval framework connecting AI-generated designs to local marketplace

5.3.3 Academic Contribution

The project makes a technical contribution through:

- A novel two-plane GPU inference architecture using zrok tunneling, applicable beyond interior design
- A documented ablation study methodology for selecting VLMs (Qwen3-VL 2B/4B/8B) in taxonomy generation tasks

- An end-to-end reference implementation of ComfyUI API integration within a production Node.js/React stack
- Demonstration of zero-shot visual retrieval (RF-DETR + SigLIP2) for e-commerce product matching

5.4 Remaining Implementation Roadmap

The remaining project phases will follow a structured implementation plan, building upon the completed AI pipeline infrastructure.

5.4.1 Phase 2: Frontend and User Experience (Next Sprint)

Priority Tasks:

- Conduct usability testing on the finalized frontend admin and marketplace tools
- Implement before/after comparison visualization
- Deploy user engagement features (gamification during generation)
- Finalize design generation workflow UI connecting to completed ComfyUI pipelines
- Integrate recommendation engine results into the frontend product display

5.4.2 Phase 3: Marketplace Integration and Scale

Tasks:

- Connect furniture recommendations to live marketplace
- Refine billing system, subscription lifecycle management, and invoice tracking
- Scale out recommendation engine instances to support higher concurrency
- Complete plugin marketplace integration

5.4.3 Phase 4: Testing and Quality Assurance

Comprehensive testing will ensure platform quality:

- **Unit Testing:** Individual component validation
- **Integration Testing:** End-to-end workflow verification
- **Performance Testing:** Load testing with simulated concurrent users

- **User Acceptance Testing:** Feedback collection from Egyptian user focus groups
- **Security Testing:** Vulnerability assessment

5.4.4 Phase 5: Deployment and Documentation

Final phase activities:

- Production environment configuration and deployment
- Performance optimization and scaling preparation
- User documentation and tutorial content creation
- Final project report compilation
- Graduation project presentation preparation

5.5 Future Enhancements

Beyond the graduation project scope, IntelliRoom has potential for significant feature expansion:

5.5.1 Short-Term Enhancements (Post-Graduation)

1. **Mobile Native Applications:** iOS and Android applications providing optimized mobile experience with camera integration
2. **AR Furniture Preview:** Augmented reality feature allowing users to visualize recommended furniture in their actual spaces using device cameras
3. **Professional Tier:** Expanded features for interior designers including client collaboration tools, project management, and portfolio showcasing
4. **3D Room Modeling:** Full 3D scene generation from 2D floor plans with realistic lighting and materials

5.5.2 Medium-Term Enhancements (6-12 months)

1. **Retailer Partnership Program:** Direct integration with Egyptian furniture manufacturers and retailers for real-time inventory and pricing
2. **AI Design Consultation:** Chat-based AI assistant providing personalized design recommendations and answering user questions

3. **Smart Home Integration:** Recommendations for smart home devices compatible with generated designs

5.5.3 Long-Term Vision (1-3 years)

1. **Regional Expansion:** Localized versions for other regional markets with country-specific furniture retailers
2. **B2B Platform:** Enterprise features for real estate developers, hospitality companies, and furniture manufacturers
3. **Design Marketplace:** Platform for professional designers to sell custom style presets and design templates
4. **Video Generation:** AI-generated walkthrough videos of redesigned spaces

5.6 Concluding Remarks

The IntelliRoom project represents a significant step toward democratizing interior design through AI technology, directly connecting spatial redesign to the Egyptian furniture marketplace. The completed backend infrastructure, authentication system, ComfyUI integration, and AI vector recommendation engine provide a comprehensive foundation for the platform's core value proposition.

Key accomplishments:

- Comprehensive market analysis validating the Egyptian e-commerce business opportunity
- Complete backend API with 35+ endpoints including secure JWT authentication
- Full e-commerce functionality ready for marketplace integration
- Operational AI recommendation engine with vector similarity search and taxonomy-constrained metadata filtering
- ComfyUI integration enabling AI workflow execution with multiple Flux2Klein specialized workflows
- Flux2Klein 4B integration under Apache 2.0 license for commercial deployment
- Docker-based development environment for team collaboration

The project is on track to deliver a functional prototype demonstrating the core value proposition: AI-powered interior design with seamless Egyptian furniture marketplace integration. By addressing the specific needs of Egyptian users—including local furniture sourcing, Egyptian Pound pricing, and Fawaterk payment integration—IntelliRoom establishes a practical bridge between AI-powered spatial redesign and e-commerce purchase conversion.

References

- [1] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [2] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10 684–10 695, 2022.
- [3] L. Zhang, A. Rao, and M. Agrawala, “Adding conditional control to text-to-image diffusion models,” *arXiv preprint arXiv:2302.05543*, 2023, Stanford University.
- [4] B. Klein and R. Black, “Flux2klein: Distilling diffusion transformers for rapid high-resolution image synthesis,” *arXiv preprint arXiv:2602.04351*, 2026.
- [5] A. Kirillov et al., “Segment anything,” *arXiv preprint arXiv:2304.02643*, 2023, Meta AI Research.
- [6] B. Xiao et al., “Florence-2: Advancing a unified representation for a variety of vision tasks,” *arXiv preprint arXiv:2311.06242*, 2024, Microsoft Research.
- [7] R. R. Team, “Roboflow detection transformer (rf-detr): Optimized transformer-based object detection for production edge systems,” *Technical Report, Roboflow AI*, 2025.
- [8] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer, “Siglip 2: Multilingual and high-resolution sigmoid language-image pre-training,” *arXiv preprint arXiv:2502.01188*, 2025.
- [9] L. A. Gatys, A. S. Ecker, and M. Bethge, “Image style transfer using convolutional neural networks,” *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2414–2423, 2016.
- [10] A. Yang et al., “Qwen3-vl: Visual-language models with enhanced spatial and structured reasoning,” *arXiv preprint arXiv:2412.01248*, 2024.

Appendix A

Market Analysis and Competitive Positioning

This appendix provides comprehensive market research and competitive intelligence that informed the strategic direction of the IntelliRoom platform. The analysis employs a top-down approach, beginning with global market dynamics, narrowing to the Egyptian target market. The competitive landscape assessment identifies key differentiators and strategic positioning opportunities.

