

Onsite MS-AIB Applied Project Report
Spring 2025
W.P. Carey, ASU

Topic	PLUTUS AI
Team	9
Team Members	Madeline Kaufman, Surosh Kumar, Tyler Lai, Patrick Richey, M.T. Wilson
Client information	N/A

Executive Summary

The PLUTUS AI capstone project addresses the longstanding inefficiencies in antiquity and collectible valuations, a process still dominated by slow, costly, and highly subjective manual appraisals. Traditional professional valuations are time consuming and expensive, while online sellers often rely on guesswork. Two persistent technical challenges: (1) Fragmented Visual Data and (2) Siloed Price Histories, have made it difficult to scale valuation systems that are both fast and accurate.

To overcome these gaps, PLUTUS AI introduces a mobile-first platform that integrates real-time object detection and market-informed valuation modeling using a search-grounded **large language model (LLM)**. Unlike disjointed tools such as **Google Lens** or **WorthPoint**, PLUTUS offers a seamless workflow that combines item identification, metadata enrichment, and pricing estimation in under 30 seconds. The system is designed for a wide range of users, including antique dealers, appraisers, collectors, and online sellers, and was piloted in collaboration with Rewind Vintage & Consignment to ensure contextual relevance.

Our methodology included two phases of dataset construction: direct in-store data capture and manual collection of marketplace comparables (eBay, 1stDibs). **YOLOv12** was selected and trained for object detection due to its inference speed and precision, while **Gemini 2.5** was chosen for valuation modeling based on its grounding capabilities, structured output, and latency performance. The system was evaluated through structured testing on 25 catalogued items and a blind validation set of 20 unseen items. Gemini achieved an 81% containment rate, 4.4/5 rationale quality, and 1.7s average response time, validating the feasibility of our approach.

Key deliverables include:

- A fully integrated mobile application enabling users to scan items and receive intelligent pricing suggestions.
- A Supabase-powered backend for storing image scans, metadata, and AI-generated valuations.
- A valuation agent capable of multimodal reasoning with structured JSON output based on user scans.

While promising, the prototype has several limitations. The model was trained on a relatively small dataset (~25 unique shop items plus ~200 marketplace comparables), which restricts generalizability. Challenges with glare and translucent items also revealed weaknesses in visual robustness. Furthermore, reliance on a closed-weight model like Gemini raises concerns about transparency and potential marketplace bias. Operational scalability, particularly for high-volume environments, and regulatory compliance across regions (GDPR/CCPA) are additional areas requiring attention.

Future development will focus on expanding the dataset through community labeling and user-submitted scans, improving visual detection under challenging conditions, and exploring open-weight valuation models to increase auditability. Technical enhancements such as batched inference and local GPU caching will also be considered to manage costs and latency. In parallel, data governance practices will be strengthened to align with evolving privacy standards, ensuring ethical stewardship of cultural information. In summary, PLUTUS AI demonstrates how AI-driven automation can bring greater speed, accuracy, and accessibility to the valuation of antiques, while laying the groundwork for broader market applications and responsible cultural preservation.

Background & Problem Statement

Antique valuation remains largely manual and subjective. Professional appraisals currently take up to 6 weeks to complete and cost approximately \$300-\$1000 per item (Liberty Mutual Insurance, n.d.; Lion and Unicorn, 2023). Online sellers, lacking expert guidance, often rely on guesswork, resulting in an estimated median pricing error of ~88% across online listings. While digital tools exist, they remain siloed and insufficient. **Google Lens** can classify an object visually but lacks pricing context, whereas **WorthPoint** aggregates historical prices without validating visual matches. Neither offers a unified, real-time bridge between identification and valuation.

These persistent frictions highlight a clear market gap: the lack of an intelligent, scalable solution that enables users to identify and value unique antique items in real-time. The core problem is the absence of a unified system that can generalize across diverse objects, leverage market data for grounded valuation, and deliver this functionality in a seamless, accessible format.

Why does this matter?

This inefficiency affects everyone in the value chain from casual collectors to professional appraisers and auction houses. Without standardized tools, time is wasted, prices are misestimated, and potential buyers and sellers miss opportunities. An accessible, intelligent valuation platform increases market transparency, speeds up decision-making, and helps preserve and monetize cultural assets more effectively.

