



Laidlaw Scholars Undergraduate Leadership and Research Programme
Leadership-in-Action (LiA) Project Report

Building Decision-Support for Migration Policy

Six weeks with the Global Initiative Against Transnational Organized Crime

Vedant Harlalka

University of Toronto

LiA Supervisor: Dr. Alex Bish, Head of Research Innovation,
North Africa and Sahel Observatory, GI-TOC

Submission date: 1 September 2026

Project Details

Field	Detail
Scholar	Vedant Harlalka, University of Toronto
Host organisation	Global Initiative Against Transnational Organized Crime (GI-TOC), North Africa and Sahel Observatory
Supervisor	Dr. Alex Bish, Head of Research Innovation
Dates	6 July – 15 August 2026
Location	Remote from Toronto (6 – 28 July); GI-TOC Malta office, Sliema (28 July – 15 August)
LiA option	Independent Leadership-in-Action programme
Project title	AI for Sahel Communities: Understanding and Responding to Smuggling and Trafficking
Focus after scoping	Human smuggling on Libyan routes; decision-support for GI-TOC field researchers
SDG alignment	SDG 16 (Peace, Justice and Strong Institutions); SDG 10 (Reduced Inequalities)

Introduction

I am a second-year electrical engineering student at the University of Toronto, and most of the software I had written before this summer was commercial: you build something, someone pays for it, and that transaction tells you whether it worked. This project was the first time I built something where success is measured by whether it leaves the people it describes better off.

GI-TOC is an independent network that researches organised crime and illicit economies, and its North Africa and Sahel Observatory produces field reporting on smuggling routes. Its approach is evidence-led and harm-reduction-led rather than enforcement-led, and that distinction shaped everything I built. My role was to design a decision-support tool that could take fragmented field reporting and help a researcher think through the first and second-order effects of an intervention before it happens. I worked remotely from Toronto from 6 to 28 July, then on-site in Sliema, Malta, from 28 July to 15 August, and delivered two working prototypes: an intervention simulator called InterventionAI, and a retrieval tool for the organisation's archive that the team came to call the GI-TOC Librarian.

Three themes shaped the summer more than anything I expected to write about. The first is that narrowing scope is a leadership act rather than an admission of failure. The second is that ethical constraints written down early are the ones that survive pressure later. The third is that a tool like this only earns its place if the people whose lives it describes stay at the centre of it, and that keeping them there is a matter of who you talk to as much as what you code.



First morning at the GI-TOC office on Bisazza Street, Sliema.

Context: The Problem and the People It Affects

Libya is a central node in what GI-TOC calls organised immigration crime, and the Libyan authorities, the EU and a range of donors are all trying to reduce irregular migration through it. In early scoping we narrowed the project to Libyan hotspots and, within that, to crime scripts for human smuggling only. Four structural features make interventions here go wrong in ways that are predictable in hindsight and rarely predicted in advance:

- **Fragmented authority.** Control is dispersed across state, militia and criminal actors, so there is rarely one lever to pull.
- **Crime adapts quickly.** Suppression tends to displace activity rather than remove it.
- **Migration and livelihood are coupled.** Migration-related income is embedded in local economies, so disrupting a route disrupts households.
- **Effects unfold over time.** Much of the harm surfaces later, in the lag between an action and its consequence.

The affected community is not an abstraction. It is people moving north through Kufra and Sebha toward the coast, many displaced by the conflict in Sudan; migrants held in detention; and Libyan households in towns along the route whose income depends on movement through their area. UNHCR recorded 658,477 refugees and asylum-seekers in Libya in 2025, and IOM's Displacement Tracking Matrix recorded 147,382 internally displaced people as of 31 May 2024. Both figures sit in the header of the tool, so that anyone using it sees the scale of the affected population before anything else.

Malta is the other half of that picture, and living there changed how I understood it. The island sits directly on the Central Mediterranean route, and the arrival statistics in a policy paper are also somebody's neighbourhood and somebody's argument at a dinner table. The office is in Sliema for a reason.



Valletta, looking north across the harbour. Malta sits directly on the Central Mediterranean route.

Listening First: The People Behind the Project

The most useful decision I made all summer was to spend my first week in Malta talking rather than coding, under a rule borrowed from the Mom Test: never describe the tool, never ask whether someone would use it, and never ask a question beginning with “would you”. Ask what they did last week, and what took the longest.

