TRACK 3 — COMPUTER VISION

Problem Statements

HTH-CV-01: Prioritized Road-Safety Violation & Enforcement Advisor

Category	Details
Industry	Public Safety / Traffic Enforcement
Scenario	Traffic police have limited enforcement capacity and cannot act on every detected violation equally.
Problem Statement	Detect safety violations (no helmet, wrong-way, red-light, multiple riders) from video, then prioritize a limited daily enforcement/review queue by violation severity and repeat-offense likelihood (e.g., recurring plate).
What Teams Should Build	Object/violation detection pipeline, severity scoring, capacity-constrained enforcement queue (e.g., only 30 cases reviewable/day).
24-Hour MVP	Public traffic-violation video/image datasets, detect 2–3 violation types, produce a ranked, evidence-attached enforcement report.
Data/Input	Public road-safety datasets (e.g., open traffic camera datasets) or synthetic/staged footage.
Key Constraints	Review capacity fixed and small; must attach evidence frame per flagged case for auditability.
Expected Output	Ranked violation report with evidence frames and severity.
Evaluation Criteria	Detection accuracy, prioritization logic, auditability of evidence.
Bonus Features	Repeat-offender tracking across multiple video sessions.

HTH-CV-02: Industrial PPE Compliance & Risk-Zone Monitor

Category	Details
Industry	Manufacturing / Occupational Safety
Scenario	Safety officers can't monitor every zone of a large facility at once; incidents are often discovered after the fact.
Problem Statement	Detect PPE compliance (helmet, vest, gloves, goggles) per worker per zone, and prioritize supervisor attention by combining missing-PPE severity with zone risk level (e.g., near heavy machinery = higher priority).
What Teams Should Build	Multi-class PPE detection, per-zone risk weighting, prioritized alert dashboard.
24-Hour MVP	Public PPE detection datasets, detection across simulated multi-zone footage with risk-weighted alert ranking.
Data/Input	Public PPE/safety-gear detection datasets (e.g., open PPE detection datasets on Kaggle/Roboflow Universe).
Key Constraints	Must weight by zone risk, not treat all violations equally; must handle occlusion/multiple workers per frame.
Expected Output	Risk-ranked compliance dashboard with per-worker evidence snapshots.
Evaluation Criteria	Detection accuracy, zone-risk-weighted prioritization, real-time feasibility.
Bonus Features	Trend report showing compliance improvement/decline over a shift.

HTH-CV-03: Shelf-Availability Loss-Prevention System

Category	Details
Industry	Retail Operations
Scenario	Out-of-stock shelves cause direct revenue loss, but store staff can't check every aisle continuously.
Problem Statement	Detect empty/low-stock shelf sections and misplacement from shelf images, then prioritize restocking tasks by estimated revenue-loss-per-hour of the affected SKU/section, not just detection confidence.
What Teams Should Build	Shelf-gap/product detection model, stock-level estimation, revenue-loss-weighted restocking task ranker.
24-Hour MVP	Public retail shelf datasets, detect gaps/misplacement, generate a restocking task list ranked by estimated $ impact (using assumed/synthetic SKU price data).
Data/Input	Public shelf-monitoring datasets (e.g., SKU110K, Grocery Shelves) plus a synthetic price/margin table.
Key Constraints	Ranking must use business impact (price × estimated lost sales), not raw detection score; must handle cluttered/varied shelf layouts.
Expected Output	Restocking task list ranked by revenue impact with shelf-image evidence.
Evaluation Criteria	Detection accuracy, business-impact-based ranking, robustness across shelf types.
Bonus Features	Planogram-compliance check (correct product in correct location).

