



Red Light Camera Site Selection

AI-Powered Video Analytics for Pedestrian Safety & Traffic Enforcement in Hartford, CT

Kwadwo Osarfo-Akoto '26 | François Carpe-Elías '26 | Alejandro Sanchez '26 | Advisor: Professor Kenneth Kousen

INTRODUCTION & MOTIVATION

The City of Hartford is pursuing legislation to implement red light cameras at high-risk intersections. Without data-driven analysis, officials must rely on intuition to select camera sites, leading to suboptimal placement and missed safety opportunities.

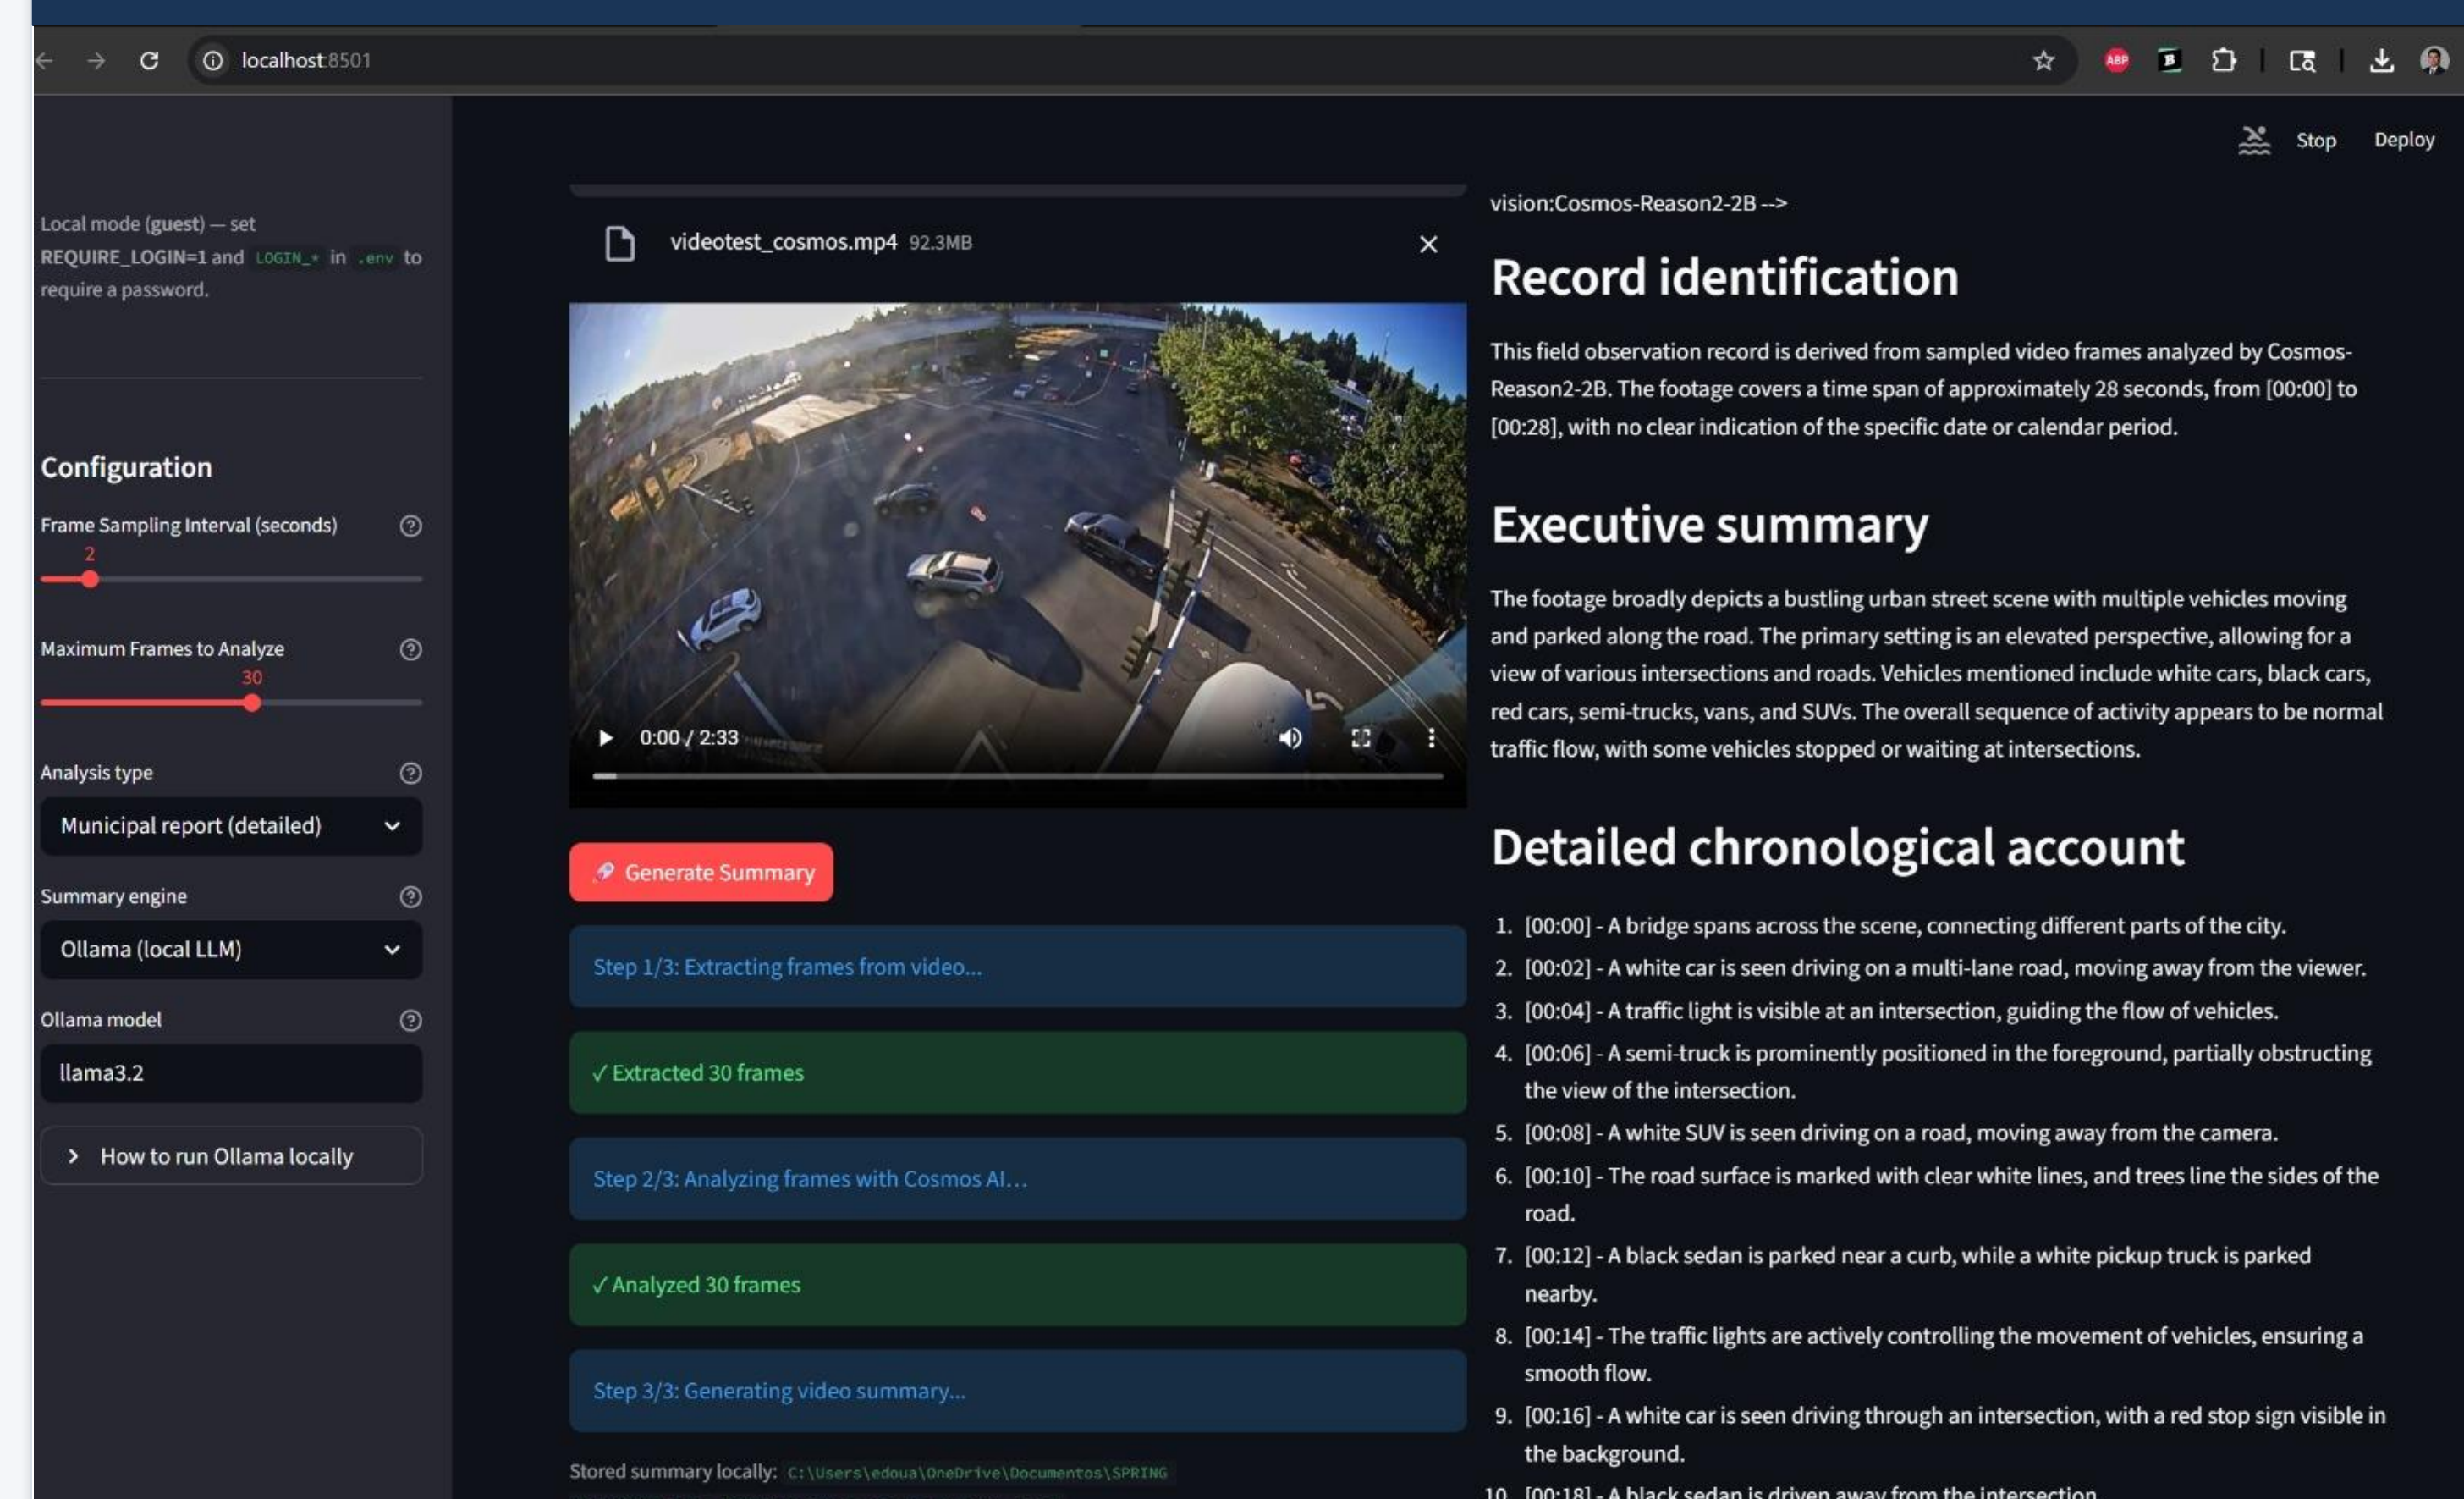
PROJECT OBJECTIVES

- Recognize traffic behavior patterns across intersections and streets in Hartford using video-derived observations of vehicles and pedestrian activity.
- Assisting Hartford PD generating comprehensive municipal reports on busy downtown intersections.
- Designing a program blueprint and foundation for the city of Hartford to implement planned smart city infrastructure, creating incentive for future funding.

APP FEATURES

- Video Uploading and Processing: Supports multiple formats (MP4, AVI, MOV, MKV)
- AI-Powered Frame Analysis: Uses Nvidia Cosmos-Reason2-8B to analyze extracted frames and identify meaningful scene-level events.
- Flexible Summaries: Summaries can be generated in unique styles: detailed, concise, bullet points and municipal reports.
- Customizable Analysis Settings: Users can configure frame sampling interval and maximum frames analyzed.
- Exportable results: Download generated summaries as text files for sharing and reposting.