Inside the office that meant structured conversations with Alex, Matt, and Tasnim about how they really work — where a case starts, which sources they trust on which subject, what slows them down. Almost nobody described their problem as prediction. What came up repeatedly was time lost to finding things: the report from eighteen months ago with the detail they now needed, the thread where a colleague had already worked out who controlled which stretch of road. That single finding redirected the project and produced the second prototype.

I deliberately did not stop at the office door. I spoke with a historian in Mdina, who walked me through how the island’s fortifications were built in response to arrivals from across the same stretch of water GI-TOC now studies, and who was blunt about how long Malta has argued with itself about who belongs. I had a version of the same conversation with a guide at Fort St Elmo, and others in cafes, on buses and at the dive school in Sliema, with people whose views on migration were more textured and more generous than anything in the policy literature.

None of that is data, and I have not treated it as such. What it produced was calibration. It is easy, building software about migration from a library in Toronto, to hold the subject at a comfortable distance, and much harder once you have argued about it with someone whose town is on the map. Community engagement in this project meant letting those conversations sit in front of the design decisions rather than behind them, and routing migrant voice itself through material GI-TOC had already gathered under proper ethical protocols, including a set of 71 interviews with migrants the team had conducted.



Working through the smuggling question with the team in Malta.

What I Built

1. InterventionAI — a scenario check for proposed interventions

The tool runs in four steps. A researcher describes a situation in plain language, checks the machine’s understanding of it, chooses a response from a ranked set, and sees what is likely to follow. Every claim links back to its source in one click, and the footer reads “decision-support only”.

Input can be a PDF, a report, or two sentences. The case I used throughout development was: “Boats are departing the Tobruk–Derna coastline toward Crete, fed by Sudanese arrivals moving north through Kufra.” Live UNHCR and IOM figures sit across the top, so the situation is framed against current data rather than the model’s memory.

A screenshot of the InterventionAI Scenario Check web application. The interface has a light beige background. At the top left, it says "InterventionAI SCENARIO CHECK" and at the top right "V1 WORKSPACE". The main heading is "What is happening?". Below this is a progress bar with four steps: "Step 1 - Case" (highlighted with a red circle), "Step 2 - ACCP", "Step 3 - Response", and "Step 4 - Impact". Below the progress bar, there are two live data feeds: "LIVE UNHCR - refugees & asylum-seekers in Libya (2025): 650,477" and "LIVE IOM DTM - IDPs in Libya (2024-05-31): 147,382". To the right of these is a button labeled "HEURISTIC BANK" and a status indicator "ACLED - offline". Below the data feeds is a large text input area with the placeholder "Describe the situation — who is moving, where, what changed...". Below the input area is a button labeled "Analyse". Below the input area, there is a section titled "PARSED & MATCHED AGAINST THE EVIDENCE BASE" which contains three bullet points: "Boats are departing the Tobruk-Derna coastline toward Crete, fed by Sudanese arrivals moving north through ...", "Smuggling fees from Sebha to the coast have doubled in three months: migrants take on debt and forced-labo...", and "A single armed group has taken control of the southern corridor and taxes every convoy at its checkpoints." At the bottom, there is a link "OR START FROM A DOCUMENTED CASE". At the very bottom, a small footer reads "Decision-support only - every claim traces to a cited source".

Figure 1. Step 1 — case entry, with live UNHCR and IOM figures for Libya above the input.

The system then performs an ACCP analysis, breaking the case into Actors, Context, Content and Processes. The framework comes from ETH Zurich’s conflict-analysis material, chosen because it was the simplest thing a person could verify by hand; before writing any code I worked through two GI-TOC case studies on paper with sticky notes, and the card board is that exercise turned into software. Elements are tagged as

connectors or dividers, following Mary Anderson’s Do No Harm framework, and nothing proceeds until a human confirms or edits the cards.