A.1 Research Methodology

The market analysis conducted for IntelliRoom synthesizes data from multiple sources to establish a comprehensive understanding of the opportunity landscape. Quantitative market data was gathered through extensive online research across industry publications, e-commerce platforms, social media analytics, and publicly available market reports from Deep Market Insights, IMARC Group, Euromonitor International, and Grand View Research. Competitive intelligence was developed through systematic platform trials and feature analysis of existing interior design solutions.

Qualitative insights regarding user needs, cultural preferences, and pain points were gathered through informal discussions with friends, relatives, colleagues, and architectural engineering students at Ain Shams University who provided perspectives on Egyptian interior design preferences, furniture purchasing behaviors, and technology adoption patterns.

The analysis framework follows the standard TAM-SAM-SOM model (Total Addressable Market, Serviceable Addressable Market, Serviceable Obtainable Market) to establish realistic market sizing based on available industry data.

A.2 Global Market Opportunity

The global interior design technology market presents substantial growth potential, driven by accelerating digitalization of home improvement processes and increasing consumer comfort with AI-powered design tools. Industry projections indicate the market will expand from USD 18.32 billion in 2024 toward USD 184 billion by 2032, representing a compound annual growth

rate exceeding 30%. This explosive growth reflects fundamental shifts in consumer behavior: the COVID-19 pandemic accelerated remote work adoption, increasing homeowner investment in residential spaces, while simultaneously normalizing digital-first purchasing behaviors.

The technological landscape enabling this growth includes advances in computer vision (particularly image segmentation and object detection), generative AI models capable of photorealistic rendering, and cloud computing infrastructure that makes sophisticated processing accessible to consumer applications. Major technology companies including Google, IKEA, and Wayfair have validated the market through significant investments in AR furniture placement and AI design assistance features.

However, analysis of existing solutions reveals a critical gap: current platforms predominantly serve Western aesthetic preferences, with limited understanding of cultural design traditions beyond European and North American contexts. This represents a strategic opportunity for platforms capable of delivering culturally-aware design intelligence.

A.2.1 Market Size Analysis

Table A.1: Market Opportunity Sizing Framework

Market Level	Value	Timeline	Strategic Focus
TAM (Global)	USD 4.8B	By 2030	Interior design technology sector
SAM (Egypt)	USD 940M	Current	Egyptian interior design and e-commerce sector
SOM (Egypt)	USD 50M	Launch	Initial market for product-market fit

The Total Addressable Market (TAM) of USD 4.8 billion represents the global interior design software and AI-powered design assistance sector. The Serviceable Addressable Market (SAM) of USD 940 million focuses specifically on the Egyptian interior design and furniture market. The Serviceable Obtainable Market (SOM) of USD 50 million represents a realistic penetration target for IntelliRoom’s initial launch phase in Egypt, assuming capture of 5-7% of the addressable Egyptian market within three years.

A.3 Egyptian Market Analysis

Egypt represents the strategic launch market for IntelliRoom, selected based on multiple converging factors: substantial population scale, growing digital penetration, underserved interior design needs, and the project team’s deep cultural understanding and local networks.

A.3.1 Digital Infrastructure and Demographics

The Egyptian digital landscape has matured significantly, creating favorable conditions for technology-driven consumer services. With 113.5 million population, Egypt provides scale advantages, while 75% internet penetration (85.2 million users) ensures adequate digital reach. Mobile connectivity is near-universal at 97% (110 million connections), and critically, 58.5% of users exhibit mobile-first behavior, accessing internet services primarily through smartphones rather than desktop computers. This mobile preference directly influences IntelliRoom’s responsive design strategy and mobile-optimized user experience requirements.

Social media penetration reaches 46% (52.25 million users), providing established channels for viral marketing and community engagement features. Payment infrastructure continues to evolve, with local digital payment providers (Fawry, InstaPay) gaining adoption alongside international card networks.

Table A.2: Egyptian Digital Market Indicators (2025)

Metric	Value
Total Population	113.5 million
Internet Users	85.2 million (75%)
Mobile Connections	110 million (97%)
Social Media Users	52.25 million (46%)
Mobile-First Users	58.5%

A.3.2 Interior Design Market Dynamics

The Egyptian interior design services market demonstrates robust growth trajectory, valued at USD 940 million in 2024 and projected to reach USD 1.51 billion by 2033, representing a 6% compound annual growth rate. This growth is driven by expanding middle class demographics, increasing urbanization rates, and rising homeownership among millennials entering peak household formation years (ages 28-40).

However, traditional interior design services remain economically inaccessible to most Egyptian households. Professional design consultations typically cost 500-2,000 EGP per room, representing 5-20% of average monthly household income. This affordability gap creates substantial unmet demand, with cost being a primary barrier preventing many Egyptian homeowners from accessing professional design assistance.

Furniture e-commerce presents the most dynamic growth segment, expanding at 27.5% CAGR from USD 116 million (2024) to projected USD 267 million by 2028. This acceleration reflects increasing consumer comfort with online furniture purchasing, improved logistics infrastructure in major cities, and emergence of Egyptian e-commerce platforms (Homzmart, Furnex) offering competitive pricing and convenient payment options.

A critical insight emerges from e-commerce research: online furniture retailers report cart abandonment rates exceeding 90%, significantly higher than typical e-commerce averages (70-75%). Industry research indicates that inability to visualize furniture in actual living spaces represents a major purchase barrier for online furniture shoppers. This visualization gap represents IntelliRoom’s core value proposition opportunity.




Figure A.1: Egyptian Home Improvement Market Growth Trajectories (2024-2033)

Table A.3: Egyptian Home Improvement Sector Projections

Sector	2024 Value	Projected	CAGR	Source
Interior Design	USD 940M	USD 1.51B (2033)	6.0%	Deep Market Insights
Home Decor	USD 2.73B	USD 3.98B (2033)	4.28%	IMARC Group
Furniture E-commerce	USD 116M	USD 267M (2028)	27.5%	ECDB
Smart Home Tech	USD 565M	Growing	28.1%	Grand View Research

The convergence of these sectors creates IntelliRoom’s strategic opportunity: users seeking interior design guidance (USD 940M market) increasingly purchase furniture online (USD 116M market growing at 27.5%), but lack visualization tools to bridge the decision gap. Platforms that solve this visualization challenge while respecting Egyptian cultural aesthetics can capture value across both markets.

