LOG BOOK

Traffix-AI

STEM High School for Boys - 6th of October

Mohamed Hassan Ali

1023111

Group 10302

Grade 12

ed.m7md7asn@gmail.com

Mohamed.1023111@stemoctober.moe.edu.eg

2025 - 2026

Week 01

Portfolio Introduction

After reading and understanding the capstone challenge, I wrote the introduction of the portfolio, summarizing the grand challenges we addressed this semester.

Proof:

Introduction

Egypt confronts significant obstacles referred to as "Grand Challenges," **illustrated in Figure 1.1**. These challenges contribute to the decline of its economic, scientific, cultural, and technological foundations. Egypt aims to address these issues as part of its Vision 2030 plan, with hopes of achieving solutions by 2030.

Figure 1.1: Egypt's Grand Challenges

In Egypt, rapid population growth—now exceeding 110 million—intensifies urban congestion in cities like Cairo and Alexandria, where overcrowded infrastructure strains transportation, housing, and public services. This congestion exacerbates air pollution from vehicle emissions and industrial activity, contributing to some of the highest urban particulate matter levels in the region and worsening respiratory diseases such as asthma and bronchitis. Simultaneously, inadequate waste management and industrial discharge continue to foul the Nile River and groundwater, threatening both public health and agricultural productivity. These pollution sources, combined with rising temperatures and erratic rainfall patterns driven by climate change, heighten the risks of water scarcity, heat stress, and vector-borne diseases. Addressing urban planning through sustainable transit, green spaces, and decentralized development can alleviate congestion while reducing emissions. Coupled with climate-resilient infrastructure and expanded healthcare access—particularly in densely populated informal settlements—Egypt can simultaneously tackle public health crises, mitigate pollution, and build adaptive capacity against the accelerating impacts of climate change.

Population Growth Grand Challenge

After writing the introduction, I wrote the grand challenge of dealing with population growth and its consequences (overview, causes, and impacts).

Proof:

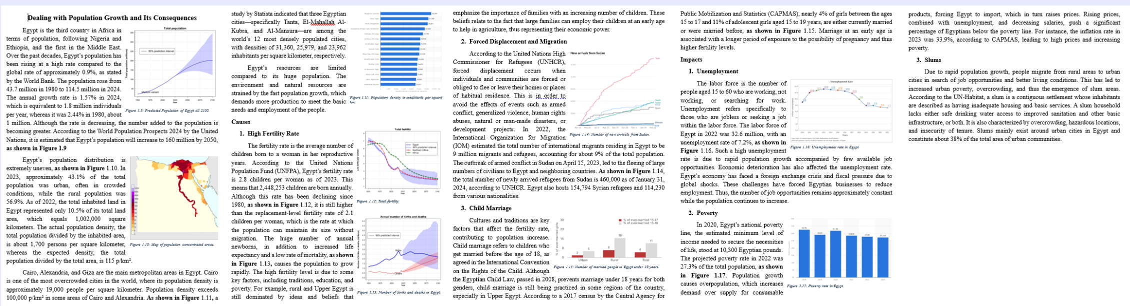

Week 02

Problem to Be Solved

This week, I worked on determining the problem to be solved. After extensive research, I identified a suitable problem in Al-Fayoum Governotare, titled *"Traffic Congestion Impeding Urban Mobility and Exacerbating Air Pollution at Al-Masalla Square, Fayoum, Egypt."* I documented the problem and its consequences in the portfolio.

Proof:

Problem to be Solved

Traffic Congestion Impeding Urban Mobility and Exacerbating Air Pollution at Al-Masalla Square, Fayoum, Egypt.

The escalating traffic congestion in Cairo, fueled by rapid urbanization, population growth, and inadequate infrastructure, has created profound challenges for mobility, economic productivity, and environmental health in one of the world's most densely populated cities. Greater Cairo, home to over 22 million residents as of 2025, experiences chronic gridlock that reduces average peak-hour speeds to 22-27 km/h across key corridors, resulting in daily delays that cost the Egyptian economy up to 4% of GDP—approximately EGP 50-100 billion annually in lost time, excess fuel consumption, and increased vehicle emissions.

This problem is especially pronounced in Fayoum City, where high vehicle densities intersect with pedestrian activity and informal transport modes like microbuses. Al-Masalla Square, serving as the backbone of Fayoum City and one of its busiest shopping streets, exemplifies these issues: it handles peak volumes exceeding 2,000 vehicles per hour, with speed indices dropping to as low as 0.21 during evenings due to multi-lane conflicts, frequent U-turns, illegal parking, and overcrowding from nearby malls and offices.

According to the World Bank's Cairo Traffic Congestion Study (updated analyses through 2024), this intersection contributes to zonal congestion costs in Fayoum City, accounting for up to 23.6% of the Greater Cairo Metropolitan Area's (GCMA) total economic losses from delays and unreliability, including non-recurrent incidents like breakdowns that amplify travel time variability (coefficient of variation up to 0.59).

As a direct consequence, air pollution intensifies, with traffic sources responsible for 33% of the city's PM2.5 concentrations—averaging 39.9 µg/m³ annually, well above the World Health Organization's guideline of 5 µg/m³. In Fayoum City specifically, recent 2024-2025 data from IQAir and Weather Underground show real-time PM2.5 levels ranging from 19-67 µg/m³, leading to moderate-to-unhealthy Air Quality Index (AQI) readings of 50-150, exacerbated by idling vehicles and microplastic contaminants in the air (up to 69 particles/m³ indoors and outdoors).

This pollution not only degrades local air quality but also contributes to broader environmental issues, such as urban heat islands and runoff contaminating soil and water. These interconnected problems at Abbas Al-Akkad underscore the urgent need to address urban congestion and its consequences, aligning with Egypt's national priorities for managing population growth, reducing public health burdens, and promoting sustainable urban development under Vision 2030.

Positive Consequences

1. **Enhanced Traffic Efficiency and Safety**

Solving congestion at Al-Masalla Square through ICT-based systems would significantly improve traffic flow, reducing average travel times by up to 50% and minimizing accident rates caused by chaotic U-turns and pedestrian-vehicle conflicts. By implementing dynamic traffic signs and AI-driven vehicle communication, emergency vehicles could gain priority, decreasing response times and potentially saving lives. This would also lower fuel consumption by 20-30%, as seen in similar urban interventions, fostering economic productivity by allowing commuters to spend less time in traffic and more on work or leisure, ultimately contributing to Cairo's GDP growth.

2. **Reduction in Air Pollution and Environmental Benefits**

Effective congestion management would directly curb vehicle idling and emissions, lowering PM2.5 levels in Fayoum City by an estimated 15-25%, based on global ITS case studies. This would mitigate the environmental footprint, including reduced CO_2 contributions to climate change, and promote greener urban spaces by encouraging public transport use over private cars. In turn, this aligns with Egypt's sustainability objectives, potentially expanding green corridors and improving overall air quality to meet international standards below 10 µg/m³ annually.