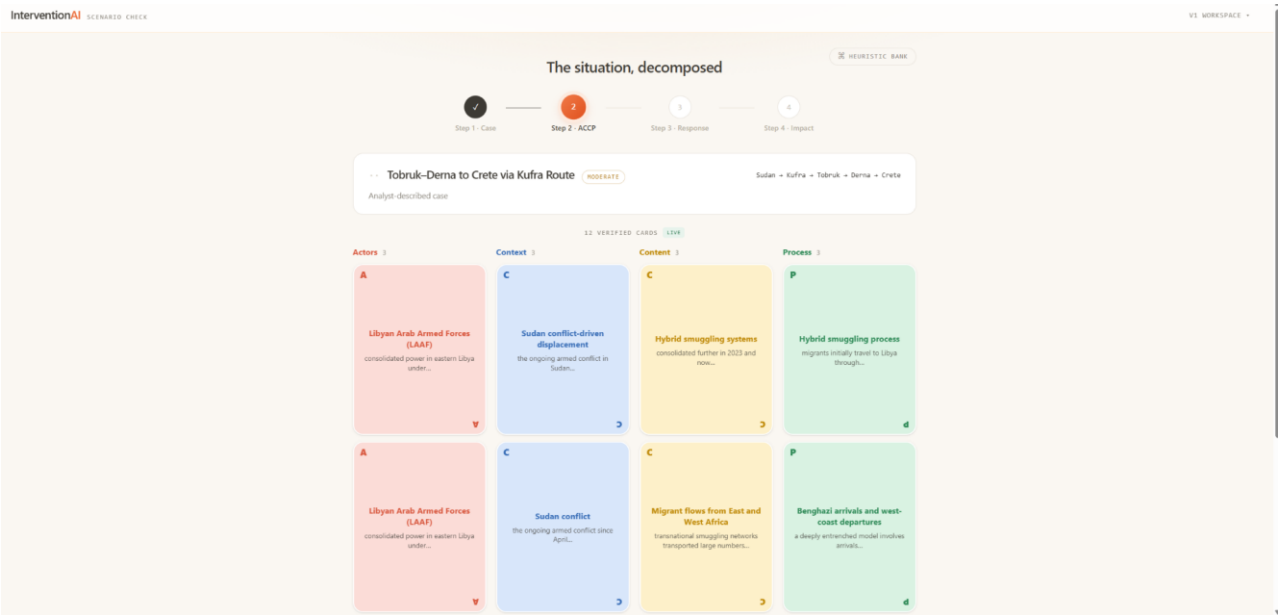


Figure 2. Step 2 — twelve verified ACCP cards for the Tobruk–Derna to Crete route via Kufra.

Seven intervention families are then ranked against the case, from security and law enforcement through to migrant protection and legal pathways, each carrying a feasibility, effectiveness and ethics score. The researcher chooses. The tool does not choose for them.




Figure 3. Step 3 — ranked intervention families with feasibility, effectiveness and ethics indicators.

The final screen is the one the whole project exists for. For the top-ranked option here, enforcement at medium intensity, the tool returns feasibility 3, effectiveness 3 and ethics 2 out of 5, and raises a gate warning that enforcement at this intensity increases migrant-harm exposure and should be paired with protection measures.

The two columns underneath are the argument of the project in one image. In the first year, irregular departures fall moderately, which is what the intervention is meant to achieve and what would be reported as success. In the same column, migrant harm rises, smuggling prices rise, and the detention population

risers. Three to five years out, the fall in departures has gone and the harm has not, while network strength has risen slightly. The benefit expires and the costs do not, and those costs land on the people described earlier. The displacement panel names where activity moves rather than stops, the western coast around Zawiya and Zuwara, overland toward Salloum and Egypt, adjacent hubs, lower-signature methods, and a divider flag notes that security support to contested eastern actors would entrench their position.