A.3.3 User Segmentation and Target Personas

Market research identified three primary user segments, each with distinct needs, pain points, and willingness-to-pay profiles:

The **Homeowners & DIY** segment (45%) represents the largest opportunity and initial launch target. These users typically allocate 15,000-50,000 EGP for room renovation projects

Table A.4: Target User Segment Analysis

Segment	Share	Characteristics and Needs
Homeowners & DIY	45%	Age 25-45, middle class, first-time homeowners, mobile-first behavior, seeking affordable design solutions that respect traditional aesthetics
Professional Designers	43%	Interior designers, architects, firms of 2-20 employees, seeking efficiency tools to reduce manual rendering time, cost-sensitive pricing models
Small Businesses	12%	Cafes, restaurants, retail stores, offices; budget-constrained, need quick design turnaround for commercial space planning

but cannot justify 2,000-5,000 EGP for professional design consultations. They exhibit high social media engagement (average 3.2 hours daily on Facebook and Instagram) and actively seek inspiration from home improvement content. However, they struggle to translate Western design ideas to Egyptian contexts, creating frustration and abandoned projects.

Professional Designers (43%) present a high-value segment with different value drivers. These users seek efficiency gains rather than design inspiration: manual 3D rendering consumes 6-12 hours per client project, and clients frequently request multiple revision iterations. AI-powered rapid visualization can reduce this cycle from days to hours, enabling designers to serve more clients with existing staff. Pricing sensitivity remains high due to competitive pressure in the Egyptian design services market.

Small Businesses (12%) represent a smaller but strategically valuable segment. Commercial spaces (cafes, retail stores) undergo more frequent redesigns (every 18-24 months) compared to residential spaces (every 5-7 years), creating recurring revenue potential. These users prioritize speed and cost over design sophistication, making them ideal candidates for template-based design workflows.

A.3.4 Seasonal Demand Patterns and Launch Timing

Understanding renovation seasonality is critical for launch timing and marketing resource allocation. Egyptian home improvement activity exhibits pronounced seasonal fluctuations driven by religious calendar, weather patterns, and cultural practices:

The April peak (demand index 92) corresponds to pre-Ramadan and Eid renovation surge, when Egyptian families traditionally refresh their homes for religious celebrations and family gatherings. This represents the single highest-value marketing window, suggesting platform launch should target February-March to capture planning phase engagement. The September-November period (demand index 75-82) reflects back-to-school momentum and winter prepa-




Figure A.2: Egyptian Home Renovation Seasonal Demand Pattern (Indexed to Peak)

ration, providing a secondary acquisition window.

Summer months (June-August, demand index 55-65) exhibit the lowest activity due to extreme heat discouraging construction work and widespread vacation travel. However, this low-demand period can be strategically utilized for platform development sprints, user research activities, and infrastructure scaling without opportunity cost of foregone user acquisition.

A.4 Regional Expansion Context: MENA Markets

While Egypt serves as the strategic launch market, broader MENA region dynamics inform the long-term platform roadmap. The Middle East furniture and home decor market reached USD 24.8 billion in 2024, with interior design technology adoption accelerating across wealthy Gulf Cooperation Council (GCC) nations.




Figure A.3: Middle East Furniture Market Share Distribution by Country (2024)

Egypt’s 18% market share provides sufficient scale for product-market fit validation while maintaining manageable operational complexity. However, GCC markets (35% combined) represent the primary expansion opportunity post-launch, offering significantly higher purchasing power (UAE GDP per capita: USD 44,000 vs Egypt: USD 4,000) and technology adoption rates (UAE smartphone penetration: 98% vs Egypt: 97% but with far higher average device capabilities).

The linguistic commonalities across regional markets, including Arabic language, traditional design aesthetics, and furniture preferences, enable platform localization learnings from Egypt to transfer efficiently to expansion markets. However, important regional variations exist: GCC markets exhibit greater Western influence in design aesthetics and higher willingness to pay for premium features, while North African markets (Morocco, Tunisia) more closely align with Egyptian price sensitivity and traditional design preferences.

Strategic expansion sequencing should prioritize UAE (Q3 2027 target), Saudi Arabia (Q1 2028), followed by Morocco and Jordan based on IntelliRoom’s cultural design model refinement and marketplace partnership development timelines.

A.5 Technology Adoption Readiness Assessment

Egyptian professional designers and tech-savvy homeowners demonstrate varying adoption rates across design technologies, creating a nuanced picture of market readiness for AI-powered design platforms:




Figure A.4: Professional Designer Technology Adoption Rates in Egypt (2025)

The data reveals important strategic insights: traditional 3D modeling tools (85% adoption) and mobile design apps (73% adoption) achieve mainstream acceptance, demonstrating professional designers’ comfort with digital workflows. However, emerging technologies show

adoption gaps that represent IntelliRoom's market entry opportunity.

AI-powered design assistance (38% adoption) and AR furniture preview (42% adoption) cluster together at early-adopter phase penetration. This positioning is strategically favorable: adoption rates are low enough that IntelliRoom enters without facing entrenched competitive solutions, yet high enough (crossing the 30% adoption threshold identified by Rogers' Diffusion of Innovations theory) to indicate market validation and readiness for mainstream adoption.

Analysis of existing AI design tool adoption patterns reveals barriers IntelliRoom must address: lack of furniture visualization tools integrated with local product purchasing, concerns about output quality compared to manual design work, and pricing uncertainty around cost-benefit tradeoffs. These observations directly inform IntelliRoom's e-commerce integration strategy and freemium pricing.

A.5.1 E-commerce Category Performance Analysis

Analysis of Egyptian furniture e-commerce transaction data reveals consumer purchasing patterns that inform IntelliRoom's marketplace integration strategy:




Figure A.5: Egyptian Furniture E-commerce Category Distribution by Revenue

Home & Living dominates with 49% revenue share, encompassing sofas, beds, wardrobes, and living room furniture, precisely the categories where visualization uncertainty drives cart abandonment. This category also exhibits the highest average order value (2,500-8,000 EGP) compared to Decor Items (200-600 EGP), maximizing affiliate commission potential for IntelliRoom's marketplace partnerships.

The data validates IntelliRoom's strategic focus on residential living spaces rather than commercial or outdoor applications: the combined Home & Living plus Decor Items segments represent 64% of total market value and directly align with the platform's room redesign core functionality.

A.6 Competitive Landscape Analysis

The interior design technology market exhibits competitive intensity with both established players and emerging startups. However, comprehensive competitive analysis reveals that no existing platform combines AI-powered room redesign with direct integration into a local Egyptian furniture marketplace, and this gap represents IntelliRoom’s strategic positioning opportunity.

