



# Fall 2025 University Innovation Impact Report

# CONTENTS

- 03 Acknowledgments
- 04 Introduction
- 05 Engagement Metrics Snapshot
- 06 Technology Snapshot List
- 07 Comprehensive Table Part 1
- 08 Comprehensive Table Part 2
- 09 Carnegie Mellon University – Corporate Startup Lab (capstone)
- 10 Carnegie Mellon University – Heinz College (capstone)
- 11 Carnegie Mellon University – Information Networking Institute (capstone)
- 12 Carnegie Mellon University – Masters of Science in Product Management (capstone)
- 13 Carnegie Mellon University – Masters of Science in Product Management (capstone)
- 14 The Ohio State University - Electrical & Computer Engineering (capstone)
- 15 The Ohio State University – Hack OHIO (hackathon)
- 16 The Ohio State University – Center for Innovation Strategies (capstone)
- 17 University of California, Berkeley – Data Science Discovery (capstone)
- 18 University of California, San Diego – DSC (capstone)
- 19 University of California, San Diego – DSC (capstone)
- 20 University of California, San Diego – DSC (capstone)
- 21 University of Colorado, Boulder – Masters of Science in Data Science (capstone)
- 22 Smith College - Data Clinic (capstone)
- 23 Smith College - Design Clinic (capstone)
- 24 Conclusion

# ACKNOWLEDGMENTS

We want to extend our heartfelt thanks to everyone at 99P Labs, Honda Research Institute, and all of our university partners. Your dedication and collaborative spirit have been instrumental in making these student innovation projects possible.

Thank you all for your hard work, creativity, and commitment to bridging industry and academia. Your contributions have made a real impact on our mission and will continue to inspire future projects. We're excited about what we're building together.

# INTRODUCTION

Welcome to 99P Labs' Fall 2025 University Innovation Impact Report. This document serves as a window into the university engagements we participated in during the 2025 fall semester.

The report begins with the Engagement Metrics Snapshot section. This section includes metrics from our Medium blog and LinkedIn activities, as well as the engagements themselves. It provides a visual representation of our outreach and impact, making it easy to track the progress and effectiveness of our efforts directly from the report.

Following the engagement metrics, we've included a snapshot of the technologies explored and applied across our projects. This section highlights the breadth of tools and approaches investigated, offering a quick sense of the technical landscape shaping our work this semester.

Next, the report features a comprehensive table view of all the engagements. This table provides a glimpse of each engagement, including the problem statement, the domain, the associated university or program, key numbers, and links to related assets. These assets, which include blogs, videos, posters, and other materials, offer a deeper dive into each program and are designed to spark your interest.

Each engagement is also presented in greater detail later in the report, offering a look into the challenges we are striving to address.

**This report is not just a testament to our commitment towards fostering academic collaborations, but also an invitation. We are reaching out to all those interested in these projects, seeking your valuable feedback and potential collaboration.**

**For further information, queries, or feedback, please feel free to contact Rajeev Chhajer at [rajeev\\_chhajer@honda-ri.com](mailto:rajeev_chhajer@honda-ri.com) and/or Ryan Lingo at [ryan\\_lingo@honda-ri.com](mailto:ryan_lingo@honda-ri.com). We are eager to connect with you and explore opportunities for collaboration**

Thank you for your interest in our work. We look forward to hearing from you. You can also connect with our community by checkout out our [website](#), following us on [Medium](#) and [LinkedIn](#).

# ENGAGEMENT METRICS SNAPSHOT FOR FALL 2025

## Engagements

**Student Engagement Projects:** 15

**University Partnerships:** 6

**Unique Programs:** 12

**Students Engaged With:** 179

## Medium

**Number of Blogs:** 9

**Number of Blog Views:** 359

## Outcomes

**Technologies Utilized/ Researched:** 88

**Prototypes / Demos:** 44

**Repos / Decks:** 50

**Videos:** 38

## LinkedIn

**Number of LinkedIn Posts:** 9

**LinkedIn Organic Impressions:** 8,584

**LinkedIn Engagement Rate Avg:** 8.08%

# TECHNOLOGIES SNAPSHOT FOR FALL 2025

List of the 88 technologies utilized/researched across Fall 2025 projects, highlighting the breadth of student engagement and technical exploration.

## Intelligence + AI

- OpenAI API
- ChatGPT
- Google Gemini
- Anthropic Claude
- GPT-4
- GPT-5
- GPT-5-Nano
- GPT-4o-mini
- text-embedding-3-small
- LightGBM
- Prophet
- scikit-learn
- PointNet
- LSTM
- Qwen-VL
- PaddleOCR
- LangGraph
- Model Context Protocol (MCP)
- ReAct
- Reflexion
- Self-Refine
- TF-IDF
- ROUGE
- BLEU
- METEOR
- BERTScore
- MMLU
- MT-Bench
- Chatbot Arena
- HELM
- MLE-Bench
- ContextEval
- NOMAD
- UseVille
- Arize Phoenix
- Langfuse
- Agent Lightning
- Memori
- CoT
- Retrieval-Augmented Generation (RAG)

## Cloud / Data / Communication

- Docker
- Kubernetes
- Redis
- PostgreSQL
- PostGIS
- SQLite
- Kafka
- MQTT
- Vector database
- SQL database
- arXiv API
- USPTO API
- WebSearch

## Software

- Opencode
- REST API
- React
- React Native
- Redux
- TypeScript
- JavaScript
- HTML
- Python
- Go
- Node.js
- FastAPI
- Flask
- VSCode
- Jupyter Notebook
- Google Colab
- JSON
- JSONL
- Regex-based parsing
- React Native
- Redux
- TypeScript
- JavaScript
- HTML
- Python
- Go
- Node.js
- FastAPI
- Flask
- VSCode
- Jupyter Notebook
- Google Colab
- JSON
- JSONL
- Regex-based parsing

## Hardware

- RealSense 99ACXA
- Raspberry Pi
- TI IWR6843 mmWave radar

# Fall 2025 University Engagements Summary Table Part 1

Innovation Prompt	University / Program	Key Numbers	Link(s)
How might we leverage Honda's robotics expertise and partnerships to systematically identify, validate, and commercialize new robotics opportunities beyond factory automation?	 Carnegie Mellon University Corporate Startup Lab	2 students 12 technologies 14 weeks	<a href="#">Blog</a> <a href="#">Process Book</a> 
How might we leverage open data and digital twin concepts to help LA28 organizers and city planners proactively model and mitigate high-impact mobility, energy, and societal challenges for the 2028 Olympic Games?	 Carnegie Mellon University Heinz College	5 students 17 technologies 14 weeks	<a href="#">Blog</a> <a href="#">Final Video</a> 
How might we create a sandbox environment for emulating networking between devices in constrained environments, with considerations on distance, scale and real-world factors (weather, crowding, etc.)?	 Carnegie Mellon University Information Networking Institute	5 students 11 technologies 14 weeks	<a href="#">Blog</a> <a href="#">Final Video</a> 
How might we use AI to help organizations detect emerging trends early and turn them into strategy-ready insights?	 Carnegie Mellon University Master of Science in Product Management	3 students 2 technologies 14 weeks	<a href="#">Blog</a> <a href="#">Final Video</a> <a href="#">Process Book</a>  
How might we use AI to unlock new innovation opportunities from dormant patents by making them easier to understand, connect, and evaluate at scale?	 Carnegie Mellon University Master of Science in Product Management	3 students 3 technologies 14 weeks	<a href="#">Blog</a> <a href="#">Final Video</a> <a href="#">Process Book</a>   
How might we detect and model group coordination or disruption in shared spaces using anonymous, edge-processed motion signals?	 Ohio State University Electrical & Computer Engineering	3 students 5 technologies 14 weeks	<a href="#">Mind-Report Video</a> <a href="#">Deck</a> 
How might we reduce everyday friction in how people move, coordinate, and make decisions in shared spaces?	 Ohio State University Hack OHIO	120 students 2 days	<a href="#">Blog</a> <a href="#">1st Place</a> <a href="#">2nd Place</a>  
How might we conceptualize a proactive assistant experience that anticipates needs and offers timely help while remaining transparent, privacy-first, and always under the user's control?	 Ohio State University Center for Innovation Studies	8 students 4 technologies 10 weeks	<a href="#">AI Guardian Deck</a> <a href="#">Variable Tech Deck</a> 
How might we conceptualize an interface that adapts its complexity and interaction modes to different people and situations so advanced technology feels intuitive and low-friction?			

# Fall 2025 University Engagements Summary Table Part 2

Innovation Prompt	University / Program	Key Numbers	Link(s)
How might we transform the research ideation process from a solitary activity into a dynamic, collaborative dialogue between a researcher and an intelligent agent?	 Data Discovery	4 students 5 technologies 14 weeks	<a href="#">Blog</a> <a href="#">Poster</a> <a href="#">LinkedIn</a>
How might we evaluate an LLM agent's context strategy so we can reliably improve performance while reducing token and time cost?	 Halıcıoğlu Data Science Institute (HDSI)	4 students 13 technologies 14 weeks	<a href="#">1<sup>st</sup> Quarter Report</a> <a href="#">2<sup>nd</sup> Quarter Proposal</a>
How might we define explanation quality dimensions that are predictive of learning outcomes and can be measured?	 Halıcıoğlu Data Science Institute (HDSI)	4 students 7 technologies 14 weeks	<a href="#">1<sup>st</sup> Quarter Report</a> <a href="#">2<sup>nd</sup> Quarter Proposal</a>
How might we evaluate LLM generated information in a way that is scalable, reliable, and aligned with human judgment?	 Halıcıoğlu Data Science Institute (HDSI)	4 students 5 technologies 14 weeks	<a href="#">1<sup>st</sup> Quarter Report</a> <a href="#">2<sup>nd</sup> Quarter Proposal</a>
How might we design a debates agent to run debates using multiple roles and evidence-based arguments?	 Master of Science in Data Science	5 students 15 technologies 14 weeks	<a href="#">Blog</a> <a href="#">LinkedIn</a>
How might we transform massive, static archives of slide decks into an interactive and reliable knowledge base?	 Data Clinic	5 students 13 technologies 14 weeks	<a href="#">Blog</a> <a href="#">LinkedIn</a>
How might we design a compact, efficient, and low-noise vertical axis wind turbine that performs reliably in turbulent rooftop conditions and integrates seamlessly with urban architecture?	 Design Clinic	4 students 7 technologies 14 weeks	<a href="#">Concept Generation Memo</a> <a href="#">TFM 1</a> <a href="#">TFM 2</a>

## Innovation Prompt

How might we leverage Honda's robotics expertise and partnerships to systematically identify, validate, and commercialize new robotics opportunities beyond factory automation?

## Description

This project developed an evidence-based approach for assessing whether emerging robotics opportunities can succeed at scale in real-world operations. Rather than producing a one-time market ranking, it synthesized market research, investor and founder insights, and case evidence into a practical evaluation toolkit: an 8-stage commercialization maturity model, a 4+1 ROI lens (speed of value, profit engine, learning loop, hidden costs, plus team adaptability), and standardized templates for comparing opportunities side-by-side. The result is a decision framework that helps leadership cut through hype, surface operational and integration risks early, and prioritize robotics investments with a clear path to scalable deployment and sustainable economics.

## Engagement Highlights

### Four Levers of ROI



A financial viability lens that evaluates robotics opportunities through speed of value, profit engine, learning loop, and hidden costs, highlighting the factors that determine whether deployments succeed at scale.

They produced these standardized "startup baseball card" profiles for a number of companies to compare metrics, maturity, and commercialization signals side-by-side.

