

A Multi-Agent System to Identify Trends using LLMs

MSPM Honda B Capstone Project – Process Book



Ankit Shukla

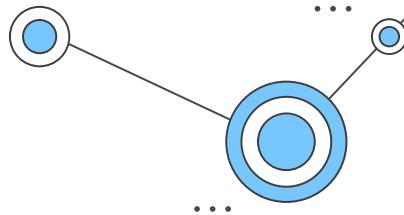


Kritika Rastogi



Himakar YV

About the Team



- ❖ **Ankit Shukla** has over 6 years of experience in Computer Vision and AI from his time at Samsung Research in the Camera Team, and he recently worked on Agentic and Gen AI platforms in the financial services solutions space at ServiceLink.
- ❖ **Kritika Rastogi** holds an MBA and has over 8 years of experience in product marketing and management across five industries and four countries. She has led initiatives in strategy, user research, and product development, and most recently worked in Palo Alto on AI-driven HR Tech SaaS, building workflows with agentic platforms.
- ❖ **Himakar Yanamandra** has over 6 years of AI engineering experience, with dual degrees in Computer Science and Artificial Intelligence and published research in the field. In addition, he pursues creative work as a photographer and content creator, managing multiple social media pages.

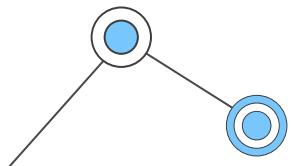


Table of Contents



A Brief Introduction

Setting The Stage



Our Approach

End-to-End Project Overview



Research Plan

Understanding the Landscape



Identifying User Persona & Pain Points

From Interviews to Core Challenges



Brainstorming

Critical Thinking Sessions



Scoping and Prioritization

Research Insights to Solution Priorities



Proposed Solution

Trend Identification Pyramid



Building a Prototype

System Architecture and Demo



Future Scope

Evaluation and Next Steps



Project Deliverables

Final Deliverables and Artefacts



Key Learnings

What We Learned as a team



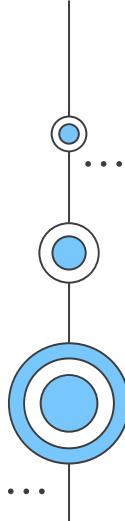
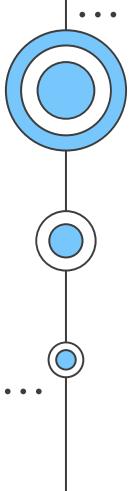
Conclusion

Closing Remarks

01

A Brief Introduction

Setting The Stage



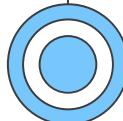


How many of us truly know where the industry is heading?

We are living in a time where technology shifts at an unprecedented pace.

- We try Googling.
- We skim reports.
- We ask GPT.

But we end up with endless tabs, scattered notes, and zero clarity. There is a lot of information out there, but it's difficult and time consuming to make sense of it and connect the dots.



Project Background and Objectives

Context

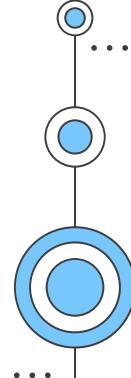
Honda Research Institute/99P Labs is seeking a systematic process to identify external innovation trends and transform them into structured, actionable knowledge outputs.

Problem

Today, while vast amounts of trend-related information already exist, the challenge lies in making sense out of it: the current path to gather, curate, and synthesize this information is either nonexistent or painfully manual.

Opportunity

With the rise of LLMs and AI-driven tools, there is now an opportunity to build an automated pipeline that can bridge this gap, turning unstructured, distributed information into curated knowledge that fuels innovation insights.



Why this Matters for HRI and 99P Labs

There is no current, repeatable process at 99P Labs/HRI to define, discover, and communicate trends relevant to HRI's "beyond next-gen" bets. The team needs a lightweight, evidence-based way to identify trends from public sources and synthesize them into usable outputs (reports/briefs) for downstream tasks. No current process exists at HRI/99P Labs for systematically identifying or quantifying trends.

Downstream tasks:

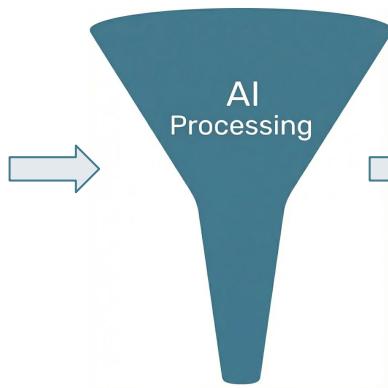
- ❖ **Investment decisions:** Executives need evidence-based insights (e.g., "Which robotics segments are growing fastest?") to guide where HRI should place long-term bets.
- ❖ **Research prioritization:** 99P researchers can quickly spot white spaces or hot areas to shape their project pipeline (e.g., batteries vs AI vs robotics).
- ❖ **External positioning:** 99P Labs can publish trend briefs or LinkedIn posts to establish thought leadership and attract partners.
- ❖ **Internal knowledge management:** Large troves of PDFs/reports become usable trend summaries instead of being buried in drives.
- ❖ **Innovation scouting:** Serves as an early radar to catch emerging technologies before competitors act.

PROBLEM STATEMENT



How might we define a process for using AI to research trends and generate high quality content?

Lot of Information



Automobiles

AI

Energy

Investment Decisions

Research Prioritization

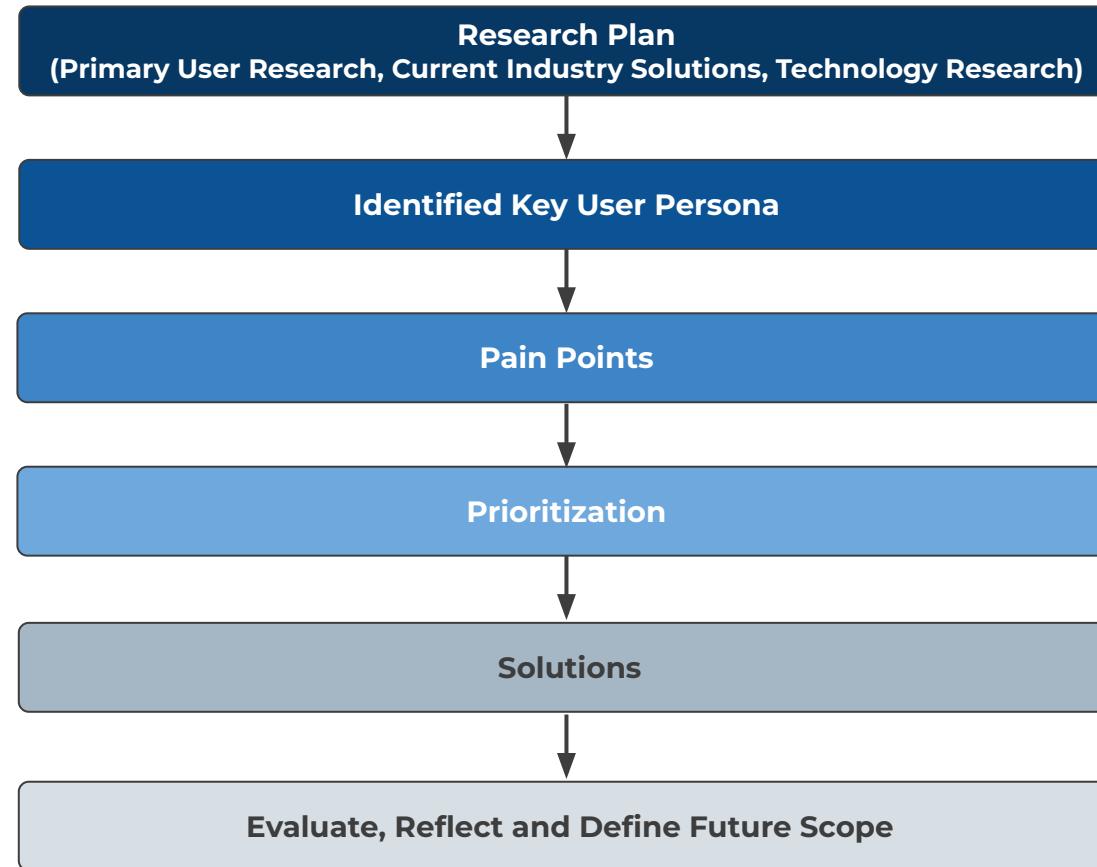
External Positioning

02

Our Approach

End-to-End Project Overview

Our Approach



Our Approach In Detail

We approached this problem by applying product thinking principles.

1. **Research Plan:** Conducted user interviews, researched industry frameworks and explored existing tools to map how trends are identified today.
2. **Identified Key User Persona:** Defined a user persona (Monica Geller) with clear responsibilities and downstream tasks.
3. **Pain Points :** Consolidated recurring challenges into five core pain points around fragmentation, overload, memory, trust and verification.
4. **Prioritization:** Used the MoSCoW framework to decide which capabilities the prototype must deliver first and what can be left for the future roadmap.
5. **Design and Build the Solution:** Designed the Trend Identification Pyramid and implemented a multi-agent system using OpenCode platform.
6. **Evaluate, Reflect and Define Future Scope:** Designed an evaluation framework, stated current limitations and outlined future scope for Honda Research Institute/99P Labs.



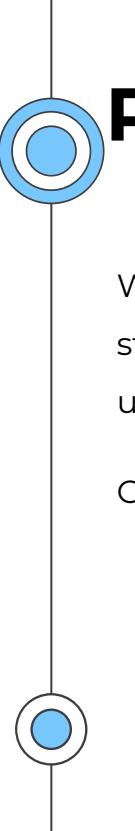
03

Research Plan

Understanding the Landscape

Research Plan and Methodology

- ❖ **Stage 1 – Primary Research - Expert Interviews**
 - Conduct semi-structured interviews with researchers and executives.
 - Map their workflows, sources and decision points.
 - Understand their definition of a trend.
- ❖ **Stage 2 – Current Industry Solutions**
 - Study analyst frameworks for trend identification and technology scanning.
- ❖ **Stage 3 – Technology Research**
 - Evaluate existing LLM-based tools and platforms.



Primary User Research and Interviews

We conducted **interviews with nine experts**, including HRI researchers, CMU faculty and PhD students across artificial intelligence, robotics, automotive technology and neuroscience to understand their customer journey from question to insight

Our goals were to understand:

- ❖ How experts currently **discover and monitor trends**
- ❖ Which **sources they trust** for different kinds of questions
- ❖ Where they **experience friction** and redundancy

Initial learnings from User Interviews: How Different Fields Define Trends



Neuroscience / Biomedical

Slow-moving; trends stabilize only after reproducibility across multiple labs and extensive peer review to confirm maturity.



AI / Language Technology

Very fast cycles; massive post-paper explosion leads to high noise. Hard to separate hype from fundamental impact.



Robotics & Engineering

Requires sustained, reproducible methods. Labs pivot only when core scientific assumptions are actively challenged.



Industry & Market

Action-oriented; prioritizes continuity over time and visible investment signals from major competitors.

Every domain defines 'trend' differently—some chase novelty, others wait for proof—but all struggle to connect evidence over time.

Quotes from User Interviews

“

“Labs don’t pivot because of a single spike on Twitter/Google News. Labs pursue core problems. We’d only reconsider if a core assumption is broken.”

”

“

Example: “If my assumption is ‘robots must work with humans’ and someone demonstrates robots that don’t need humans, that would make me rethink.”

”

“

“A self-driving crash is a news trend, not a research trend.”

”

Current (Manual) Trend Extraction Process

Literature Review & Search

Relying heavily on standard library searches (Google, Scholar) and manual database queries.

Intuition & Internal Synthesis

No standardized process; experts rely on personal experience to validate and synthesize signals.



Conference Scanning

Following key speakers, monitoring accepted papers, and attending high-value industry events.