A.6.1 Primary Competitors and Market Positioning

Table A.5: Competitive Landscape Overview

Platform	Target ment	Seg-	Core Strengths	Strategic Weaknesses
Coohom	Professionals		Arabic UI, 20K+ 3D models, established Egypt presence	Generic cultural features, no local marketplace, Western-focused design library
Planner 5D	DIY (90M users)		Massive user base, accessible interface, mobile-first	Basic AI capabilities, Western-centric aesthetics, limited professional features
Houzz Pro	Professionals		Comprehensive CRM, invoicing, business tools	Expensive (\$65/month), minimal AI design assistance, US-focused
IKEA Place	DIY		Best-in-class AR (98% accuracy), completely free	Limited to IKEA catalog only, no design customization, Swedish aesthetics
IntelliRoom	Egypt		Flux2Klein-powered redesign, gamification, local marketplace, affordable EGP pricing	New market entrant, building brand awareness

Coohom represents the most direct competitive threat due to existing Egyptian market presence and Arabic interface. However, platform trials reveal that despite Arabic UI, the design output remains Western-focused: style recommendations emphasize Scandinavian minimalism, industrial lofts, and mid-century modern aesthetics with minimal understanding of traditional Egyptian furniture layouts or local spatial arrangements (e.g., sink placement near windows). Their 20,000+ model library contains few models representing local furniture styles.

Planner 5D’s 90 million user base demonstrates market validation for accessible 3D design tools. However, their AI capabilities remain rudimentary (basic room type detection without style intelligence), and user reviews consistently cite the lack of localized furniture options as

a pain point among regional users. Their freemium model (free with limitations, \$10/month premium) establishes pricing expectations that IntelliRoom’s strategy must consider.

IKEA Place provides strategic insights into AR furniture preview adoption: their 98% object placement accuracy rate sets the technical benchmark users expect. However, their catalog limitation (IKEA products only) creates frustration for users seeking local furniture options, and IKEA’s Scandinavian design language inherently conflicts with traditional Egyptian aesthetic preferences.

A.6.2 Feature-by-Feature Competitive Analysis

Table A.6: Detailed Feature Comparison Matrix

Feature Category	Coohom	Planner 5D	Houzz Pro	IntelliRoom
AI Design Intelligence	Yes	Limited	Limited	Advanced
3D Visualization Quality	Advanced	Good	Basic	Advanced
Arabic Language UI	Yes	No	No	Yes
Cultural Style Library	Limited	No	No	Deep
Local Marketplace Integration	No	No	No	Yes
Gamification & Community	No	No	No	Yes
Local E-Commerce Integration	No	No	No	Yes
Free Tier Availability	Yes	Yes	No	Yes
Mobile-First Design	Partial	Yes	No	Yes

The feature matrix reveals IntelliRoom’s differentiation strategy: while competitors offer superior features in individual categories (Coohom’s 3D quality, Planner 5D’s mobile experience), none deliver the integrated value proposition of Flux2Klein-powered spatial redesign connected to a local Egyptian furniture marketplace, specifically optimized for Egyptian users.

A.6.3 Strategic Gap Analysis

Market gap analysis identifies four critical opportunity areas where existing solutions fail to meet Egyptian market needs:

The **E-Commerce Integration** gap represents IntelliRoom’s most defensible competitive advantage. Connecting AI-generated room designs directly to a purchasable Egyptian furniture catalog requires: (1) the Flux2Klein + RF-DETR + SigLIP2 + Qwen3-VL + ChromaDB pipeline producing accurate furniture crop embeddings, (2) curated retailer relationships with local pricing agreements, (3) Fawaterk payment integration supporting Fawry and mobile wallets, and (4) a complete order management workflow. These requirements create substantial barriers to replication by international competitors lacking local commercial relationships.

The **User Engagement** gap addresses a universal pain point: AI generation requires 30-60 seconds processing time, during which users passively wait. IntelliRoom’s gamification ap-

Table A.7: Strategic Market Gap Analysis

Gap gory	Cate-	Current Market State	IntelliRoom Solution	Defensibility
E-Commerce Integration		Generic links to inter- national retailers	Flux2Klein gener- ative engine + RF- DETR/SigLIP2/VDB retrieval pipeline con- necting to Egyptian retailer catalog; EGP pricing and Fawaterk checkout	Deep technical pipeline required
User Engage- ment		Passive AI processing (30-60s wait times)	Gamified voting during generation, community-driven model improvement	Novel approach
Local Com- merce		External links to inter- national retailers	Direct Egyptian retailer partnerships, local pay- ment integration	Requires on- ground relation- ships
Professional Accessibility		Expensive tools (\$65+/month) or basic free tools	Powerful professional tier at accessible Eryp- tian pricing	Value positioning

proach transforms this latency into engagement opportunity, allowing users to vote on design elements, participate in style preference surveys, and contribute to community model training, activities that simultaneously improve platform intelligence while reducing perceived wait times.

A.7 Business Model and Revenue Strategy

IntelliRoom adopts a phased monetization strategy aligned with product-market fit validation timelines and graduation project constraints.

A.7.1 Phase 1: Free Beta (Current Graduation Project Scope)

The initial platform launch operates as a **Free Beta** accessible to all Egyptian users without payment requirements. This approach serves multiple strategic objectives: (1) maximizing user acquisition to demonstrate market traction, (2) gathering qualitative feedback to refine cultural design models, (3) validating technical architecture under real usage patterns, and (4) establishing initial Egyptian furniture retailer partnerships without revenue-sharing pressure.

During the beta phase, the platform implements "soft" monetization validation through tracking of key conversion indicators: furniture marketplace click-through rates, premium feature engagement metrics, and user feedback patterns. This data informs post-graduation pric-

ing optimization without requiring actual payment infrastructure during the academic project phase.

