DIGIMITH

From idea to software that runs

CASE STUDY

AI-DRIVEN PRODUCT PLACEMENT PIPELINE

Intelligently placing products onto stands through synthetic views and deterministic AI pipelines.

CLIENT

Emperia

EMPERIA · PYTHON/FASTAPI/COMPUTER VISION · GPT-4O · GEMINI PRO

01 PROJECT OVERVIEW & CLIENT CONTEXT

Emperia is the global standard-bearer for virtual commerce, specialising in high-fidelity, immersive retail environments for the world's most prestigious luxury brands. To maintain the "pixel-perfect" reputation of these virtual storefronts, the production team required a more agile method for updating product inventory within 3D-integrated spaces.

- **Key Achievement:** Successfully replaced a 48-hour manual 3D modelling workflow with a near-instantaneous AI-driven pipeline, resulting in a 90% efficiency gain in production delivery.

PROJECT METADATA

FIELD	DETAILS
Timeline	2 Months
Team Size	Lead Developer — Muhammad Zeeshan (Sole Developer)
Client	Emperia
Status	Internal Tool — Production-Ready

02 THE CHALLENGE: BUSINESS VS. TECHNICAL BARRIERS

2.1 BUSINESS CHALLENGES: THE CREATIVE BOTTLENECK

The legacy workflow relied heavily on manual intervention from the 3D production team. According to Charlie Adams (Director of 3D Production), the manual pipeline required constant artist supervision to "assist with obstacles" and navigate the friction of Figma-to-3D translation. Tomas Turevicius (3D Environment Artist) noted that even minor updates required artists to "update different functionalities" across the virtual web experience, consuming up to 48 hours per environment.

- **Operational Friction** — Production timelines were tethered to artist availability rather than client demand.
- **Inventory Lag** — Rapidly changing client requirements led to a minimum one-week turnaround for product updates.

2.2 TECHNICAL CONSTRAINTS: THE RESOLUTION BARRIER

From an architectural perspective, the primary hurdle was the resolution and context window limitations of multimodal Large Language Models (LLMs).

- **Pixel Density Loss** — Modern vision models such as Gemini Pro have maximum input resolution thresholds. When a high-fidelity 4K panorama is fed directly into these models, the system downscales the image, resulting in a catastrophic loss of detail for the product.
- **Non-Native 4K Support** — Generative AI does not natively support 4K output. Maintaining the visual integrity required for luxury retail necessitated a custom programmatic mitigation strategy to bypass these hardware-level constraints.

03 THE SOLUTION: DETERMINISTIC AI PIPELINE

The solution involved engineering a stateless API orchestration that leverages deterministic coordinate-based logic to place products accurately while maintaining 4K fidelity. By moving away from purely "generative" haphazardness toward a structured pipeline, the system ensures that synthetic views of products integrate seamlessly with their 3D stands.

04 TECHNOLOGY STACK

CATEGORY	TECHNOLOGIES
Languages & Frameworks	Python, FastAPI
Computer Vision	OpenCV, PIL
LLMs & Multimodal	GPT-4o, Gemini Pro

05 TECHNICAL IMPLEMENTATION DETAILS

The architecture follows a "Crop-Process-Stitch" workflow — a mitigation strategy engineered to preserve pixel density while utilising the full reasoning capabilities of multimodal LLMs.

THE "CROP-PROCESS-STITCH" WORKFLOW

1. **Coordinate-Based Slot Identification** — Rather than relying on visual "guesses," the system uses deterministic coordinate logic to identify pre-defined placement slots within the 4K environment.
2. **Strategic Segmentation** — The 4K panorama is programmatically cropped into smaller, high-resolution segments. This ensures that the input fed to the AI maintains the necessary pixel density for high-end retail.
3. **Multimodal Reasoning & Placement** — Using Gemini Pro and GPT-4o, the system reasons about the product's orientation and scales the synthetic view onto the designated stand within the segment.
4. **Deterministic Re-stitching** — Processed segments are re-integrated into the original 4K background using OpenCV and PIL, ensuring the final output remains a seamless, high-resolution panorama.

> Architectural Optimisation: To handle enterprise-scale experiences (e.g., 13+ products in a single environment), the pipeline utilises parallel processing via Python's asyncio within FastAPI, ensuring near-instant execution regardless of product volume.

06 IMPLEMENTATION RESULTS & IMPACT

INSTANT WORKFLOW DURATION (WAS 48H)	**90%** PRODUCTION EFFICIENCY GAIN
1 DAY PRODUCT UPDATE (WAS 1 WEEK)	**13+** PRODUCTS IN SINGLE PASS

The transition to a deterministic AI pipeline decoupled production capacity from artist headcount, allowing Emperia to scale its virtual retail offerings without throughput constraints.

METRIC	MANUAL PROCESS (BEFORE)	AI PIPELINE (AFTER)
Workflow Duration	48 Hours	Near-Instant
Product Change Turnaround	1 Week	1 Day
Output Quality	Manual 3D Modelling	Deterministic Synthetic Views
Scalability	Limited by Artist Headcount	Infinite via Horizontal API Scaling

> Volume Note: The system has been stress-tested to handle 13 individual product placements within a single high-resolution experience in one pass.

07 PROJECT TIMELINE

- **Month 1 — Logic & MVP** — Focused on the core placement logic and coordinate-based slot identification. Rigorous testing was conducted to ensure the synthetic product views aligned perfectly with the 3D backgrounds.
- **Month 2 — Architecture & Integration** — The system was migrated to a robust FastAPI-based architecture. APIs were exposed to the internal production team, allowing them to integrate the tool directly into their daily creative workflows.

08 KEY LESSONS LEARNED

LESSON 1: DECOUPLING CREATIVE FROM REPETITIVE LABOURS

The most significant ROI was achieved not by replacing artists, but by automating high-friction, repetitive tasks. By removing the 48-hour 3D modelling bottleneck, the production team can now focus on high-level environment design and brand strategy, while the AI handles the "heavy lifting" of inventory placement.

LESSON 2: MASTERING MULTIMODAL REASONING

Success in production-grade AI requires moving beyond simple generation. This project demonstrated that when LLMs are used for multimodal reasoning within a deterministic pipeline — constrained by OpenCV and PIL — they can deliver the reliability and "pixel-perfect" quality required by global luxury brands.

DIGIMITH

From idea to software that runs

hello@digimith.com · www.digimith.com · +92 306 1234567

Johar Town, Lahore, Pakistan