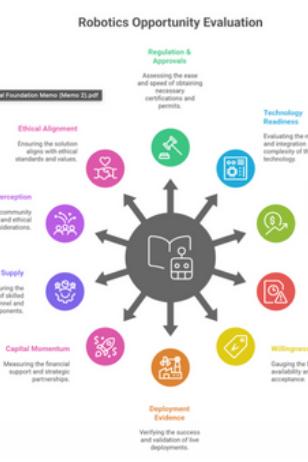


## Key Learnings

**Market size is not a proxy for "can we win."** Top-down sector maps and funding trends describe what exists, but they do not explain the step-by-step evidence of how robotics ventures actually progress from prototype to scalable operations.

**Operational reality dominates technical novelty.** The biggest commercialization blockers are often real-world constraints (integration effort, deployment friction, training burden, infrastructure and safety constraints), not model performance in controlled settings.

**Evidence can be reconstructed with a structured lens.** Treating company signals (especially funding announcements and public milestones) as "archaeological" evidence, then mapping them to maturity stages and ROI levers, enables more reliable comparisons across opportunities than narrative claims.



A multi-factor robotics opportunity evaluation framework that screens ideas across deployment evidence, economics, readiness, buyer demand, regulatory and ethical fit, and go-to-market feasibility.



An eight-stage commercialization maturity model that maps robotics ventures from lab prototype to production deployments, scaled manufacturing, expansion, and platform evolution.

## Innovation Prompt

How might we leverage open data and digital twin concepts to help LA28 organizers and city planners proactively model and mitigate high-impact mobility, energy, and societal challenges for the 2028 Olympic Games?

### Description

This project was a research-driven exploration of how Los Angeles can manage extreme, time-compressed mobility demand in preparation for the LA 2028 Olympics. The team synthesized lessons from prior Olympics, conducted user and stakeholder research, and built data-informed models to design a structured “node-to-node” mobility network that improves throughput and predictability through tiered pickup hubs, pooling, and transit connections. The work combined system design with demand forecasting, routing and fleet optimization, and operational monitoring concepts, and outlined simulation-based validation to stress-test performance under peak surges and disruptions.

### Engagement Highlights

#### UN SDG & Corporate Sustainability Alignment

Direct integration with Honda's ESG reporting framework and global sustainability commitments creates measurable impact across four critical UN Sustainable Development Goals.



##### SDG 11 – Sustainable Cities

Decongests venues and urban zones through intelligent routing and dedicated mobility infrastructure, reducing traffic by an estimated 30-40% in Olympic corridors.



##### SDG 12 – Responsible Consumption and Production

Shared EV fleet model eliminates thousands of private vehicle trips, promoting efficient resource utilization and reducing per-capita transportation footprint.



##### SDG 13 – Climate Action

Achieves 65%+ CO<sub>2</sub> emission reduction compared to traditional Olympic transport, with verified carbon accounting for corporate disclosure requirements.



##### SDG 9 – Innovation & Infrastructure

Smart mobility infrastructure serves as proof-of-concept for future urban deployments, demonstrating scalable solutions for sustainable city development.



Aligned the node-to-node mobility system with key UN Sustainable Development Goals, connecting design choices to sustainability outcomes. Supported executive and partner reporting by translating system performance into measurable impact.

**Honda LA28 Node-to-Node MaaS: Our High-Level Solution**

Our proposed solution is a Shared EV Mobility Network designed specifically for LA28. This system will connect people between fixed “nodes” strategically placed near residential neighborhoods, existing public transit hubs, and Olympic venues. This approach prioritizes clarity and efficiency.

- Node-to-Node Routing**: Provides clarity and predictability, simplifying user journeys and optimising fleet deployment.
- Dynamic Pooling**: Aims for 3-4 riders per EV, significantly reducing the total number of vehicles on the road.
- Exclusive Access**: Leverages exclusive Olympic lane access and real-time official road closure data for efficient travel.
- Fully Electric Fleet**: Utilises Honda's EV fleet, directly supporting LA28's sustainability targets and reducing environmental impact.

Summarized the proposed LA28 node-to-node MaaS concept and the key system capabilities (routing, pooling, access control, and EV-first operations). Set the foundation for the project by clearly defining the solution approach and what it was designed to deliver.

### Key Learnings

#### Structure beats flexibility.

Fixed mobility nodes and clear pickup flows turn chaotic, peak-demand movement into predictable throughput, reducing wait times and operational breakdowns compared to door-to-door pickup.

#### Design for segments, not averages.

Treating visitors, residents, commuters, workforce, and credentialed personnel as distinct user groups leads to clearer product decisions and fewer failure modes than a one-size-fits-all transport experience.

#### Model, simulate, then operate.

Demand forecasting, routing and fleet optimization, and scenario-based simulation provide the evidence needed to size fleets, place nodes, and plan contingencies before real-world surges expose weaknesses.

### Node Strategy: Three-Tier Hierarchy



Presented the three-tier node strategy, showing how venue, neighborhood, and connector nodes structured pickups and transfers across Los Angeles. This hierarchy was critical to achieving predictable throughput at scale by concentrating demand into well-defined, manageable locations.

### Possible Next Steps

#### 1. Simulation and scenario testing.

Build a SUMO-based model of the node network and run “what if” cases (overcapacity events, transit delays, road closures) to quantify throughput and bottlenecks.

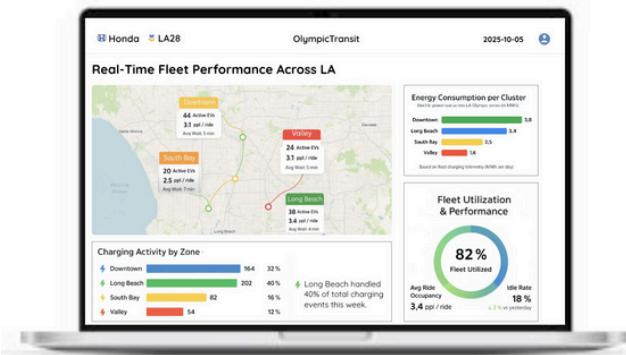
#### 2. Node placement and coverage optimization.

Use GIS + optimization to refine where nodes go and how many are needed, trading off walk distance, safety, roadway access, and expected demand.

#### 3. Operations policy experiments.

Test strategies for pooling rules, pricing and incentives, vehicle staging, and charging schedules to see which policies reduce waits and smooth peaks.

### Dashboard UI - Fleet Performance



Showcased the fleet performance dashboard used to monitor real-time vehicle status, utilization, and charging across the network. It was important because it gave operators the visibility needed to make fast decisions and keep service reliable during peak surges.

## Innovation Prompt

How might we create a sandbox environment for emulating networking between devices in constrained environments, with considerations on distance, scale and real-world factors (weather, crowding, etc.)?

### Description

This capstone project explored how to evaluate future critical digital infrastructure by creating a repeatable way to test network behavior under real-world conditions. The team prototyped Omen, a scenario-driven testbed that turns a single experiment description into automated runs and clear, interactive results, so researchers can compare scenarios and understand when and why connectivity degrades. The work focused on practical research questions like how to represent mobility over time in a repeatable way, how to approximate environmental effects, and how to convert simulation output into evidence that supports decision-making, while also mapping the tool landscape and identifying which backends are viable for scaling beyond the prototype.

### Engagement Highlights

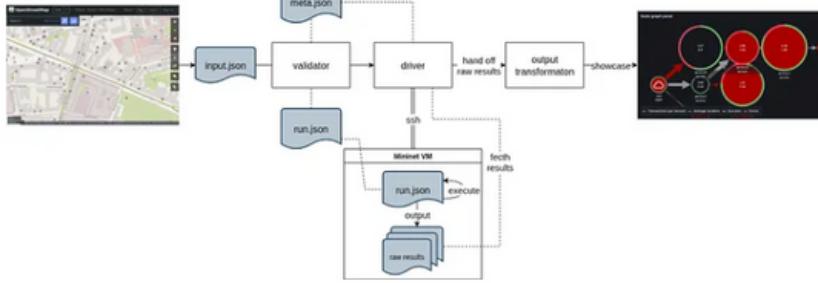


Figure 1: An early implementation of the prototype

An early Omen pipeline diagram showed JSON scenario inputs flowing through validation and a driver to execute tests on a Mininet VM, then transform outputs into dashboards. It mattered because it proved an end-to-end, repeatable path from scenario definition to measurable results.

### Key Learnings

#### Repeatability comes from a pipeline approach.

Defining each study as “JSON in, results out” made experiments easier to rerun, compare, and share, while revealing practical stack limitations early.

#### Mobility needs an explicit time structure.

Using discrete timeframes for movement, followed by a connectivity test each timeframe, enabled clean comparisons across scenarios.

#### Insights require standardized outputs and visualization.

Converting Mininet stdout into CSVs and Grafana dashboards made it easy to spot when and where a topology degrades.

Category	Mininet	ns-3	OMNeT++	OMNeT++ (with OMNet-Cloud)	OpenWRT	OpenWRT (with OMNet-Cloud)	OpenWRT (with OMNet-Cloud, with OMNet-Cloud)	OpenWRT (with OMNet-Cloud, with OMNet-Cloud, with OMNet-Cloud)	OpenWRT (with OMNet-Cloud, with OMNet-Cloud, with OMNet-Cloud, with OMNet-Cloud)
Time	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.	Time-based, discrete time steps. Best for testing real applications, large TCP/IP flows, latency measurements, and real-world environments.
Environment	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.	Environment is not considered. Focus is on the basic network metrics rather than real-world effects.
Scalability	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).	Scalability is limited. Results are only generated for a small number of nodes (e.g., 10).
Usability	Usability is high. Focus is on the basic network metrics rather than real-world effects.	Usability is high. Focus is on the basic network metrics rather than real-world effects.	Usability is high. Focus is on the basic network metrics rather than real-world effects.	Usability is high. Focus is on the basic network metrics rather than real-world effects.	Usability is high. Focus is on the basic network metrics rather than real-world effects.	Usability is high. Focus is on the basic network metrics rather than real-world effects.	Usability is high. Focus is on the basic network metrics rather than real-world effects.	Usability is high. Focus is on the basic network metrics rather than real-world effects.	Usability is high. Focus is on the basic network metrics rather than real-world effects.

A scorecard the team made that compared candidate simulators across key criteria such as wireless fidelity, mobility, scalability, and usability. It mattered because it informed the backend recommendation and de-risked future scaling decisions.

### Possible Next Steps

#### 1. Add a second simulation backend.

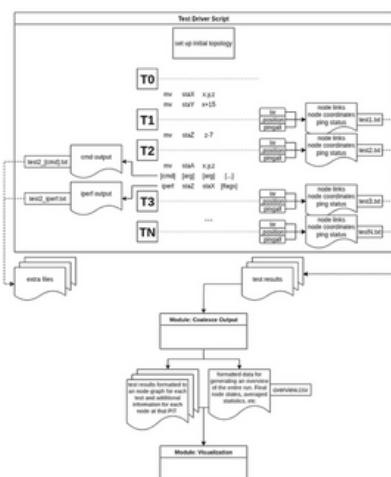
Implement a swappable backend (for example, ns-3) and benchmark fidelity, scale, and runtime using the same JSON scenarios.

#### 2. Increase realism without losing repeatability.

Strengthen time and mobility handling and extend environmental modeling so scenarios better match real conditions while staying comparable.