A.7.2 Phase 2: Freemium Monetization (Post-Graduation)

Following successful product-market fit validation, IntelliRoom transitions to a freemium subscription model optimized for Egyptian purchasing power and payment preferences:

Table A.8: Planned Subscription Tier Structure (Post-Graduation Phase)

Tier	Pricing	Target Segment	Key Features
Free	Permanent	User Acquisition	3 designs/month, basic cultural styles, watermarked outputs, community gallery access
Premium	~199 EGP/month	Homeowners	Unlimited designs, all cultural styles, 4K resolution exports, no watermarks, AR preview (future)
Professional	~350 EGP/month	Interior Designers	All Premium features + team collaboration, client presentation mode, PDF export, project management tools
Business	Custom	Design Firms	Team accounts (5+ users), usage analytics, white-label options, API access, priority support

Pricing strategy balances three factors: (1) Egyptian median household income (approximately 6,000 EGP/month), (2) competitive benchmarks (Planner 5D charges \$10/month = 500 EGP, Coohom charges \$29/month = 1,450 EGP), and (3) perceived value relative to traditional designer consultations (500-2,000 EGP per room). The Premium tier at 199 EGP represents approximately 3% of median household income, positioning it as accessible for middle-class homeowners undertaking renovation projects.

The freemium model deliberately maintains accessible free tier limitations to enable users to complete typical single-room projects without payment, reducing conversion friction while creating upgrade incentives for users planning multi-room renovations or frequent design experimentation. Premium tiers offer significantly more generous limits and enhanced features.

A.7.3 Marketplace Revenue Model

Beyond subscription revenue, IntelliRoom captures value through affiliate partnerships with Egyptian furniture retailers and manufacturers. When users click "Shop This Look" or add rec-

ommended furniture to shopping carts, IntelliRoom earns commission on completed purchases (negotiated rates: 5-12% depending on product category and order value).

Preliminary partnership discussions with Homzmart (Egypt's largest online furniture retailer) indicate willingness to offer 8% commission on referrals, with higher rates (10-12%) available for exclusive partnership arrangements. Given average furniture order values of 2,500-8,000 EGP, successful conversions generate 200-960 EGP per transaction, comparable to multiple months of Premium subscription revenue from a single marketplace transaction.

The marketplace model creates aligned incentives: retailers gain qualified traffic (users who have already visualized furniture in their spaces), users receive convenient purchasing options with local payment support, and IntelliRoom earns performance-based revenue without inventory or fulfillment risks.

A.8 Risk Assessment and Mitigation Strategies

Comprehensive risk analysis identifies five primary threat categories that could impact IntelliRoom's market success, along with corresponding mitigation strategies:

Competitive Response Risk (High probability, High impact): If IntelliRoom demonstrates market traction, well-funded competitors (particularly Coohom with existing Egyptian presence) may rapidly add Egyptian cultural design features. Mitigation requires building defensible advantages through: (1) exclusive partnerships with Egyptian furniture manufacturers offering special pricing unavailable to international platforms, (2) continuous refinement of cultural design models through local user feedback loops that competitors cannot easily replicate, and (3) community network effects where users generate cultural design content that improves platform intelligence.

AI Quality Issues (Medium probability, High impact): Users expect photorealistic outputs that respect structural constraints (walls, windows remain intact) while applying cultural styles authentically. Poor quality outputs drive immediate abandonment. Mitigation strategies include rigorous testing with diverse room types, partnership with Ain Shams University design faculty for qualitative validation, and transparent "beta" labeling that manages expectations during quality refinement phases.

User Acquisition Costs (Medium probability, High impact): Customer acquisition cost (CAC) in Egyptian digital markets averages 50-150 EGP per user through paid channels, creating unsustainable economics given free beta positioning. Mitigation emphasizes zero-cost acquisition channels: viral social media campaigns showcasing before/after transformations, university partnerships leveraging team members' ASU networks, and referral programs offering credit rewards for successful friend invitations.

Table A.9: Strategic Risk Assessment Matrix

Risk Category	Probability	Impact	Mitigation Strategy
Competitive Re-sponse	High	High	Build defensible cultural expertise moat; unique gamification mechanics; establish exclusive local partnerships that global competitors cannot easily replicate
AI Quality Issues	Medium	High	Partner with Ain Shams University for academic validation; curate high-quality training datasets; implement continuous user feedback loops for model refinement
User Acquisition Costs	Medium	High	Leverage team's marketing expertise; establish university campus partnerships; deploy viral social media strategy; utilize community-driven growth
Infrastructure Costs	Medium	Medium	Employ pay-per-use AI APIs; implement generation limits on free tier; continuously optimize model efficiency to reduce per-request costs
Payment Integration	Low	Medium	Partner with established Egyptian payment providers (Fawry, InstaPay); support local payment methods through marketplace partners

A.9 Success Metrics and Validation Framework

The beta phase validation framework establishes quantitative and qualitative success criteria across three evaluation timeframes:

Table A.10: Beta Phase Validation Milestones and Success Criteria

Timeline Phase	Success Criteria and Key Metrics
3-Month Initial Validation	500-1,000 registered users, positive qualitative feedback on cultural authenticity (NPS >30), establishment of 2-3 Egyptian furniture retailer partnerships, technical infrastructure stability (<5% error rate)
6-Month Product-Market Fit	2,000-5,000 active users, demonstrated user retention (>40% 7-day retention, >20% 30-day retention), feature usage validation (>60% users try AI redesign, >30% explore marketplace), generation quality satisfaction (>70% positive ratings)
12-Month Scale Readiness	10,000-20,000 total users, proven viral growth coefficient (k >0.5), demonstrated premium feature interest through user engagement patterns, marketplace conversion tracking (>5% click-to-purchase rate), technical readiness for monetization phase

Critical Success Indicators: Beyond raw user counts, product-market fit validation emphasizes qualitative metrics indicating cultural resonance: user testimonials explicitly mentioning "Egyptian style" or "matches my culture," designer testimonials validating time savings, and organic social media sharing of generated designs without platform prompting (indicating genuine user satisfaction rather than incentivized behavior).

Failure Indicators and Pivot Triggers: If 6-month metrics fall significantly below targets (<1,000 active users, <20% 7-day retention, <50% design satisfaction), the project roadmap includes defined pivot options: (1) narrow focus to professional designer segment only, (2) expand the AI redesign feature set to validate technical capabilities before adding full marketplace integration, or (3) pivot to B2B model serving furniture retailers directly rather than consumer-facing platform.

A.10 Strategic Recommendations and Conclusions

This comprehensive market analysis validates substantial opportunity for AI-powered interior design platforms in Egypt, representing the optimal launch geography. Key strategic conclusions inform IntelliRoom’s development roadmap:

Conclusion 1: E-Commerce Integration Represents Defensible Differentiation. Competitive analysis confirms no existing platform delivers AI-powered room redesign integrated

with a local Egyptian furniture e-commerce marketplace. This gap creates first-mover advantage opportunity, provided IntelliRoom executes reliable Flux2Klein generative quality and maintains the RF-DETR/SigLIP2/VDB retrieval pipeline with an expanding retailer catalog.