Peer & Competitor Discussion

Informal discussions combined with following funding announcements and competitor investment signals.



Key Learnings



Key patterns from these interviews included:

- ❖ Workflows are **manual and fragmented**. Experts routinely combine conference talks, academic papers, patents, internal reports and ad hoc web searches.
- ❖ Time horizons differ widely. Some signals relate to near term commercialization, while others relate to long running research agendas measured in years.
- ❖ Each researcher maintains their own tools and notes. There is no shared pipeline or common vocabulary for what constitutes a trend.



A recurring theme was that **information volume was not the problem**. The missing piece was a **connected and explainable structure** for moving from raw evidence to trend level statements that can be shared and defended.

Industry Frameworks

In parallel, we analyzed reports and frameworks from leading firms such as Gartner and McKinsey to understand how trend identification is handled in the broader market.

These frameworks:

- ❖ Provide strong high-level structure for technology lifecycles and adoption curves
- ❖ Offer useful language for describing trends and drivers

But they are not optimized for HRI's needs because:

- ❖ They target broad, market-level views rather than **domain-specific research questions**.
- ❖ They do not provide a reusable, **LLM-ready workflow** that can be run frequently within a research organization.

Gartner

BCG

McKinsey
& Company

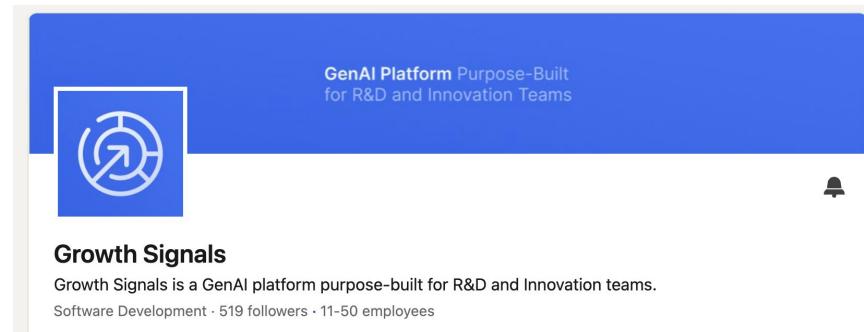
Deloitte.

pwc

Industry Research - Growth Signals

We also explored adjacent work in the trend-intelligence space. As part of this, we met with the team at **Growth Signals**, who are building a research platform focused on surfacing market and technology developments for innovation teams. They provided us access to their tool and walked us through how it aggregates external signals into dashboards and alerts.

This exploration was useful to see how others structure research workflows, but the platform was not directly usable for our specific goal of defining a transparent, LLM-driven workflow for identifying and scoring trends at HRI. It reinforced our decision to focus on designing the underlying process and architecture rather than adopting an existing tool.



Growth Signals

The screenshot shows the homepage of the Growth Signals website. The header features the Growth Signals logo and the URL www.growthsignals.com. The main content area has a light gray background with a central callout box. The callout box contains two sections: 'Research' and 'Concept'. The 'Research' section includes icons for a folder and a magnifying glass, and lists tasks like 'Analyze your docs & source material', 'Discover scholarly articles & publications', 'Explore the patent database', and 'Search the entire web'. The 'Concept' section includes an icon of a briefcase and lists tasks like 'Generate venture concepts from data', 'Define initial pitch & business model', 'Assess market size and opportunities', and 'Identify potential partners & suppliers'. A blue call-to-action button in the center of the callout box reads: 'Let's see how the Growth Signals GenAI platform can accelerate your research and innovation workflows.' Below the callout box, there is a blue button with a right-pointing arrow. The left sidebar contains navigation links for Home, Source, Research, Concept, Pipeline, and Feedback, each with a corresponding icon. The 'Home' link is currently selected, indicated by a gray background and a white icon.

www.growthsignals.com

Growth Signals

Home

Source

Research

Concept

Pipeline

Feedback

Venture diligence for internal innovation.

Research

- Analyze your docs & source material
- Discover scholarly articles & publications
- Explore the patent database
- Search the entire web

Concept

- Generate venture concepts from data
- Define initial pitch & business model
- Assess market size and opportunities
- Identify potential partners & suppliers

Let's see how the Growth Signals GenAI platform can **accelerate** your research and innovation workflows.

→

Technology and LLM Tool Landscape

We explored general-purpose LLM tools and deeper research platforms to assess their capabilities in trend identification and analysis.

Findings:

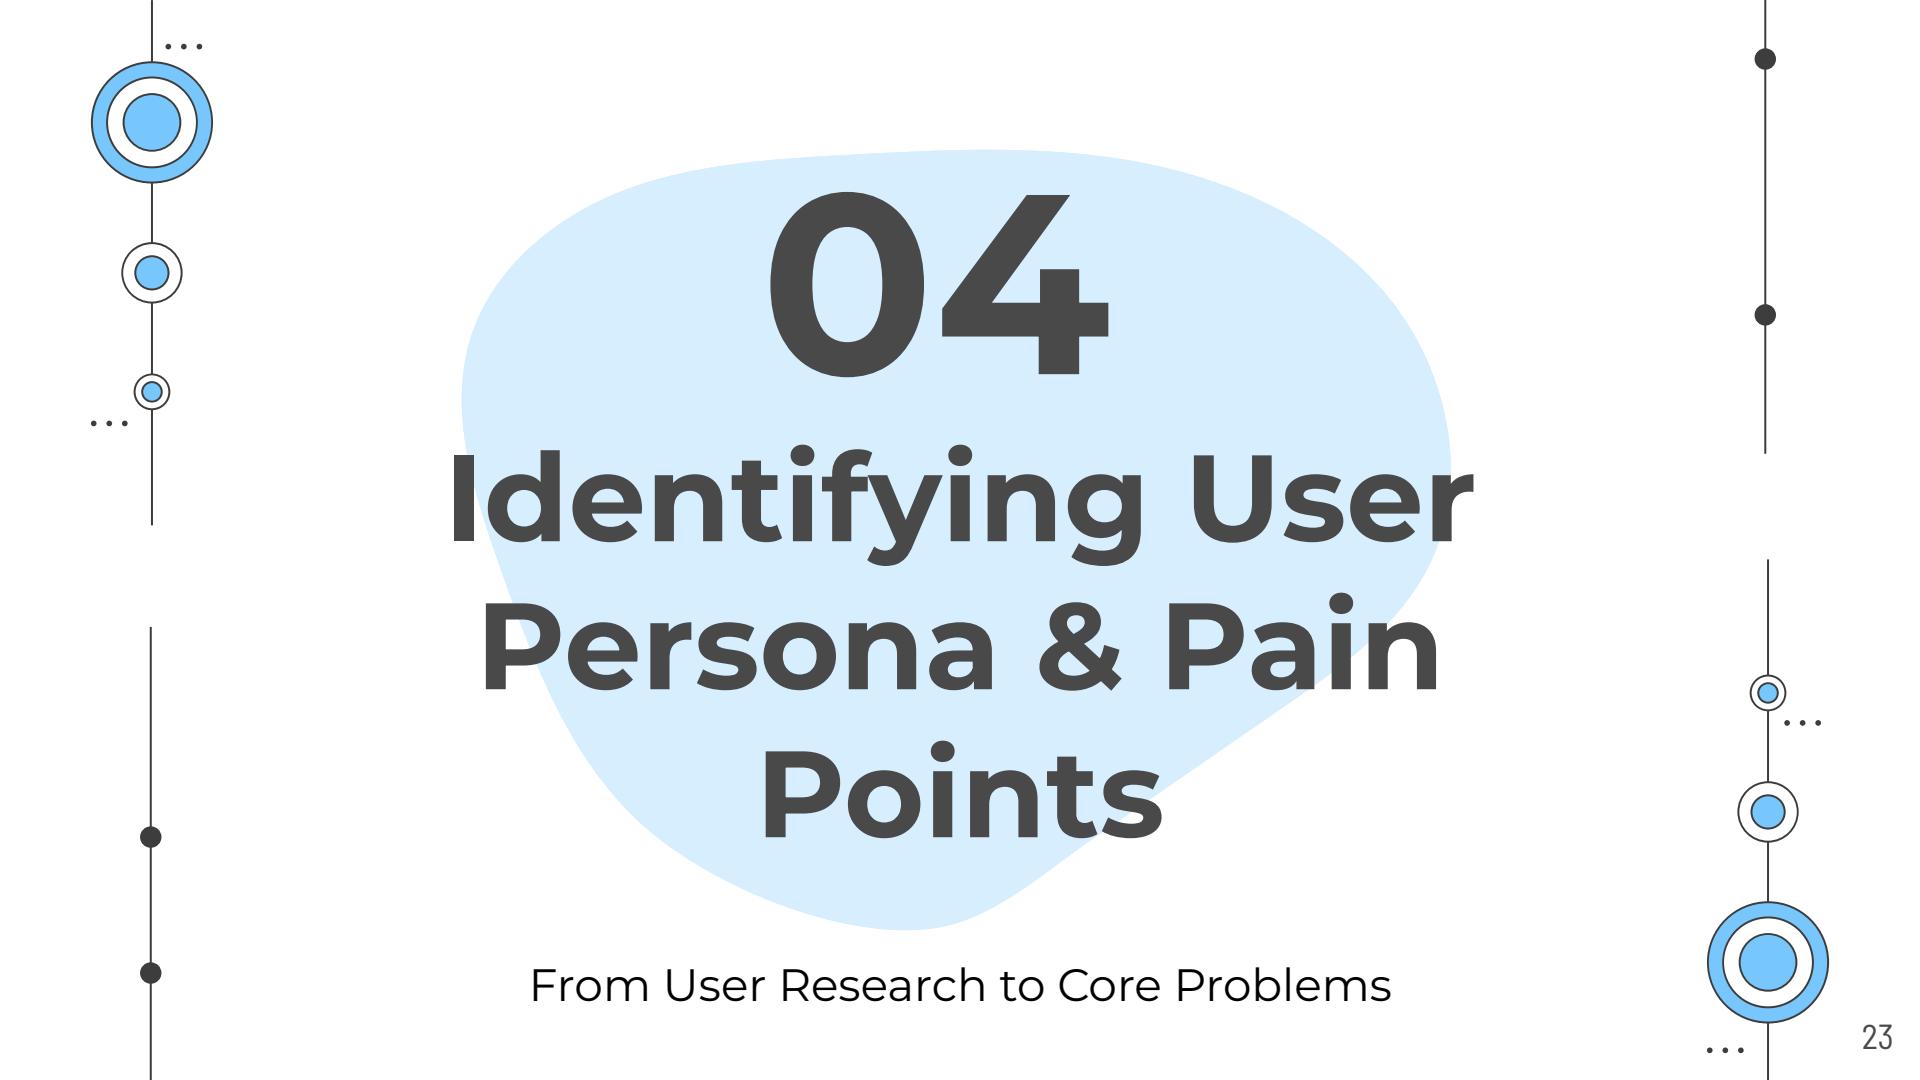
- ❖ LLMs excel at **summarizing individual documents**, answering questions and generating narratives.
- ❖ Out-of-the-box, they do **not**:
 - Enforce consistent scoring criteria across sources
 - Maintain historical memory of how topics evolve
 - Provide transparent, decision-grade ranking of signals

This reinforced the need to architect a **structured pipeline around the models**, rather than treating prompting alone as the solution.



ChatGPT





04

Identifying User Persona & Pain Points

From User Research to Core Problems

Key User Persona - Monica Geller from HRI



"I spend hours researching trends across scattered sources. It's manual, overwhelming, and I often can't connect the dots."

From our interviews, we developed a primary persona. Meet **Monica Geller** who represents the core user of the system:

- A senior executive at HRI with 15+ years of experience
- Leads a team of 20 analysts

This persona is heavily inspired by a real stakeholder at HRI and kept the project anchored to a concrete decision-maker.