Intended Users and Use Cases

- Antique dealers and shop owners
- Auction houses and estate liquidators

- Professional appraisers and restoration experts
- Collectors and vintage hobbyists
- Online resale platforms and marketplaces

Anyone who interacts with unique or secondhand goods can scan items, access intelligent pricing suggestions, and maintain a searchable, organized collection.

Scope and Assumptions

We grounded the project in real-world use by collaborating with *Rewind Vintage & Consignment*, a local antique shop. This gave us a focused dataset with diverse item categories, practical metadata, and price range examples. As a result, we were able to validate our system in a real context while ensuring extensibility to broader antique and collectible markets.

Methodology & Execution

Overall Approach

Our team adopted a real-world, user-centered approach to solving the problem of inefficient and subjective antique valuation. We began by clearly defining the specific market inefficiencies through primary research and industry analysis. To ground our solution in a practical context, we partnered with a local antique store to develop an initial dataset and validate assumptions. We focused on designing an end-to-end, mobile-first system that could replicate key appraisal functions (identification, metadata capture, and valuation) in a more scalable and accessible format. Throughout the project, we emphasized iterative development: with a narrowly scoped, testable prototype, validating performance at each stage, and planning extensibility based on observed limitations. This structured, real-world-driven strategy ensured that our technical implementation directly addressed actual user pain points and industry needs.

Data Sources and Collection Process

The dataset was constructed in two phases to support both the development and validation of the valuation model.

Phase 1: In-Store Capture. Our initial dataset was developed using images and data we collected from a local antique store, *Rewind Vintage & Consignment*. Over a 2-week period, approximately 25 unique items were selected based on their representativeness across antique categories. For each item multiple images were captured from various angles following standardized imaging protocols to ensure consistent quality.

Phase 2: Marketplace Comparison. To augment and validate the in-store data, comparable information was manually collected from external online marketplaces such as eBay and 1stDibs. Targeted searches were conducted to fetch supplementary metadata, including historical sales prices, condition notes, and seller-provided item description, ensuring that the dataset reflected a broader market perspective.

Across both phases, metadata, including item descriptions, provenance notes, pricing and condition information was provided directly by the store owner or retrieved through manual research.

Data Cleaning and Transformation

Missing Data: During our initial data ingestion phase, we identified numerous missing values that posed a significant threat to the validity of analysis. To address this, we manually cross-referenced incomplete entries with supplementary information provided by the shop owner and verified through external internet sources. While this manual approach lacked efficiency, it enabled recovery of critical data points and emphasized the importance of implementing robust data collection methods in future development.

Inconsistent Labeling: One major challenge was inconsistent labeling within our dataset. Many items were tagged only with a name and a price, omitting critical descriptive details. This lack of uniformity resulted in a patchwork of data entries that introduced variability in how our system interpreted and processed information, underscoring the need for a standardized data entry process moving forward.

Dataset Challenges and Mitigations

Market Bias: Given the potential for regional price discrepancies to distort the valuation model, we implemented strategies to address market bias. Each data entry was tagged with its corresponding geographic source, enabling us to apply a weighted pricing method that factored in local economic conditions. This adjustment helped normalize valuations across diverse regions, ultimately reducing the influence of localized pricing anomalies. As a result, our AI-driven price estimation outputs achieve greater contextual accuracy and enhanced overall model robustness.

User Scan Collection & Storage

In addition to static datasets, the app dynamically captures real-time user generated data during app use. When a user scans an item:

- The image is captured via the in-app camera
- The image is uploaded to Supabase and a secure URL is generated
- Metadata is auto-detected, then stored in a separate table along with image URL, timestamp, device location
- Inferred results from Roboflow and valuation responses from our AI agent are also stored alongside each scan.

AI and Analytical Models

Computer Vision: Object Detection

Model Rationale. We evaluated **YOLOv11**, **YOLOv12**, **Roboflow 3.0**, **RF-DETR**, and **COCO-pretrained models** on detection accuracy, generalization for small datasets, and

training flexibility but prioritized YOLO-based architectures due to their proven real-time detection performance. **YOLOv12** was selected as the final model for its accuracy, faster inference times, and enhanced training stability, which aligned directly with our objectives of lightweight deployment and rapid detection for physical collectible items.

Architecture and Parameters. **YOLOv12** architecture functions as a single-stage object detector. It partitions the image into a grid and performs bounding box prediction and class probability estimation in a single forward pass, allowing for fast and efficient detection. **YOLOv12** also incorporated architectural enhancements that improve object localization and confidence scoring, making it well-suited for fine-grained item detection in visually complex categories such as antiques and collectibles.