HTH-CV-04: Document Intake Quality Gate for Enterprise OCR Pipelines

Category	Details
Industry	BPO / Financial Services (Document Processing)
Scenario	Enterprises processing thousands of scanned documents (KYC forms, invoices) waste downstream OCR/NLP effort on poor-quality scans that will fail anyway.
Problem Statement	Build a vision system that scores incoming scanned documents for quality issues (blur, skew, missing sections, low contrast) before they enter an OCR pipeline, routing bad scans to re-scan and good ones forward — reducing downstream failure cost.
What Teams Should Build	Multi-defect quality classifier/scorer, routing logic (pass/reject/needs-review), corrective suggestion generator.
24-Hour MVP	Public/synthetic scanned-document datasets with induced defects, quality scoring + routing dashboard showing pipeline efficiency gain vs. no gate.
Data/Input	Public document datasets (e.g., RVL-CDIP, FUNSD) with synthetically injected blur/skew/noise.
Key Constraints	Must quantify downstream cost savings (e.g., "% of docs prevented from failing OCR"); must handle handwriting interference as a distinct defect.
Expected Output	Quality-gated document queue with defect type, score, and routing decision.
Evaluation Criteria	Defect-detection accuracy, demonstrated pipeline efficiency gain, practical routing logic.
Bonus Features	Auto-corrective preprocessing (deskew/denoise) for borderline cases.

HTH-CV-05: Capacity-Aware Crowd Safety Monitor

Category	Details
Industry	Event Management / Public Venues
Scenario	Venue operators must decide when and where to intervene (open new exits, halt entry) before a crowd-safety incident occurs, not after.
Problem Statement	Estimate crowd density and detect abnormal movement patterns (sudden surge, restricted-area entry) from video, and generate a tiered intervention recommendation (monitor/alert/halt-entry) based on density trend vs. venue capacity limits.
What Teams Should Build	Density estimation model, abnormal-pattern detection over video (not single frames), capacity-threshold-based intervention tiering.
24-Hour MVP	Public crowd-density video datasets, run density + anomaly detection on sample footage, generate tiered alerts against a configurable venue-capacity limit.
Data/Input	Public crowd datasets (e.g., ShanghaiTech, UCF-QNRF, Mall dataset).
Key Constraints	Must operate on video sequences, not just static images; capacity threshold must be configurable per venue.
Expected Output	Live density/anomaly dashboard with tiered intervention recommendations and timestamped evidence clips.
Evaluation Criteria	Density estimation accuracy, video-based anomaly detection quality, realistic intervention logic.
Bonus Features	Predictive early-warning (forecast density trend 5–10 minutes ahead).

HTH-CV-06: Severity-Triaged Crop Disease Advisory System

Category	Details
Industry	Agriculture / AgriTech
Scenario	Smallholder farmers need to know not just "what disease" but how urgently to act and what it will cost them if they don't.
Problem Statement	Detect plant disease and severity from leaf images, estimate potential yield-loss if untreated, and prioritize treatment recommendations by cost-of-inaction across multiple crops/fields.
What Teams Should Build	Disease classification + severity estimation model, yield-loss estimation logic (using assumed crop-loss curves), multi-field treatment prioritization.
24-Hour MVP	Public plant-disease image datasets (e.g., PlantVillage), classification + severity + prioritized treatment list across a simulated multi-field farm.
Data/Input	Public PlantVillage or similar open crop-disease datasets.
Key Constraints	Must estimate urgency/cost-of-inaction, not just classify disease; must handle multiple fields competing for limited treatment resources/budget.
Expected Output	Prioritized treatment plan across fields with severity, estimated loss, and recommended action.
Evaluation Criteria	Classification accuracy, severity/urgency estimation logic, resource-prioritization quality.
Bonus Features	Region-highlighting (localize affected leaf area) and spread-risk estimate to neighboring plants.

HTH-CV-07: Budget-Constrained Road Maintenance Prioritization System

Category	Details
Industry	Municipal Infrastructure
Scenario	City road-maintenance departments have far more damaged road segments than budget to fix, and must decide what to repair first.
Problem Statement	Detect and classify road damage (potholes, cracks, surface wear) and severity from images/video, then generate a prioritized repair plan given a fixed maintenance budget and traffic-volume-weighted urgency.
What Teams Should Build	Damage detection + severity classification model, geo-tagging/mapping, budget-constrained repair-priority optimizer weighted by traffic volume/road class.
24-Hour MVP	Public road-damage datasets, detect + classify severity on sample footage, generate a ranked repair plan under a synthetic fixed budget across multiple road segments.
Data/Input	Public road damage datasets (e.g., RDD2022).
Key Constraints	Budget is fixed and insufficient for all repairs — must show explicit trade-offs; higher-traffic roads should weigh higher in priority.
Expected Output	Interactive map/report of ranked repair priorities with severity, cost estimate, and traffic weighting.
Evaluation Criteria	Detection/severity accuracy, budget-aware prioritization logic, map/report clarity.
Bonus Features	Before/after simulation showing network-risk reduction if top-N repairs are made.

