

Dorna Raj Gyawali

📍 Butwal, Nepal | [Email](#) | [Github](#) | [linkdein](#) | [PortFolio](#)

Education

Tribhuvan University

Bachelor in Information Communication and Technology

GPA - 3.55/4.0 (First Year Average)

Relevant coursework: Probability & Statistics, Web Technology, DBMS, OS and OOPS

Sep 2024 - June 2028

Kathmandu, Nepal

Projects

Quark - AI RAG System with Persistent Memory | [GitHub Link](#)

April 2026

TypeScript, Python, ElysiaJS, Redis, BullMQ, Qdrant, AWS S3, Supabase, React, TanStack Query, Docker

- Built a RAG-based **document intelligence system** supporting persistent multi-session memory using Redis + Qdrant hybrid architecture.
- Designed BullMQ-based ingestion pipeline to asynchronously process embeddings, **reducing API latency** under concurrent uploads.
- Designed a persistent AI memory layer using Redis for short-term state and MemO for long-term memory compression, enabling **context-aware interactions** across sessions while reducing token usage.
- Improved retrieval accuracy using **embedding + reranking pipeline**, reducing irrelevant responses in deep document queries.
- Designed a scalable S3-compatible storage system with a metadata sync layer ensuring consistency between uploads and database state under concurrent processing.

SDR-9 - P2P File Transfer System | [GitHub Link](#) | [Live](#)

March 2026

React, Express, WebRTC, Socket.io, Redis, BullMQ, Resend, Docker

- Built a browser-based P2P file transfer system **eliminating server bandwidth usage** via WebRTC-based architecture.
- Enabled 1GB+ file transfers using chunked streaming with ArrayBuffer and Uint8Array for **memory-efficient transmission**.
- Designed a **real-time signaling layer** with Socket.io for low-latency peer discovery and connection orchestration.
- Implemented Redis-backed worker queues to **offload background jobs** like email invites and notification processing.

PX47 | Distributed Media Processing Pipeline | [GitHub Link](#)

Dec 2025

Javascript, Node.js, Redis, BullMQ, FFmpeg, AudioWaveForm, AWS S3, Prisma, MongoDB

- Built a distributed media processing pipeline for asynchronous handling of **large-scale video/audio workloads**.
- Leveraged **pre-signed S3 uploads** to bypass backend bottlenecks and enable direct client-to-storage transfers.
- Designed a worker-based **processing engine** for transcoding, waveform extraction, and metadata generation at scale.
- Orchestrated **multi-stage media transformation** workflows using FFmpeg for scalable and fault-tolerant processing.

Open Source Contributions

GreedyBear (IntelOwl Ecosystem) | [PR](#) | [Live](#)

Jan 2026

Python, Django, Django Rest Framework, Celery, Docker, Elasticsearch, PostgreSQL, Machine Learning, React.js

- Contributed to a **threat intelligence pipeline** for IOC analysis and automated security data processing
- Integrated **GeoIP and ASN** enrichment to enhance **contextual intelligence** for threat attribution.
- Improved **CI reliability** by refactoring test suites and adding **robust database migration** validation tests.
- Added feature importance logging to improve interpretability of **ML-based threat classification** models.

Technical Skills

Languages: Python, TypeScript, C/C++, GO, HTML/CSS

Backend & Frameworks: Node.js, ElysiaJs, Django, FastAPI, Redis, BullMQ, Docker, Tailwind, React.js

Cloud & DevOps: AWS(s3), Docker, CI/CD, Nginx

Database & Baas: MongoDB, PostgreSQL, QdrantDB, Appwrite, Supabase

Tools & Testing: Git, Linux, Postman, Github Actions, Pytest, Vitest/jest