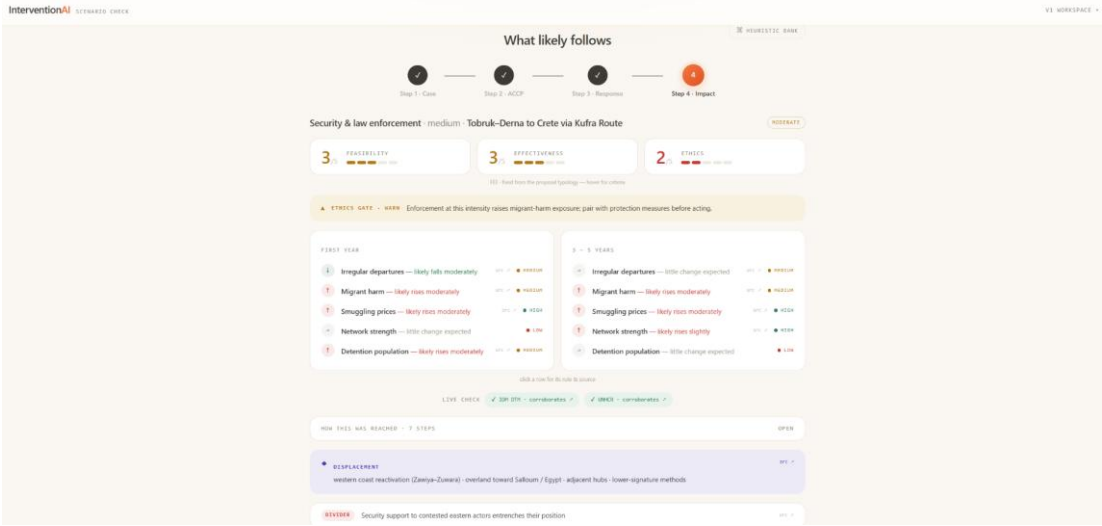


Figure 4. Step 4 — first-year and three-to-five-year effects, with the ethics gate warning and displacement note.

An instruction from my supervisor stayed with me all summer: there are no rules, only assumptions about how the world works, so state your assumptions. The reasoning trace is my attempt at that. Any output opens to show the seven steps that produced it, down to which rules fired and from which report — and where a number is not honestly computed, it says so.

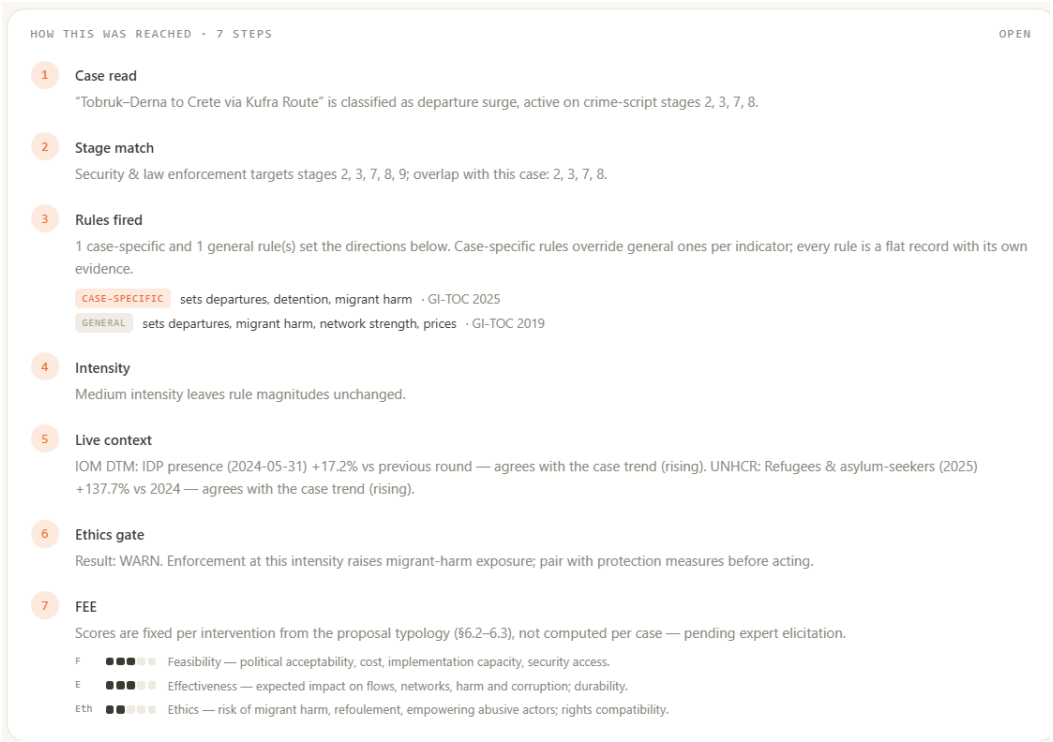


Figure 5. The reasoning trace. Every claim is traceable to a rule and a cited source.