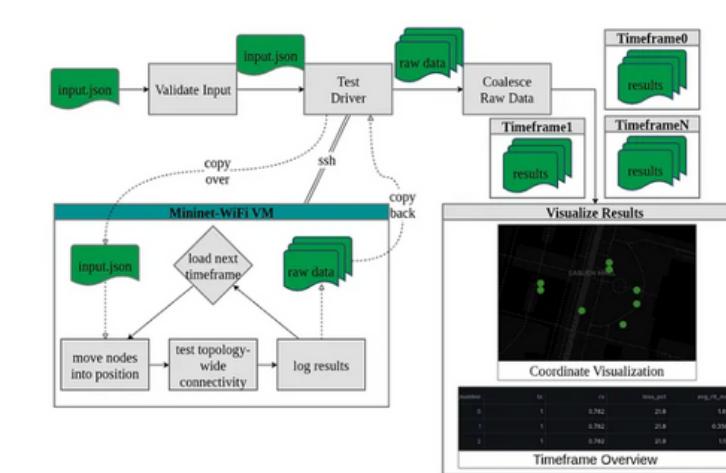
#### 3. Automate scenario exploration and reporting.

Generate controlled scenario variants and produce automatic summaries to highlight breakpoints and mitigation options quickly.



A workflow diagram showed how the test driver executed time-stepped mobility and connectivity checks in Mininet-WiFi, then coalesced outputs into summary files for visualization. It mattered because it demonstrated how experiments were made repeatable and how raw logs became analysis-ready results.

A workflow diagram showed how the test driver executed time-stepped mobility and connectivity checks in Mininet-WiFi, then coalesced outputs into summary files for visualization. It mattered because it demonstrated how experiments were made repeatable and how raw logs became analysis-ready results.



A final architecture diagram showed Omen validating a JSON scenario, executing timeframed runs on a Mininet-WiFi VM, and converting raw outputs into per-timeframe results and visualizations. It mattered because it confirmed an end-to-end, repeatable path from scenario definition to analysis-ready insights.

## Innovation Prompt

## How might we use AI to help detect emerging trends early and turn them into strategy-ready insights?



# MASTER OF SCIENCE IN PRODUCT MANAGEMENT (CAPSTONE)

## Description

This CMU MSPM capstone project investigated how LLM-based multi-agent systems can make innovation trend analysis more systematic, transparent, and repeatable. From a research perspective, it reframed trend identification as a structured pipeline that defines scope, gathers diverse signals across papers, patents, news, and social sources, applies an explicit scoring rubric (credibility, frequency, recency, relevance), and synthesizes evidence into decision-ready insights with traceable provenance. The work contributes an agentic architecture, the Trend Identification Pyramid, that prioritizes explainability and verification to reduce information overload and improve trust in AI-assisted trend sensemaking.

## Engagement Highlights



Interviews across AI, robotics, and neuroscience grounded the project in real workflows and surfaced the biggest pain points behind manual trend analysis.

## Key Learnings

## Decomposition improves reliability.

Breaking trend analysis into specialized agents (scope, sourcing, scoring, synthesis) reduces drift and makes results more consistent over time.

## Trust requires explicit, visible criteria.

Users trust AI outputs more when the system shows why a trend ranked highly using clear metrics like credibility, frequency, recency, and relevance.

## Signal diversity changes the trend picture.

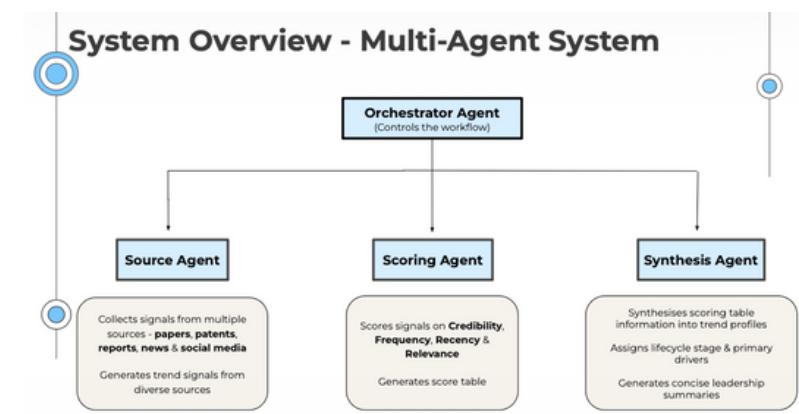
Combining papers, patents, news, and social signals reveals different stages of maturity and helps distinguish hype-driven attention from evidence-backed trends.

## Possible Next Steps

- 1. Longitudinal evaluation of trend accuracy.** Track detected trends over time and compare them against expert assessments and real-world outcomes to measure precision, recall, and time-to-detection.
- 2. Domain-adaptive scoring calibration.** Learn or tune the scoring weights (credibility, frequency, recency, relevance) per industry and time horizon, then test whether calibration improves ranking quality and user trust.
- 3. Provenance and verification UX studies.** Experiment with different ways of presenting citations, evidence trails, and counter-evidence, and measure how they affect user confidence, verification time, and decision quality.



In-person working session in Pittsburgh with the team to review early research insights and shape the direction of the capstone.



System overview image of a multi-agent architecture where an Orchestrator Agent connects to Source, Scoring, and Synthesis agents to collect signals, score them, and generate trend profiles and summaries.

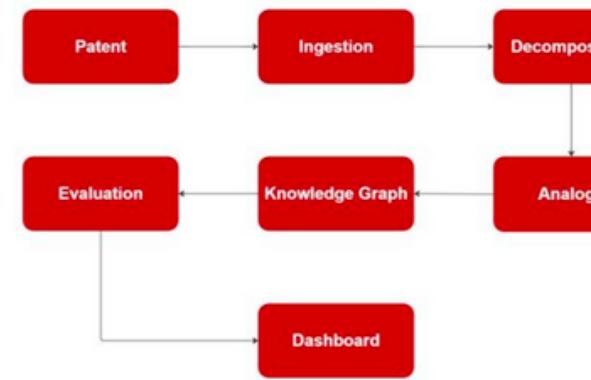
## Innovation Prompt

How might we use AI to help detect emerging trends early and turn them into strategy-ready insights?

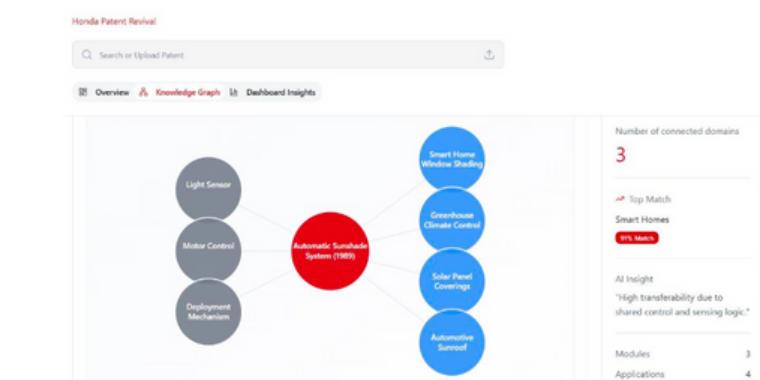
### Description

This project studies an agentic AI framework that converts dormant patents into structured modules, finds cross-domain reuse opportunities via semantic retrieval plus LLM analogical reasoning, and prioritizes them with explainable scoring. It combines patent ingestion and summarization, functional decomposition, a patent-module-domain knowledge graph, and an evaluation layer, demonstrated on two case studies (sunshade and active vehicle control) to show how module-level abstraction can surface non-obvious applications.

### Engagement Highlights



High-level workflow of the patent revival system, showing how a patent moves through ingestion, decomposition, analogy generation, knowledge graph mapping, and evaluation before results appear in the dashboard.



Sunshade knowledge graph linking patent modules to cross-domain applications through similarity-based connections.

### Key Learnings

#### Module-level representations are the right unit for transfer.

Decomposing patents into functional modules enables meaningful cross-domain matching that whole-document similarity often misses.

**Best results come from hybrid “retrieve + reason” analogies, not retrieval alone.** Using semantic similarity to shortlist candidates, then LLM-based analogical reasoning to generate and justify transfers, produces higher-quality, more interpretable opportunity hypotheses.

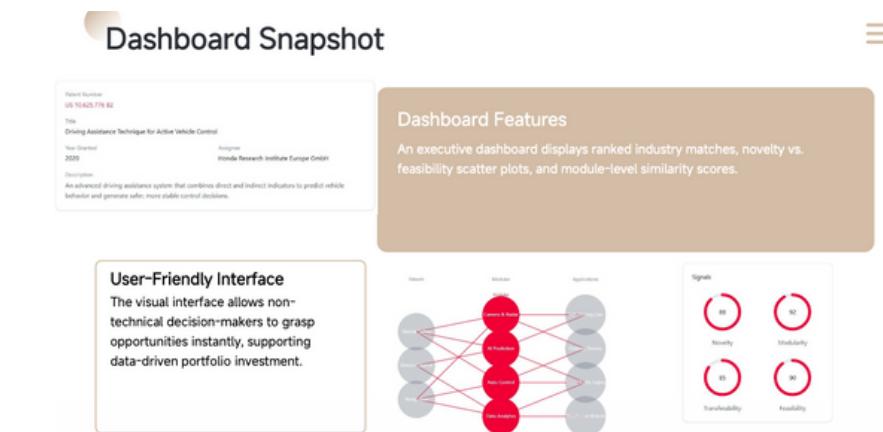
**Trust requires traceability and evaluation hooks.** Prompt templates, chain-of-thought style structured reasoning, prompt versioning/A-B testing, and validation checks (expert review and consistency testing) are essential to make outputs stable enough for decision support.

### Possible Next Steps

**1. Scale and stress-test generalization.** Run the full pipeline on a much larger, more diverse patent set (across domains and writing styles) to measure consistency of module decomposition, stability of analogies, and how performance changes with volume.

**2. Build stronger evaluation benchmarks and human-in-the-loop protocols.** Define gold standards for module boundaries and analogy quality, add inter-rater reliability with domain experts, and quantify scoring calibration (novelty, feasibility, transferability) so rankings are comparable across patents.

**3. Improve representations and reasoning constraints.** Experiment with richer module schemas and knowledge-graph features (typed relations, provenance, citations back to claims/sections) and add structured prompting to reduce hallucinations while keeping explanations actionable.



In-person working session in Pittsburgh with the team to review early research insights and shape the direction of the capstone.



Capstone poster summarizing the agentic AI patent revival framework, showing the five-agent pipeline and case-study dashboard results that surface and score cross-domain opportunities from dormant IP.

## Innovation Prompt

How might we detect and model group coordination or disruption in shared spaces using anonymous, edge-processed motion signals?

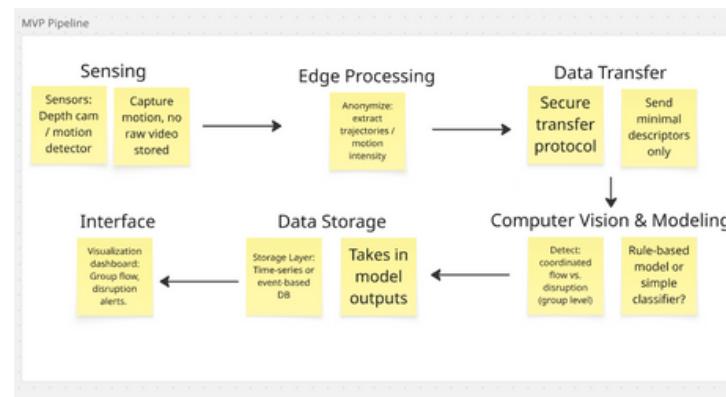


## OSU ELECTRICAL & COMPUTER ENGINEERING (CAPSTONE)

### Description

This project will continue in the spring, building on a first semester that was mostly about getting the problem definition right and designing a feasible approach. The team set out to understand how we can measure and interpret group coordination in shared spaces without relying exclusively on cameras, aiming for a privacy-preserving system that can still capture meaningful motion patterns. Over the fall, we clarified what “coordination” should mean at a high level, aligned on key requirements (accuracy, secure handling of data, and the ability to scale to multiple groups), and mapped out an end-to-end system concept using edge-based sensing and on-device processing.