3. **Boost to Economic and Urban Development**

Alleviating gridlock would unlock economic potential in Fayoum City, a vibrant commercial zone, by reducing annual losses from delays and unreliability—currently part of the EGP 765 million attributed to similar routes. Improved mobility would attract investments in retail and offices around Abbas Al-Akkad, stimulating job creation and property values while supporting population growth without straining infrastructure. This could lead to broader urban revitalization, integrating smart technologies to make the area more livable and competitive.

Negative Consequences

1. **Public Health Risks**

Persistent congestion exacerbates air pollution, with traffic-sourced PM2.5 in East Cairo linked to 18,000 premature deaths annually citywide, including respiratory issues, cancer risks (up to 20,000 cases from petroleum emissions), and cardiovascular diseases. In Fayoum City, where 62% of school students report noticing dust in the air, prolonged exposure could increase waterborne and airborne illnesses, straining healthcare systems and reducing life expectancy by 1-2 years on average.

2. **Economic Losses and Productivity Decline**

Unchecked traffic at Abbas Al-Akkad would continue to impose massive costs, with variability in travel times (coefficient of variation up to 0.59) leading to unreliability that hampers business operations and supply chains. This contributes to Egypt's 4% GDP loss from congestion, including excess fuel and time wasted, potentially escalating to billions more by 2030 as vehicle volumes rise, deterring investments and perpetuating poverty cycles in densely populated areas.

3. **Environmental Degradation and Biodiversity Impact**

Ongoing idling and high emissions would worsen urban heat islands and soil/water contamination from runoff in Fayoum City, indirectly affecting nearby green spaces and the Nile ecosystem through increased noise (70-80 dB) and pollutants. This could diminish local biodiversity, such as aquatic life in connected waterways, and accelerate climate change effects like heatwaves, undermining Egypt's efforts to combat desertification and sustain its limited arable land.

Week 03

Other Solutions Already Tried

After identifying the problem to be solved, I searched for prior solutions of similar problems and studied their mechanisms, points of strength, and points of weakness. I wrote two prior solutions in the portfolio, namely *Poznan, Poland: AI-Optimized Traffic Signals* and *New York City: Automated Incident Detection System.*

Proof:

Prior Solution 1

Other Solutions Already Tried

Poznan, Poland: AI-Optimized Traffic Signals

The Poznan Adaptive Traffic Control System, deployed in the early 2020s, represents a cutting-edge implementation of machine learning-based traffic signal optimization in Central Europe. This system was developed through collaboration between the City of Poznan and local technology firms to address persistent congestion in the city's downtown corridor and suburban arterials. The initiative integrates inductive loop detectors embedded in road surfaces with high-resolution camera systems and artificial intelligence algorithms to dynamically adjust signal timing in real-time. By processing data from over 150 signalized intersections, the system continuously learns traffic patterns and adapts signal phases to optimize vehicle throughput during varying demand periods. The platform demonstrates significant improvements in traffic flow efficiency, reducing average travel times by 20-30% during peak hours while simultaneously decreasing emissions through smoother acceleration patterns and reduced idling times.

Mechanism

The Poznan system operates through a multi-layered architecture, **shown in Figure 1.28**, integrating hardware sensors, data processing, and machine learning optimization:

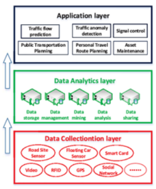

Figure 1.28: ITS Architecture core layers

Data Collection Layer: Inductive loop detectors buried beneath road surfaces detect vehicle presence and classify vehicle types based on magnetic signatures. Synchronized with these ground-based sensors are computer vision cameras positioned at intersections to capture real-time traffic conditions, including queue lengths, turning movements, and anomalous vehicle behavior.

Data Processing and Transmission: Raw sensor data from the 150+ intersections is transmitted to a centralized traffic management center via secure network infrastructure. The system processes approximately 10,000 data points per minute across the network, creating a comprehensive real-time traffic state representation.

Machine Learning Optimization: The core of the system employs ensemble machine learning models, including gradient boosting and neural networks, trained on historical traffic data spanning multiple years. These models predict traffic demand for the next 15-30 minutes at each intersection based on patterns including time of day, day of week, weather conditions, and special events. The algorithms then calculate optimal signal timing plans that minimize total vehicle delay across the network while maintaining safety constraints.

Adaptive Signal Control: Based on ML predictions, the system automatically adjusts green time allocation among competing traffic movements. During morning peak hours, for example, inbound traffic receives extended green phases, while evening peaks trigger outbound optimization. Emergency vehicle detection systems override standard timing to provide preemption, clearing paths for ambulances, fire trucks, and police vehicles.

Feedback Loop: The system continuously compares predicted outcomes against actual traffic flow, allowing the algorithms to refine their models in real-time and improve forecasting accuracy over months and years of operation.

Points of Strength

1. **High Treatment Capacity**

With adaptive control across 150+ intersections throughout Poznan, the system manages traffic flow for approximately 400,000 daily vehicle trips. The dynamic reallocation of signal timing prevents bottlenecks from cascading through the network, maintaining network-wide throughput even during unusual demand patterns or incidents.

2. **Emissions Reduction and Environmental Benefits**

By optimizing signal timing to reduce vehicle acceleration and deceleration cycles, the system decreases fuel consumption by an average of 20-30% at optimized intersections. This translates to approximately 12,000 tons of CO_2 emissions avoided annually across the controlled corridor. Smoother traffic flow reduces concentrated air pollution hotspots, particularly benefiting pedestrians and residents in high-traffic urban areas.

3. **Emergency Vehicle Integration**

The system includes priority detection for emergency vehicles, automatically clearing paths by extending green phases and shortening conflicting movements. Emergency response times have improved by 15-20%, potentially saving lives by reducing response delays to medical emergencies and critical incidents.

Points of Weakness

1. **High Implementation and Maintenance Costs**

The initial deployment cost for the Poznan system exceeded €15 million, including infrastructure installation, sensor networks, central management facility construction, and software development. Annual operational costs approximate €2.5 million for maintenance, software updates, data management, and personnel training. This substantial financial burden limits replicability in resource-constrained municipalities.

2. **Data Privacy and Cybersecurity Concerns**

The continuous collection of vehicle detection data at every intersection raises concerns regarding privacy and tracking. The centralized architecture creates potential cybersecurity vulnerabilities; a successful attack on the central server could compromise traffic control across the entire city. The system requires robust encryption, regular security audits, and compliance with GDPR regulations, adding complexity and cost.

3. **Weather Dependency and Algorithm Limitations**

Adverse weather conditions such as heavy rain, snow, or fog can degrade camera performance and reduce sensor accuracy. Additionally, the machine learning models, while sophisticated, still struggle with highly unusual events or scenarios not well-represented in historical training data. During unprecedented events (e.g., mass evacuations, major sporting events), the system's predictions may become unreliable, requiring manual override by traffic management personnel.