Monica Geller: Motivation



"I spend hours researching trends across scattered sources. It's manual, overwhelming, and I often can't connect the dots."

Motivation

- Provide **quarterly reports** on key technologies & trend shifts
- Monitor **competitors** and summarise findings for executive management
- Define **research direction** and strategy for the coming year

Monica Geller: Key Tasks



"I spend hours researching trends across scattered sources. It's manual, overwhelming, and I often can't connect the dots."

Key Tasks:

- ❖ Scan reports, patents, and news
- ❖ Summarize trends & shifts manually and using LLM's
- ❖ Compare old vs. new signals
- ❖ Uses LLMs as helpers but still relies on manual work for synthesis
- ❖ Identify the **"next big thing"**

Pain Points

For Monica and her team, the core challenge is not finding information; **it is making sense of it in a way they can trust**. Here are the major pain points that her team is facing:



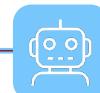
Fragmented
Sources



Information
Overload



No Historical
Memory



Low Trust in
Automation



Manual
Verification

These pain points provide a structured view of where the current process breaks down and where a new system must focus to create value: connecting sources, reducing overload, adding memory, increasing trust and lowering the verification burden.

Pain Point 1: Fragmented Sources

- ❖ Critical information lives across scattered sources—papers, patents, regulatory reports, news, and social content.
- ❖ These sources rarely speak to one another, making it difficult to form a cohesive, cross-referenced understanding of emerging trends.
- ❖ 78% of respondents said they regularly miss important sources because nothing lives in one place.
- ❖ Analysts spend time jumping between tools instead of seeing a unified picture.



Pain Point 2: Information Overload

- ❖ The volume of papers, reports, and updates is growing exponentially.
- ❖ Researchers are overwhelmed, with limited tools to filter signal from noise or extract actionable context from the flood of content.
- ❖ 80% of respondents reported feeling overwhelmed by the volume of emerging research.
- ❖ Large parts of the week are spent just “keeping up” with newspapers, patents and articles.
- ❖ Less time is available to interpret what these signals actually mean for HRI.



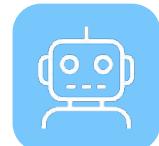
Pain Point 3: No Historical Memory

- ❖ There is no easy way to trace how an idea is originated, evolved, or gained momentum.
- ❖ Without a structured historical view, researchers struggle to understand when a trend truly began or how it has shifted over time.
- ❖ Analysts reported spending 10+ hours per week retracing old insights to connect the dots.
- ❖ Past work is hard to surface and reuse, so every new project feels like starting from scratch.
- ❖ Knowledge exists, but it is not organized as an accessible institutional memory.



Pain Point 4: Low Trust in Automation

- Only 2 in 9 respondents said they trust AI outputs without additional validation.
- According to PwC's Global CEO Survey, around two-thirds of senior leaders still require human oversight before acting on AI-generated recommendations.
- Most users still feel they must double-check model outputs before using them in slides or decisions.
- Lack of transparency in how results are produced reduces confidence in AI-generated insights.
- They want explanations backed by transparent citations, traceable evidence, and clear provenance without which trust in synthesized insights remains low.



Pain Point 5: Manual Verification

- Every important claim requires manual checking.
- There is no clear, shared way to verify whether a “trend” is legitimate.
- Analysts reopen sources, re-check citations and validate numbers by hand before sharing outputs.
- Constant rechecking results in delays, inconsistencies, and significant resource expenditure.
- Even with AI in the loop, the final mile of verification remains fully manual.



From Pain Points to Requirements

- **Fragmented sources:** The proposed system must bring signals from multiple sources into a single, consistent workflow.
- **Information overload:** The proposed system must help filter, focus and highlight what truly matters for decisions.
- **No historical memory:** The proposed system must help preserve and reuse prior signals and analyses over time.
- **Low trust in automation:** The proposed system must provide explainable, evidence-backed reasoning that users can inspect.
- **Manual verification:** The proposed system must reduce the verification burden by structuring evidence and citations upfront.

Together, these pain points defined where the new proposed system must create value.

05

Brainstorming

Critical Thinking Sessions

Whiteboarding and Early Exploration (1/3)

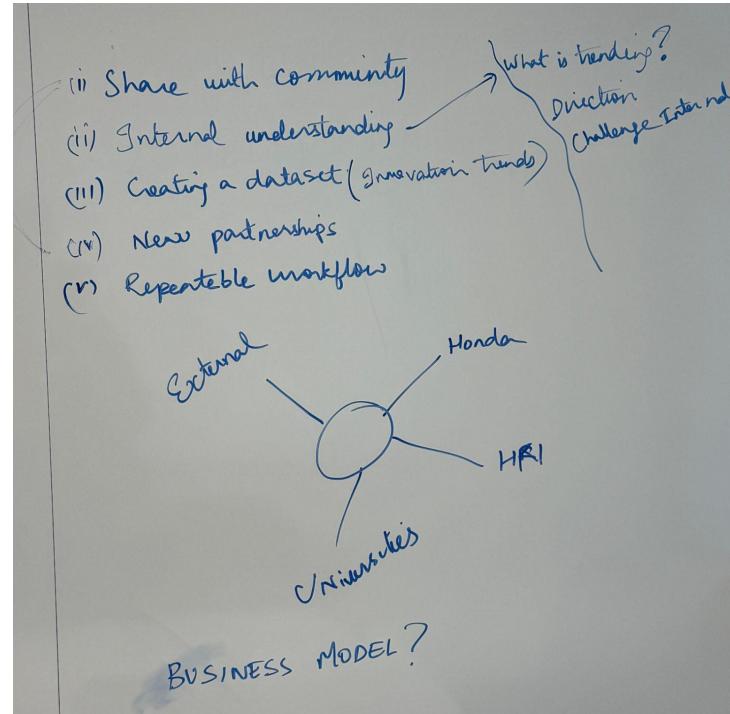
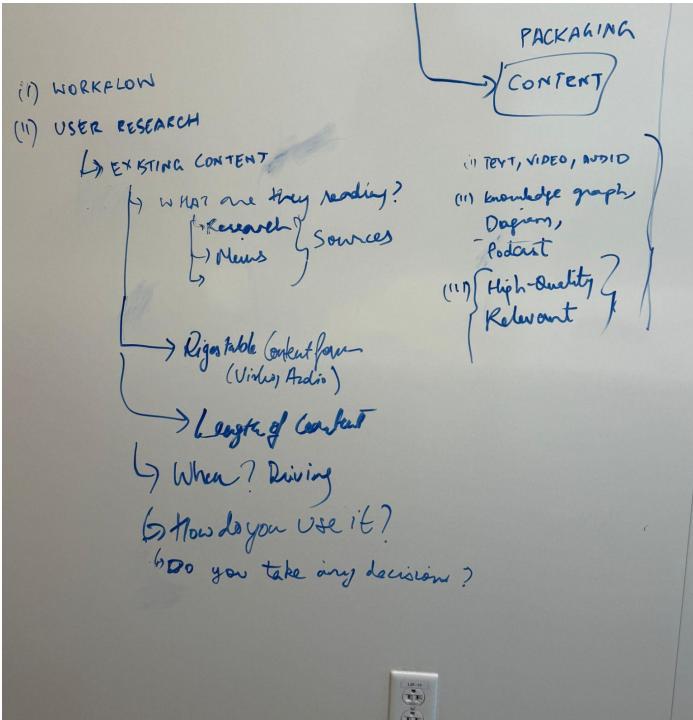
To make sense of the research, we used several rounds of whiteboarding sessions which helped us to:

- Map **all relevant source types** used by researchers, from academic literature to internal documents and social signals.
- Sketch possible **flows of information**, from initial questions through to final outputs.
- Explore different **degrees of automation** at each step

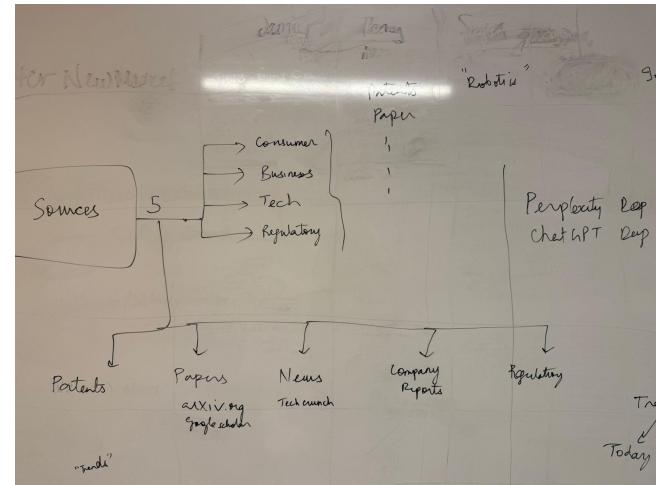
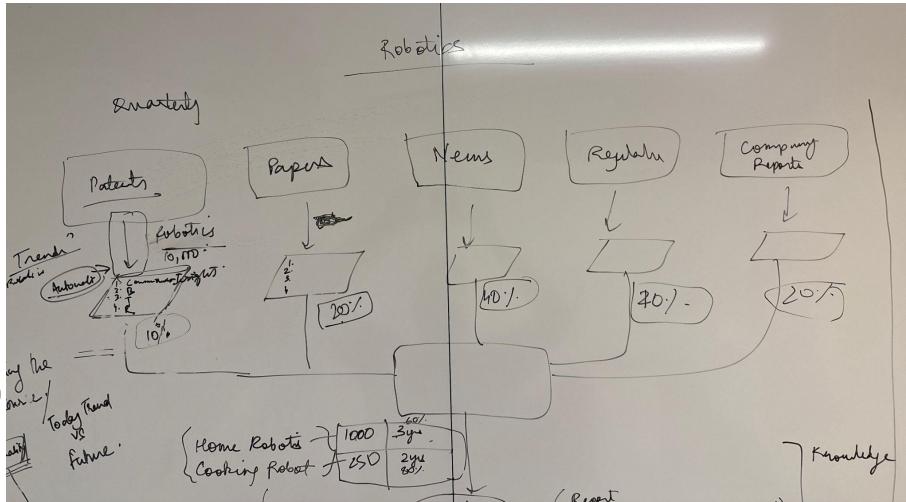
By visualizing these options on the board, we were able to:

- Separate what the system **must own end to end** from what should remain as human judgment.
- Identify where LLM capabilities could be applied safely and where they required guardrails.

