TRACK 2 — MACHINE LEARNING & DEEP LEARNING

Problem Statements

HTH-ML-01: Actionable Student-Risk Intervention Engine

Category	Details
Industry	EdTech (Institutional Operations)
Scenario	Institutions know which students are "at risk" too late to intervene meaningfully.
Problem Statement	Move beyond prediction: predict at-risk students, explain top contributing factors per student, then prioritize which students need intervention first given limited counselor capacity, and recommend specific interventions.
What Teams Should Build	Risk-prediction model, SHAP/feature-importance explanation layer, capacity-constrained prioritization (e.g., only 5 counselor slots/week), intervention recommender mapped to root cause.
24-Hour MVP	Synthetic/public student performance dataset, risk scores + explanations + ranked intervention list under a fixed counselor-capacity constraint.
Data/Input	Public open education datasets (e.g., UCI student performance) or synthetic equivalents.
Key Constraints	Counselor capacity is fixed and small — model must prioritize, not just list everyone; interventions must map to explained cause (attendance vs. assessment vs. engagement).
Expected Output	Prioritized intervention queue with per-student rationale and recommended action type.
Evaluation Criteria	Explainability, prioritization logic under constraint, actionability over raw accuracy.
Bonus Features	Simulate intervention impact ("if we act on top 5, expected outcome change").

HTH-ML-02: Explainable Real-Time Fraud Triage System

Category	Details
Industry	Banking / FinTech
Scenario	Fraud teams are overwhelmed by alert volume and cannot investigate every flagged transaction with a small human team.
Problem Statement	Detect fraud in highly imbalanced transaction data, produce fraud probability with explanation, and prioritize a limited daily investigation queue (e.g., only 50 cases/day) by expected loss × probability.
What Teams Should Build	Imbalanced-data classifier (SMOTE/class-weighting), explanation layer (SHAP/LIME), expected-loss-based queue ranking under a fixed investigator capacity.
24-Hour MVP	Public fraud dataset (e.g., Kaggle credit card fraud used only as raw signal source, not as the whole solution), imbalance handling, ranked investigation queue with rationale, not just raw scores.
Data/Input	Public financial fraud datasets (e.g., IEEE-CIS, credit card fraud) or synthetic transaction generator.
Key Constraints	Investigation capacity is fixed and much smaller than alert volume — must rank by business cost, not just probability; extreme class imbalance must be explicitly handled.
Expected Output	Ranked investigation queue (case, probability, expected loss, explanation) + summary dashboard.
Evaluation Criteria	Handling of imbalance, business-cost-aware ranking (not just AUC), explanation quality.
Bonus Features	Feedback loop simulating analyst verdicts updating future rankings.

HTH-ML-03: Congestion-Aware Route & Signal Advisory System

Category	Details
Industry	Urban Mobility / Smart City
Scenario	City traffic managers need to decide where to deploy limited traffic-control interventions (temporary signal changes, diversions) during predictable congestion windows.
Problem Statement	Predict congestion by location/time, explain contributing factors (time, weather, events), and recommend a prioritized set of interventions given a limited number of traffic officers/signal changes available per day.
What Teams Should Build	Time-series congestion model, factor explanation, constrained intervention recommender (e.g., only 3 locations can be actively managed at once).
24-Hour MVP	Public traffic dataset (city open data or simulated), congestion prediction + top-factor explanation + ranked intervention list under officer-capacity constraint.
Data/Input	Public city traffic open-data portals or synthetic traffic simulator.
Key Constraints	Officer/resource capacity fixed and small; must incorporate at least one external factor (weather or event calendar).
Expected Output	Congestion forecast map + prioritized intervention plan with expected impact.
Evaluation Criteria	Forecast usefulness, constraint-aware prioritization, explanation quality.
Bonus Features	Event-calendar API integration for spike prediction.