Figure 6. The situational picture, showing hotspots and route connections across the Libyan corridor.

2. The GI-TOC Librarian — an unplanned second project

The interviews surfaced a problem the office genuinely had: time lost searching an archive of roughly two to three terabytes. This was not in my outline. I volunteered because it was real, and because it was something I could give back to people who had been giving me their time. I produced an options assessment covering Onyx Cloud, Google Vertex AI Search, Glean and Azure OpenAI, weighing retrieval quality, OCR, cost and implementation difficulty, and recommended building the Onyx MVP first, running user acceptance testing to confirm the need, and only then assessing enterprise vendors.

The privacy analysis mattered more than the technical comparison. The corpus holds open-source material but also confidential data, including witness and informant information, and every credible provider on my shortlist is US-owned and subject to subpoena. That is not a footnote for an organisation whose sources have taken personal risk to speak to researchers. I recommended that open-source material could reasonably move to cloud retrieval while sensitive holdings need a separate answer, and flagged the prototype's limitations plainly on handover: no OCR, a workaround for login, difficulty packaging it as an agent, and unstructured testing.

Progress Against the Three SMART Goals

The project outline set three goals. Two were met and one was partly met.

Goal	Planned	Outcome
1. Data and insight generation	Aggregate field data and conduct 5–10 structured interviews with analysts, policymakers and community-informed stakeholders.	Met. Structured interviews with the GI-TOC research team in Malta, supported by conversations across the island — a historian in Mdina, a guide at Fort St Elmo, and residents in Sliema and Valletta. Four sources integrated: UNHCR, IOM DTM, ACLED and GI-TOC Libya reporting.
2. Framework usefulness	Structure data into system maps and generate recommendations validated with analysts.	Met. Twelve verified ACCP cards, seven ranked intervention families and a full reasoning trace, reviewed by my supervisor and revised across

Goal	Planned	Outcome
		three iterations. Source-grounding is close to complete across the interface.
3. Real-world impact (pilot test)	Apply and test recommendations in real scenarios, with qualitative evidence of improved decision-making.	Partly met. The tool was applied to a documented Libyan case and reviewed by GI-TOC staff, and the method held up. Six weeks was not long enough to evidence a change in a live decision, and independent testing surfaced deployment issues that remain on the list.

On the third goal I decided not to round up. When the prototypes were tested independently they hit a session error and a loading problem I had not seen on my own machine. The fair version is that the method works and has been reviewed by the people it was built for, and that making it run reliably in someone else's hands is the next piece of work.

Challenges Faced

Scope, and the reason it kept expanding

My original outline covered the North Sahel broadly and proposed to validate recommendations in practice. Reading it back, ambition was doing work a decision should have been doing. Partway through I emailed Alex with two extra ideas on top of the agreed project, and his reply contained an observation I have thought about since: there is a risk of spreading yourself too thin across multiple ideas, and it is better to focus on one priority and execute it well. Narrowing then happened in stages, from the Sahel to Libya, from all illicit flows to human smuggling, from prediction to orientation, and each cut made the project better.

Building before listening

I spent three weeks in Toronto building before speaking to a single end user, and the tool that emerged was aimed at a problem the researchers did not have. The ordering was backwards because building is comfortable for me and asking people questions is not. Fixing that ordering in Malta is what rescued the project.

Partway through I also wrote a candid assessment of my own work. The FEE scoring felt arbitrary and needed a real method underneath, such as multi-criteria decision analysis, rather than a model rating things on a scale. The impact page, which projects five years ahead, was really matching the current case to the nearest of roughly eighteen GI-TOC reports and presenting that as precedent; it could not triangulate across sources, because that requires knowing the weighting between them and the method a researcher uses in their head. The output as a whole, I wrote, felt more like a librarian than a seasoned analyst. I meant that as the central criticism. It turned out to be the finding, and it became the second prototype.

Two broader limits are worth recording. The grounded context available to the system is well short of what an informed judgement requires, and closing that gap is a large collection problem in itself. And this sits in the same territory as organisations that have spent decades modelling how one variable propagates through a political or economic system, so competing on general capability is not realistic. A tool like this earns its place by being scoped down far enough to solve something specific, which is the direction the project took.