### Engagement Highlights



MVP pipeline diagram showed the end-to-end flow from sensing and edge anonymization through secure transfer, modeling, storage, and the dashboard. It reduced implementation risk by clarifying required components and handoffs.

### Key Learnings

**Privacy and trust drive the sensing approach.** The team learned that understanding group interaction has to be balanced with strong privacy constraints, which pushes the design toward anonymous, edge-processed sensors instead of traditional camera-based methods.

**Off-the-shelf work gets you partway, but not to “groups.”** Radar-based activity recognition (for example, RadHar) supports the choice of mmWave radar because it can capture rich motion signals while preserving privacy, but it is largely centered on individuals, so “group coordination” needs additional modeling and definitions.

**Even with a feasible end-to-end edge setup (mmWave radar streaming point-cloud and velocity over UART to a Raspberry Pi), preprocessing is a first-class problem.** Early progress showed that reliable ingest and preprocessing have to come before ML.

### Team Requirements Based on Sponsor Needs

Needs	Requirement	Units	Range	Ideal
Accurately capture motion data	Accuracy	Percentage	80-100%	100%
Securely encrypt all data-at-rest	Encryption	Percentage	100%	100%
Smoothly handle increase in number of tracked groups	Scalability	Groups	2-5 Groups	5 Groups
Easily read the interface	Readability	Seconds	TBD	TBD

Sponsor requirements table summarized target metrics for accuracy, encryption-at-rest, scalability, and interface readability. It aligned the team on measurable success criteria and guided design tradeoffs throughout development.

### Possible Next Steps

#### 1. Deploy the full edge pipeline on hardware.

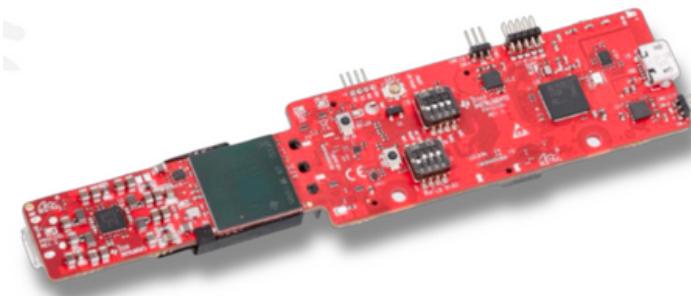
Stream mmWave point-cloud and velocity data into the Raspberry Pi and run real-time preprocessing and inference end to end.

#### 2. Train and validate a model for group coordination.

Implement the proposed two-layer approach and evaluate how well it recognizes coordination, neutral, and disruption states from point-cloud sequences.

#### 3. Harden the system for real-world use.

Build the session-summary output and ensure requirements like encrypted data-at-rest and scaling to multiple tracked groups are met.



Type: IWR6843

Hardware selection slide showed the TI IWR6843 mmWave radar chosen as the primary sensor for privacy-preserving motion capture. It de-risked the build by locking in a sensor capable of producing the point-cloud data needed for downstream modeling.



Figure 1: Example of 3D human body segmentation.

(a) Generated Point Cloud as Input

(b) Segmented Result [1]

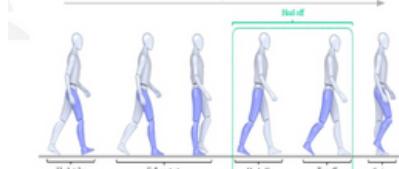


Figure 2: Four gait sub-phases according to

3-ch FSR measurement results [2]

Layer 2 software concept slide outlined a PointNet+LSTM approach to classify collaborative movements (walking, sitting, leaning) from radar point-cloud sequences. It provided a concrete modeling plan that linked raw sensor data to the higher-level behaviors needed for coordination detection.

# Innovation Prompts

How might we reduce everyday friction in how people move, coordinate, and make decisions in shared spaces?

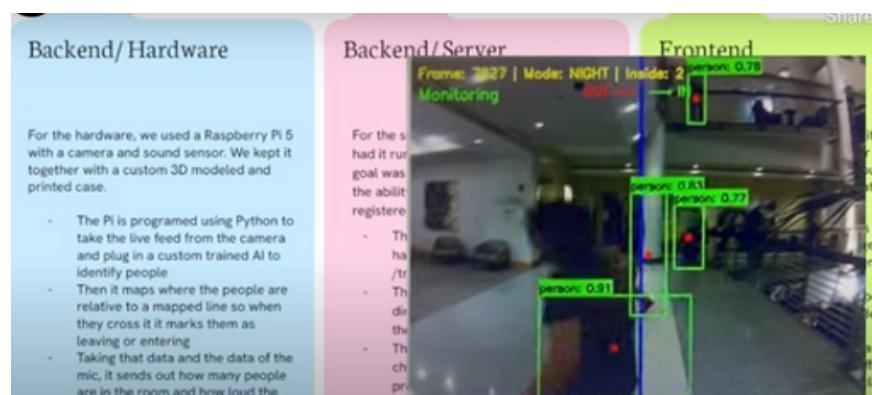


## OSU HACK OHI/O (HACKATHON)

### Description

Hack OHI/O 2025 brought students from diverse disciplines to The Ohio State University for a weekend of building, experimentation, and problem solving, resulting in working prototypes that tackled real challenges in daily life and campus mobility. As a sponsor and partner, 99P Labs hosted the Friction Finder Challenge, inviting teams to identify common moments of movement and coordination friction, like crowding, parking search, and route uncertainty, and design solutions that made these experiences smoother and more efficient. Many teams leveraged modern AI tools to accelerate development and improve polish, producing more complete, practical demos than in prior years and highlighting how accessible AI is reshaping student innovation.

### Engagement Highlights



Optimap took first place by combining low-cost sensing with AI-driven prediction to estimate and forecast crowding and noise in campus spaces, helping students choose less congested places in real time.

### Key Learnings

#### AI raised the prototype baseline.

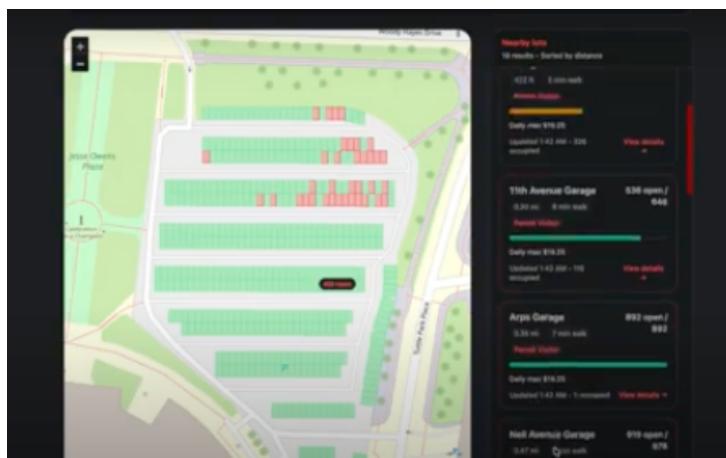
With AI tools helping with coding, design, and analysis, teams shipped more complete, polished demos in a weekend and spent more time on impact and usability.

#### Mobility friction is everywhere and highly relatable.

The strongest projects started with a specific, observable pain point (crowding, parking search, route uncertainty) and translated it into a clear “before vs. after” improvement.

#### End-to-end thinking wins.

Teams that combined solid technical execution with user experience, data flow, and safeguards (reliability, privacy, clarity of communication) stood out over isolated model performance or novelty.



OSU SmartParking earned second place by using privacy-aware computer vision to estimate real-time parking availability and display it on a simple, filterable map to reduce the time spent searching for open spots.

### Possible Next Steps

#### 1. Launch a rapid follow-on research sprint.

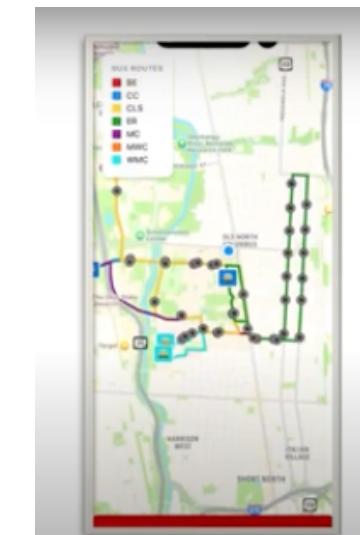
Offer a short, structured program after the event where promising teams get mentorship and a clear evaluation plan to turn demos into validated findings.

#### 2. Create “challenge starter kits” for reproducible work.

Provide lightweight templates (data schema, baseline model, evaluation script, reporting format) so teams can produce results that are comparable and easy to build on next year.

#### 3. Build a campus testbed network for mobility experiments.

Partner with a few willing sites to share safe data feeds or pilot locations, letting teams design with real constraints and generate real-world evidence quickly.



gOSU earned third place by turning live campus bus data into an intelligent transit assistant that automatically recommends the fastest route between two locations, reducing guesswork and wait-time friction.



99P Labs proudly sponsored Hack OHI/O 2025, supporting student innovation through the Friction Finder Challenge focused on reducing everyday mobility friction.

# Innovation Prompts

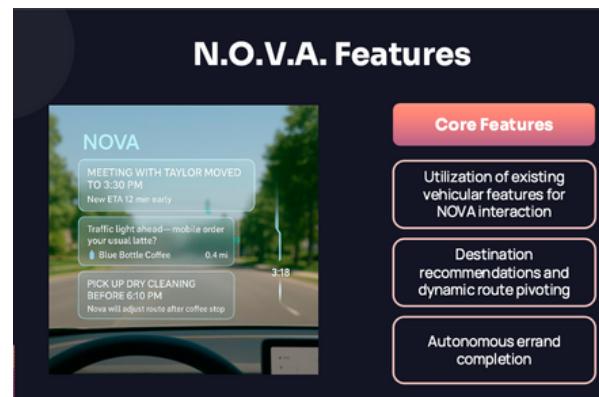
How might we conceptualize a proactive assistant experience that anticipates needs and offers timely help while remaining transparent, privacy-first, and always under the user's control?

How might we design an interface that adapts its complexity and interaction modes to different people and situations so advanced technology feels intuitive and low-friction?

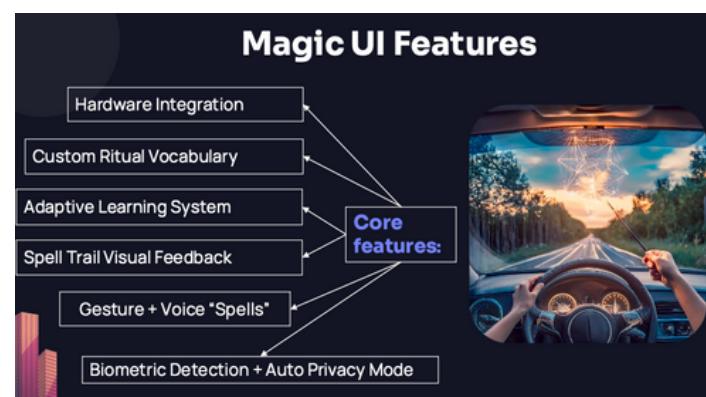
## Description

This capstone project explored how emerging AI and interface design can better support people as technology becomes more present, complex, and personalized. It included two subprojects: AI Guardian, which investigated the concept of an assistive “guardian” that proactively helps users with everyday needs while prioritizing trust, privacy, and safety; and Variable Tech UI, which examined concepts around adaptable user interfaces that can flex across different user preferences, contexts, and comfort levels to reduce friction and make advanced capabilities feel intuitive.