Conclusion 2: Freemium-to-Marketplace Revenue Model Aligns with Egyptian Economics. Direct subscription revenue faces headwinds from Egyptian purchasing power constraints and payment infrastructure limitations. The hybrid model (generous free tier for user acquisition, marketplace affiliate revenue for monetization) better aligns with local market dynamics while deferring payment friction until users demonstrate high purchase intent.

Conclusion 3: Launch Timing Critical for Seasonal Capture. February-March 2026 launch positioning enables capture of April peak renovation demand (Ramadan/Eid preparation), the single highest-value acquisition window representing 15-20% of annual market activity. Delayed launch missing this window would sacrifice 3-4 months of optimal user acquisition conditions.

Conclusion 4: Community and Gamification Enhance Competitive Moat. Network effects created through community design sharing and gamified model training create switching costs and data advantages that pure-technology competitors cannot easily replicate. Prioritizing these features in initial launch (rather than deferring to "future enhancements") strengthens defensibility from day one.

Recommended Strategic Actions: (1) Finalize minimum 2 furniture retailer partnerships pre-launch to enable marketplace functionality from day one, (2) establish Ain Shams University design faculty advisory relationship for ongoing design quality validation, (3) develop viral social media content strategy showcasing before/after room transformations to drive zero-cost user acquisition.

Appendix B

Use Case Documentation

This appendix documents the core use cases for the IntelliRoom AI platform, focusing on the three primary user workflows that define the platform's value proposition: AI-powered design generation, intelligent room analysis, and integrated furniture shopping.

B.1 Overview


IntelliRoom serves three primary user groups: homeowners seeking affordable design solutions, professionals requiring efficient client presentation tools, and furniture retailers connecting with motivated buyers. The platform's core functionality revolves around AI-powered room transformation with integrated marketplace features.

B.2 Use Case Diagram

B.3 Core Use Case Specifications

This section provides detailed specifications for the three most critical use cases that define IntelliRoom's core value proposition.

B.3.1 UC-C04: Upload & Analyze Room

**Figure B.1:** Professional Designer Workflows


**Figure B.2:** Administrator and System Management Workflows

Table B.1: UC-C04: Upload & Analyze Room Photo - Detailed Specification

Attribute	Description
Use Case ID	UC-C04
Name	Upload & Analyze Room Photo
Actors	Consumer/Professional User, AI Analysis System (Florence-2, SAM2)
Preconditions	User authenticated, navigated to design workflow, AI pipeline operational
Basic Flow	1. User uploads image or captures photo 2. System validates and uploads to backend storage 3. Florence-2 analyzes: room type, furniture items, architectural features, lighting 4. SAM2 generates: pixel-level masks (walls/floors/furniture), depth map, replacement regions 5. System stores metadata in MongoDB and displays preview 6. User confirms or requests re-analysis, then proceeds to style selection
Alternative Flows	• A1: Invalid format/low quality - System prompts correction • A2: Detection failure - Manual room type selection offered • A3: Timeout (>30s) - Retry or manual selection
Postconditions	Image and Florence-2 metadata stored; SAM2 masks saved for ControlNet; style selection enabled
Includes	None
Extends	UC-C05 (Choose Style), UC-C06 (Generate AI Design)
Technical Notes	Florence-2 runs on ComfyUI backend or via Replicate API. SAM2 segmentation provides structural guidance for ControlNet to preserve room geometry during style transfer. Analysis typically completes in 5-10 seconds.

B.3.2 UC-C06: Generate AI Design

Table B.2: UC-C06: Generate AI Design - Detailed Specification

Attribute	Description
Use Case ID	UC-C06
Name	Generate AI Design
Actors	Consumer/Professional User, AI Generation System (ComfyUI, Flux2Klein, ControlNet)
Preconditions	User authenticated, room analyzed (UC-C04), style selected (UC-C05), generation quota available, ComfyUI operational
Basic Flow	1. User confirms parameters; system validates quota 2. ComfyUI pipeline executes 4 stages: • ControlNet: Preserves structure using SAM2 masks • Flux2Klein: Applies photorealistic style transformation with structure conditioning • Post-processing: Composes result images and generates furniture recommendations 3. During processing, user votes on design pairs, earns credits 4. System stores design with metadata, presents before/after comparison 5. User can view variations, save, download, or regenerate
Alternative Flows	• A1: Quota exhausted - Upgrade prompt • A2: Timeout (>120s) - Retry or email notification • A3: Low quality - Auto-regeneration
Postconditions	Design stored with unique ID, quota decremented (free tier), added to gallery, voting data sent for training
Includes	UC-C10 (Gamified Voting During Generation), UC-C14 (View Free Tier Limits)
Extends	UC-C15 (Purchase Premium) for 4K resolution, UC-C07 (View Variations) for unlimited options
Technical Notes	ControlNet ensures structural fidelity by using depth/edge maps from SAM2.

B.3.3 UC-M02: Shop This Look

Table B.3: UC-M02: Shop This Look - Detailed Specification

Attribute	Description
Use Case ID	UC-M02
Name	Shop This Look (AI-Powered Furniture Matching)
Actors	Consumer User, Product Matching System (RF-DETR, SigLIP2, Qwen3-VL), Egyptian Furniture Retailers
Preconditions	User has generated design (UC-C06), clicks "Shop This Look", retailer catalog populated
Basic Flow	1. RF-DETR detects furniture items and extracts visual crops 2. SigLIP2 extracts high-dimensional visual embeddings for vector search matching 3. System displays interactive overlay with clickable hotspots 4. Side panel shows 3-5 products per item: name, price (EGP), retailer, similarity, stock 5. User clicks hotspot to view full details: specs, images, reviews, alternatives 6. User adds items to cart (UC-M08); system tracks conversion for analytics
Alternative Flows	• A1: No matches or low confidence (<0.6) - Manual browsing or category suggestions • A2: Out of stock - Notify option and alternatives • A3: User filters - Re-run matching with constraints
Postconditions	Products linked to design, cart updated, analytics logged, affiliate tracking attached
Includes	UC-M01 (Browse Egyptian Furniture Catalog) for product data access
Extends	UC-M05 (Compare Prices), UC-M08 (Add to Cart), UC-M10 (Place Order)
Technical Notes	RF-DETR provides object detection, SigLIP2 enables visual similarity embedding retrieval, and Qwen3-VL generates taxonomy tags for metadata-based filtering.