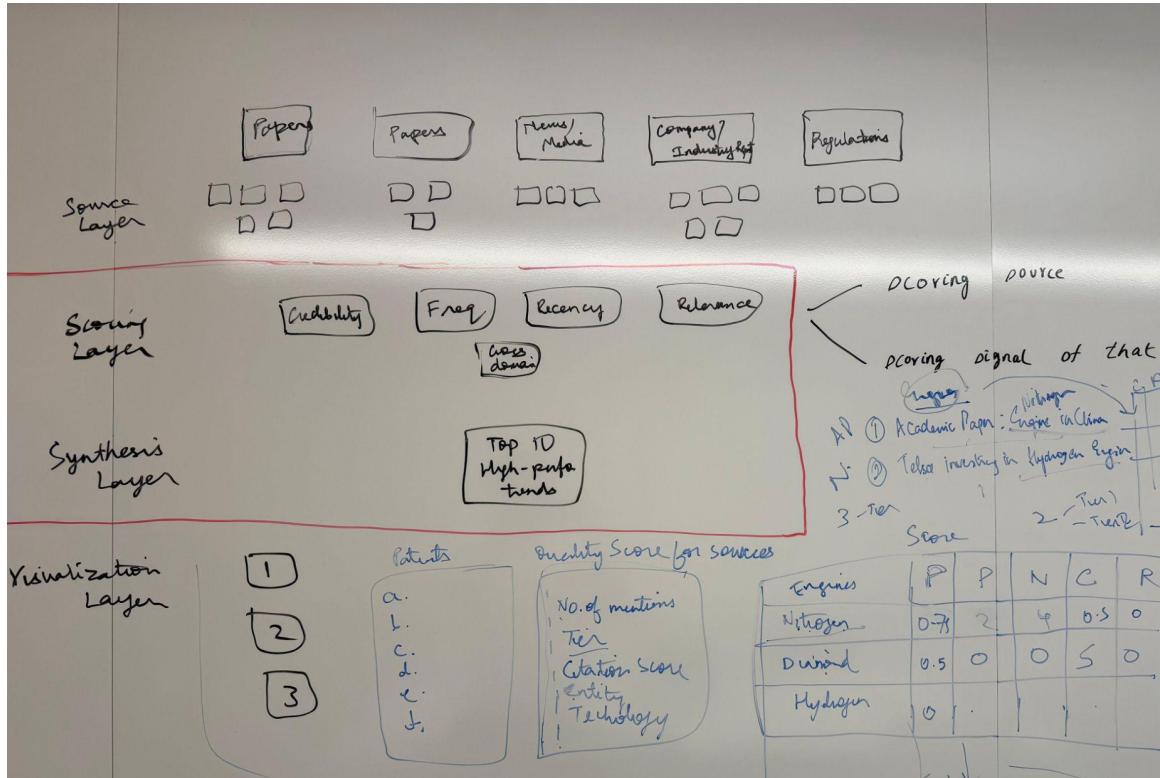
Whiteboarding and Early Exploration (2/3)



Whiteboarding and Early Exploration (3/3)



Initial WorkFlow Draft



Miro Boards and Clustering of Insights

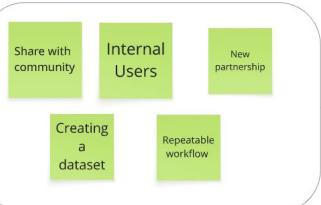
We captured key interview quotes, workflow sketches and early solution ideas on shared Miro boards. This allowed us to:

- ❖ Cluster similar pain points
- ❖ Map potential solution components back to specific problems
- ❖ Maintain traceability from user insight to design decision

The Miro boards became both a **bridge between raw interview data and structured requirements** and a shared collaboration space for the team. We could work through ideas freely and quickly align. They also made it easier to walk stakeholders through our thinking in a visual, transparent way.

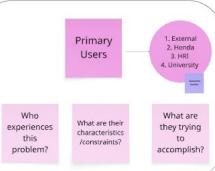
Brainstorming Problem Statement

Why?



1. What's the motivation?
2. How did you come up with this Problem Statement?
3. Why now?
4. What do you think is missing that you need something like this?

Who?



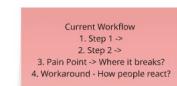
What?



HRI Internal

1. Is Honda going to use the tool/content directly or HRI?
2. What does Honda want? - New Techniques/Insights/Models/IP/etc.
3. How frequently do you interact with Honda today?
4. Inform Product Roadmap - Explain? - How it happens today?

How?



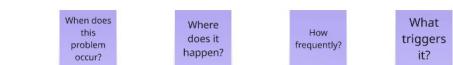
Root cause analysis



Honda

1. Is Honda going to use the tool/content directly or HRI?
2. What does Honda want? - New Techniques/Insights/Models/IP/etc.
3. How frequently do you interact with Honda today?
4. Inform Product Roadmap - Explain? - How it happens today?

When & Where?



Clustering Sources

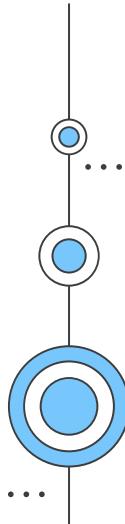
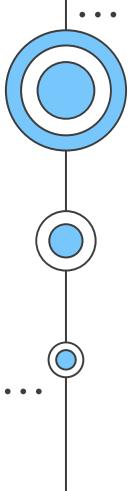




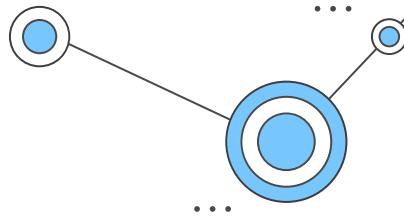
06

Scoping and Prioritization

Research Insights to Solution
Requirements



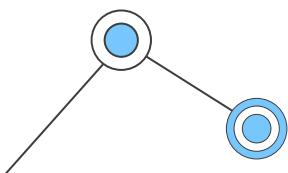
Solution Requirements



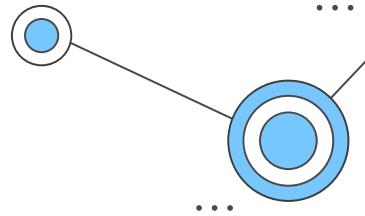
From the persona and pain points, we derived solution requirements:

The proposed system must:

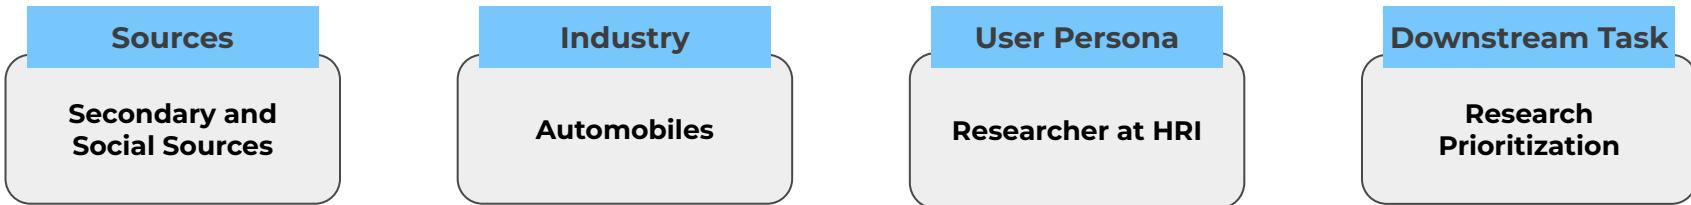
- Consolidate signals from multiple external sources into a single structured workflow
- Provide **objective, transparent scoring** of signals and trends
- Produce **consistent, explainable outputs** that can be trusted in high-stakes discussions
- Operate across the sources and time horizons relevant to HRI's current decisions



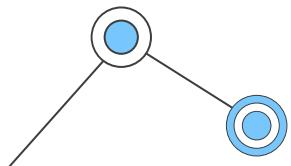
Scoping



For this prototype, we scoped the problem to automotive **industry**, with Monica Geller as the **primary user** and research prioritization as the main **downstream task**. The system focuses on the sources this user actually relies on: papers, patents, industry and regulatory reports, and social signals.



This focused scope ensures the prototype directly aligns with our core user's pain-points.



Prioritization - MoSCoW Framework

What is it?

Must Have

These are essential capabilities. Without them, the solution does not achieve its purpose

Should Have

These are very important capabilities that improve outcomes. Version 1 can still function without them

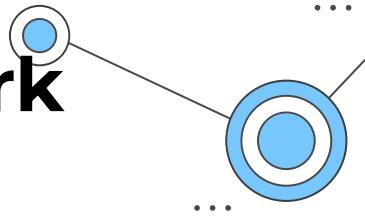
Could Have

These are nice-to-have capabilities that add depth or usability. They are not core to the solution

Won't Have (for now)

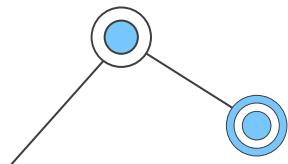
These are capabilities that are explicitly out of scope for this phase.

Prioritization - MoSCoW Framework



We performed prioritization using the MoSCoW framework.

Must Have	Should Have	Could Have	Won't Have
<ul style="list-style-type: none">• Handle Fragmented Sources• Score Signals Objectively• Produce Consistent & Explainable Insights	<ul style="list-style-type: none">• Identify lifecycle stage• Identify what is driving the trend (Primary Driver)• Executive-Ready Summary	<ul style="list-style-type: none">• Visual dashboards	<ul style="list-style-type: none">• Real-time monitoring or alerts• Expansion beyond automobiles (Future Scope)



Must Haves - Core Of the System (1/3)

- ❖ Capabilities without which the system does **not** solve Monica's problem.
- ❖ If any Must Have is missing, the prototype cannot be used for real research decisions.

Our Must Haves

1. Handle Fragmented Sources

- Problem:** Monica's work is scattered across papers, patents, news and internal tools.
- Decision:** The system must bring signals from multiple external sources into one structured pipeline.
- Effect:** Turns "fragmented" into connected, so she no longer has to manually stitch sources together.



Must Haves - Core Of the System (2/3)



2. Score Signals Objectively

- a. **Problem:** Today, trend ranking is often based on intuition or whoever is in the room.
- b. **Decision:** The system must score signals using objective framework/criteria (credibility, frequency, recency, relevance).
- c. **Effect:** Turns “inconsistent” into scored, giving HRI a transparent way to compare signals and address low trust in automation.

3. Produce Consistent and Explainable Insights

- a. **Problem:** Outputs are often manual and hard to reuse; each analyst formats and explains trends differently
- b. **Decision:** The system must generate repeatable, citation-backed summaries that follow a consistent structure
- c. **Effect:** Turns “manual” into assisted, so Monica’s team can reuse outputs without rebuilding every slide from scratch



Should Haves (1/2)



Capabilities that are very important, but not strictly required for the first working version. By treating them as **“Should Haves,”** we improve the quality of strategic decisions while keeping the core scope focused.



1. Identify Lifecycle Stage

- a. Classifies each trend as inception, growth, peak, maturity or decline.
- b. Helps Monica judge timing and decide whether a topic is emerging, established or already fading.
- c. The core workflow can still connect, score and summarize trends without this label, which is why it is a Should Have rather than a Must Have.

2. Identify Primary Driver

- a. Indicates whether a trend is driven mainly by technology, regulation, consumers, markets or a combination.
- b. Guides the type of response that teams should consider, such as R&D investment, regulatory engagement or product changes.
- c. The system can already provide useful trend summaries without this classification, so it strengthens decisions but is not essential for basic operation



Should Haves (1/2)



3. Executive Ready Summaries:

- a.** Provides concise, leadership friendly narratives that highlight what the trend is, why it matters and what decisions it informs.
- b.** Reduces extra work between research output and boardroom discussion and makes it easier to share insights across teams.
- c.** Analysts can still use the core summaries if this layer is missing, which is why it is treated as a Should Have for early versions.



Could Haves

Capabilities that make the system easier to explore and present, but do not change the underlying value of the workflow. By keeping them as “**Could Haves**,” we recognise their importance without diluting the focus of the prototype.

Visual Dashboards

- a. Dashboards, heat maps and charts that display trends by score, domain, timeframe or source type.
- b. Help researchers and leaders quickly see patterns, compare trends and drill into areas of interest visually.
- c. All core functions of connecting sources, scoring signals and generating summaries work without these views, which is why they are enhancements rather than core requirements.

Won't Haves (Out of Scope for Now)

Capabilities that are valuable, but **explicitly excluded** from this phase. By parking them in future scope, we stayed focused on delivering a reliable, decision-ready prototype for today, while providing a clear path for tomorrow.