## Engagement Highlights



NOVA is one of their concepts: a predictive vehicle assistant that uses signals like location and biometrics to anticipate needs, adjust routes in real time, and complete tasks hands-free with low-effort personalization.



Magic UI is another concept: a personalized interaction system that replaces traditional menus with user-defined gestures and voice commands, using biometric recognition to load a driver's custom “vocabulary,” adapt to who's present, and enable automatic privacy protection.



## Key Learnings

**Trust and control drive adoption.** Across both concepts, people consistently worried about privacy and unwanted access to personal data, and they wanted clearer controls and transparency over what the system knows, shares, and does.

**Personalization matters, but it must stay lightweight.** Users value customization and self-expression, yet they also want experiences that work “out of the box” with minimal setup and effort.

**Reduce cognitive load with flexible interaction modes.** People want tech that makes tasks faster and easier (not more distracting), and they prefer efficient, context-aware interaction options, leaning toward physical interaction with voice as an optional fallback when hands are busy.

## Possible Next Steps

### 1. Launch a rapid concept validation sprint.

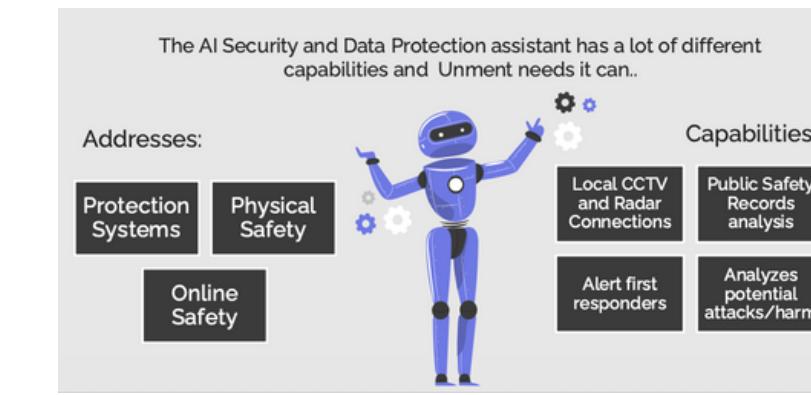
Run short, structured user studies with scenario walkthroughs and lightweight prototypes to confirm top use cases, perceived value, and key adoption blockers.

### 2. Define a trust and control framework.

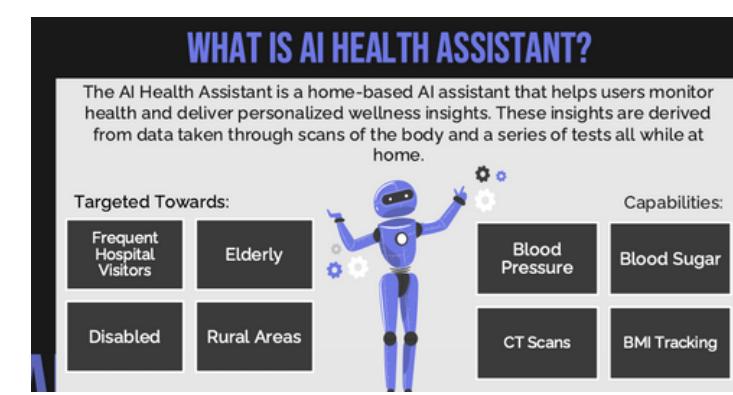
Specify what data is used, where it lives, how actions are approved, and how users can inspect and override behavior, then translate that into consistent UX patterns.

### 3. Build a narrow end-to-end MVP for each concept.

Implement one core workflow per concept to demonstrate feasibility and measure usability, trust, and value with a concrete demo.



One of the AI Guardian concepts was a safety and security assistant that uses nearby sensors and public safety data to detect potential threats and trigger timely alerts or responder support.



Another AI Guardian concept was a home-based health assistant that helps people monitor key metrics and receive personalized wellness insights from at-home scans and tests.

# Innovation Prompt

How might we transform the research ideation process from a solitary activity into a dynamic, collaborative dialogue between a researcher and an intelligent agent?

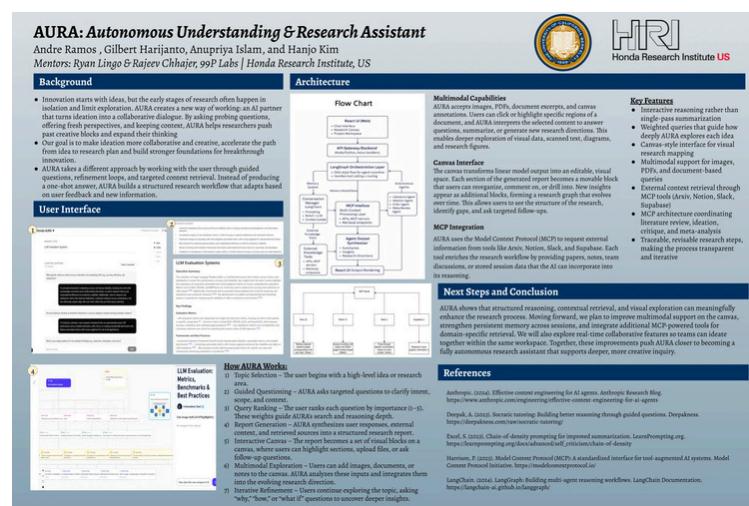


UNIVERSITY OF CALIFORNIA, BERKELEY (CAPSTONE)

## Description

AURA (Autonomous Understanding & Research Assistant) is a retrieval-first research ideation system that turns early-stage curiosity into prioritized, evidence-grounded research directions. It models ideation as an iterative workflow: it asks targeted clarifying questions to surface goals and constraints, generates and ranks candidate research queries, retrieves relevant sources (for example, from arXiv), and synthesizes findings into a structured, citation-backed report. AURA then supports deeper exploration through a visual canvas that maps themes, links evidence to claims, and enables context-aware follow-ups, aiming to reduce hallucination risk and make research exploration more transparent, systematic, and actionable.

## Engagement Highlights



Research poster showing an overview of AURA, including the problem, retrieval-first architecture, workflow stages, and key outcomes.

## Key Learnings

### Ideation is a multi-stage process, not a prompt.

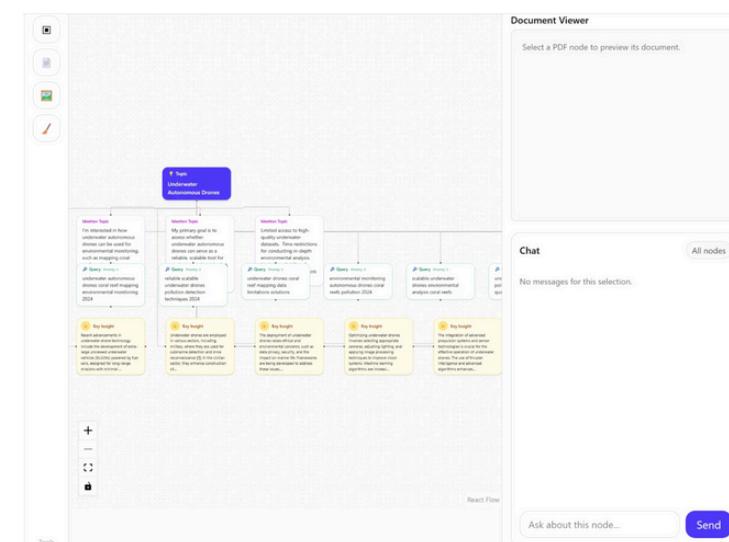
The highest leverage comes from structured steps (clarifying questions → query expansion → prioritization) that turn vague interests into testable, scoped directions.

### Retrieval-first beats “model memory” for research reliability.

Grounding outputs in external sources and keeping citations tied to claims is essential to reduce hallucinations and make results auditable.

### Curation matters as much as retrieval.

Relevance ranking, filtering stale/low-signal material, and explicitly handling source disagreement improves both output quality and user trust compared to “more context is better.”



Canvas view showing a visual workspace where AURA maps and clusters ideas, links evidence to themes, and supports node-based follow-up exploration.

SET PRIORITY WEIGHTS (1-5)

Assign a priority weight to each answer. Higher weights (5) indicate more important aspects to focus on during research.

What specific aspect of “Underwater Autonomous Drones” interests you most?

I’m interested in how underwater autonomous drones can be used for environmental monitoring, such as mapping coral reefs or detecting pollution.

Priority: 4 1 2 3 4 5

What is your primary goal or outcome for this research?

My primary goal is to assess whether underwater autonomous drones can serve as a reliable, scalable tool for environmental monitoring, with a focus on coral reef mapping and pollution detection.

Priority: 5 1 2 3 4 5

Are there constraints (time, data, tools) we should consider?

Limited access to high-quality underwater datasets. Time restrictions for conducting in-depth environmental analysis. Limited availability of simulation or testing tools for underwater drones.

Priority: 3 1 2 3 4 5

Back to Questions Continue with These Weights

Refine and respond screen showing the step where users edit and answer AURA’s clarifying questions to shape the research scope and intent.

RESEARCH REPORT

1 question • 81 sources • 12/11/2025, 7:30:33 AM

Regenerate Download MD Copy MD Download JSON Copy JSON

**Underwater Autonomous Drones**

**Executive Summary**

Underwater autonomous drones, also known as Unmanned Underwater Vehicles (UUVs), are revolutionizing marine exploration, military applications, and environmental monitoring. These drones are capable of operating autonomously, performing tasks such as surveillance, data collection, and even submarine hunting. The development of these technologies is driven by advancements in robotics, artificial intelligence, and sensor integration. Key insights from recent literature highlight the importance of optimizing underwater vision systems, addressing ethical and environmental concerns, and implementing robust control strategies for effective deployment.

**Key Findings**

**Technological Advancements**

Recent advancements in underwater drone technology include the development of extra-large uncrewed underwater vehicles (KLUUV) powered by fuel cells, designed for long-range missions with minimal supervision [1]. The integration of autonomous underwater gliders with UAVs and USVs for deployment and recovery is also a significant development, enhancing the operational capabilities of these drones [2].

**Applications and Use Cases**

Underwater drones are employed in various sectors, including military, where they are used for submarine detection and mine reconnaissance [3]. In the civilian sector, they enhance construction site assessments and environmental monitoring by providing detailed underwater imagery and data [4].

Sample report showing AURA’s structured research output with an executive summary, key findings, literature references, and actionable next steps grounded in sources.

## Innovation Prompt

How might we evaluate an LLM agent's context strategy so we can reliably improve performance while reducing token and time cost?

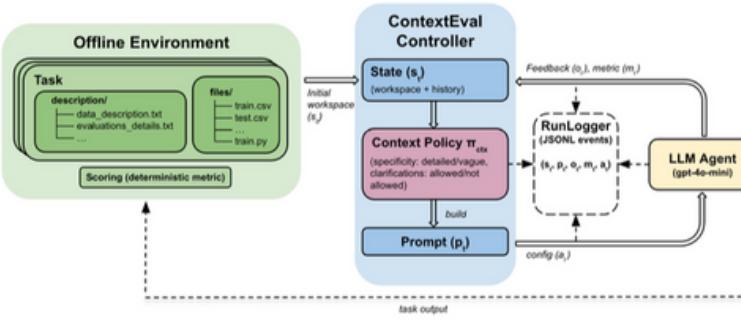


UNIVERSITY OF CALIFORNIA, SAN DIEGO (CAPSTONE)

### Description

This quarter, the team developed ContextEval, an evaluation layer for LLM-based systems that measures how context policy choices (what task information, history, and tool guidance the model can use) affect reliability, performance, and cost in an end-to-end workflow. The framework pairs an offline ML environment with a controller that builds prompts under different policies, runs a fixed LLM agent through iterative steps, and logs structured JSONL traces to compare outcomes, latency, and behaviors across settings. Early results show clear trade-offs across short vs. long context and agentic vs. controller modes. The project will be continued next quarter with a larger controlled study across more policies and seeds and deeper trace-based analysis to produce actionable guidance.