HTH-ML-04: Cost-Optimized Predictive Maintenance Scheduler

Category	Details
Industry	Manufacturing / Industrial Operations
Scenario	Unplanned equipment failure is costly, but maintenance crews and downtime windows are limited — over-maintaining is also expensive.
Problem Statement	Predict failure probability from sensor data, then schedule maintenance across multiple machines given a limited number of maintenance slots and downtime cost, prioritizing by expected-failure-cost, not just probability.
What Teams Should Build	Failure-prediction model (sensor time-series), maintenance-scheduling optimizer under slot constraints, cost-based prioritization.
24-Hour MVP	Public/synthetic sensor datasets (e.g., NASA CMAPSS turbofan), failure prediction + constrained maintenance schedule across 10+ simulated machines with only 3 available slots/day.
Data/Input	Public NASA CMAPSS or synthetic sensor-degradation dataset.
Key Constraints	Maintenance capacity strictly limited; false negatives (missed failures) must be weighted more heavily than false positives; downtime cost varies by machine.
Expected Output	Ranked maintenance schedule with cost-benefit rationale per machine.
Evaluation Criteria	Predictive quality, cost-aware scheduling logic, realistic constraint handling.
Bonus Features	Dynamic re-scheduling when a new urgent failure signal appears mid-demo.

HTH-ML-05: Retention-ROI Churn Prioritization Engine

Category	Details
Industry	Subscription / Telecom / SaaS
Scenario	Retention teams have a limited budget for retention offers and cannot target every at-risk customer.
Problem Statement	Predict churn probability, segment customers by churn driver, and prioritize a retention-offer budget (e.g., $10,000 total) across customers by expected-value-saved, not just churn probability.
What Teams Should Build	Churn model, driver-based segmentation, budget-constrained offer allocator (customer LTV × churn probability × offer cost).
24-Hour MVP	Public telecom churn dataset (e.g., IBM Telco), churn prediction + segmentation + budget-constrained offer allocation with ROI projection.
Data/Input	Public telecom/SaaS churn datasets.
Key Constraints	Fixed retention budget; must show which customers are NOT worth retaining (negative ROI cases).
Expected Output	Budget allocation table with projected ROI, plus segment-level retention strategy notes.
Evaluation Criteria	ROI-awareness (not just accuracy), segmentation insight, realistic budget handling.
Bonus Features	A/B simulate two allocation strategies and compare projected outcomes.

HTH-ML-06: Inventory-Constrained Demand Forecasting & Allocation

Category	Details
Industry	Retail / Supply Chain
Scenario	A retailer has limited warehouse capacity and must decide not just how much to forecast but how to allocate limited stock across multiple stores/SKUs.
Problem Statement	Forecast demand per SKU/store considering seasonality/promotions, then allocate a fixed total inventory budget/capacity across stores to minimize stockouts and overstock simultaneously.
What Teams Should Build	Time-series forecasting model, multi-store allocation optimizer under a total capacity constraint, stockout/overstock cost trade-off logic.
24-Hour MVP	Public retail demand dataset (e.g., M5/Walmart), forecast + constrained allocation across 5–10 simulated stores with a fixed total unit budget.
Data/Input	Public retail forecasting datasets (M5 Forecasting, Rossmann, or similar).
Key Constraints	Total inventory is fixed and insufficient to satisfy all stores fully — must trade off; holiday/promo effects must be incorporated.
Expected Output	Store-level allocation plan with projected stockout/overstock risk.
Evaluation Criteria	Forecast accuracy, allocation logic under scarcity, business trade-off clarity.
Bonus Features	Simulate a promo event mid-demo and show re-allocation.

HTH-ML-07: Cold-Start-Aware Recommendation Engine with Business Guardrails

Category	Details
Industry	E-commerce / Media
Scenario	New users/items get poor recommendations, and pure engagement-maximizing recommenders can promote low-margin or low-quality items.
Problem Statement	Build a recommender that handles cold-start via content/metadata signals, and applies business guardrails (e.g., margin, inventory availability) so recommendations balance relevance with business value, not engagement alone.
What Teams Should Build	Hybrid collaborative + content-based model, cold-start fallback strategy, guardrail/re-ranking layer incorporating a secondary business metric.
24-Hour MVP	Public interaction dataset (MovieLens/e-commerce), demonstrate cold-start handling for a synthetic new user/item, and guardrail-based re-ranking with a toggle to compare "pure engagement" vs "business-aware" output.
Data/Input	Public recommendation datasets (MovieLens, Retailrocket, or similar).
Key Constraints	Must explicitly handle zero-history users/items; guardrail metric must be configurable, not hardcoded.
Expected Output	Ranked recommendations with a visible "why recommended" explanation and guardrail impact comparison.
Evaluation Criteria	Cold-start handling quality, business-guardrail integration, explainability.
Bonus Features	Diversity/fairness metric alongside relevance.