Leadership Skills Applied and Developed

Communication

The most reliable thing I did all summer was turn up prepared. Written documents, a diagram and a list of open questions consistently produced better guidance than turning up with enthusiasm, and the sharpest feedback always came in response to my most specific documents. Being concrete is a way of handing someone else the ability to tell you that you are wrong. The harder lesson was volunteering problems rather than waiting until they were solved: partway through I sent a message admitting I had not reached a point I was happy to present, and the working relationship became easier immediately.

Problem-solving and choosing the right problem

The most valuable technical decision I made was not technical. Holding to the interview discipline — asking what people did rather than what they wanted — redirected the project and produced the Librarian. I did break my own rule near the end of that week, opened the laptop and showed the prototype, and the conversation shifted from what they do to whether my thing was good. The quality of the information dropped in real time, which was a useful thing to feel.

Cultural humility

I arrived assuming the work in front of me was a modelling problem that happened to be set in North Africa, and that with enough data an experienced regional researcher's reasoning could be reconstructed. Time in the office made clear how wrong that was. What the researchers hold is not a larger dataset. It is judgement built from years of relationships, language, political context and knowing which sources to trust on which subject.

The moment I keep returning to is one where I proposed a private site letting local intervention groups in Malta see the roll-off effects of interventions, and suggested it might help inform their questioning of migrants. I wrote it quickly and thought it was helpful. Alex's response was that migration to Malta now involves very small numbers and, more importantly, that supporting security forces in this context carries significant ethical considerations that are not obvious from the outside. He was right. I had proposed handing an enforcement actor a tool for questioning vulnerable people and framed it as an improvement, while my own requirements document already said enforcement actors are represented by the model, not served by it. There was nothing malicious in the idea, which is what makes it worth writing down: distance is what makes well-intentioned proposals dangerous in this field. What I practised afterwards was slower, listening more than proposing, and hearing "we have ongoing relationships here" as what it is, which is people's safety.



Mdina. Conversations here about how Malta has long argued with itself over arrivals reshaped how I framed the project.

Ethical Considerations

The requirements document I wrote in Week 3 fixed four constraints: the tool is pre-intervention only and never an operational or real-time targeting system; there is no black-box output, and every result traces to source data and an explicit line of reasoning; a human validates all outputs before use; and behaviour derives from field data and historical precedent rather than intuition.

Writing those down while nobody was pushing against them was the most useful ethical decision of the summer, because pressure arrived later and from an unexpected direction. While struggling to justify the tool's value, I proposed a use case aimed at people who make a high volume of decisions quickly. I was not thinking about enforcement or about migrants; I was looking for a market. The constraint written weeks earlier is what caught it, not my judgement in the moment. That is the practical case for writing ethics down early: your future self, under pressure to show the work matters, is not a reliable ethical actor.

The second set of decisions was about the product itself. The impact page was the most impressive-looking thing I built and the least defensible, so it now presents nearest-precedent matches as exactly that rather than as a five-year forecast. The FEE panel states that its scores are fixed per intervention type and pending expert elicitation. The ethics gate raises a warning where a response increases migrant-harm exposure. Each makes the demonstration weaker and the tool more honest, and I have come to think that trade-off is the whole job.

Collaboration and Team Dynamics

The working relationship with my supervisor was the spine of this project. We set it up deliberately in Week 1 — weekly calls, Signal and email in between, GitHub for code, a shared Drive split by purpose, and short iterations rather than disappearing for a month and returning with something finished. That structure felt like overhead at the time. It is the reason the project could change direction three times without falling apart.

Almost every correction I received depended on context I did not have: that Malta receives very few arrivals now, that supporting security forces here is ethically complicated, that a name in a report can carry risk.

Expertise in this field is not mostly analytical technique. It is knowing what is true about a place, and the only way to acquire it is from people who already do.

In the Malta office, Matt, Tasnim, Mark and Shanti gave up hours of their working weeks to sit with an undergraduate asking basic questions about work they had spent years on, and not one of them made me feel like an inconvenience. Somebody had worked out where I should get lunch before I asked. Somebody else spent forty minutes explaining a case with no benefit to their own week. That generosity is the reason the Librarian project exists, and it is what I will remember long after the code stops running.