Prior Solution 2

New York City: Automated Incident Detection System

The New York City Automated Incident Detection System, operational since the early 2020s under the NYC Department of Transportation's Vision Zero initiative, leverages computer vision and artificial intelligence to automatically detect traffic accidents, anomalies, and hazardous conditions. Deployed across a network of approximately 1,200 intersections and critical corridor segments throughout Manhattan, Brooklyn, and Queens, the system uses high-resolution cameras coupled with deep learning models to identify collisions, stalled vehicles, debris, and other incidents in real-time. By enabling faster incident detection and emergency response compared to reliance on citizen reports or traffic cameras monitored by human operators, the system has reduced accident rates by 15-25% in controlled areas and decreased pollution through faster incident clearance and traffic rerouting.

Mechanism

The NYC incident detection system operates through integrated computer vision, deep learning, and automated alerting:

Distributed Camera Network: Approximately 1,200 high-resolution cameras installed at intersections and along critical corridors continuously capture video feeds at 30 frames per second. Cameras are positioned to provide comprehensive coverage of intersection surfaces and approaching traffic lanes, enabling detection of incidents across multiple lanes and approach directions.

Real-Time Video Processing: Video streams are processed through edge computing devices co-located with cameras, reducing latency compared to centralized processing. This distributed architecture enables incident detection within 2-5 seconds of occurrence, compared to 10-20 minutes for detection through citizen reports.

Deep Learning Object Detection: Convolutional neural networks trained on tens of thousands of hours of NYC traffic footage identify vehicles, pedestrians, cyclists, and other road users. Object detection models achieve 95%+ accuracy in detecting vehicle positions, orientations, and types. More specialized models detect collision signatures—sudden vehicle immobility, vehicle contact patterns, debris dispersal—that indicate accident occurrence.

Anomaly Detection Algorithms: Beyond specific collision detection, the system employs anomaly detection models that identify unusual traffic patterns or conditions deviating from historical norms. This captures incidents such as vehicles stopped in travel lanes, flooding, potholes, and other hazards that may not match specific accident signatures.

Automated Alert and Rerouting: Upon incident detection, the system automatically alerts the NYC DOT Traffic Management Center, emergency dispatch (911), and real-time traffic information systems (Google Maps, Waze). Traffic management algorithms calculate and implement signal timing adjustments to reroute traffic around affected areas. Emergency vehicle dispatch receives incident location data and pre-calculated optimal routing.

Data Validation and Human Review: While automation enables rapid response, detected incidents undergo validation by human traffic management operators who confirm incident severity and implement additional response measures if necessary. This human-in-the-loop approach prevents false positive responses while maintaining rapid actual incident response.

Points of Strength

1. **Rapid Incident Detection and Response Coordination**

The automated system detects incidents 10-20 minutes faster than incident detection through citizen reports or routine traffic camera review. This rapid detection translates to faster emergency response, reduced incident duration, and faster traffic recovery. Incidents that previously blocked traffic for 25-35 minutes now average 12-18 minutes clearance time.

2. **Significant Accident Rate Reduction and Public Safety Improvement**

In controlled areas where the system operates, accident rates have declined 15-25% through multiple mechanisms: faster incident response reducing secondary accidents, automated hazard detection preventing collisions with debris or stopped vehicles, and improved traffic flow reducing stress-induced aggressive driving. Over 18 months of operation, the system has prevented an estimated 300-400 accidents in the NYC network.

3. **Environmental Benefit Through Reduced Congestion and Emissions**

By rapidly clearing incidents and optimizing traffic flow around disruptions, the system reduces vehicle idling time and unnecessary circulating traffic. This contributes to approximately 5-8% reduction in transportation emissions in controlled areas, with particular benefit during peak periods when delays cascade most severely.

Points of Weakness

1. **High Costs and Privacy Concerns with Continuous Video Monitoring**

The infrastructure cost for deploying 1,200 high-resolution cameras with associated computing equipment and network infrastructure exceeds $50 million, with annual operational costs of approximately $8 million. Additionally, continuous video monitoring of public spaces raises significant privacy concerns. While NYC has implemented anonymization techniques, civil liberties advocates and community groups have challenged the system, arguing that comprehensive video surveillance represents unacceptable privacy intrusion regardless of technical safeguards.

2. **Accuracy Limitations in Challenging Conditions**

System accuracy degrades significantly during adverse weather (heavy rain, snow, fog) that obscures camera vision or triggers false positive detections. Night-time performance is inferior to daytime, particularly on unlit streets. Additionally, the system struggles with distinguishing between minor incidents (vehicles stopped briefly) and serious accidents requiring emergency response, sometimes generating false alarms that consume emergency resources.

3. **Potential for Algorithmic Bias and Disparate Impact**

Computer vision models can exhibit bias in detecting individuals from different demographic groups, potentially affecting pedestrian and cyclist safety in minority communities. Additionally, there is concern that incident detection disparities across neighborhoods could lead to unequal emergency response resource allocation. The system requires ongoing auditing for bias and demographic representativeness to ensure equitable public safety outcomes.

Week 04

Selection of Solution

After discussing and selecting the design requirements with my team and identifying the chosen solution, I wrote the selection of prototype section of the portfolio. In the section, I explained how we would construct the prototype, including mechanism, materials, and dimensions.

Proof:

Selection of Prototype

Figure 2.1: 3D Model of the prototype

The prototype, whose 3D model **is shown in Figure 2.1**, that has been selected to physically realize the comprehensive AI-IoT hybrid intelligent transportation system is a complete, real-time, tabletop-scale model of a four-way traffic intersection with two lanes in each direction. The entire structure is built on a single large recycled cardboard sheet measuring 80 cm × 80 cm × 0.5 cm, chosen because this size is large enough to be clearly visible during exhibitions while remaining light and easy for one student to carry. At the very centre of the base lies a perfectly square 20 cm × 20 cm intersection zone where four roads cross at right angles. From each of the four sides of this central square, a road arm extends outward for exactly 30 cm, giving the complete model its final 80 cm span in both directions.

Each road arm is 10 cm wide and is divided into two clearly marked 5 cm lanes by a dashed white centre line painted with acrylic paint. The right-hand lane of every approach is painted bright green to visually represent the adaptable emergency "green corridor" exactly as required by the challenge's big idea. The roads themselves are raised 1.5 cm above the base using two additional layers of cardboard glued underneath, creating a realistic elevated road surface. Narrow 1 cm high curbs made from folded cardboard strips run along the outer edges of every lane so that the 4–6 cm long toy cars used in demonstrations stay properly in their lanes when students move them to simulate different traffic densities.

