

Colin Joy Lobo

Karnataka, India • +91 8867731318 • colinjoylobo16@gmail.com • colinjoylobo.vercel.app • linkedin.com/in/colinjoylobo • github.com/colinjoylobo

EXPERIENCE

AI Engineer , <i>Collective Artists Network (Galleri5)</i>	Nov 2025 – Present
- Built agentic AI systems across the three-layer generative-video stack (storyboard / generation / orchestration), in production on the AI series <i>Mahabharat</i> (JioHotstar) and the upcoming films <i>Hanuman</i> and <i>Krishna</i>. - Designed and shipped six production multi-agent systems on Azure AI Foundry covering script writing, shot breakdown, asset generation, visual breakdown, BG-plate and video composition. MCP-style tool calling, RAG over internal vector stores, 4-level LLM-as-judge eval harness for output QC. - Fine-tuned open-source foundation models (Flux, Wan, Qwen, SDXL) and trained custom LoRAs with reference-image control to lock character likeness and art style across long-running shows. - Built Carve, an internal VFX segmentation and image-matting tool around SAM 2 / SAM 3 and ViTMatte; trained a clean-license ViTMatte-S replacement that shipped as the production matte model; delivered EXR shot mattes (SAM 3.1 tiled, CLaHE QC) on live film projects. - Productionised generative-video pipelines on Seedance 2.0, Bernini (ByteDance) and Nanobanana for OTT and film clients, including Hansal Mehta's <i>Khana Dil Se</i>. - Authored internal references and benchmarked new open-source models (Wan 2.7, LTX, HunyuanVideo, SAM 3.1, LaMa, IC-Light, FBA Matting) to pick the right fit per production.	
Trainee QA Executive , <i>Winman Software Pvt. Ltd.</i>	Mar 2025 – Oct 2025
Manual testing of flagship desktop ERP / tax software; root-cause analysis of functional defects across financial data modules and validation of fixes with dev + support.	
Machine Learning Intern , <i>Suvidha Mahila Mandal NGO, Walni</i>	Dec 2023 – Jan 2024
Built churn-prediction models lifting accuracy from 70% to 85%; automated reporting and cut reporting time by 30%.	
Data Science Intern , <i>BrainoVision Solutions India Pvt. Ltd.</i>	Jun – Aug 2023
Curated NLP datasets and co-designed text-classification models, improving classification accuracy by 20%.	

SKILLS

Agentic AI: Multi-agent system design, Azure AI Foundry hosted agents, Model Context Protocol (MCP), tool / function calling, retrieval-augmented generation (RAG), LLM-as-judge eval harnesses, agent QC, prompt engineering.
Model Training: LoRA training, foundation-model fine-tuning (Flux, Wan, Qwen, SDXL), reference-image control for character consistency, dataset curation, PyTorch.
Open-Source Models: Flux, Wan, Qwen, SDXL, Stable Diffusion, SAM 2 / SAM 3, ViTMatte, LaMa, IC-Light, OWLv2, FBA Matting, Hermes, Gemma.
Generative Video & Image: Three-layer pipeline (storyboard / generation / orchestration), Seedance 2.0, Bernini (ByteDance), Nanobanana, gpt-image-2, ComfyUI workflows, video-to-video upscaling, multi-character lipsync.
Infrastructure / LLMOps: Azure AI Foundry, Azure H100 GPUs, FastAPI, serverless inference endpoints, vector stores, Gradio, systemd services, Cloudflare tunnels, Git.
3D / VFX: Blender Python API, retopology, texture baking, UV unwrapping, EXR / ACES colour pipelines.
Languages: Python, SQL, R; English, Hindi, Kannada (all fluent).