1. Real Time Trend Monitoring

- a. Real time streaming, live score updates and alerts would demand new infrastructure, monitoring and operational support
- b. This requires significant infrastructure and operational maturity beyond our scope
- c. The system can still be fully evaluated and used in batch mode, so this was intentionally deferred

2. Expansion Beyond Automobiles

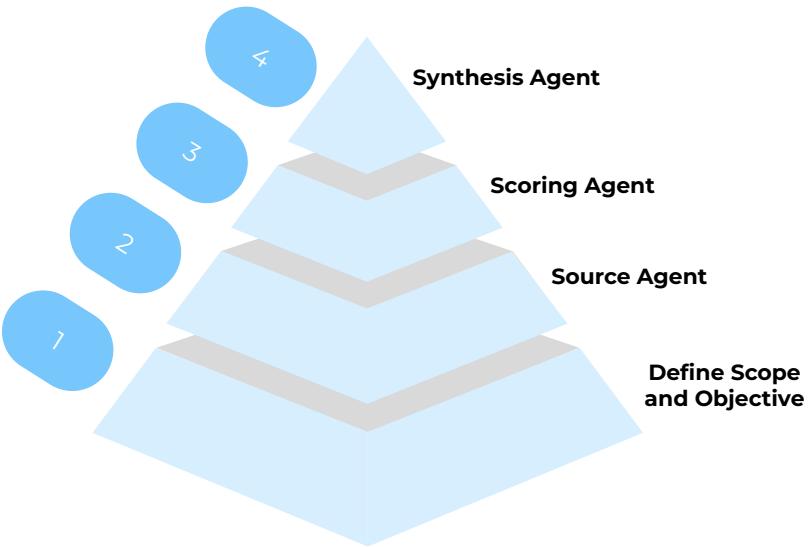
- a. Proving the workflow in one focused context is more important than spreading it thinly across many domains.
- b. Starting with automotive makes it possible to validate the approach deeply, while leaving clear room to generalize later.

07

Proposed Solution

Trend Identification Pyramid

Solution Overview : Trend Identification Pyramid



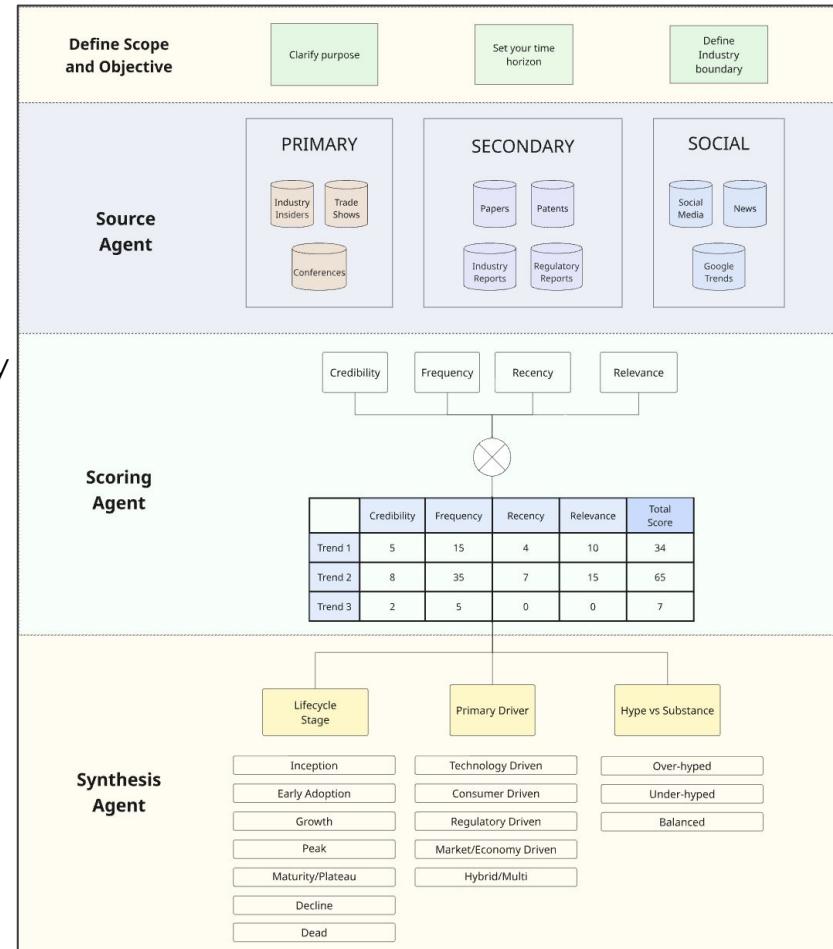
We designed a multi-agent system structured as a layered pyramid, where each stage builds on the one beneath it — from defining the scope, to sourcing signals, to scoring them, and ultimately synthesizing insights.

This architecture enables information to be aggregated and interpreted at scale, forming what we call the Trend Identification Pyramid.

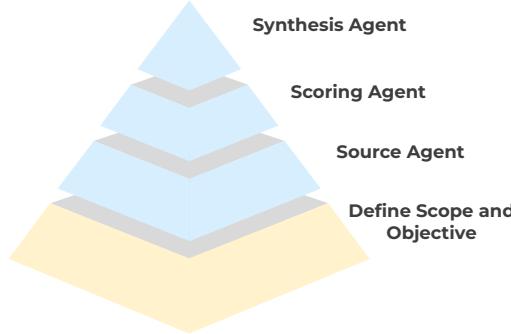
Solution: High Level Workflow

We designed a multi-agent system that contains four layers and transforms fragmented information into structured, evidence-based insights.

- **Define Scope & Objective** : Establish a clear purpose, time horizon, and industry boundary to align the entire analysis.
- **Source Agent**: Collects signals from primary, secondary, and social sources to ensure comprehensive, diverse coverage.
- **Scoring Agent**: Evaluates each signal using Credibility, Frequency, Recency, and Relevance to generate objective, comparable scores.
- **Synthesis Agent**: Converts scored signals into actionable insights by determining lifecycle stage, primary driver, and hype vs. substance.



Solution Layer 1: Define Scope and Objective



Users told us that trend analysis today feels scattered: different people use different definitions, different periods, different interpretations. By formalizing the scope upfront, we create a single shared lens for the entire workflow. This layer is critical because it prevents the analysis from drifting into irrelevant noise.

For example:

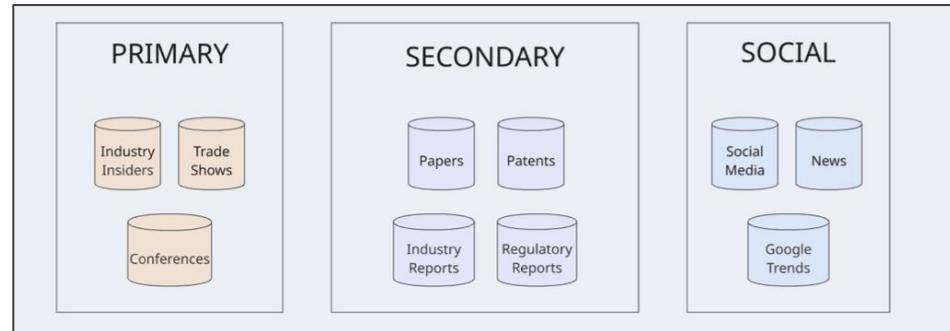
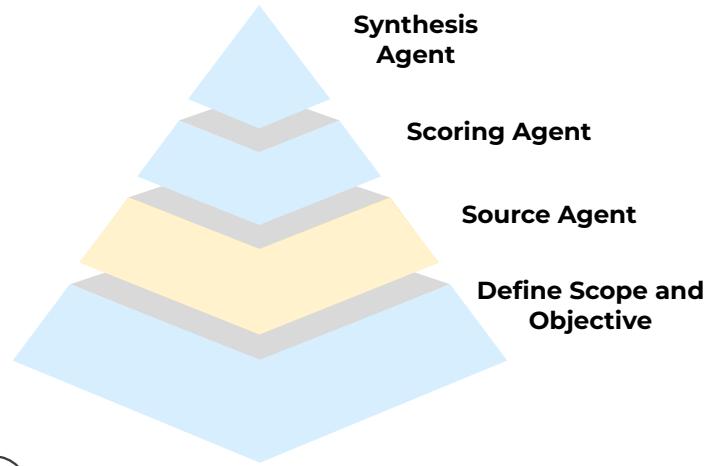
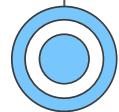
- Are we scouting technology for R&D investment?
- Are we identifying long-term industry shifts?
- Are we prioritizing short-term opportunities?

Each of those use cases drastically changes what counts as a “trend.”

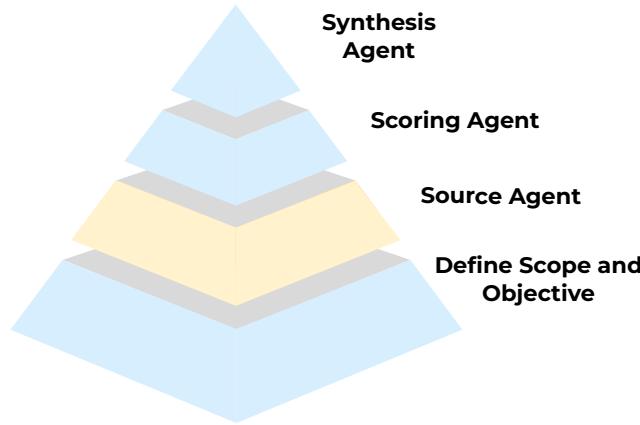
So Layer 1 forces clarity around:

- Purpose: the strategic decision being supported.
- Time Horizon: 1 year or 3-5 years or 10 years, each implying different maturity levels and risk profiles
- Industry Scope: because insights must be relevant to a specific domain to be actionable

Solution Layer 2: Source Agent



Solution Layer 2: Source Agent

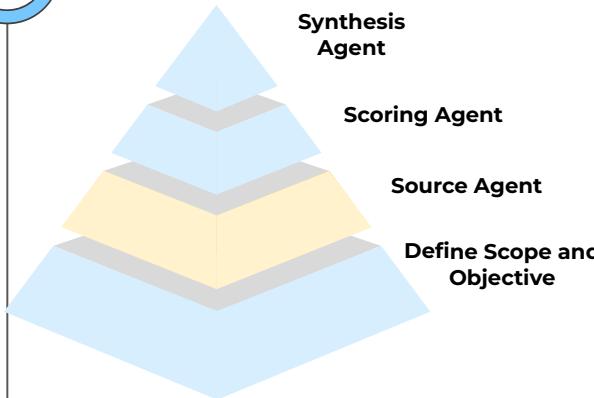


Users told us they spend hours scraping papers, news, reports, and social feeds, but the information never comes together cohesively.

The result is lot of information and blind spots. Some users rely too heavily on academic material, while others rely on fast signals like news or social media.

A new technology may appear first as a patent, then in papers, then in industry news, and finally in social signals. If you only look at one type of source, you misread the timeline and ur not getting the full picture.