B.4 Supporting Use Case Summary

The following table summarizes additional use cases that support the platform's complete functionality. Detailed specifications are available in the requirements documentation.

Table B.4: Supporting Use Cases Summary

Use Case ID	Description
Consumer Account & Access	
UC-C01: Register Account	User creates account with email, password, profile information. Email verification required.
UC-C02: Log In / Log Out	User authenticates with credentials to access personalized features. Session management.
UC-C03: Manage Profile	User updates personal information, design preferences, notification settings.
Consumer Design Workflow	
UC-C05: Choose Style/Theme	User selects a design style preset (Modern, Minimalist, Traditional, Mediterranean) before generation.
UC-C07: View Design Variations	User reviews multiple design options with before/after comparison slider.
UC-C08: Save/Favorite Design	User saves generated designs to personal gallery for later access.
UC-C09: Download Design Image	User downloads design in available resolution (1080p free, 4K premium). Watermark on free tier.
Consumer Engagement	
UC-C10: Gamified Voting	User ranks design pairs during AI generation (30-60s), contributes to model improvement.
UC-C11: Earn Credits	User earns credits through voting, sharing, daily streaks. Credits unlock premium features.
UC-C12: Browse Community Gallery	User explores designs shared by other users, filtered by style and room type.
UC-C13: Share to Social Media	User shares designs to Facebook, Instagram, WhatsApp with customizable templates.
Marketplace Shopping	
UC-M01: Browse Furniture Catalog	User explores curated Egyptian furniture catalog organized by category and style.

Use Case ID	Description
UC-M03: Filter by Category/Style	User filters furniture products by category, room type, or style tag.
UC-M05: Compare Prices	User compares pricing for similar items across multiple Egyptian retailers side-by-side.
UC-M08: Add to Cart	User adds products to shopping cart. Cart persists across sessions.
UC-M09: Select Payment Method	User selects credit card, Fawry, or InstaPay mobile payment.
UC-M10: Place Order	User confirms order. Order routed to retailer for fulfillment with tracking.
Professional Designer Tools	
UC-P01: Register Professional Account	Designer creates professional account with portfolio, credentials, business information.
UC-P04: Create Client Project	Designer creates project with client information, requirements, timeline. Multi-room support.
UC-P08: Priority Queue Generation	Professional users get priority AI processing for faster client turnaround.
UC-P10: Batch Process Rooms	Designer processes multiple rooms simultaneously for commercial projects (cafes, offices).
UC-P17: Share Client Preview Link	Designer generates read-only link for client to view designs without account.
UC-P20: Export PDF Report	Designer generates branded PDF with before/after comparisons and shopping lists.

Appendix C

Business Model Canvas

This appendix presents the Business Model Canvas for IntelliRoom, providing a strategic view of the platform's value creation, delivery, and monetization approach.

C.1 Business Model Overview

The IntelliRoom Business Model Canvas captures the key elements of our business strategy for the Egyptian interior design and e-commerce market.

C.1.1 Key Partners

- Egyptian furniture manufacturers (Mobica, Divano, Danube Home)
- Payment gateways (Fawaterk, supporting Fawry, credit cards, and mobile wallets)
- Cloud providers (Vercel, Railway, Supabase)
- AI model communities and APIs (Replicate, RunPod)
- Interior design influencers
- Ain Shams University architectural engineering department

C.1.2 Key Activities

- ComfyUI workflow development and AI pipeline maintenance (Flux2Klein, ControlNet, RF-DETR)
- Platform development and maintenance
- Retailer partnership onboarding and management
- Cultural content curation with architectural partners
- Community management and moderation
- Marketing and user acquisition campaigns

C.1.3 Key Resources

- AI infrastructure (SAM2, ControlNet, Florence-2, Flux2Klein 4B, RF-DETR, SigLIP2, Qwen3-VL, ComfyUI)
- Egyptian furniture catalog database
- Cross-disciplinary development team
- GPU compute resources (pay-per-use APIs)
- University partnership network

C.1.4 Value Propositions

- AI-powered room redesign powered by Flux2Klein for Egyptian e-commerce integration
- Integrated Egyptian furniture marketplace
- Affordable alternative to traditional professional design services
- Full Arabic language support
- Local payment methods (Fawry, InstaPay)
- Community inspiration gallery
- Gamified engagement during AI processing

C.1.5 Customer Relationships

- Self-service platform with intuitive UX
- Tutorial content and onboarding
- WhatsApp Business support
- Community forums and galleries
- Email notifications and engagement
- Gamified loyalty program (credits, streaks)

C.1.6 Customer Segments

- Egyptian homeowners (45% of target market)
- Young professionals (25-34 age group)
- Families renovating homes
- Interior design professionals (43% of target market)
- Small businesses (cafes, offices, clinics) (12% of target market)
- Real estate agents and property developers

C.1.7 Channels

- Web Application (PWA for mobile-first experience)
- Social Media (Instagram, TikTok, Facebook)
- WhatsApp Business for support and marketing
- Egyptian furniture retailer partnerships
- Influencer marketing collaborations
- University and architecture school partnerships

C.1.8 Revenue Streams

Initial Launch Phase - Free Beta:

The IntelliRoom platform will launch as a **completely free beta** for all users during the initial deployment phase (estimated 6-12 months post-graduation). This allows the team to:

- Validate product-market fit with real Egyptian users
- Collect user feedback for feature refinement
- Build brand awareness and user base
- Establish retailer partnerships and marketplace operations
- Train and optimize AI models with real-world usage data

Future Phase - Monetization Strategy:

Following successful beta validation, IntelliRoom will transition to a freemium subscription model:

- **Free Tier:** Limited AI generations per month, basic styles, watermarked downloads (continues indefinitely for user acquisition)
- **Premium Tier:** Significantly more generous generation limits, full style preset library, 4K downloads, no watermarks
- **Professional Tier:** Unlimited generations plus collaboration tools, client presentation features, priority queue, PDF exports
- **Business Tier:** To be determined based on enterprise needs, featuring team accounts, analytics, white-label options
- **Marketplace Revenue:** Affiliate commissions on furniture sales referred to Egyptian retailers (percentage-based, negotiated per partner)
- **Future Opportunities:** Featured retailer placements, B2B API access, white-label partnerships with real estate developers

The pricing strategy emphasizes affordability for the Egyptian market while remaining sustainable. All pricing remains flexible and subject to adjustment based on beta phase learnings and market conditions.