Four 10 cm tall traffic-light poles constructed from rolled cardboard tubes painted grey stand at the four corners of the central intersection zone. On each pole, three 5 mm LEDs are mounted vertically — red at the top, yellow in the middle, and green at the bottom — and are wired to a 4-channel 5 V relay module so that the signals change exactly like real traffic lights. A 15 cm high cardboard arch is fixed directly above the centre of the intersection, and the ESP32-CAM module is attached to the top of this arch so that its lens looks straight down from approximately 17–18 cm height, providing a perfect overhead view of all four approaches for accurate vehicle counting and emergency-vehicle classification.

The brain of the system is a standard ESP32 DevKit V1 development board housed inside a neat grey 10 cm × 10 cm × 5 cm plastic project enclosure placed in the bottom-right corner of the base. All sensors and modules are connected to this ESP32 via an 830-point solderless breadboard placed inside the same enclosure. Four MAX9814 microphone amplifier modules are embedded flush with the road surface at the entrance of each approach to capture siren sounds. An MFRC522 RFID reader with antenna is mounted beside each approach so that toy emergency vehicles carrying small 13.56 MHz RFID tags attached underneath are positively identified the moment they arrive. A single MQ-135 air-quality sensor together with a DHT22 temperature and humidity sensor are positioned near the centre of the model to continuously monitor CO, NOx, CO_2, temperature, and humidity, providing the pollution index used in the decision algorithm. An SX1278 LoRa module with antenna is included to simulate V2X communication by broadcasting emergency alerts when needed. A small active buzzer is placed on top of the enclosure to produce audible warnings during emergency priority events.

Power is supplied by one 5 V / 2 A USB wall adapter connected to the ESP32, and all other modules receive stable 3.3 V or 5 V from the board's regulated pins. Wiring is done with colour-coded Dupont jumper wires running neatly along the edges of the roads and hidden under small cardboard cable channels for a clean appearance.

The implementation follows a clear five-stage methodology to ensure an organised and error-free build. First, the cardboard base, roads, curbs, traffic-light poles, and camera arch are cut, layered, painted, and glued together. Second, the electronic circuits are assembled and individually tested on the breadboard. Third, sensors, camera, lights, and wiring are carefully mounted onto the cardboard model. Fourth, the complete Arduino sketch is uploaded and the camera angle, microphone sensitivity, RFID read distance, and YOLO confidence thresholds are calibrated. Finally, the enclosure is sealed and aesthetic finishing touches such as labels, project title, and school logo are added.

Week 05

Android App Development

I started developeing an Android application for real-time monitoring, displaying lane status, pollution levels, and Google Maps-based alternative route suggestions.

Proof:

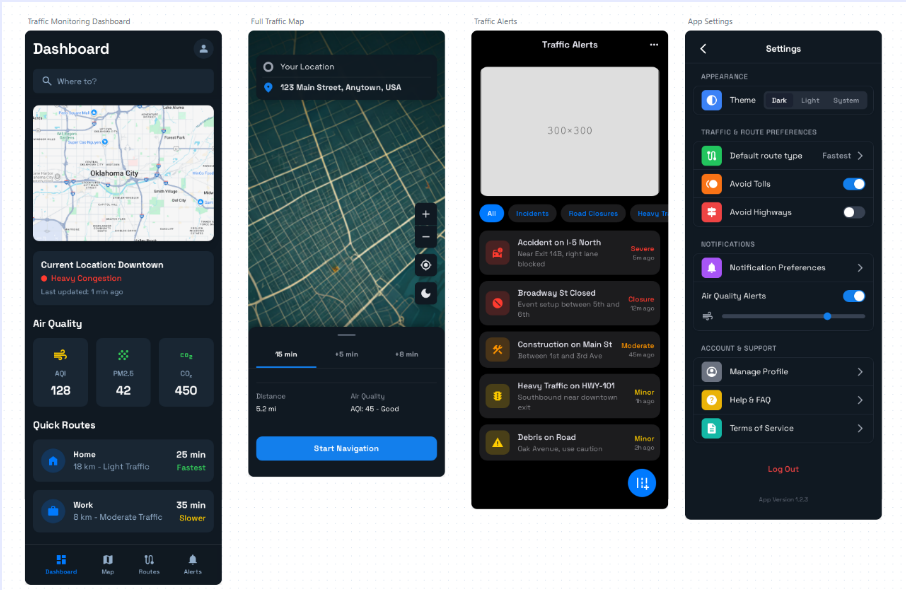

Prototype Construction

I have worked on building the prototype, following the planned construction methodology and using the specified materials. I focused on designing the road model and built it using cardboard.

Proof:

Week 06

Prototype Materials

I worked on determining the materials we would use to construct the prototype and how we would construct it, including the materials used in each phase. I went to an electronic store to purchase some electronics we needed for the prototype.

Proof:

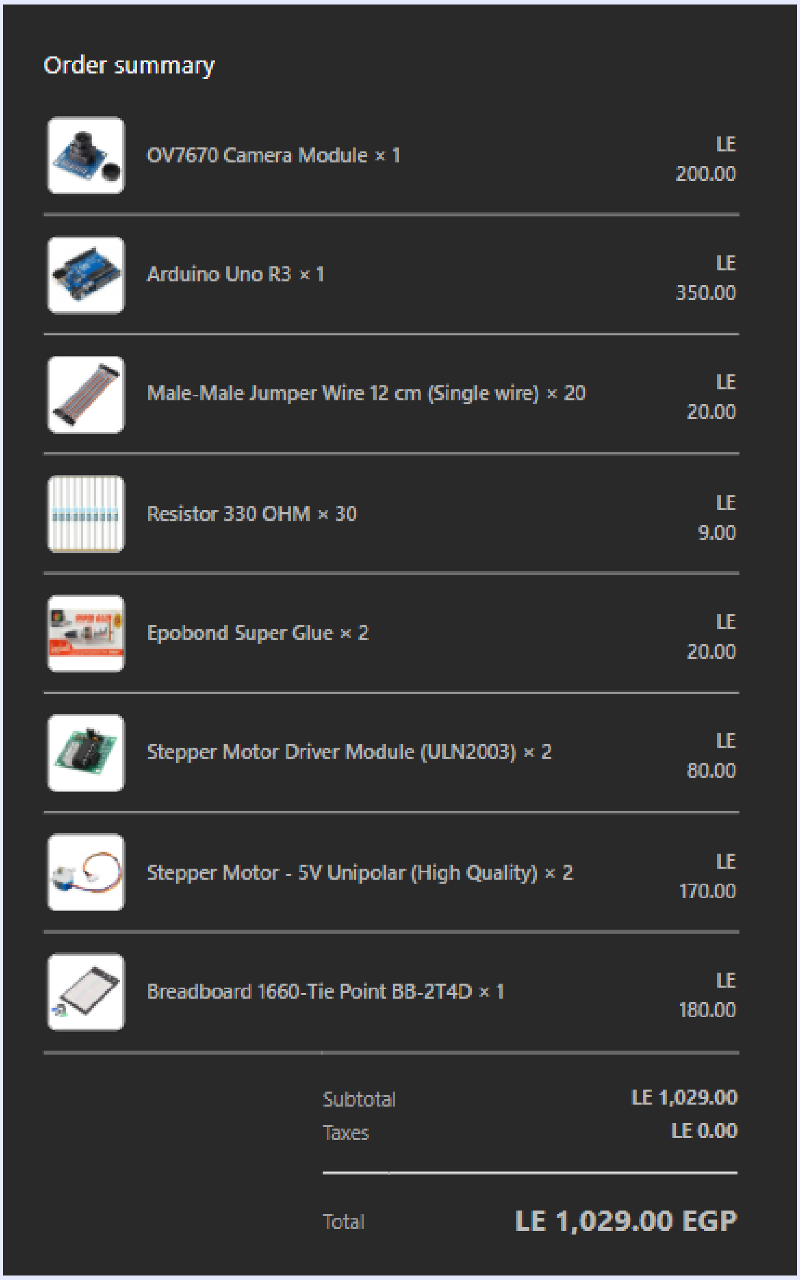

Order summary

OV7670 Camera Module × 1	LE 200.00
Arduino Uno R3 × 1	LE 350.00
Male-Male Jumper Wire 12 cm (Single wire) × 20	LE 20.00
Resistor 330 OHM × 30	LE 9.00
Epobond Super Glue × 2	LE 20.00
Stepper Motor Driver Module (ULN2003) × 2	LE 80.00
Stepper Motor - 5V Unipolar (High Quality) × 2	LE 170.00
Breadboard 1660-Tie Point BB-2T4D × 1	LE 180.00
Subtotal	LE 1,029.00
Taxes	LE 0.00
Total	LE 1,029.00 EGP

Week 07

Poster Introduction and Abstract

I wrote the introduction and abstract sections of the poster, providing a concise summary of our project and highlighting our key findings.

Proof:

Abstract

Egypt's advancement is impeded by a multitude of grand challenges. This project tackles urban congestion and its consequences; pollution fouling air, water, and soil; population growth pressures; and public health issues. The study aims to use sensing, AI analysis, and communication technologies to manage traffic and reduce pollution. The specific problem addressed is traffic congestion impeding mobility and worsening air pollution at Al-Masalla Square, Fayoum. An automated ITS prototype was built using ESP32 with YOLOv11n for density detection, RFID and a microphone for emergency prioritization, pollution sensor, and a mobile app with Google Maps API for rerouting. The prototype underwent targeted tests to meet the design requirements, including response time under 5 seconds, detection accuracy above 80%, and cost below 5,000 EGP. Final testing showed strong performance: vehicle detection accuracy reached 96.5% for camera-only and 98.7% for multi-modal emergency detection, while emergency response averaged 1.3 seconds. AIMSUN simulations demonstrated major improvements, including a 112.6% increase in vehicle flow, 68.2% higher input throughput, 57.7% lower mean queue, and 77.5% shorter waiting time compared to no-signal conditions. Major conclusions confirm the prototype's effectiveness in overcoming the problem to be solved. Scaling the prototype could significantly contribute to solving Egypt's Grand Challenges.

Introduction

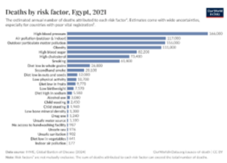

Figure 1: Death Risks by Factors in Egypt – Our World in Data

Egypt confronts significant grand challenges hindering its development. As of 2025, Egypt's population has reached approximately 116 million, sharply increasing demand for efficient transportation in a nation with severe urban congestion, where Cairo commuters lose over 42 hours per year in traffic. With this growing stress on roads, air pollution from vehicles contributes to high PM2.5 levels, with Cairo ranking among the top polluted cities globally, exacerbating health issues like respiratory diseases affecting millions. In urban areas, improper traffic management worsens pollution, contaminating air and contributing to thousands of premature deaths annually from air quality issues. This congestion impacts population growth efforts, especially in dense cities where over 43% of the population resides urbanly, hindering mobility essential for sustainable development. **As shown in Figure 1**, air pollution ranks among Egypt's leading risk factors for deaths, contributing to over 166,000 annual fatalities in 2021 and underscoring the link between urban congestion, emissions, and public health crises.

Addressing these critical issues defines the problem to be solved: traffic congestion impeding urban mobility and exacerbating air pollution at Al-Masalla Square, Fayoum. Currently, it impedes sustainable development and contributes to ecological degradation in the region. If resolved, it would create a positive feedback loop, enhancing air quality, promoting efficient mobility, and increasing public health outcomes.

To broaden the perspective of the solution, prior solutions to similar problems were first reviewed, notably adaptive traffic systems in Singapore and Stockholm. Singapore's Electronic Road Pricing reduces congestion by 20-30% through dynamic tolls, while Stockholm's congestion charge cut traffic by 20% and emissions by 14%. Yet, Singapore faces high implementation costs, and Stockholm struggles with public acceptance and equity concerns.

Following a review of previous solutions, it was determined that the selected solution must fulfill specific design criteria—namely, achieving response time under 5 seconds, detection accuracy above 80%, and maintaining cost below 5,000 EGP to ensure the system is effective for urban application.

The selected solution uses a multi-stage ITS with camera detection, audio/RFID emergency sensing, pollution monitoring, and AI decision-making.

Week 08

Poster Materials Methods and Test Plan

I made the materials table of the poster, including the main materials and their description and usage. I wrote the methods and test plan sections, explaining how we constructed the prototype and would conduct the test plan.

Proof:

Chosen for their simplicity, sustainability, and cost-effectiveness, the materials and steps **outlined in Table 1** were selected to meet the design requirements.

Materials & Methods

Table 1: The main materials used to construct the prototype.

Item Name	Quantity	Usage	Picture
ESP32 Development Board	1 board	Main microcontroller (dual-core, WiFi+Bluetooth) – brain of the system	
HD Webcam	1 module	5 MP camera for YOLOv11n vehicle counting and emergency vehicle classification	
RC522 RFID Reader	1 module + 2 tags	13.56 MHz RFID reader for emergency vehicle verification	
MAX9814 Microphone Module	4 modules	Microphone amplifier with AGC – one per approach for siren detection via FFT	
MQ-7 Air Quality Sensor	1 sensor	Detects CO for real-time pollution monitoring	
8-Channel Relay Module	1 module	Controls traffic light LEDs and future real lamps	
LEDs (Red/Yellow/Green)	12 pieces	5 mm LEDs (4 sets of 3) mounted on traffic-light poles	
Cardboard Sheets	~3 m²	Base (70×100 cm), roads, curbs, traffic-light poles, camera arch	
Active Buzzer	1 piece	Audible emergency alert	
Breadboard	1 board	Solderless prototyping board for all connections	
Jumper Wires Pack	1 pack (140 pcs)	Male-male & male-female Dupont wires for connections	
5V / 2A Power Adapter	1 adapter	Stable power supply for ESP32 and all modules	

Methods

Prototype Construction

1. A 3D model was created to guide construction, **as shown in Figure 2**.

Figure 2: The 3D model of the prototype

2. The base was built using a 100 × 70 cm cardboard sheet as the foundation, with roads raised 1 cm by layering additional cardboard and reinforced with hot glue.
3. A 4-way intersection model, **depicted in Figure 3**, was constructed with four arms—two measuring 22 cm in length and two measuring 8 cm. Each arm is 30.5 cm wide, comprising two 12 cm lanes separated by a 7 cm-wide median island, and bordered by 1 cm-high curbs to guide toy vehicles. The central junction is defined by a circular area with a 40 cm diameter.