HTH-CV-08: Multi-Waste-Stream Contamination & Recycling Yield Optimizer

Category	Details
Industry	Waste Management / Sustainability
Scenario	Recycling facilities lose revenue when contaminated waste streams reduce the value/purity of recyclable batches.
Problem Statement	Classify waste materials from images/video on a conveyor-style feed, detect multiple objects per frame, and estimate contamination rate per batch to recommend whether a batch should proceed, be re-sorted, or be rejected.
What Teams Should Build	Multi-object waste classification model, batch-level contamination-rate aggregation, decision logic (accept/re-sort/reject) with estimated value impact.
24-Hour MVP	Public waste-classification image datasets, multi-object detection on sample batches, batch decision report with estimated recyclable-value impact.
Data/Input	Public waste classification datasets (e.g., TrashNet, TACO).
Key Constraints	Must handle multiple objects per frame, not single-item classification; must translate contamination rate into a business decision, not just a percentage.
Expected Output	Batch-level contamination report with accept/re-sort/reject recommendation and estimated value impact.
Evaluation Criteria	Multi-object detection accuracy, business-decision logic quality, practical batch-level aggregation.
Bonus Features	Real-time conveyor simulation with live batch decisions.

HTH-CV-09: Accessibility-First Sign Language Communication Bridge

Category	Details
Industry	Accessibility Tech / Public Services
Scenario	Public service counters (banks, hospitals, government offices) often lack real-time interpretation support for Deaf/hard-of-hearing visitors.
Problem Statement	Recognize a defined vocabulary of sign-language gestures across varied hand positions/orientations/backgrounds and convert to text/speech in near real-time to support a live service-counter interaction.
What Teams Should Build	Gesture recognition model robust to orientation/background variance, text output, real-time inference pipeline, optional text-to-speech.
24-Hour MVP	Public sign-language gesture datasets, recognize 15–20 vocabulary signs reliably across at least 2 different backgrounds/lighting conditions, live demo translating to text/speech.
Data/Input	Public sign-language datasets (e.g., WLASL, ASL Alphabet).
Key Constraints	Must demonstrate robustness across backgrounds/orientations, not a single clean condition; latency must support "conversation-like" use, not batch processing.
Expected Output	Real-time gesture-to-text (and optionally speech) interface.
Evaluation Criteria	Accuracy across varied conditions, real-time feasibility, practical UX for a service-counter scenario.
Bonus Features	Two-way support (speech-to-simplified-sign prompts on screen).

HTH-CV-10: Store-Level Loss-Prevention & Suspicious-Behavior Flagging System

Category	Details
Industry	Retail Security
Scenario	Retail loss-prevention teams review hours of footage manually and miss most incidents; they need computer vision to surface only the highest-likelihood events.
Problem Statement	Analyze in-store video to flag behaviorally suspicious patterns (loitering near high-value shelves, concealment gestures, unusual dwell time) and rank flagged clips by confidence and estimated loss-value of the nearby merchandise for human review.
What Teams Should Build	Behavior/pattern detection over video sequences, dwell-time and zone analysis, value-weighted ranking of flagged clips for a limited daily review capacity.
24-Hour MVP	Public retail/behavior-analysis video datasets (or staged footage), detect 1–2 suspicious behavior patterns over sequences, generate a ranked review queue with clip evidence.
Data/Input	Public video-anomaly/behavior datasets (e.g., UCF-Crime subset, staged synthetic footage) — avoid real surveillance of identifiable individuals; use staged/public data only.
Key Constraints	Must work over video sequences (temporal), not single frames; review capacity is limited, so ranking by estimated value-at-risk matters, not just confidence; must be explicit that this is a decision-support tool for human review, not an automated accusation system.
Expected Output	Ranked review queue with video clip evidence, confidence, and estimated nearby merchandise value.
Evaluation Criteria	Temporal detection quality, responsible framing (human-in-the-loop), ranking logic.
Bonus Features	False-positive feedback loop where reviewer verdicts retrain ranking weights.