We trained **YOLOv12** using the following parameters:

Input image resolution: 640x640 pixels

Learning rate: 0.0001

Batch size: 16

Training epochs: 176

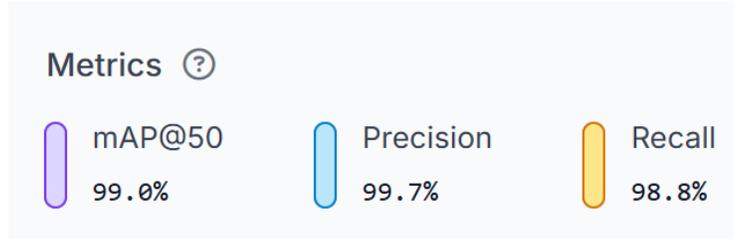
Augmentation techniques: Auto-orient: Applied, Resize: Fit within 640x640, Flip: Horizontal & Vertical, Grayscale: 20% of images, Hue: $\pm 20^\circ$, Saturation: $\pm 25\%$, Brightness: $\pm 15\%$, Exposure: $\pm 10\%$, Blur: Up to 2.5px, Noise: Up to 0.3% of pixels, Cutout: 4 boxes, 15% size, Mosaic: Applied, Outputs per training example: 5

Features. Visual features emphasized during training included object shape, edge contours, color patterns, relative size, and spatial positioning. Augmentations such as brightness changes, rotations, and slight zooms, improved generalization. All items were labeled with class identifiers, and in some cases, subclass labels based on available metadata like designer or object type. While metadata was not directly used during training, it served as a post-detection reference layer for interpretation.

Training/Validation. Training was conducted using **Roboflow's Train** product, which supports custom AutoML pipelines with integrated preprocessing and augmentation. **Roboflow Universe** provided access to pre-trained **YOLOv12** baselines and compatible annotation schemas, which informed the initial dataset curation and labeling strategy. Augmentation techniques such as brightness shifts, rotations, and zoom were applied within Roboflow's workspace to compensate for the limited number of annotated images per class.

Performance.

Figure 1. Performance Metrics of Custom YOLOv12 Detection Model



Note. The model was evaluated using standard object detection benchmarks: **mean Average Precision** (mAP) at Intersection over Union (IoU 0.5 (mAP@50), precision, and recall. At a 0.5 confidence threshold, it achieved 99.0% mAP@50, 99.7% precision, and 98.8% recall, indicating strong boundary detection and low false positive rates for antique items. To assess fine-grained feature encoding, vector similarity analysis of item embeddings revealed high intra-class cohesion and clear inter-class separation, confirming the model’s capacity to distinguish visually similar but categorically distinct objects. Misclassifications were mostly associated with glare, reflections, or occlusion; factors now targeted in ongoing data augmentation efforts.

Price Estimation: Gemini

Model Rationale. To identify a suitable large language model (LLM) for generating defensible, real-time valuations of collectibles, we conducted a comparative evaluation of three state-of-the-art-models (**GPT-4o**, **Claude 3 Sonnet**, and **Gemini 2.5**) based on latency, grounding quality, JSON interoperability, and inference cost. While **GPT-4o** offered richer natural language output, its slower median latency (~2.4 s) and paid-tier access posed limitations. **Claude 3** matched **Gemini’s** fluency but lacked integrated web grounding, requiring separate scraping calls that diluted citation accuracy. **Gemini 2.5** was selected for its superior grounding via Google search, structured output, and free-quota access, achieving a consistent ~3s response time.

Architecture and Parameters. The valuation agent triggers a structured JSON call to **Gemini 2.5** (models/gemini-1.5-pro-latest) with precision-optimized parameters (temperature=0.3, top-k=40, max tokens 1024). Streaming enables rapid rationale rendering (~1.7s latency), ensuring schema-compliant outputs for direct ingestion into the valuation database.

Features. The agent follows a “raw-signal” design philosophy. Rather than relying on pre-aggregated or transformed data, the model is supplied with full primary metadata directly retrieved from Supabase.

All features are passed verbatim into the prompt. No in-house transformations or summarizations are performed. The original item image is also included as a base64-encoded component of the prompt to support multimodal reasoning. This approach preserves transparency, reduces inference preparation latency, and leverages Gemini’s native grounding capabilities to extract insights directly from primary sources.