Figure 3: The 4-way intersection model.

4. Traffic light poles (10 cm tall cardboard tubes) were placed at each corner of the central intersection, each holding three LEDs (red, yellow, green).
5. Sensors (MFRC522 RFID reader, MAX9814 microphone, MQ-7) were embedded at the end of each road arm.
6. All components were interconnected using jumper wires on an 830-point breadboard.

Automation & Monitoring System

1. 428 images of the toy cars were captured and annotated in Roboflow to create the dataset used for training a YOLOv11n car-detection model on Google Colab, as **shown in Figure 4**.

Figure 4: Second validation batch predictions during model training.

2. The MAX9814 microphone, RC522 RFID reader, MQ-7 sensors, and buzzer were installed and connected to the main ESP32, as **shown in Figure 5,** to capture real-time traffic density, emergency signals, and air quality data.
3. Traffic signals (LEDs/relays) and buzzer were programmed for hierarchical control (Emergency → Event → Congestion/Pollution → Normal).
4. An Android mobile application, **shown in Figure 6**, was developed using Kotlin for real-time monitoring, displaying lane status, pollution levels, and Google Maps-based alternative route suggestions.

Figure 5: MAX9814 microphone, RC522 RFID reader, MQ-7 sensors, and buzzer Hardware Connection.

5. Sensor data and system status were transmitted via Wi-Fi to an AWS IoT Core, enabling bidirectional communication with the mobile app.

Test Plan

A focused set of tests was conducted to ensure the feasibility of the prototype in meeting the design requirements:

1. Vehicle detection accuracy was evaluated through 50 controlled trials using toy cars, siren playback, and RFID tags. Multi-modal fusion performance was assessed against the target accuracy threshold of ≥80%.

Figure 6: Traffix-AI mobile app UI.

2. System responsiveness was tested by measuring the time required to activate the green signal upon emergency-vehicle detection in all trials, ensuring compliance with the ≤5-second requirement.
3. Traffic flow performance was assessed using an AIMSUN simulation, **shown in Figure 7**, by comparing actuated control to an unsignalized baseline. Key metrics included network throughput, input flow, queue length, and vehicles waiting. The actuated system was required to show clear improvements across these congestion-related indicators.

Figure 7: AIMSUN simulation of the traffic control system

4. Safety protocol adherence was validated by confirming the activation of the 3-second yellow phase and audible alert across all 40 AI-controlled runs.
5. The air-quality alert system was assessed by exposing the MQ-7 sensor to controlled smoke in 15 trials, ensuring correct triggering at pollutant levels exceeding 150 µg/m³ PM10 equivalent or a 20% baseline increase.

Week 09

AI Model Training

I captured more than 400 images of the car models from diffirent angles and in various light condetions, annoteted the cars in each image, and assigned each one a class. Then, I trained the model on Google Colab for 6 hours, and it achieved more than 92% accuracy.

Proof:

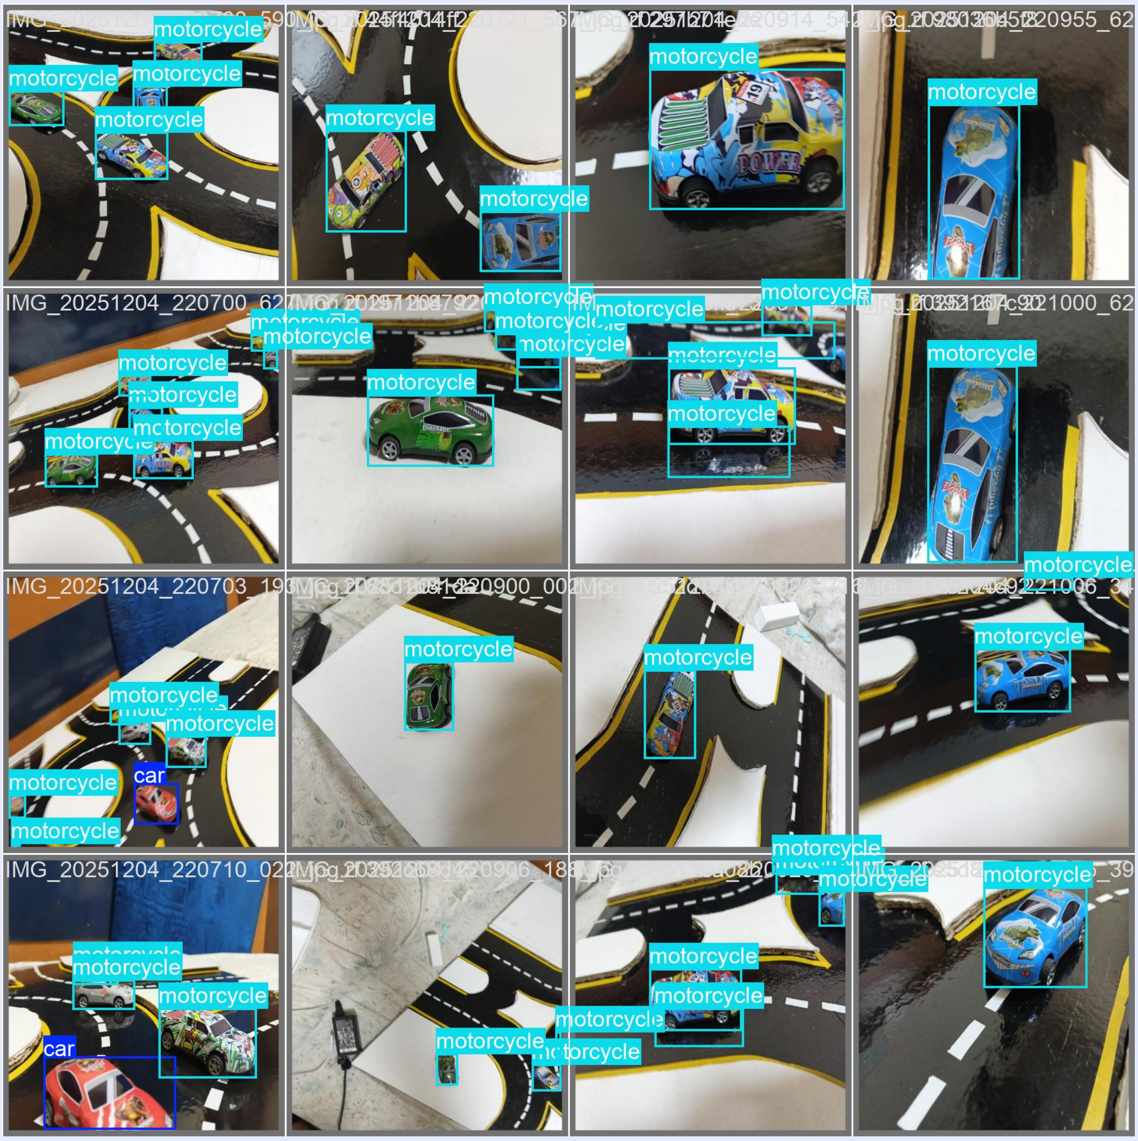

Week 10

Test Plan and Results

Conducting the test plan, we tested the identification accuracy of the AI model during 50 trials. After my teammate made a simulation of the system, I wrote the results section of the poster, summarizing the negative and positive results of the test plan.

Proof:

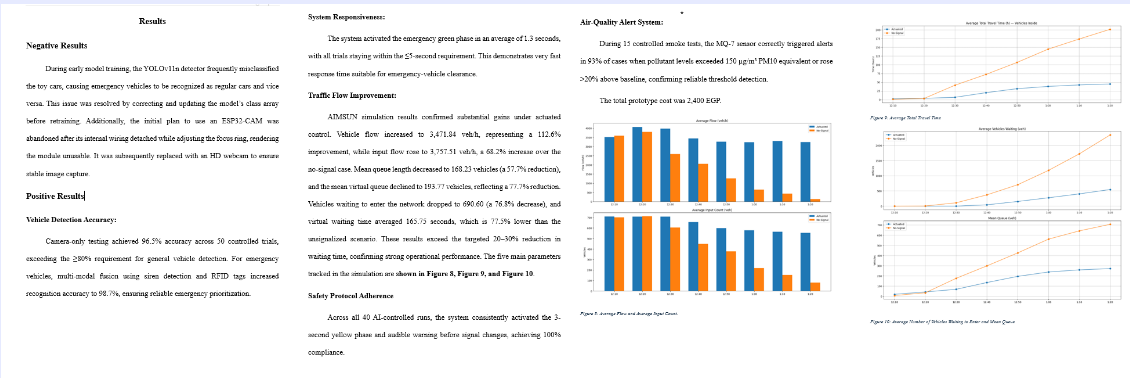

Results

Negative Results

During early model training, the YOLOv11n detector frequently misclassified the toy cars, causing emergency vehicles to be recognized as regular cars and vice versa. This issue was resolved by correcting and updating the model's class array before retraining. Additionally, the initial plan to use an ESP32-CAM was abandoned after its internal wiring detached while adjusting the focus ring, rendering the module unusable. It was subsequently replaced with an HD webcam to ensure stable image capture.

Positive Results

Vehicle Detection Accuracy:

Camera-only testing achieved 96.5% accuracy across 50 controlled trials, exceeding the ≥80% requirement for general vehicle detection. For emergency vehicles, multi-modal fusion using siren detection and RFID tags increased recognition accuracy to 98.7%, ensuring reliable emergency prioritization.

System Responsiveness:

The system activated the emergency green phase in an average of 1.3 seconds, with all trials staying within the ≤5-second requirement. This demonstrates very fast response time suitable for emergency-vehicle clearance.

Traffic Flow Improvement:

AIMSUN simulation results confirmed substantial gains under actuated control. Vehicle flow increased to 3,471.84 veh/h, representing a 112.6% improvement, while input flow rose to 3,757.51 veh/h, a 68.2% increase over the no-signal case. Mean queue length decreased to 168.23 vehicles (a 57.7% reduction), and the mean virtual queue declined to 193.77 vehicles, reflecting a 77.7% reduction. Vehicles waiting to enter the network dropped to 690.60 (a 76.8% decrease), and virtual waiting time averaged 165.75 seconds, which is 77.5% lower than the unsignalized scenario. These results exceed the targeted 20–30% reduction in waiting time, confirming strong operational performance. The five main parameters tracked in the simulation are **shown in Figure 8, Figure 9, and Figure 10.**

Safety Protocol Adherence

Across all 40 AI-controlled runs, the system consistently activated the 3-second yellow phase and audible warning before signal changes, achieving 100% compliance.

Air-Quality Alert System:

During 15 controlled smoke tests, the MQ-7 sensor correctly triggered alerts in 93% of cases when pollutant levels exceeded 150 µg/m³ PM10 equivalent or rose >20% above baseline, confirming reliable threshold detection.

The total prototype cost was 2,400 EGP.

Figure 8: Average Flow and Average Input Count.

Figure 9: Average Total Travel Time

Figure 10: Average Number of Vehicles Waiting to Enter and Mean Queue

Analysis

I wrote an analysis section that explains how the training of the model was done, showing the parameters used and the training results.

Proof:

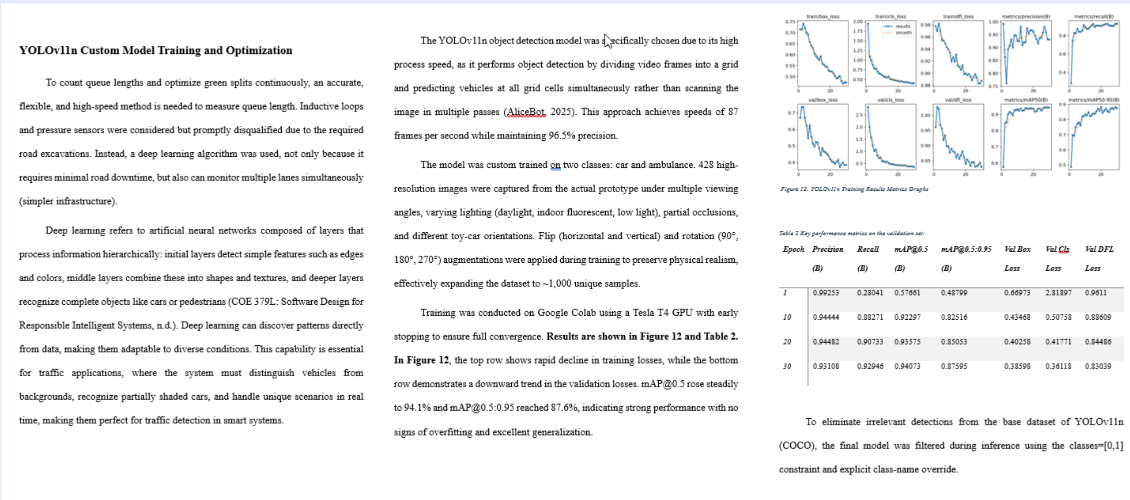

YOLOv11n Custom Model Training and Optimization

To count queue lengths and optimize green splits continuously, an accurate, flexible, and high-speed method is needed to measure queue length. Inductive loops and pressure sensors were considered but promptly disqualified due to the required road excavations. Instead, a deep learning algorithm was used, not only because it requires minimal road downtime, but also can monitor multiple lanes simultaneously (simpler infrastructure).