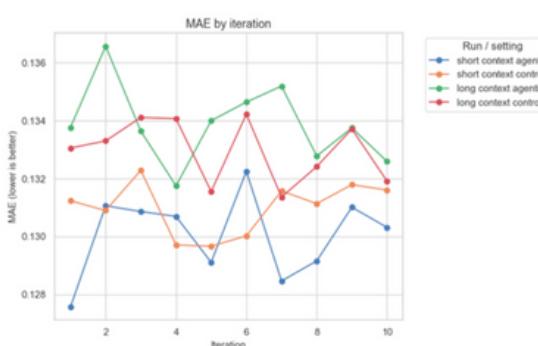
### Engagement Highlights



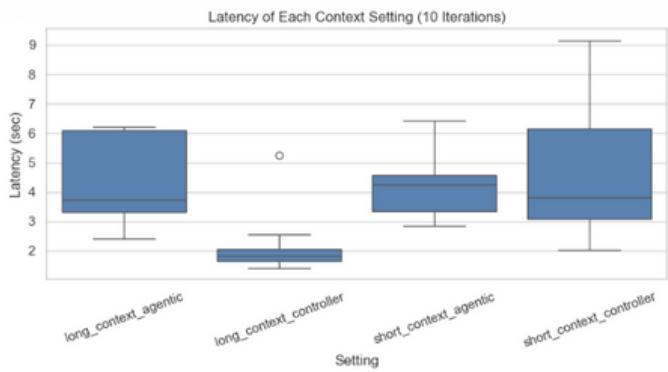
ContextEval architecture showing an offline ML environment that scores agent outputs, a controller that applies a configurable context policy to build prompts for a fixed LLM agent, and a JSONL run logger that records state, actions, and feedback for analysis.

Field	Description
run_id	Unique identifier for the run (episode)
event_type	Type of event, e.g., run.start, op.train, step.summary
step_idx	Iteration index $t$ (integer) or null for run-level events
timestamp	UTC timestamp in ISO-8601 format
task_id	Logical task name (e.g., "toy_tabular", "nomad")
dataset_id	Dataset name; often matches task_id
agent_id	Identifier for the LLM agent configuration (e.g., gpt-40-mini)
details	Event-specific payload (see below)

Top-level JSONL event schema for run traces, including run and step identifiers, timestamps, task and dataset metadata, agent configuration, and an event-specific details payload.



Validation MAE over 10 NOMAD iterations comparing short vs. long context policies and agentic vs. controller modes.



This is an image of a boxplot showing the distribution of per-iteration wall-clock latency on NOMAD across four context-policy and reasoning-mode settings.

## Innovation Prompt

How might we define explanation quality dimensions that are predictive of learning outcomes and can be measured?

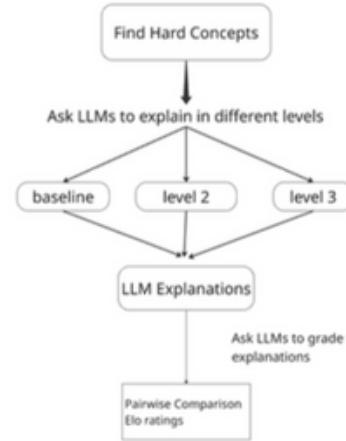


UNIVERSITY OF CALIFORNIA, SAN DIEGO (CAPSTONE)

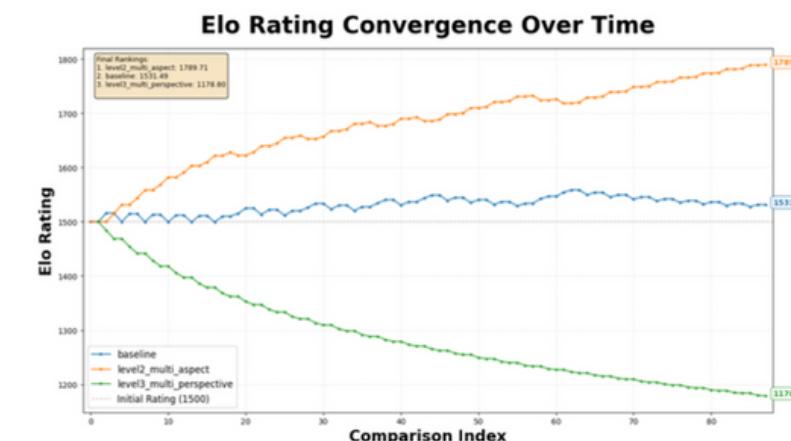
### Description

This UC San Diego capstone investigates whether automated evaluation can serve as a reliable instrument for studying explanation quality in large language models. It treats prompt design as an experimental variable, generates explanations under systematically different prompting regimes, and uses pairwise preference modeling (Elo-style ranking) to estimate which strategies produce explanations that align with human judgments. The project will continue next quarter by scaling the benchmark, adding rigorous human evaluation to validate and calibrate the LLM-as-judge, and testing whether critique-guided iteration yields reproducible gains under human preference rather than only automated scores.

### Engagement Highlights



This image shows the end-to-end pipeline: select hard concepts, generate explanations with three prompt levels, then rank them via pairwise LLM judging and Elo scores.



This image shows Elo ratings converging over pairwise comparisons, with the level2\_multi\_aspect prompt steadily rising to the top, the baseline staying near the middle, and the level3\_multi\_perspective prompt declining to the lowest score.

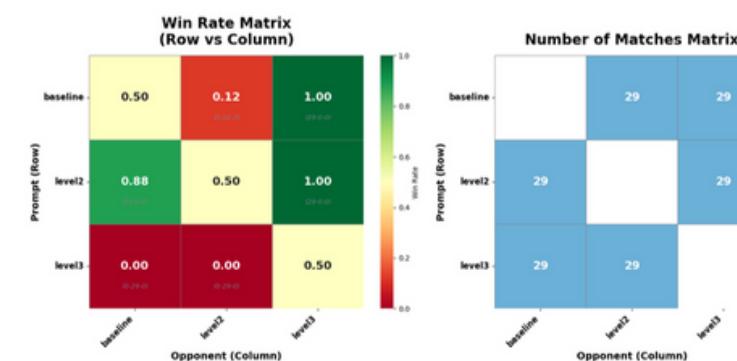
### Key Learnings

**Moderate structure beats heavy structure**, as flexible prompts significantly outrank rigid ones. Light scaffolding improves accessibility, while complex instructions reduce interpretability. Programs help train with less data.

**Pairwise preference with Elo is a practical evaluation method** that uses reverse-order controls to minimize bias. This yields stable rankings without requiring absolute ground-truth labels.

**LLM-as-judge must be validated against humans** because automated metrics can diverge from learner needs. Alignment requires human leaderboards and rubric-based ratings to verify results.

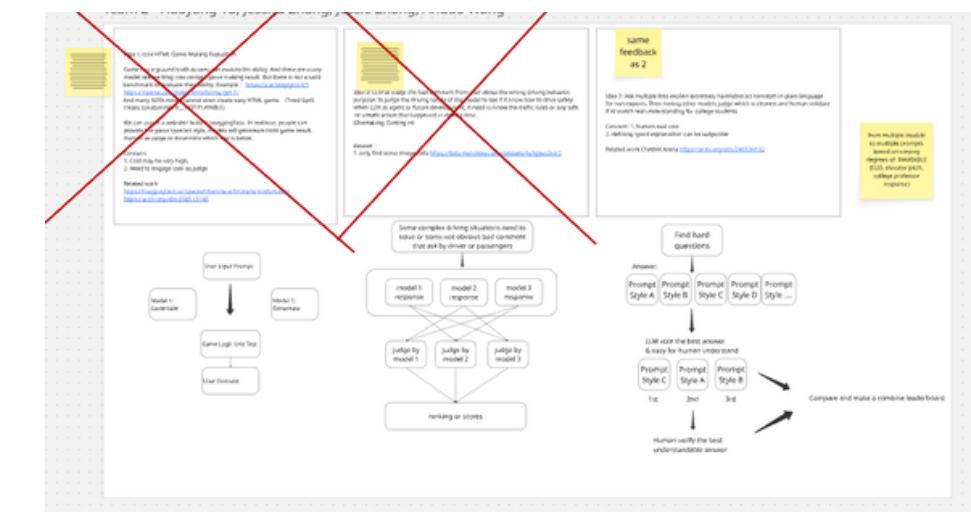
Rank	Prompt	Elo Rating
1	level2_multi_aspect	1789.71
2	baseline	1531.49
3	level3_multi_perspective	1178.80



Final rankings and head-to-head outcomes: level2\_multi\_aspect leads in Elo, baseline follows, and level3\_multi\_perspective trails, with heatmaps showing win rates and match counts for each prompt pair.

### Possible Next Steps

- Scale it up.** Expand from the 30-concept pilot to a few hundred concepts, using the same glossary-driven selection and difficulty filtering. This will tell you if the prompt ranking holds up beyond the small sample
- Bring in more humans.** Run rubric-based human scoring (explainability, complexity, familiarity) and publish a human leaderboard. Then compare it directly to the LLM-as-judge Elo rankings to see where the judge is right or off.
- Iterate with critique.** Use critique to rewrite or refine prompts, regenerate explanations, and rerun the judge pipeline. Track whether those edits actually move the Elo scores and improve agreement with humans.



## Innovation Prompt

How might we evaluate generated information in a way that is scalable, reliable, and aligned with human judgment?



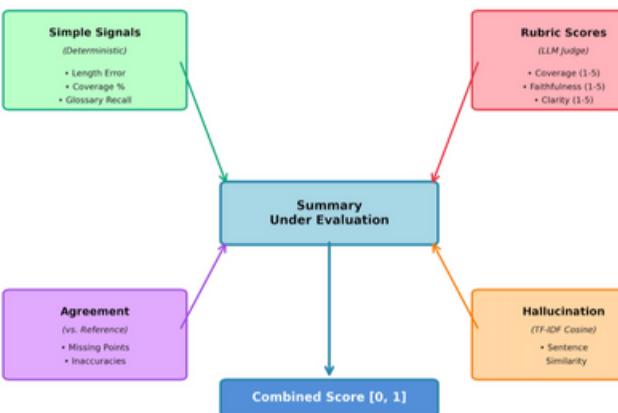
UNIVERSITY OF CALIFORNIA, SAN DIEGO (CAPSTONE)

20

### Description

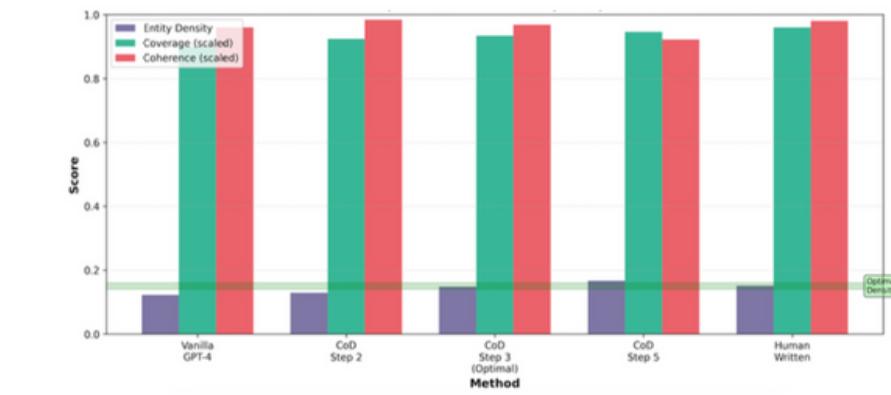
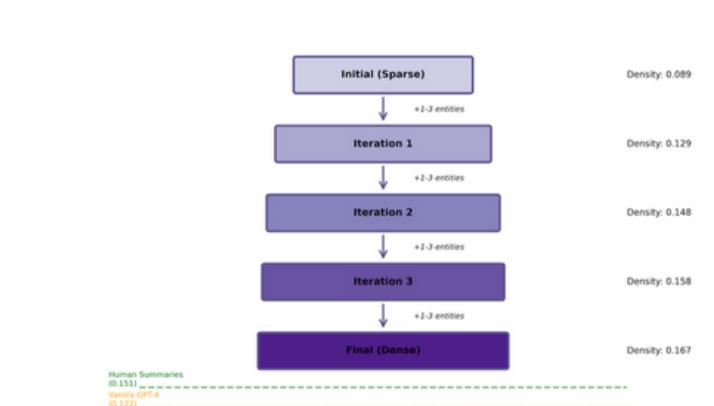
This UC San Diego capstone project is building a scalable system to evaluate the quality of generated summaries using a rubric-driven "LLM-as-judge" alongside lightweight automatic checks, with the goal of producing scores that are reliable, align with human judgment, and provide actionable feedback for improvement. The pipeline supports repeatable benchmarking across inputs and uses evaluation outputs to guide iterative refinement of summaries and reduce errors like missing key points or unsupported claims. The work will be continued next quarter (spring quarter).