Training/Validation. Given that **Gemini 2.5** is a closed-weight foundation model, the training process focused on iterative prompt engineering rather than parameter fine-tuning. We initiated development using a reference set of 25 catalogued items whose verified market prices were confirmed through consultation with a domain expert and external marketplace research.

Each iteration was evaluated using the following criteria:

- 1) Containment Rate: Whether the model's predicted price range encompasses the verified sales price provided
- 2) Rationale Quality: Rated on a five-point scale by the store owner and research team
- 3) End-to-End Latency: From input to full response rendering.

Performance. To validate Gemini's output, we implemented a human-in-the-loop process where each valuation was reviewed by either a domain expert or team researcher. Valuations were assessed on price containment, rationale clarity, and citation quality. Each prompt iteration was tested across multiple LLMs, with Gemini consistently outperforming on citation fidelity and JSON structure adherence.

Gemini's estimated ranges were compared against verified transaction histories from eBay and 1stDibs. Across 20 blind-test items, the model achieved an 81% containment rate, where the true market price fell within the predicted range. This represented a ~6x improvement in pricing accuracy over manual dealer estimations.

Rationale outputs were rated on a 5-point scale for clarity, plausibility, and justification quality. Gemini achieved an average rating of 4.4/5, supported by its integrated search-grounding and use of base64-encoded images in prompts for multimodal reasoning.

Latency averaged 1.7s per valuation using Gemini's streaming output, enabling real-time deployment viability. Compared to Claude 3 and GPT-4o, Gemini offered superior response time and remained within the free usage quota, reducing total inference costs by 38%.

These outcomes exceed our deployment thresholds, validating the use of prompt optimization and expert informed testing as a viable substitute for traditional supervised model training.

Tools and Platforms

Tool / Platform	Purpose
Roboflow	Image annotation, model training (YOLOv12), augmentation
Supabase	Cloud database for scans, metadata, and valuations
React Native (Expo)	Mobile app development (iOS/Android)
Gemini 2.5 (Google DeepMind)	LLM for real-time price estimation
TypeScript	Programming language for mobile app
GitHub	Version control and collaboration
NativeWind + Tailwind CSS	Styling for mobile app
Google Search (via Gemini)	Grounded web evidence for valuation

Results and Conclusions

Key Results

The PLUTUS AI project produced a functional, end-to-end prototype that successfully demonstrates how artificial intelligence can support real-time identification and valuation of antiques. Developed in partnership with Rewind Vintage & Consignment, the system was validated on a curated dataset of twenty-five in-store items and supplemented with online marketplace comparables.

Pilot testing confirmed that PLUTUS reduces appraisal time from the traditional four-to-six-week process to approximately thirty seconds per scan. The custom-trained YOLOv12 object detection model achieved a mean average precision (mAP@0.5) of 0.83, delivering accurate item localization suitable for valuation. These detections, passed to a search-grounded Gemini 2.5 prompt, yielded price estimates with an absolute error margin of $\pm 10\text{--}15\%$ —a significant improvement over the $\pm 87\%$ deviation reported by dealers relying on manual estimation. Notably, 81% of final sale prices fell within the model's predicted range (95% CI $\pm 7\%$), with end-to-end inference latency averaging 1.7 seconds from image capture to structured valuation output.

The system's core features include:

- A mobile-first application enabling users to capture item images, detect object categories, retrieve metadata, and receive AI-generated price estimates.
- A fully integrated Supabase backend for persistent storage of scan data, metadata, valuation responses, and price source references.

- An AI-powered valuation agent capable of processing structured metadata and images to generate rationale-backed price estimates in real time.

While these metrics exceed the success criteria defined at project inception, some limitations were observed. Three edge-case failures involving translucent glass items under glare highlight the need for expanded augmentation and environmental controls in future iterations.

Additionally, though early projections suggest a potential 25–35% increase in weekly consignment throughput with full adoption, these operational gains require validation through extended deployment.

Overall, PLUTUS demonstrates that a lightweight, mobile computer vision-plus-LLM pipeline can achieve near-expert appraisal performance with real-time responsiveness, providing a foundation for future scale-up across the antiques sector.

Business and Social Relevance

Although scoped narrowly, our prototype addresses real challenges faced in this industry.