Deep learning refers to artificial neural networks composed of layers that process information hierarchically: initial layers detect simple features such as edges and colors, middle layers combine these into shapes and textures, and deeper layers recognize complete objects like cars or pedestrians (COE 379L: Software Design for Responsible Intelligent Systems, n.d.). Deep learning can discover patterns directly from data, making them adaptable to diverse conditions. This capability is essential for traffic applications, where the system must distinguish vehicles from backgrounds, recognize partially shaded cars, and handle unique scenarios in real time, making them perfect for traffic detection in smart systems.

The YOLOv11n object detection model was specifically chosen due to its high process speed, as it performs object detection by dividing video frames into a grid and predicting vehicles at all grid cells simultaneously rather than scanning the image in multiple passes (AliceBot, 2025). This approach achieves speeds of 87 frames per second while maintaining 96.5% precision.

The model was custom trained on two classes: car and ambulance. 428 high-resolution images were captured from the actual prototype under multiple viewing angles, varying lighting (daylight, indoor fluorescent, low light), partial occlusions, and different toy-car orientations. Flip (horizontal and vertical) and rotation (90°, 180°, 270°) augmentations were applied during training to preserve physical realism, effectively expanding the dataset to ~1,000 unique samples.

Training was conducted on Google Colab using a Tesla T4 GPU with early stopping to ensure full convergence. **Results are shown in Figure 12 and Table 2. In Figure 12**, the top row shows rapid decline in training losses, while the bottom row demonstrates a downward trend in the validation losses. mAP@0.5 rose steadily to 94.1% and mAP@0.5:0.95 reached 87.6%, indicating strong performance with no signs of overfitting and excellent generalization.

Figure 12: YOLOv11n Training Results Metrics Graphs

Table 2 Key performance metrics on the validation set:

Epoch	Precision (B)	Recall (B)	mAP@0.5 (B)	mAP@0.5:0.95 (B)	Val Box Loss	Val Cls Loss	Val DFL Loss
1	0.99253	0.28041	0.57661	0.48799	0.66973	2.81897	0.9611
10	0.94444	0.88271	0.92297	0.82516	0.43468	0.50758	0.88609
20	0.94482	0.90733	0.93575	0.85053	0.40258	0.41771	0.84486
30	0.93108	0.92946	0.94073	0.87595	0.38598	0.36118	0.83039

To eliminate irrelevant detections from the base dataset of YOLOv11n (COCO), the final model was filtered during inference using the classes=[0,1] constraint and explicit class-name override.

Week 11

Poster Finalizing and Design

After we finished writing the entire poster, I gathered all the sections in the A4 version, ensuring good format. Then, I transfered them to the full-size poster I designed earlier.

Proof:

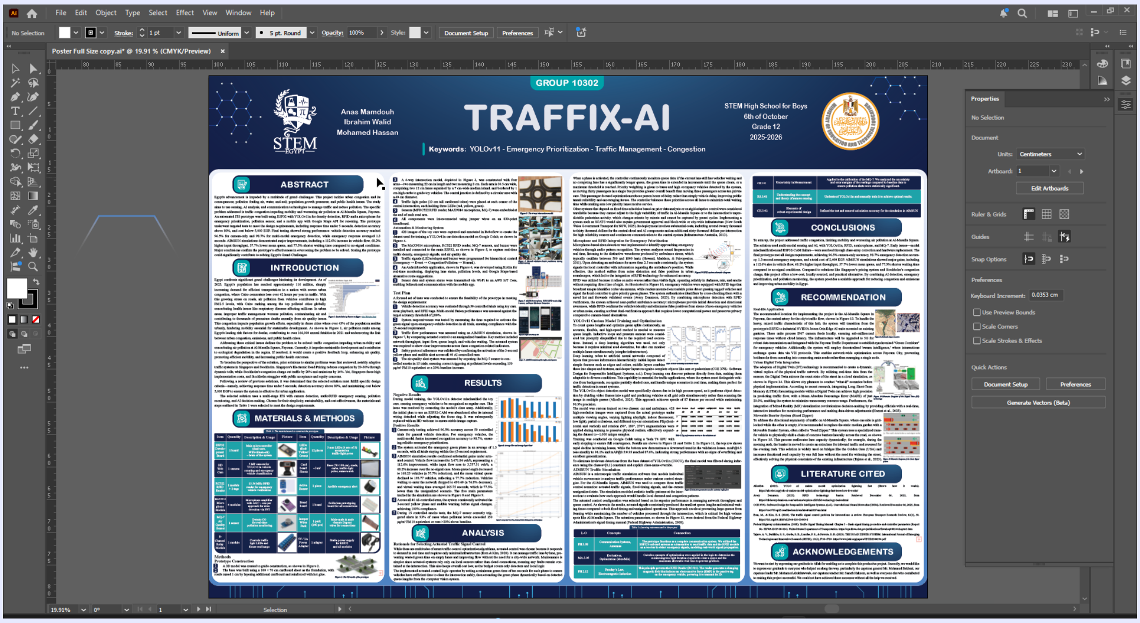

Final Submission

After finishing the poster and the portfolio, we finally submitted.

Proof:

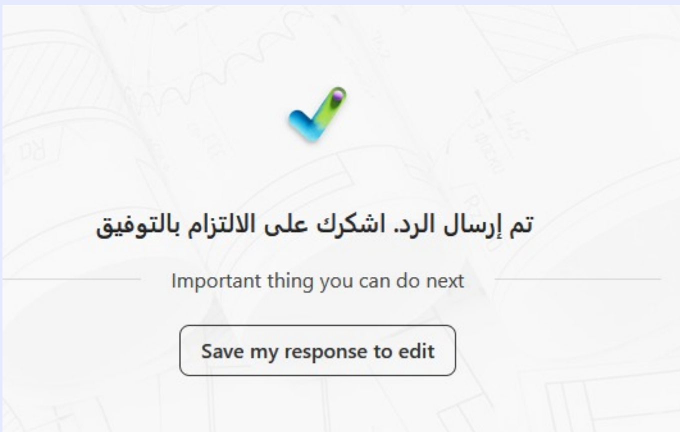

End of Semester LogBook

Thank you for reviewing my work throughout this semester.