### Engagement Highlights



Summary evaluation overview where heuristics, reference agreement, LLM rubric scores, and hallucination checks are combined into a single normalized score.

Chain-of-density summarization increases information density over successive revisions while keeping length fixed, approaching the density of human-written summaries.



Comparison of summarization methods showing chain-of-density increases entity density while maintaining high coverage and coherence, approaching human-written performance.

Feature	Chain of Density	Our System
Iterative refinement	✓	✓
Fixed-length constraint	✓	✗
Entity tracking	✓	✗
Ensemble evaluation	✓	✓
Deterministic metrics	Limited	Comprehensive
Target density	0.15 entities/token	Not enforced
Number of iterations	5 fixed	3 fixed
Domain	News articles	Lecture slides

Feature comparison between the chain-of-density approach and our system, highlighting shared iterative refinement and key differences in constraints, metrics, and target domain.

# Innovation Prompt

How might we design a debates agent to run debates using multiple roles and evidence-based arguments?

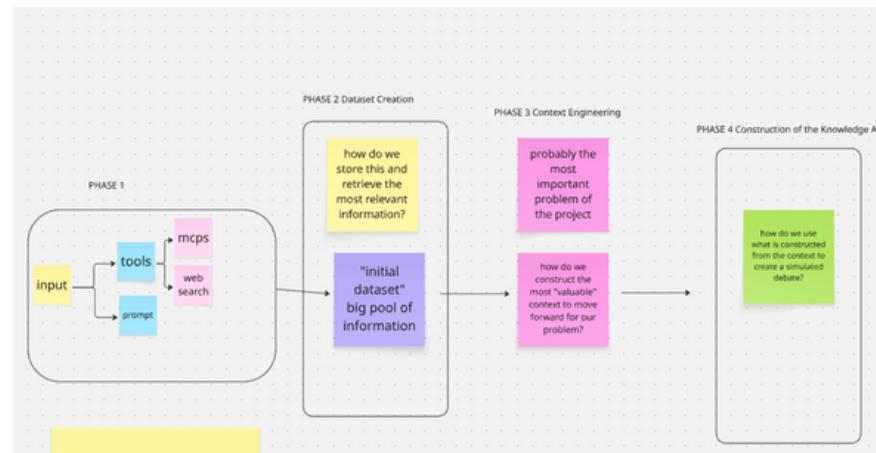


UNIVERSITY OF COLORADO BOULDER (CAPSTONE)

## Description

This project investigates how structured, multi-agent debate can make LLM reasoning more transparent, adversarially robust, and measurable. We built a real-time platform that orchestrates Pro and Con agents over multiple rounds, grounds their arguments with retrieval from a curated document corpus (and optional web sources), and applies a Judge agent with an explicit rubric to score factuality, responsiveness, logic, and evidence use each round. By logging full transcripts, evidence packs, prompts, and per-round scores, the system functions as an experimental testbed for studying debate dynamics, prompt strategies, context management, and failure modes such as repetition, drift, and evaluation format instability.

## Engagement Highlights



Early ideation diagram outlining phases from input/tools to dataset creation, context engineering, and using that context to generate a simulated debate.

## Key Learnings

### Multi-agent systems are mostly state management.

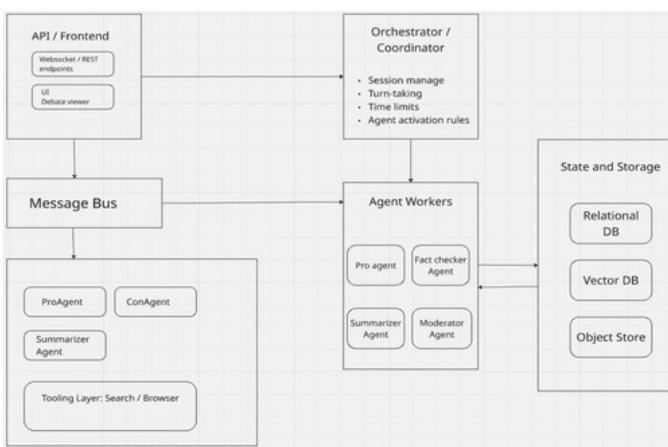
The debate feels coherent only when you maintain shared context, roles, and round structure across turns.

### Prompting and evaluation need engineering, not vibes.

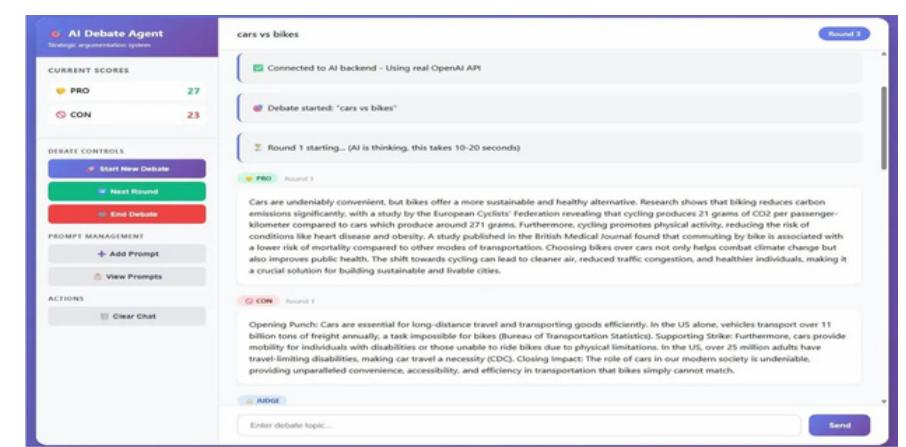
Reliable scoring required strict judge formatting plus robust parsing, retries, and fallbacks when outputs drift.

### Evidence and observability change everything.

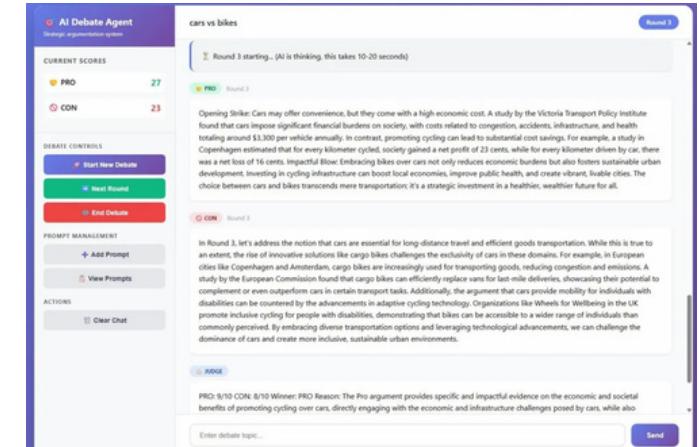
Retrieval-grounded arguments and structured logs made debates auditable and turned transcripts into data for analysis and iteration.



System architecture for the debate platform: frontend to orchestrator/message bus to agent workers, backed by relational, vector, and object storage.



Screenshot of the debate platform UI showing live Pro/Con chat, round status, scores, and debate controls.



Screenshot of the debate UI during Round 3, showing Pro and Con closing arguments, the live score sidebar, and the Judge's scored verdict (with winner and rationale).

## Possible Next Steps

### 1. Real-Time Claim Verification and Calibration

Research a fact-checker agent that audits each argument against retrieved sources, flags unsupported claims, and reports confidence levels with inline evidence snippets.

### 2. Measuring Debate Quality and Robustness

Study how prompt variants, retrieval settings, and round structure affect outcomes using metrics like score stability, evidence usage, repetition rate, and judge consistency across topics.

### 3. Domain-Specific Debate Protocols and Personas

Explore specialized panels (for example compliance, safety, or scientific review) with tailored rubrics and constraints to evaluate how domain framing changes argument quality and reliability.

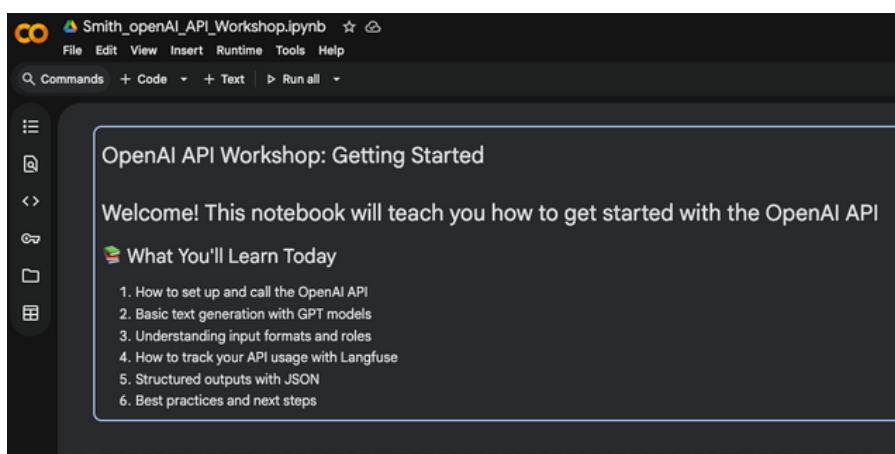
## Innovation Prompt

How might we transform massive, static archives of slide decks into an interactive and reliable knowledge base?

### Description

This capstone project developed a Retrieval-Augmented Generation (RAG) system designed to transform static PowerPoint archives into an interactive and queryable knowledge base. Leveraging Python and OpenAI's API, the pipeline ingests slide decks and converts both text and visual content into searchable embeddings indexed via FAISS. The system specifically uses GPT-4o-mini to generate text descriptions for slide images to ensure all data is captured. The final deliverable is a context-aware chatbot that allows users to accurately query and synthesize insights from historical presentations. This solution effectively automates the retrieval of institutional knowledge and eliminates the need for manual file review.