Solution Layer 2: Source Agent



The Source Agent collects data from three buckets:

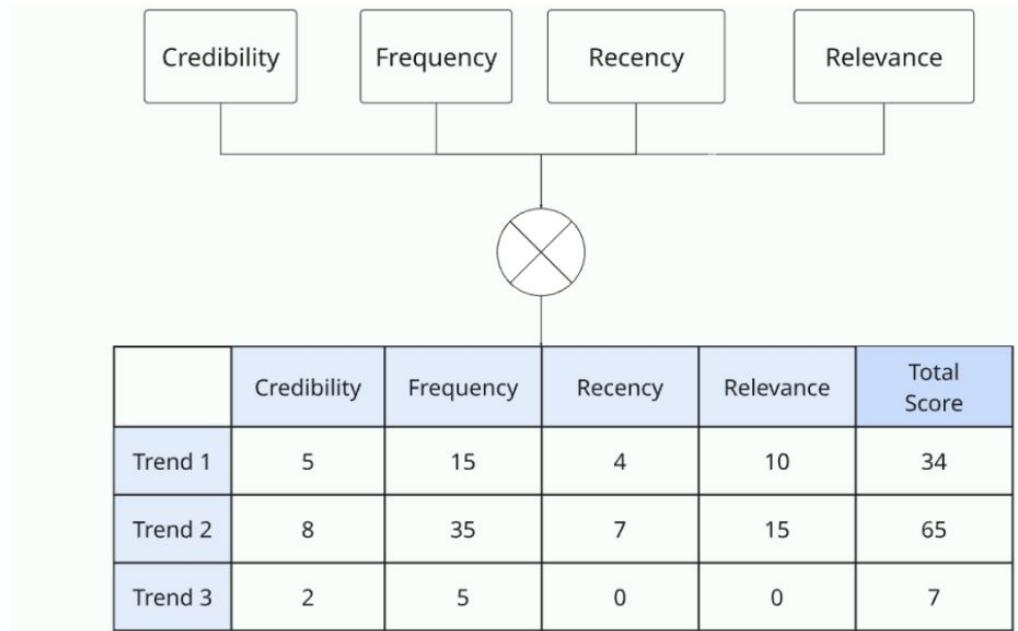
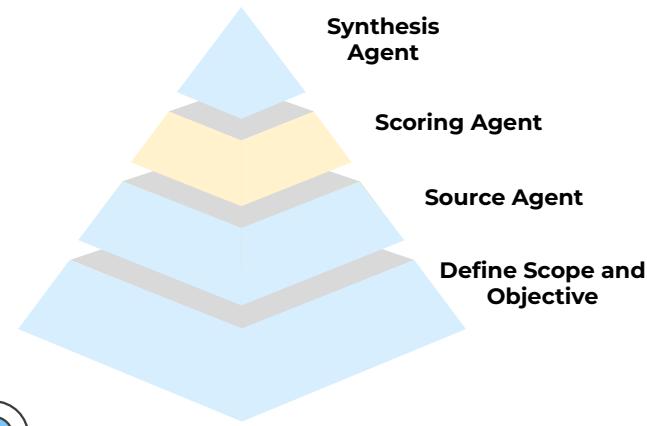
- **Primary sources:** Expert insight, Events, Conferences
- **Secondary sources:** Papers, Patents, Industry & Regulatory Reports
- **Social sources:** News, Google Trends, Social media

Together, these buckets create a diverse, cross-referenced information set that reduces blind spots and helps track how a trend evolves across the ecosystem from inception to broader adoption.

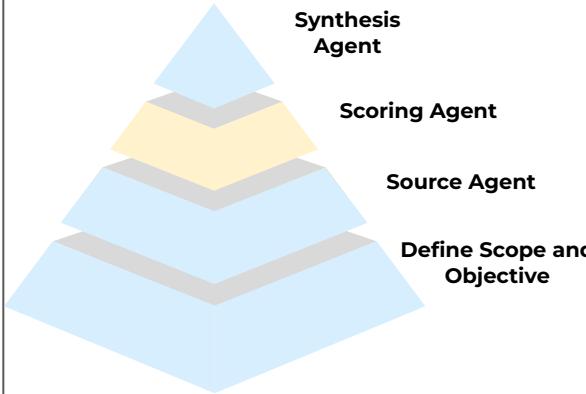
Why Combine All Three?

Trends don't emerge from a single source they evolve across ecosystems. By triangulating primary, secondary, and social signals, the Source Agent builds a richer dataset that prevents blind spots and strengthens downstream scoring and synthesis.

Solution Layer 2: Source Agent



Solution Layer 3: Scoring Agent



This is where the pain point “*Low Trust in Automation*” is addressed. We don’t just aggregate links; we score them. The system evaluates data based on four metrics to create a weighted “Total Score” for every identified trend. Also, depending on the industry and how it values sources the weighted averages differ.

The 4 core dimensions:

- **Credibility:** The trustworthiness and authority of the source.
- **Frequency:** The volume and repetition of the trend’s signal across different sources.
- **Recency:** How recently the evidence for the trend was published or observed.
- **Relevance:** How directly the trend relates to the user’s initial defined scope and industry boundary.

Solution Layer 3: Scoring Agent

To illustrate how the scoring system evaluates different trend signals, we walk through one concrete example for each of the four dimensions - Credibility, Frequency, Recency, and Relevance.

Example: Scoring Framework for Patents

Total Score = **Credibility** Score (Weight: 30%) + **Frequency** Score(Weight: 25%) + **Recency** Score (Weight: 25%) + **Relevance** Score (Weight: 20%)

- **Credibility** = $(\text{Patent Office Score} \times 0.4) + (\text{Assignee Score} \times 0.3) + (\text{Quality Indicators} \times 0.3)$
- **Frequency** = $(\text{Growth Rate} \times 0.4) + (\text{Volume} \times 0.3) + (\text{Competitive Intensity} \times 0.3)$
- **Recency** = $(\text{Timeline Score} \times 0.5) + (\text{Acceleration} \times 0.3) + (\text{Evolution Signals} \times 0.2)$
- **Relevance** = $(\text{Purpose Fit} \times 0.5) + (\text{Horizon Fit} \times 0.3) + (\text{Industry Alignment} \times 0.2)$

Solution Layer 3: Scoring Agent

Going deeper into each dimension:

1. Credibility Score = (Patent Office Score \times 0.4) + (Assignee Score \times 0.3) + (Quality Score \times 0.3)

Patent Office Score:

- Tier 1 (9-10 points): Triadic patents (filed in USPTO, EPO, and JPO simultaneously)
- Tier 2 (7-8 points): USPTO or EPO patents with examination completed
- Tier 3 (5-6 points): Single major office (USPTO, EPO, JPO, CNIPA) with granted status
- Tier 4 (3-4 points): PCT applications or pending applications in major offices
- Tier 5 (1-2 points): Single country filings, provisional patents, or applications only

Assignee Score:

- Industry Leaders (10 points): Fortune 500, recognized innovation leaders (IBM, Samsung, Google)
- Specialized Leaders (8 points): Domain leaders in specific technology (ASML for lithography)
- Mid-tier Innovators (6 points): Established companies with consistent patent portfolios
- Startups/Scale-ups (4 points): VC-backed companies with <500 employees
- Individual inventors (2 points): Non-affiliated inventors or micro-entities

Solution Layer 3: Scoring Agent

Patent Quality Indicators:

- Number of independent claims >10: +2 bonus points
- Cited by >20 subsequent patents: +2 bonus points
- Examiner citations to patent during other examinations: +1 bonus points
- Standard-essential patent (SEP): +3 bonus points
- Litigation history (survived challenge): +2 bonus points
- Licensed to multiple parties: +2 bonus points

Overall Credibility Score = (Patent Office Score \times 0.4) + (Assignee Score \times 0.3) + (Quality Indicators \times 0.3)

Solution Layer 3: Scoring Agent

Moving to the next dimension: **Frequency**

2. Frequency Score = (Growth Rate \times 0.4) + (Volume \times 0.3) + (Competitive Intensity \times 0.3)

Growth Rate Score:

- Explosive growth (10 points): >200% year-over-year increase in patent families
- Rapid growth (8 points): 100-200% YoY increase
- Strong growth (6 points): 50-100% YoY increase
- Moderate growth (4 points): 20-50% YoY increase
- Emerging (2 points): <20% YoY increase

Volume Score:

- 1000 patent families/year: +3 bonus
- 500 patent families/year: +2 bonus
- 100 patent families/year: +1 bonus

Solution Layer 3: Scoring Agent

Competitive Intensity Score:

- 50 different assignees filing patents: +2 bonus
- Top 5 assignees control <50% of patents: +2 bonus (indicates broad competition)
- New entrants (companies filing first patents in space): +1 per 10 new entrants (max +3)
- Geographic diversity (>2 countries): +1 bonus

Overall Frequency Score = $(\text{Growth Rate} \times 0.4) + (\text{Volume} \times 0.3) + (\text{Competitive Intensity} \times 0.3)$

Solution Layer 3: Scoring Agent

Moving to the next dimension: **Recency**

2. Recency Score = (Timeline Score × 0.5) + (Acceleration × 0.3) + (Evolution Signals × 0.2)

Timeline Score:

- Cutting edge (10 points): >70% of patents filed in last 18 months
- Current wave (8 points): >60% filed in last 24 months
- Recent surge (6 points): >50% filed in last 36 months
- Steady stream (4 points): Consistent filings over 3-5 years
- Mature field (2 points): Patent activity >5 years old

Acceleration Score:

- Priority claims showing shortened development cycles: +2 bonus
- Continuation-in-part (CIP) applications showing rapid iteration: +2 bonus
- Fast-track examination requests: +1 bonus
- Multiple divisional applications (indicating broad innovation): +1 bonus

Solution Layer 3: Scoring Agent

Evolution Signals Score:

- New CPC/IPC classifications created for technology: +3 bonus
- Patent offices issuing examination guidelines for field: +2 bonus
- Increase in PCT filings (international protection sought): +2 bonus
- Provisional-to-nonprovisional conversion rate >80%: +1 bonus

Overall Recency Score = (Timeline Score \times 0.5) + (Acceleration Score \times 0.3) + (Evolution Signals Score \times 0.2)

Solution Layer 3: Scoring Agent

Moving to the next dimension: **Relevance**

2. Relevance Score = (Purpose Fit \times 0.5) + (Horizon Fit \times 0.3) + (Industry Alignment \times 0.2)

Purpose Fit Score:

Direct Strategic Fit (10 points)

- Directly informs R&D investment decisions
- Solves a top research question or key decision identified in Layer 1
- Strong connection to downstream tasks (roadmap, opportunity sizing, capability building)

Strong but Indirect Fit (8 points)

- Influences adjacent decisions (materials, safety, user experience)
- Relevant to R&D but not core to the defined decision

Partial Fit (6 points): Related to automotive but weak connection to interiors or investment questions

Weak Fit (4 points): Interesting but peripheral to the defined purpose

Minimal Fit (2 points): Trend is exciting but irrelevant to the decision context

Solution Layer 3: Scoring Agent

Horizon Fit Score

(Let's assume 3-5 year horizon from Layer 1)

- Perfect Horizon Fit (10 points): Aligned with 3–5 year adoption; prototypes and early deployments visible.
- Strong Fit (8 points): Expected to mature within the 3–5 year window with scaling underway.
- Slightly Ahead or Behind (6 points): Off by 1–2 years; strong research but unclear adoption.
- Misaligned (4 points): Too speculative (10+ years) or already fully mature (0–1 year).
- Poor Fit (2 points): No realistic path to adoption in 3–5 years.

Industry Alignment Score

(Let's assume Automotive interiors from Layer 1)

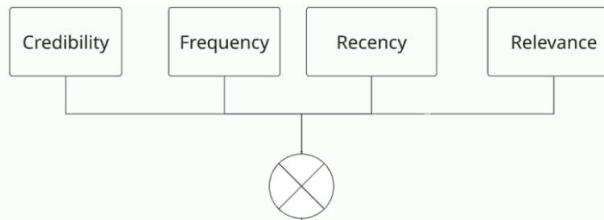
- Direct Impact (10 points): Directly affects materials, comfort, UX, safety, sensors, or cabin energy management.
- High Impact (8 points): Impacts vehicle systems with indirect or secondary effects on interiors.
- Adjacent Domain (6 points): Relevant to automotive but requires adaptation to apply to interiors.
- Industry-Neutral (4 points): Broad tech trend with limited automotive relevance.
- No Relevance (2 points): No meaningful connection to automotive or interior applications.