HTH-ML-08: Peak-Constrained Energy Optimization Advisor

Category	Details
Industry	Facilities Management / Energy
Scenario	A facility faces demand-charge penalties for exceeding peak power thresholds, not just total consumption.
Problem Statement	Predict energy consumption by hour, identify likely peak-threshold breaches, and recommend a load-shifting plan that keeps the facility under its peak limit given which systems can/cannot be shifted.
What Teams Should Build	Time-series consumption forecasting, peak-breach detection, load-shift recommender respecting which loads are shiftable vs. critical.
24-Hour MVP	Public/synthetic building energy dataset, forecast + peak-breach alerts + shift plan for 24-hour horizon under a fixed peak-demand cap.
Data/Input	Public building energy datasets (e.g., UCI household power, ASHRAE) or synthetic facility data.
Key Constraints	Some loads are non-shiftable (must be respected as hard constraints); peak cap is fixed and breach has a real cost penalty to quantify.
Expected Output	Hourly forecast, breach-risk alerts, and a load-shift schedule with projected cost savings.
Evaluation Criteria	Forecast quality, constraint respect, quantified savings.
Bonus Features	Weather-responsive re-forecasting mid-demo.

HTH-ML-09: Evidence-Grounded Misinformation Triage System

Category	Details
Industry	Media / Trust & Safety
Scenario	Content moderation teams can't manually fact-check every flagged post; they need prioritized, explainable triage.
Problem Statement	Predict likelihood an article/post is misleading using text/source/claim-consistency signals, explain the reasoning per flagged item, and prioritize a moderation queue under a limited daily review capacity.
What Teams Should Build	NLP classifier using linguistic + source-credibility + claim-consistency features, explanation generator, capacity-constrained moderation queue ranking by reach × risk.
24-Hour MVP	Public misinformation/fake-news datasets, classification + explanation + prioritized queue under a fixed reviewer capacity (e.g., 20 items/day).
Data/Input	Public datasets (LIAR, FakeNewsNet, or similar).
Key Constraints	Must output reasoning, not just a label; must incorporate estimated reach/virality into prioritization, not just confidence score.
Expected Output	Prioritized moderation queue with explanation and recommended action (review/deprioritize/escalate).
Evaluation Criteria	Explanation quality, prioritization logic, realistic moderation-capacity handling.
Bonus Features	Source-credibility trend tracking over time.

HTH-ML-10: Severity-Aware Streaming Anomaly Detector for Industrial Sensors

Category	Details
Industry	Industrial IoT / Operations
Scenario	Not all sensor anomalies are equal — a small spike may be noise, a sustained deviation may be critical — and operators need severity, not just a binary flag.
Problem Statement	Monitor streaming sensor data, detect multiple anomaly types (spike, drift, dropout, stuck-sensor), estimate severity and likely root cause, and route alerts to the right response tier (ignore/monitor/urgent).
What Teams Should Build	Streaming anomaly-detection pipeline (statistical + ML, e.g., isolation forest/autoencoder), multi-type anomaly classification, severity scoring, tiered alert routing logic.
24-Hour MVP	Simulated streaming sensor feed with injected anomaly types, live detection dashboard with severity tiers and routing decisions.
Data/Input	Public sensor datasets (e.g., NASA, SKAB) or a synthetic streaming generator with controllable anomaly injection.
Key Constraints	Single fixed threshold is explicitly disallowed — must differentiate anomaly types; must minimize alert fatigue (avoid flooding operators).
Expected Output	Live anomaly feed with type, severity, and routing tier; end-of-run summary report.
Evaluation Criteria	Multi-type detection accuracy, severity calibration, alert-fatigue reduction design.
Bonus Features	Root-cause hints (e.g., "correlated with pump #3 vibration").