### Engagement Highlights



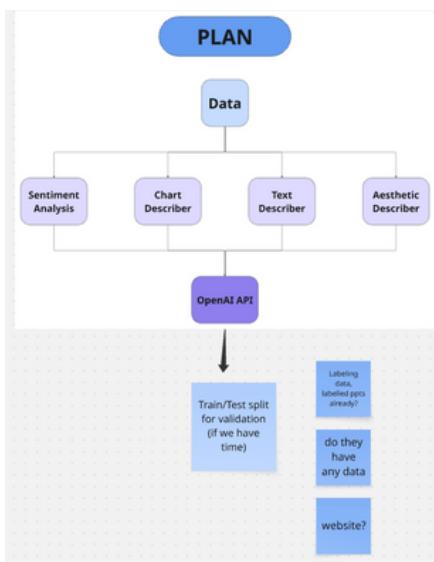
We held a workshop with the student team to show them how to get started with the OpenAI API.

### Key Learnings

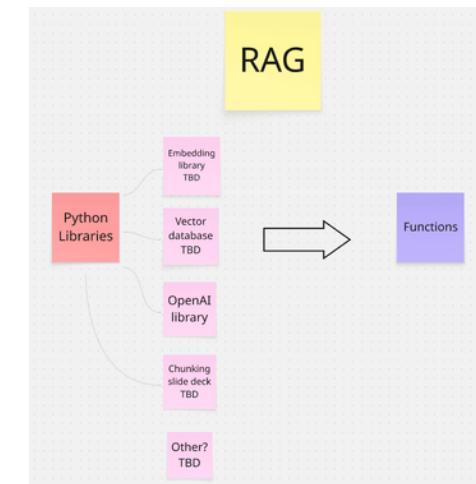
**Retrieval-Augmented Generation (RAG) is critical for grounding LLMs in specific, internal data.** By anchoring the model to retrieved slide content, the team eliminated the hallucinations common in standard, ungrounded AI responses.

**Effective indexing requires converting visual slide elements into searchable text.** The use of vision models to generate captions for charts and images ensured that no critical information was lost during the retrieval process.

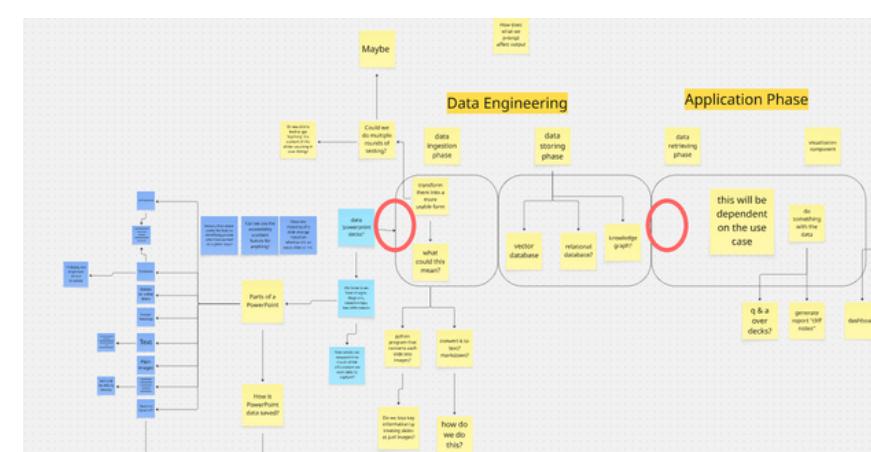
**Transitioning from notebooks to modular code is vital for system scalability.** Although Jupyter Notebooks facilitated early prototyping, the project highlighted that standalone Python files are necessary for managing complex logic and error handling.



This flowchart outlines the team's initial strategy for decomposing slide data into text and visual descriptions for processing via the OpenAI API.



This diagram maps out the preliminary technical stack and core Python libraries required to build the Retrieval-Augmented Generation pipeline.



This whiteboard snapshot details the complete data engineering lifecycle, mapping the process from ingesting raw PowerPoint decks to retrieving them in the final application.

### Possible Next Steps

#### 1. Develop a standalone web interface to replace the current Jupyter Notebook environment.

Moving the chatbot to a dedicated application would significantly improve accessibility and usability for non-technical stakeholders.

#### 2. Expand the ingestion pipeline to capture speaker notes and hidden metadata.

Including these overlooked data sources would provide the RAG system with deeper context and improve search accuracy.

#### 3. Refactor the codebase into modular scripts for improved error handling.

Transitioning from a notebook prototype to structured production code would increase system stability and ease of maintenance.

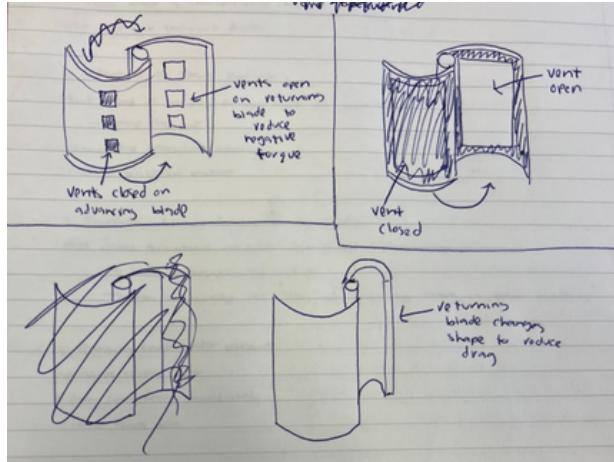
# Innovation Prompt

How might we design a compact, efficient, and low-noise vertical axis wind turbine that performs reliably in turbulent rooftop conditions and integrates seamlessly with urban architecture?

## Description

This capstone project developed an urban-ready vertical-axis wind turbine (VAWT) concept for campus and building-edge locations where winds are turbulent and multi-directional. The team combined technical modeling with a survey of commercial VAWTs and a structured concept-generation process to identify high-impact design features, set realistic energy expectations, and organize candidate Darrieus, Savonius, and hybrid approaches for evaluation. The work will continue in Spring 2026 with concept downselection, prototyping, and performance testing in realistic urban wind conditions.

## Engagement Highlights



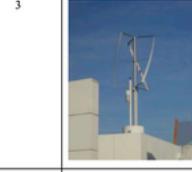
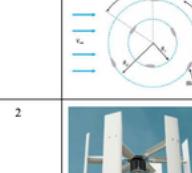
Early prototype sketch of a vented Savonius rotor with vents opening on the returning blade to reduce drag and negative torque.

## Key Learnings

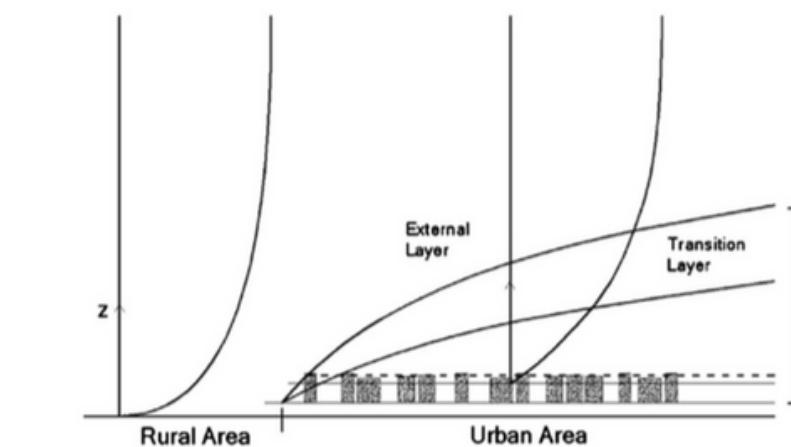
**VAWTs are better suited than HAWTs for turbulent, multi-directional urban winds.** That makes “messy wind” performance a primary design requirement for campus and building-edge deployment.

**Commercial VAWTs repeatedly use hybrid rotors and load-management concepts to improve operating range and durability.** Those patterns are strong candidates to borrow when designing for real installations.

**Small VAWT energy production needs conservative assumptions to avoid overpromising.** Using conservative assumptions when estimating annual energy helps set credible project goals and impact.

Concept	Description	Level of Development (1-3)	Image	Reference number
16	3-blade helical darrieus	3		[9]
17	Doubling the number of blades used (3 blades to 6 blades)	2		[10][11]
18	Using a horizontal Savonius profile combined with a vertical Darrieus one	2		[10]
19	Anodizing the blades	2		[12]

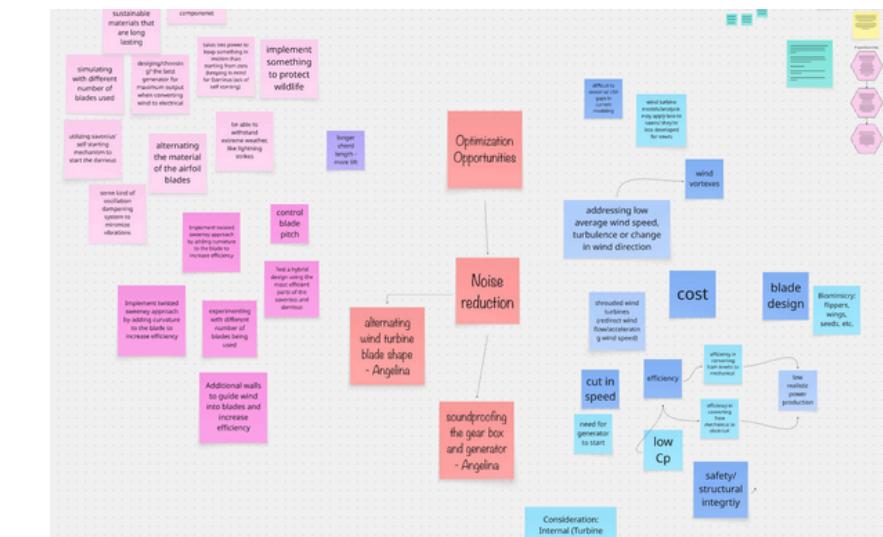
Sample from the Concept Generation Memo summarizing hundreds of wind turbine design concepts and grouping them by turbine type and maturity.



Research figure comparing rural and urban wind profiles in the urban boundary layer, illustrating increased shear and turbulence that favors VAWTs in built environments.

## Possible Next Steps

- 1. Downselect to 1-2 candidate rotor architectures for urban winds.** Use performance, manufacturability, safety, and maintenance criteria to choose a lead Darrieus, Savonius, or hybrid concept for prototyping.
- 2. Build and test a scaled prototype in representative turbulent flow.** Run wind-tunnel, fan-array, or outdoor testing to measure startup behavior, torque ripple, noise, and power output under changing wind direction.
- 3. Iterate on performance-enhancing and durability features.** Evaluate add-ons like guide vanes/deflectors or endplates for efficiency gains, and refine structural/load-management components to improve reliability and lifespan.



Early brainstorming board mapping key VAWT challenges and opportunities, including blade design, cut-in speed, low efficiency, cost, turbulence, and noise reduction.

# CONCLUSION

As we wrap up this edition of the 99P Labs' University Innovation Impact Report, we extend our sincere gratitude for your engagement and the time you have dedicated to reviewing our endeavors from the fall semester of 2025. Your interest is the cornerstone of our continued pursuit of excellence in academic and research collaborations.

This document is not only a reflection of our past efforts but also a stepping stone towards future collaborations. We invite you to share your insights or express your interest in joining us on this journey of innovation and discovery.

Should you wish to discuss the content of this report further or explore avenues for collaboration, please do not hesitate to reach out to Rajeev Chhajer at [rajeev\\_chhajer@honda-ri.com](mailto:rajeev_chhajer@honda-ri.com) or Ryan Lingo at [ryan\\_lingo@honda-ri.com](mailto:ryan_lingo@honda-ri.com).

We are excited about the potential collaborations that may arise from this report and are looking forward to the opportunity to bring these prospects to fruition together.

Thank you once again for your interest and the possibility of future collaborations. Your participation is invaluable to us, and we eagerly anticipate your thoughts and contributions.