Overall Relevance Score = (Purpose Fit Score \times 0.5) + (Horizon Fit Score \times 0.3) + (Industry Alignment Score \times 0.2)

Solution Layer 3: Scoring Agent

Summarizing the scoring logic:

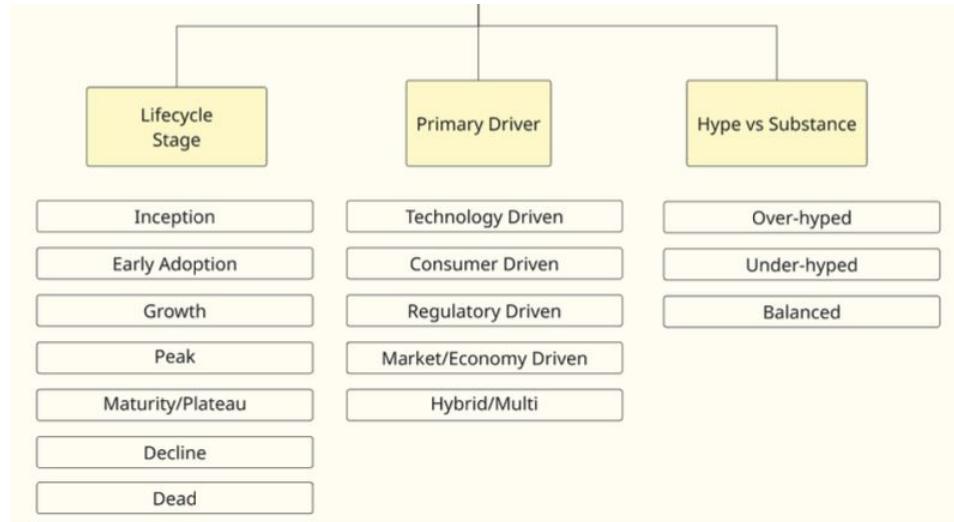
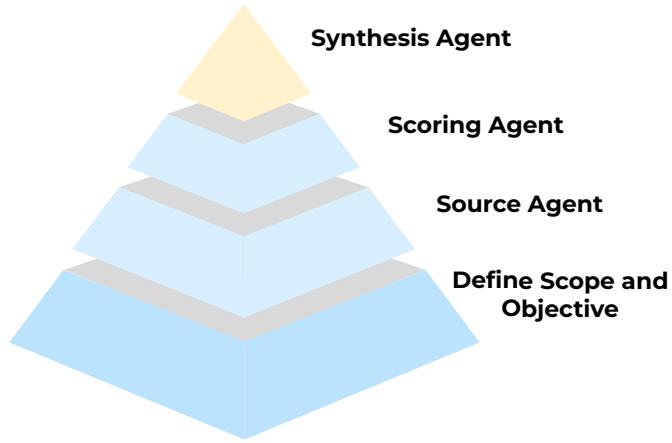
- Credibility = (Patent Office Score \times 0.4) + (Assignee Score \times 0.3) + (Quality Indicators \times 0.3)
- Frequency = (Growth Rate \times 0.4) + (Volume \times 0.3) + (Competitive Intensity \times 0.3)
- Recency = (Timeline Score \times 0.5) + (Acceleration \times 0.3) + (Evolution Signals \times 0.2)
- Relevance = (Purpose Fit \times 0.5) + (Horizon Fit \times 0.3) + (Industry Alignment \times 0.2)



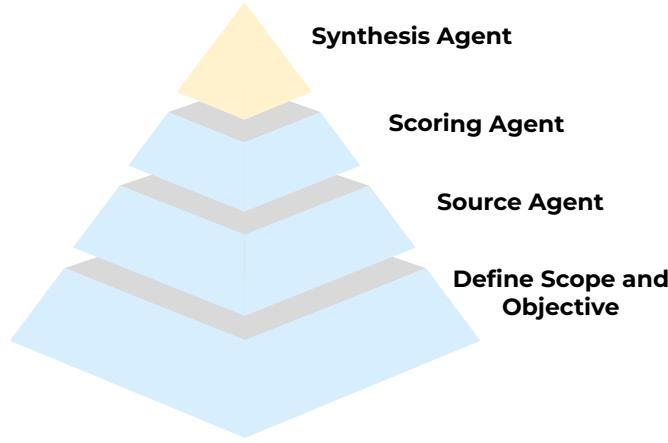
	Credibility	Frequency	Recency	Relevance	Total Score
Trend 1	5	15	4	10	34
Trend 2	8	35	7	15	65
Trend 3	2	5	0	0	7

Total Score = Credibility Score (Weight: 30%) + **Frequency** Score(Weight: 25%) + **Recency** Score (Weight: 25%) + **Relevance** Score (Weight: 20%)

Solution Layer 4: Synthesis Agent



Solution Layer 4: Synthesis Agent



Finally, the system transforms the scored data into digestible insights. Instead of a list of links, the user receives a synthesized analysis categorized by :

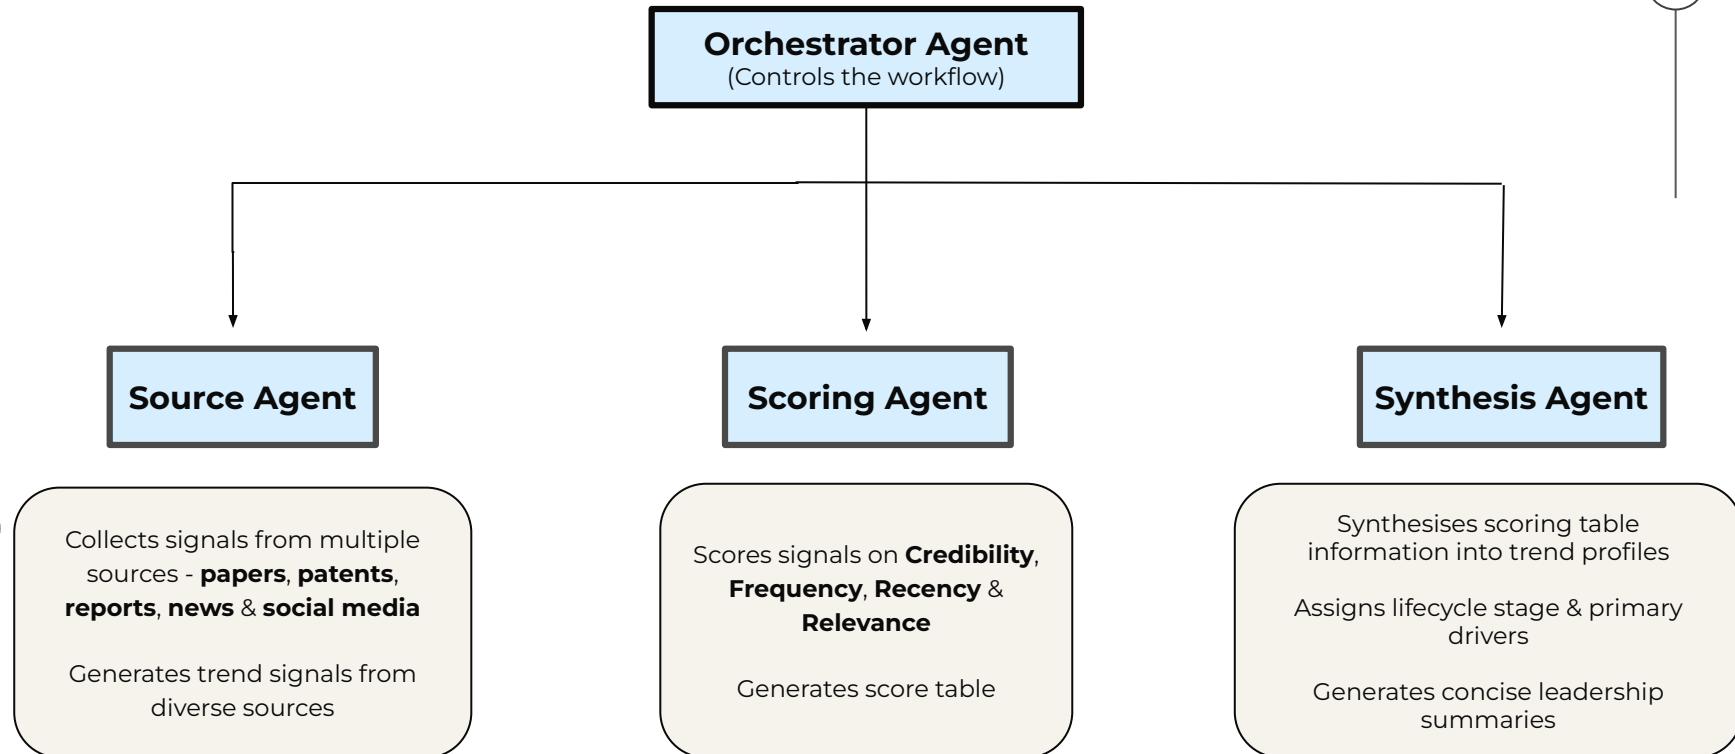
- **Lifecycle Stage:** Is the trend in Inception, Growth, Peak, or Decline?
- **Primary Driver:** Is it Technology Driven, Consumer Driven, or Regulatory Driven?
- **Hype vs. Substance:** Is the trend Over-hyped, Under-hyped, or Balanced?

08

Building a Prototype

System Architecture And Learnings

System Overview - Multi-Agent System



Prototype Layer 1: Define Scope and Objective

- **Purpose:** Scout R&D Investment opportunities for Honda
- **Time Horizon:** 3-5 years
- **Industry Scope:** Interiors in Automotives

Prompt

Run honda-orchestrator-agent to get trends in automotive industry with 3-5 years horizon and interiors as scope. Run the complete workflow and show me the results at each step and also mention clearly when the orchestrator agent is calling other agents – state it in the CLI output
ankitsmac

19,377 tokens
5% used
\$0.10 spent

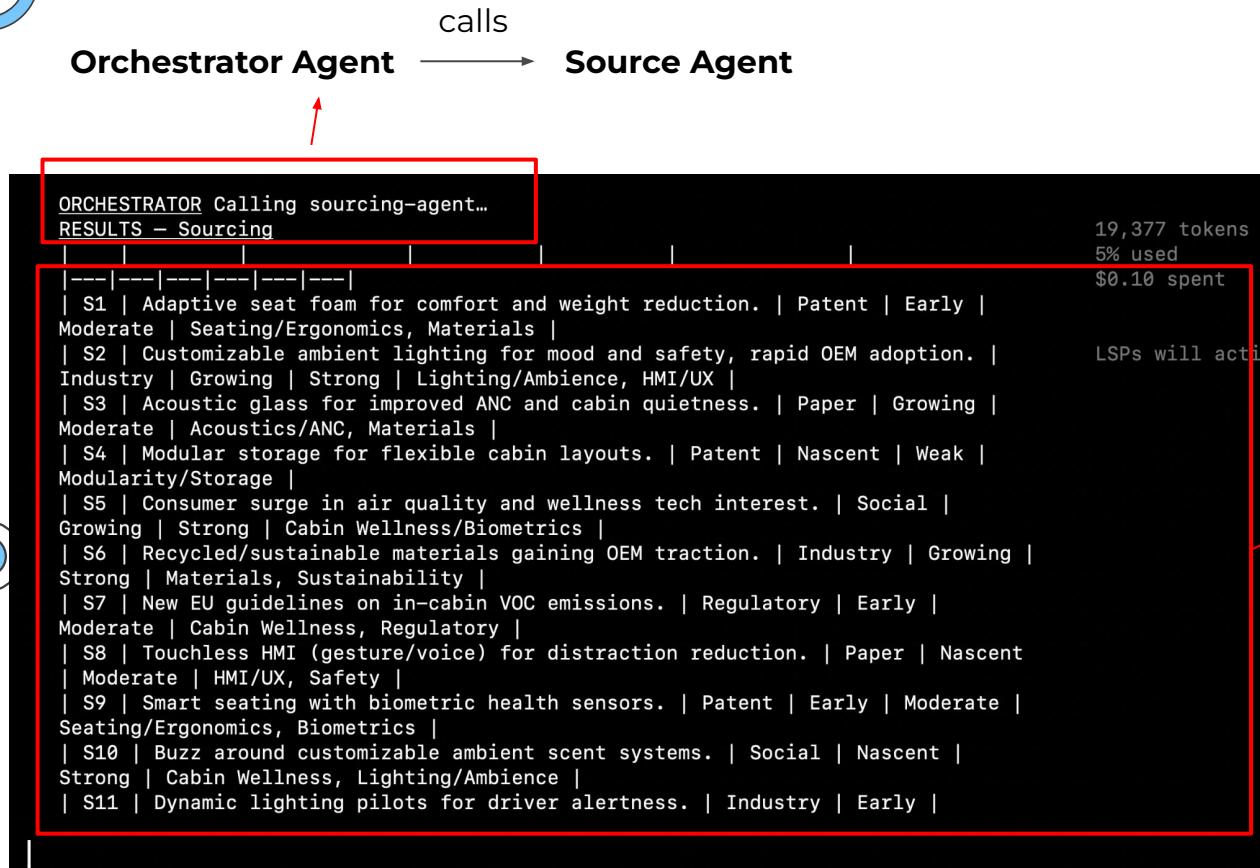
I'll run the orchestrator now and display step-by-step outputs, clearly noting each sub-agent call.