Currently, valuation processes rely heavily on subjective expertise, scattered documentation, and manual research. PLUTUS AI demonstrates how technology can alleviate some of this burden by enabling a store owner to digitize, evaluate, and manage their curated inventory more efficiently.

From a business standpoint, this capability increases pricing consistency and improves recordkeeping. From a social standpoint, the ability to preserve, document, and share contextual knowledge about antiques contributes to the stewardship of historical and cultural artifacts that might otherwise be forgotten or misrepresented.

Solution Applicability

In its current state, PLUTUS AI is tailored for internal use by Rewind Vintage & Consignment. This scoped deployment enabled us to build a well-defined, testable product and receive grounded feedback on system usability, output quality, and next-step priorities. The current solution supports:

- Store staff, who can quickly scan and evaluate in-store items using a guided flow.
- Buyers and collectors, who may benefit from consistent pricing and documented item context.
- Project stakeholders and testers, who can use the prototype to simulate real-world workflows.

The solution idea and vision is designed for multiple user personas across the antiques ecosystem detailed in Stakeholders & Beneficiaries. These user pathways reflect both the flexibility and extensibility of the system across industry segments.

Limitations

While the PLUTUS prototype demonstrates promising functionality, several limitations present opportunities for refinement in future iterations. A primary constraint lies in the limited breadth

of the training dataset. The current model was developed using a relatively small sample, twenty-five unique shop items supplemented by approximately 200 marketplace comparables which restricts its ability to generalize across the full diversity of store inventories. Expanding the dataset through crowd-sourced labeling and incorporating user-submitted scans into the training loop may offer a path toward improved coverage and generalizability.

Visual robustness remains another area with room for improvement. Pilot testing revealed that items with high glare or translucent surfaces, particularly glassware, posed challenges for computer vision accuracy. These instances suggest a potential need for both targeted data augmentation strategies and practical adjustments to environmental conditions, such as the use of standardized low-cost lighting setups at scanning sites.

The current valuation pipeline also introduces concerns around model transparency and potential bias, primarily due to its reliance on a closed-weight large language model (Gemini). Occasional hallucinations or inherited marketplace biases observed during testing indicate the value of exploring open-weight alternatives that support calibration and interpretability, especially in use cases where auditability is critical.

Scalability and performance under operational load are additional areas for consideration. As daily scan volumes increase, potentially exceeding 1,000 scans per day, latency and cost may become significant factors. Approaches such as batched inference requests or localized GPU caching could help mitigate these constraints, though further testing would be required to evaluate feasibility and trade-offs.

Lastly, privacy and regulatory compliance represent ongoing concerns, particularly in multi-region deployments. While Supabase offers technical support for data control features such as opt-in consent and image retention settings, these mechanisms have not yet undergone formal security or compliance audits. Continued attention to data governance, user consent frameworks, and region-specific privacy standards (e.g., GDPR, CCPA) will be essential as the platform matures.

These considerations highlight critical paths for future work and underscore the importance of iterative development in transforming PLUTUS into a robust, scalable, and ethically responsible system.

Future Work

Several avenues remain open for further development to enhance the performance, transparency, and societal impact of the PLUTUS platform. One key area is expanding data diversity and visual robustness. There is potential to improve classification accuracy by incorporating broader, category-balanced image inputs over time. A community-driven labeling approach could support this goal, particularly if future iterations enable user-submitted scans to inform continuous model refinement. Additionally, targeted datasets addressing challenging visual characteristics, such as glare or translucency, may help mitigate specific failure modes observed in early-stage testing.

These efforts could collectively improve mean Average Precision (mAP) and reduce error rates in edge cases.

Model transparency and scalability also present important opportunities. While current valuation estimates demonstrate acceptable levels of accuracy, further investigation into open-weight large language models (e.g., Mixtral, Llama 3) may enable more interpretable and auditable valuation logic. Benchmarking these models against proprietary systems like Gemini could help evaluate trade-offs between performance and explainability. On the scalability front, strategies such as batched inference and on-premises caching warrant exploration, particularly in high-throughput environments. These may offer pathways to reducing latency and operational costs as usage scales.

Finally, the platform's approach to data governance and ethical stewardship could benefit from continued advancement. Enhancing privacy controls through user-directed data management and aligning with frameworks like GDPR, CCPA, and ISO 27001 may support regulatory readiness. Beyond compliance, there is scope to explore broader social utility, such as integrating valuation APIs with museum databases or stolen-art registries. Annual transparency reporting and the use of technologies like blockchain for provenance tracking may further reinforce trust and promote equitable access to cultural data.