C.2 Detailed Component Analysis

C.2.1 Value Proposition Summary

IntelliRoom's core value centers on democratizing interior design for the Egyptian market through: (1) AI-powered spatial redesign via Flux2Klein and ComfyUI workflows, (2) local furniture marketplace integration with EGP pricing, (3) accessibility for middle-class families, (4) Arabic language and local payment support, and (5) community-driven inspiration.

C.2.2 Customer Segments

C.3 Growth Strategy

C.4 Competitive Advantages

IntelliRoom's sustainable competitive advantages include: (1) **Pipeline Moat** — the Flux2Klein + RF-DETR + SigLIP2 + VDB two-plane architecture is not easily replicable, (2) **Network Effects** from community-driven content and gallery, (3) **Local Partnerships** with Egyptian retailers and universities, and (4) **First-Mover** positioning in AI-powered Egyptian furniture e-commerce.

Table C.1: Target Customer Segments

Segment		Share	Key Needs
Young Professionals (25-35)		Primary	Affordable first-home design, social sharing
Egyptian Families (35-50)		Secondary	Renovation visualization, purchase confidence
Design Professionals		43%	Client presentation, rapid visualization
Small Businesses		12%	Commercial space design on budget

Table C.2: Growth Phase Summary

Phase	Timeline	Key Objectives
Free Beta	Months 1-6	1,000-5,000 users, 2-3 retailer partnerships, product-market fit
Beta Scaling	Months 7-12	10,000-50,000 users, marketplace validation, prepare monetization
Monetization	Months 12-24	Paid tiers, sustainable revenue, broader Egyptian retailer partnerships

C.5 Risk Mitigation

C.5.1 Key Risks and Mitigation Strategies

Table C.3: Risk Mitigation Matrix

Risk	Impact	Mitigation
Competitive Re- sponse	Global players add cultural features	Deepen local partnerships, accelerate innovation, lever- age network effects
User Acquisition	High acquisition costs in competitive market	Leverage team marketing ex- pertise, viral features, referral program
AI Quality Issues	Cultural authenticity failures	Ain Shams partnership vali- dation, user feedback loops, expert curation
Infrastructure Costs	Compute costs during free beta	Pay-per-use model, genera- tion limits, efficiency opti- mization
Currency Volatil- ity	Egyptian Pound depreciation	Flexible EGP-based pricing, gradual localization, USD re- serves

Appendix D

Additional User Interface Mockups

This appendix provides additional user interface mockups beyond the core screens presented in Chapter ???. These designs cover authentication flows and subscription management.

D.1 Authentication Screens

D.1.1 Login Screen

The login screen (Figure ??) provides returning users with secure access to their accounts. The interface emphasizes ease of use with social authentication options and traditional email/password login.




Figure D.1: Login Interface: Secure user authentication with social login integration (Google, Facebook).

Design considerations for the login screen:

- Clean, minimalist design with centered modal layout
- Social login buttons prominently displayed above traditional form
- Clear visual hierarchy with “Welcome Back” header
- Direct link to registration for new users
- Mobile-responsive layout with touch-friendly input fields

D.1.2 Registration Screen

The registration screen (Figure ??) guides new users through account creation with comprehensive input validation and social registration options.


A mockup of a registration screen. At the top, it says "Become an artist!" and "Sign up to start designing". Below this are two buttons: "continue with Google" and "continue with Facebook". The form fields are: "First Name" (with "badr" entered), "Last Name" (with "sayed" entered), "Email" (with "your@email.com" entered), "User Name" (with "badr21" entered), "Password" (with "*****" entered), and "Repeat Password" (with "*****" entered). A large blue "Sign up" button is at the bottom. Below the button, it says "Already have an account? Login".

Figure D.2: Registration Interface: Comprehensive sign-up form collecting user profile data with social registration alternatives.

Key features of the registration screen:

- Structured form collecting first name, last name, email, username, and password
- Password confirmation field with visual feedback
- Social registration options (Google, Facebook) for streamlined onboarding

- “Become an artist!” value proposition encouraging content creation
- Direct link to login page for returning users
- Mobile-responsive centered modal layout

D.2 Pricing and Subscription

D.2.1 Pricing Plans Screen

The pricing page (Figure ??) displays available subscription tiers with clear feature comparison, guiding users in selecting the appropriate plan for their needs.




Figure D.3: Pricing & Plans Interface: Comparison of subscription tiers (Free, Pro) denominated in Egyptian Pounds (EGP) with monthly/annual billing option.

Key features of the pricing interface:

- **Two-tier pricing structure:** Free (EGP 0/month) and Pro (EGP 23/month).
- **Usage limits:** Free plan provides up to 1,000 renders and 1,000 3D models. Pro plan provides unlimited renders and 3D models.
- **Billing interval toggle:** Supports monthly and annual billing options with a highlighted 17% savings badge for annual billing.
- **Clear calls-to-action:** Status indicator for the active subscription (“Current Plan”) and clear purchase action button (“Choose Pro”).

This interface satisfies requirement FR-27 (subscription purchase via Fawaterk).

D.3 Responsive Design Considerations

All interface screens are designed to be responsive across device sizes:

Table D.1: Responsive Breakpoints

Device	Breakpoint	Layout Adaptation
Mobile (Portrait)	320px – 480px	Single column, bottom navigation
Mobile (Landscape)	481px – 767px	Single column, expanded controls
Tablet	768px – 1024px	Two column, sidebar navigation
Desktop	1025px – 1440px	Full layout, sidebar + content
Large Desktop	1441px+	Full layout with max-width container

Key responsive design principles:

- Touch-friendly targets (minimum 44px × 44px)
- Mobile-first CSS approach
- Progressive enhancement for larger screens
- Lazy loading for gallery images
- Optimized image delivery based on viewport

D.4 Accessibility Considerations

All interfaces adhere to WCAG 2.1 AA standards (NFR-11):

- Color contrast ratios minimum 4.5:1 for text
- Keyboard navigation support throughout
- ARIA labels for interactive elements
- Focus indicators for all interactive elements
- Screen reader compatible structure
- Reduced motion options for animations

This appendix provides additional UI mockups. For core interface designs including Dashboard, Design Workspace, Marketplace, and Floor Planner, see Chapter ??.