EDUCATION

Master of Computer Application (Online) — Jain (Deemed-to-be University), Bengaluru	2024 – Present
Bachelor of Computer Application (Data Science) — St. Joseph's University, Bengaluru CGPA: 7.72	2022 – 2024

SELECTED PROJECTS AT COLLECTIVE ARTISTS NETWORK

MultiCam Control — GoPro Fleet Production App *Electron · React 18 · FastAPI · Python 3.11 · BLE · COHN · WebSockets*

Cross-platform desktop app for synchronised fleet control of 3+ GoPro Hero 11/12/13 cameras. BLE discovery, COHN multi-stream MPEG-TS live preview, parallel S3 downloads, HiLight tagging, synchronised record/stop, snapshot grids and SD-card management. Companion pipeline wires captured footage into Blender asset assembly and style-transfer + LTX-2 video output for cinematic stylisation. Installers for macOS, Windows, Linux. Public repo on GitHub.

Multi-Agent Creative Pipeline on Azure AI Foundry *Foundry Hosted Agents · Sonnet 4.6 · gpt-image-2 · MCP · RAG · Bing grounding*

Six production hosted agents (ScriptWriter, ShotBreakdown, Asset, VisualBreakdown, BGPlate, VideoComposer) that take a script through to renderable shot panels. MCP-style tool calling, retrieval-augmented grounding over internal vector stores, and a 4-level LLM-as-judge eval harness for output QC. Shot-Sequence agent renders per-shot composite grids and runs an LLM aesthetic classifier that injects style and mood metadata into every render prompt.

Cue — AI Storyboard Generation App *Gemini · Azure AI Foundry · fal.ai · gpt-image-2*

Artist-facing storyboard tool and cost-efficient successor to the agent pipeline: turns a rough director's sketch into a finished storyboard (pencil or colour) with reference and style control, camera-path annotation and batch generation. Runs on Gemini + fal.ai with no dedicated GPU; UI shaped by a round of lead feedback (export, job pause / resume, progress). In production on the upcoming *Hanuman* and *Krishna*.

LoRA Training & Foundation-Model Fine-Tuning *Flux · Wan · Qwen · SDXL · PyTorch · ComfyUI*

Trained character and style LoRAs on Flux and SDXL with reference-image control to lock likeness and art direction across long-running shows. Fine-tuned Wan and Qwen for in-house generation tasks, with workflows shipped to the artist team via ComfyUI.

Carve — VFX Segmentation, Matting & Clean-Plate *Python · FastAPI · PyTorch · SAM 2 / 3 / 3.1 · ViTMatte · BiRefNet · ProPainter*

Web tool with four production features — SAM auto multi-ID, SAM prompted, ViTMatte alpha matting, and clean-plate (paint a mask, ProPainter inpaints the object out of the whole shot) — with native EXR ingest, ACES tonemapping and SAM 2 sequence propagation; A/B-tested SAM 3.1 against tiled SAM 3 on dense-crowd footage. Trained a clean-license ViTMatte-S replacement (BiRefNet pseudo-labels on COCO, P3M-10k, CorridorKey greenscreen alphas, MANIFEST.json provenance) so mattes can ship to external clients under a clean license; v2.1 is the production matte model. Delivered mattes on live film projects.

Mahabharat (JioHotstar), Hanuman & Krishna — Production Pipelines *ComfyUI · Stable Diffusion · Blender · Ultra-Sharp Upscaler · Azure H100*

Built and operate the ComfyUI workflows and upscaler service that scale output for *Mahabharat* on JioHotstar, supporting 50+ artists, plus storyboard agents, 3D-environment automation in Blender, and creative tooling for the upcoming *Hanuman* and *Krishna* films.

Generative Video Pipelines — Seedance 2.0, Bernini, Lipsync *Python · Seedance · Bernini (ByteDance) · Nanobanana*

Built pipelines spanning static-still ingestion through 88 generated Bernini prompts for Hansal Mehta's *Khana Dil Se* Ep04, plus a multi-character lipsync workflow. Wrote the internal Seedance 2.0 reference guide (endpoints, guardrails, prompt templates) used across the team.

CERTIFICATES & RECOGNITION

Best-Performing Team, IDE Bootcamp (MIC + AICTE) • **Advanced Excel**, St. Joseph's University • **Introduction to AI**, Infosys Springboard • **Unboxing AI: Healthcare & Medical Data**, St. Aloysius College • Hosted *Data Visualisers* at college data-analytics fest *Analytica*.