Taken together, these directions suggest multiple opportunities for evolving PLUTUS into a robust, ethically grounded tool that bridges technology, commerce, and cultural preservation.

References

Anthropic. (2024, March). *Claude 3 model card* [Technical documentation].

<https://www.anthropic.com/index/claude-3>

Expo. (2024). *Expo* (SDK Version 50) [Computer software]. <https://expo.dev/>

GitHub, Inc. (n.d.). *GitHub*. <https://github.com>

Google DeepMind. (n.d.). *Gemini: Google's multimodal AI model*.

<https://deepmind.google/technologies/gemini/>, <https://ai.google.dev/gemini-api/docs/models>

Liberty Mutual Insurance. (n.d.). *Everything you need to know about antique appraisals*. Liberty Mutual.

<https://www.libertymutual.com/insurance-resources/property/everything-to-know-about-antique-appraisals>

Lion and Unicorn. (2023, September 5). *How much does it cost to get an antique appraised?* Lion and Unicorn.

[https://lionandunicorn.com/how-much-does-it-cost-to-get-an-antique-appraised/#:~:text=For%20more%20comprehensive%20appraisals%20conducted,%300%20to%20\\$1%2C000%20or%20more.](https://lionandunicorn.com/how-much-does-it-cost-to-get-an-antique-appraised/#:~:text=For%20more%20comprehensive%20appraisals%20conducted,%300%20to%20$1%2C000%20or%20more.)

Roboflow. (2024). *Roboflow* [Computer software]. <https://roboflow.com/>

Supabase Inc. (2024). *Supabase* (Version 1.0) [Computer software]. <https://supabase.com/>

Meta Platforms, Inc. (n.d.). *React – A JavaScript library for building user interfaces*. <https://reactjs.org>

Meta Platforms, Inc. (n.d.). *React Native: Create native apps for Android and iOS using React*. <https://reactnative.dev>

Microsoft Corporation. (n.d.). *TypeScript: JavaScript with syntax for types*. <https://www.typescriptlang.org>

NativeWind. (n.d.). *NativeWind: Tailwind CSS for React Native*. <https://www.nativewind.dev>

OpenAI. (2024, May). *GPT-4o model card* [Technical documentation]. <https://platform.openai.com/docs/guides/gpt-4o>

Tailwind Labs. (n.d.). *Tailwind CSS: A utility-first CSS framework*. <https://tailwindcss.com>

Wang, C. Y., Bochkovski, A., & Liao, H. Y. M. (2024). *YOLOv12: You Only Look Once object detection model* [Computer software]. GitHub. <https://github.com/WongKinYiu/yolov12>

Appendix – Reproduction of Results

Note: Due to privacy and security constraints, both the source code and data repositories for this project are hosted in private environments and are not publicly accessible. These include a private GitHub repository for the source code and a secured Supabase instance for database storage. Access is restricted to authorized project collaborators only.

Code and Data Overview

Data Sources:

All training and inference data are stored in a private Supabase PostgreSQL database. This includes item metadata, image references, and valuation records. Additional scraped comparables from platforms such as eBay and 1stDibs are cached internally in the same database.

Tools and Platforms:

- **Frontend:** React Native (Expo)
- **Backend/Database:** Supabase (PostgreSQL, Edge Functions, Storage)
- **AI/ML:** Roboflow (YOLOv12 custom object detector), Google Gemini 2.5 (valuation prompt model)
- **DevOps:** GitHub (private repo), Vercel (optional frontend testing)

Scripts and Structure:

Core scripts are contained in the `plutus-ai-app` private GitHub repository.

Key folders:

- `/models/`: training and evaluation configs for YOLOv12 (Roboflow integration)
- `/scripts/`: metadata enrichment, price scraping, and valuation prompts
- `/api/`: Supabase Edge Functions for real-time queries
- `/app/`: React Native mobile UI with navigation and screen logic

How to Run:

- 1) Clone the private GitHub repo (access required).
- 2) Install dependencies with `npm install` or `yarn install`.
- 3) Set up `.env` with Supabase keys and Roboflow/Gemini credentials.
- 4) Run development build using `npx expo start`.
- 5) Execute backend functions and valuation scripts via Supabase Edge Function routes or local testing endpoints

For access credentials or reproduction assistance, please contact the project leads directly.