Dingli Cliffs, on a weekend away from the office.

Sustainability and How Impact Will Be Measured

A contribution that only works while I am in the room is not a contribution to the organisation. Three things were done to give this work a life beyond six weeks.

- **Handover material.** Both prototypes were delivered with documentation: an implementation guide for the retrieval agent, an options assessment with a recommended sequence, and an explicit list of known weaknesses, so whoever picks the work up inherits the problems along with the code.
- **Cheap, replaceable foundations.** A Streamlit interface, one orchestrator, FAISS retrieval and a SQLite store, with the agent-based modelling and FEE scoring left as a clearly marked future seam rather than half-implemented. The recommended cloud path starts at roughly sixteen dollars per seat per month, which can be trialled without a procurement process.
- **A defined next step rather than a finished product.** Complete the MVP, run user acceptance testing to establish whether the need is real, and only then assess vendors and data-privacy arrangements — a sequence designed so the organisation can stop if the answer is no.

The measures that matter are community-level rather than technical: the prevalence of irregular migration, tensions in host communities, and effects on local economies, alongside migrant harm and displacement within criminal markets. The project-level targets — source-grounding coverage, interventions a researcher judges plausible, and recall against a known Libyan case — are proxies for those. The recall test matters most, because it is the Do No Harm principle turned into something that can fail.

Conclusion

The most useful thing I learned is that ambition and honesty pull in opposite directions, and you have to keep choosing between them. Nearly every version of this project that looked more impressive was also less true. The five-year forecast was the most impressive thing I built and the least defensible; the broad Sahel scope was more ambitious than the Libya scope and much worse. Each time, the more honest choice made the work smaller and better.

The second is that constraints written down in advance are worth more than good intentions held in the moment. I did not catch my own weakest idea through judgement on the day. I caught it because weeks earlier, with nothing at stake, I had written down what this tool is and is not for.

The third is about closeness. I came into this thinking my contribution would be technical, because that is what I am confident about. Almost everything valuable came from a conversation — with researchers who knew the routes, with a historian who knew the walls, with people who live next to the question. The tool changed direction three times, and every change came from something somebody said. If there is one habit I keep, it is going to find those people first rather than last.



Golden Bay at sunset, toward the end of the six weeks.

Acknowledgements

My sincere thanks to the Laidlaw Foundation for funding this experience and making it possible for a student to spend a summer on work with no commercial return, and to Susanna Kempe, Chief Executive of the Laidlaw Foundation, for the vision behind a programme that puts leadership and service at its centre.

To Dr. Alex Bish, my LiA project advisor, for his guidance, his patience with a very junior collaborator, and above all for the corrections, several of which are the most valuable things I took from this summer. He gave me room to be wrong and then made sure I understood why, which is a rarer kind of generosity than it sounds. I look forward to working and learning with Alex again.

To Matt, Tasnim, Mark and Shanti and the wider team at the GI-TOC Malta office for their time and generosity.

To the University of Toronto Laidlaw Scholars Programme, and to Shraddha Prasad, Elena Shik, You Jia Lee and Tyler Thom, for supporting this project from the outline stage through to submission, and for a great deal of practical help along the way.

References and Sources

Literature informing the framework

- Anderson, M. B. (1999). Do No Harm: How Aid Can Support Peace — or War. Boulder, CO: Lynne Rienner Publishers. Source of the connectors and dividers analysis.
- Bish, A. Studying Irregular Migration through Crime Science. Doctoral thesis. Agent-based modelling of the Libya–Malta route.
- Anatomy of a Route: Script Analysis of Irregular Migration, Smuggling and Harms on the Central Mediterranean Route to Europe. Crime script modelling.
- ETH Zurich, Center for Security Studies. Infographics on Actors–Context–Content–Processes (ACCP) conflict analysis.
- World Bank. Conflict Assessment Framework.
- Global Initiative Against Transnational Organized Crime. Libya and Sahel case studies and annual observatory reporting.

Live data sources integrated into the tool

- UNHCR — refugees and asylum-seekers by location (658,477 in Libya, 2025).
- IOM Displacement Tracking Matrix — internally displaced persons by country (147,382 in Libya as of 31 May 2024).
- ACLED — violent conflict event reporting, rolling one-month window.