USER INTERFACE



TECHNICAL APPROACH

1 CCTV Data Ingestion

Source historical footage from Hartford Police CCTV cameras near downtown intersections via FOIA request. Geofenced area focused on Arena/Stadium district.

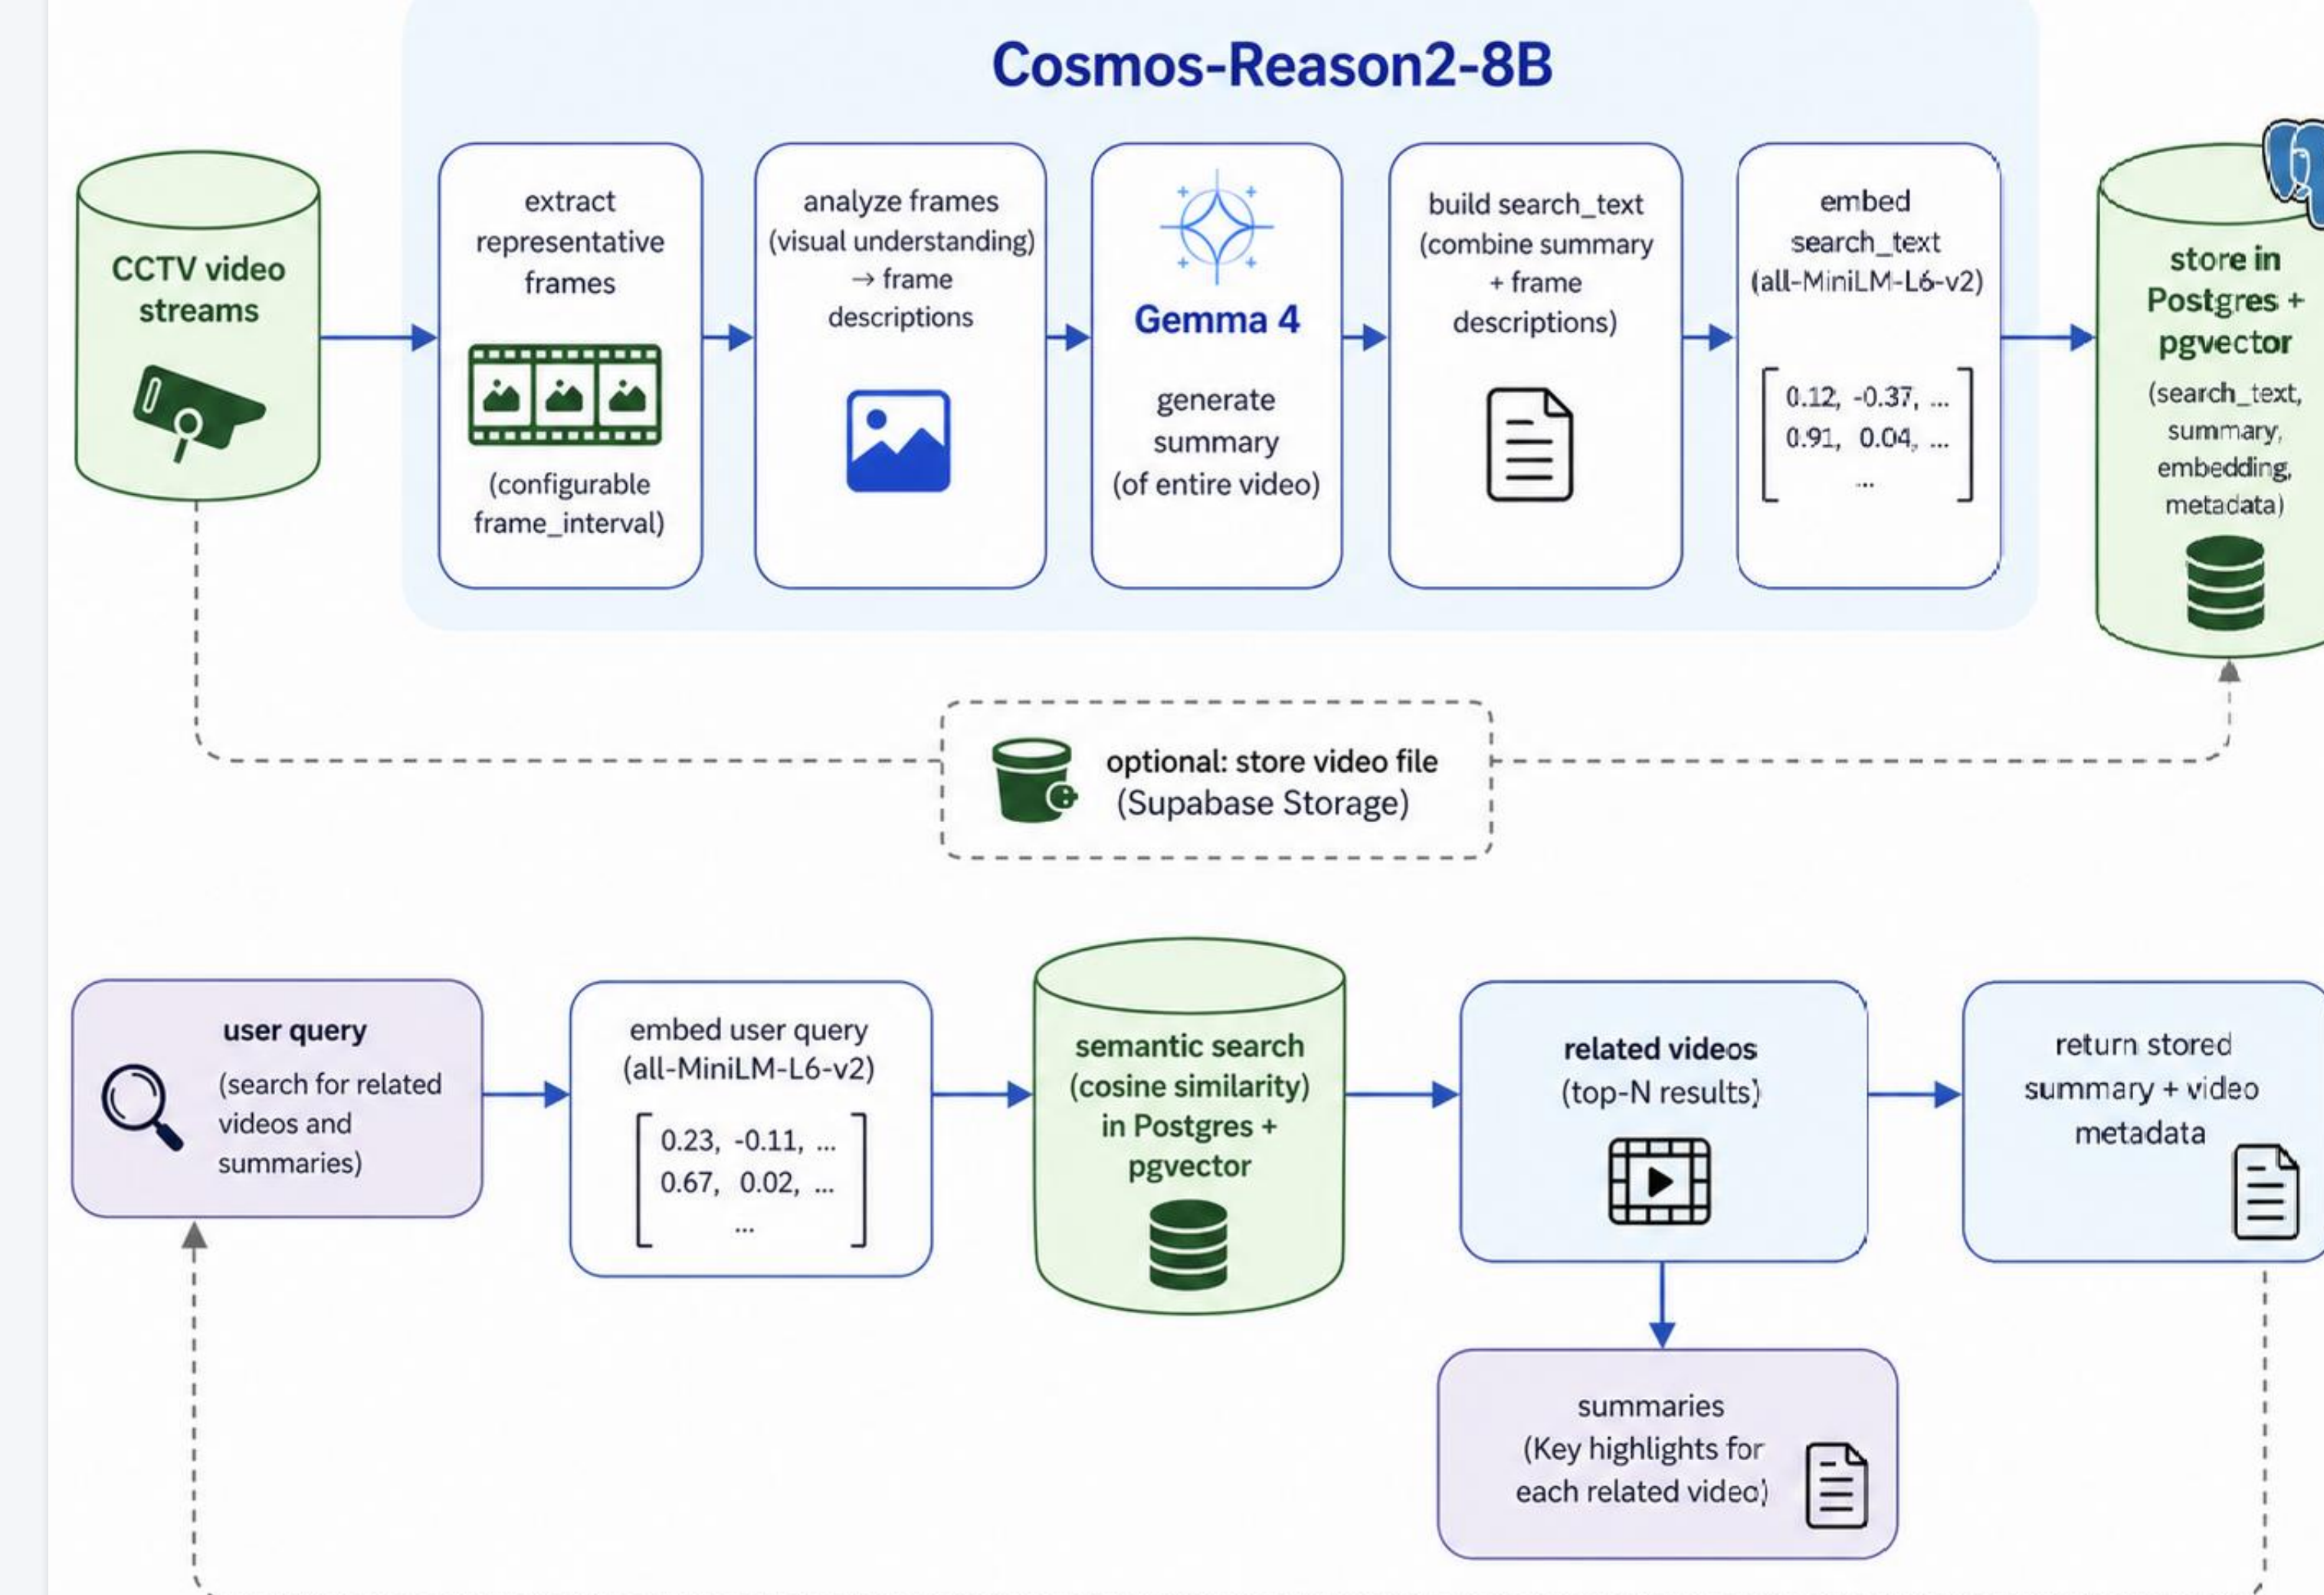
2 Video Preprocessing and Summarization

Sample representative frames from uploaded footage (configurable interval), use NVIDIA Cosmos-Reason2-8B to generate scene descriptions, then synthesize and store summary metadata and embeddings for retrieval.

3 Video Search

Search for video within vector database using summary embeddings, returning the most semantically relevant results.

PROJECT ARCHITECTURE



TECHNOLOGIES USED

- Nvidia VLM Model: Cosmos-Reason2-8B
- Summarization LLM: Gemma4
- Main language: Python
- Frontend: Streamlit
- Video and data storage: Supabase



ACKNOWLEDGEMENTS

- Mike Binko and the AIQUI Innovation Sandbox
- Professor Kenneth Kousen from the Computer Science Department
- Danny Briere and Elting Center for Innovation and Entrepreneurship
- Paul Lanzoni and Hangxin Lu from the Cisco Team
- David Cedrone and Chandra Andhe from the Vzure Team
- Donna Tadiello From Campus Safety