LSPs will activate as files are

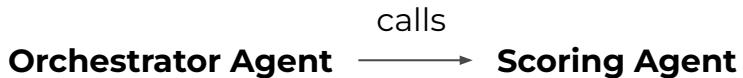
Honda-Trend-Orchestrator Task "Run orchestrator for auto interiors 3-5y"

ctrl+x right, ctrl+x left to navigate between subagent sessions

Prototype Layer 2: Source Agent



Prototype Layer 3: Scoring Agent



```
~/Documents/HondaDemo — opencode + opencode ... duments/HondaDemo/honda-b-trendly/opencode/agent — opencode ... rendly — oc | Running honda-a
ORCHESTRATOR Calling scoring-agent...
RESULTS — Scoring
---|---:|---:|---:|---|
| S1 | 4 | 3 | 3 | Patent, moderate OEM interest |
| S2 | 5 | 4 | 5 | Strong industry/consumer pull |
| S3 | 4 | 4 | 3 | Research-backed, growing OEM use |
| S4 | 3 | 2 | 1 | Nascent, few pilots |
| S5 | 4 | 3 | 5 | Social/industry convergence, strong demand |
| S6 | 5 | 4 | 5 | Sustainability top priority, regulatory push |
| S7 | 3 | 4 | 3 | Regulatory, early implementation |
| S8 | 4 | 3 | 3 | Research, growing interest |
| S9 | 4 | 3 | 3 | Patent, early pilots |
| S10 | 3 | 2 | 4 | Social buzz, tech not mature |
| S11 | 4 | 3 | 3 | Early pilots, moderate impact |
| S12 | 4 | 4 | 3 | Regulatory, growing in China |

- Insights:
  - Top sub-themes: HMI/UX, Sustainability/Materials, Cabin Wellness, Lighting/Ambience.
  - Strongest momentum: S2 (ambient lighting), S5 (wellness), S6 (sustainability).
  - Lower-confidence outliers: S4 (modularity), S10 (ambient scent).
  - Regulatory signals (S7, S12): moderate momentum, high confidence.
  - Most signals are early/growing maturity, matching 3-5y horizon.

19,377 tokens
5% used
$0.10 spent
LSPs will activate
```

Output: Score Table

Prototype Layer 4: Synthesis Agent

calls

Orchestrator Agent ————— **Synthesis Agent**

ORCHESTRATOR Step 5 – Synthesis & Executive View

Results:

Trend Profiles

- T1. Seamless, Personalized UX & HMI
 - Drivers: S2 (ambient lighting customization), S8 (touchless HMI), S15 (device integration)
 - Why 3-5y: Rapid consumer demand, strong OEM investment, maturing enabling tech
 - Implications: Intuitive, distraction-reducing interfaces; software content; safety/privacy considerations
 - Risks: Integration complexity, distraction, regulatory lag
 - Indicators: Models with advanced HMI, satisfaction scores, regulatory guidance
- T2. Sustainable & Circular Interior Materials
 - Drivers: S1 (adaptive foams), S6 (recycled materials), S13 (circular plastics)
 - Why 3-5y: Regulations, cost parity for recycled content, sustainability pledges
 - Implications: Recycled/biobased shift, new supply chains, possible cost savings
 - Risks: Supply reliability, perception, evolving standards
 - Indicators: % recycled content, supplier announcements, regulatory targets
- T3. Cabin Wellness & Comfort as Differentiators
 - Drivers: S5 (air quality/wellness), S7 (VOC regulation), S10 (ambient scent), S9 (biometric seating)
 - Why 3-5y: Post-pandemic focus, regulatory action, willingness to pay
 - Implications: Sensor integration, air quality systems, premium features
 - Risks: Proven health impact, complexity, privacy

19,377 tokens

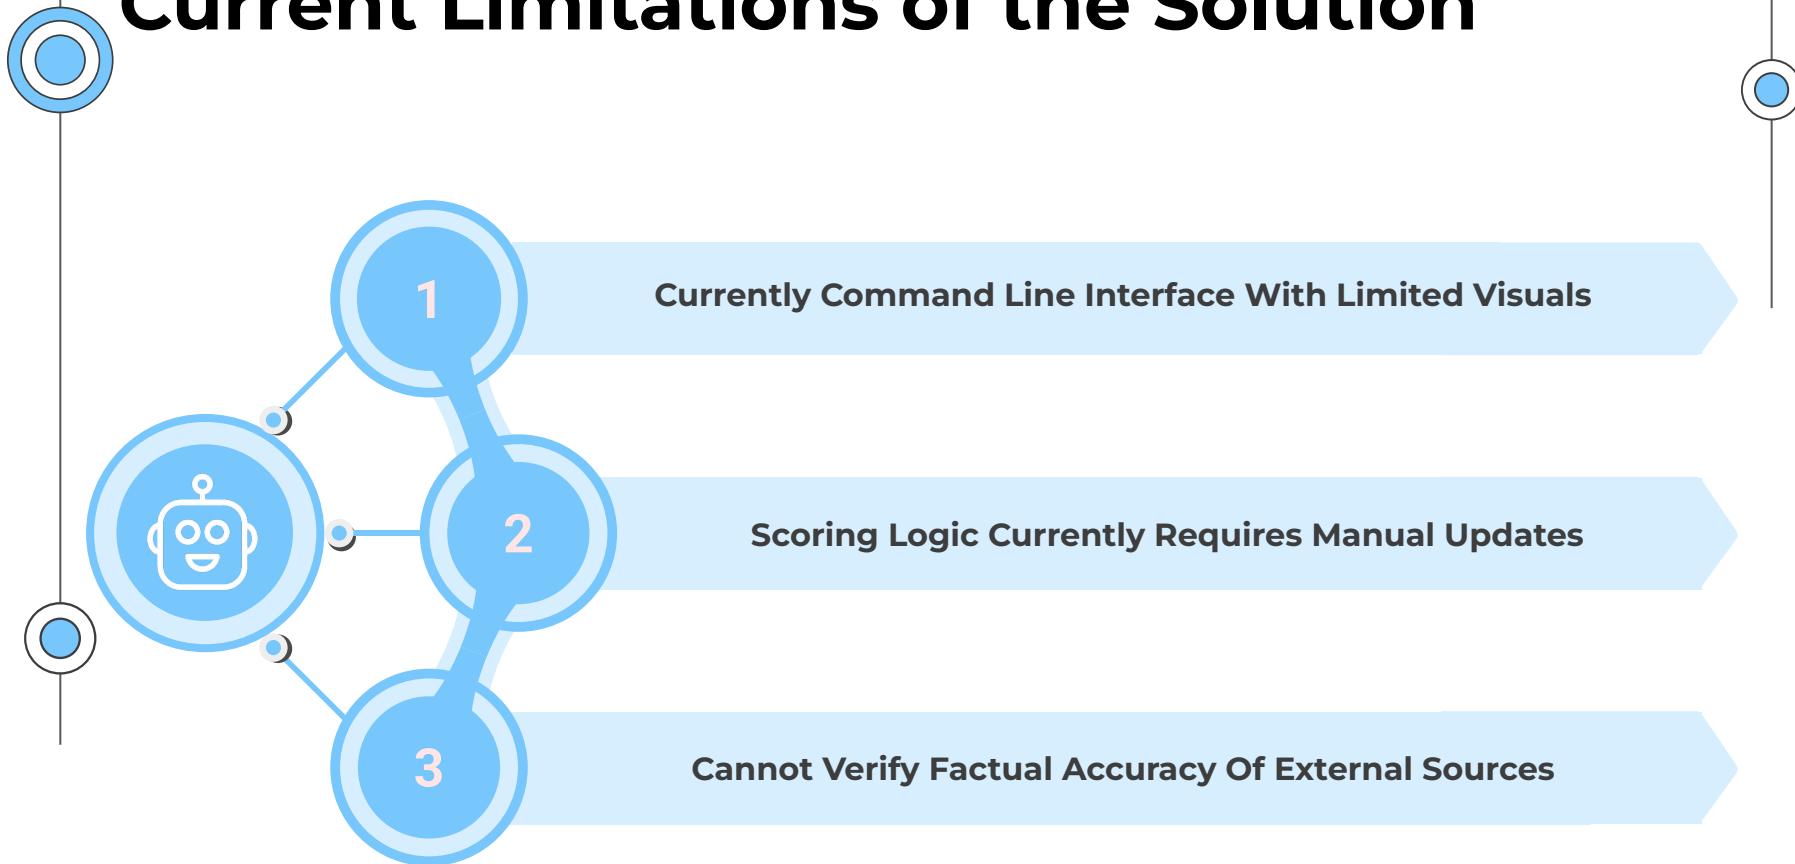
5% used

\$0.10 spent

LSPs will activate as

Output: Trend Profiles

Current Limitations of the Solution

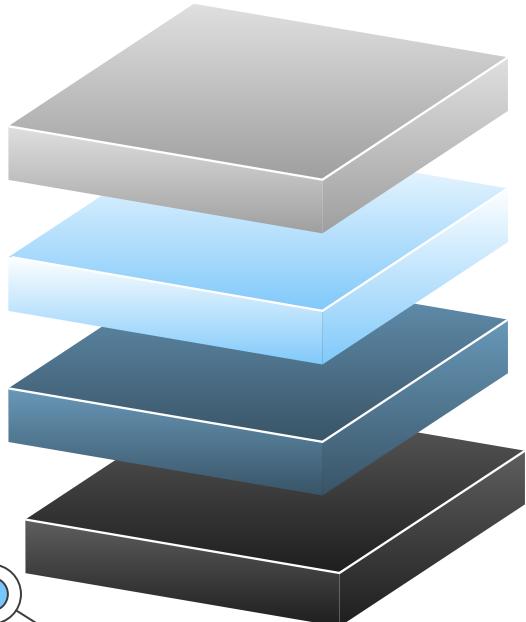
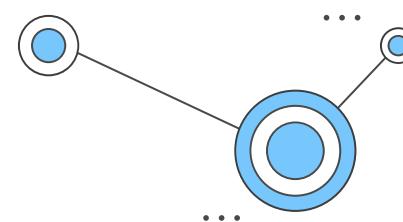


09

Future Scope

Evaluation and Next Steps

Recommended Next Steps (1/4)



1

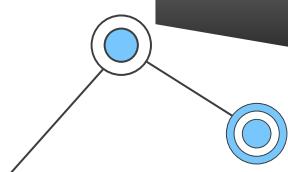
Roll out the tool internally

- Provide access to researchers so they can use the system in day-to-day work.
- Collect feedback to refine scoring logic, prompts and workflows.

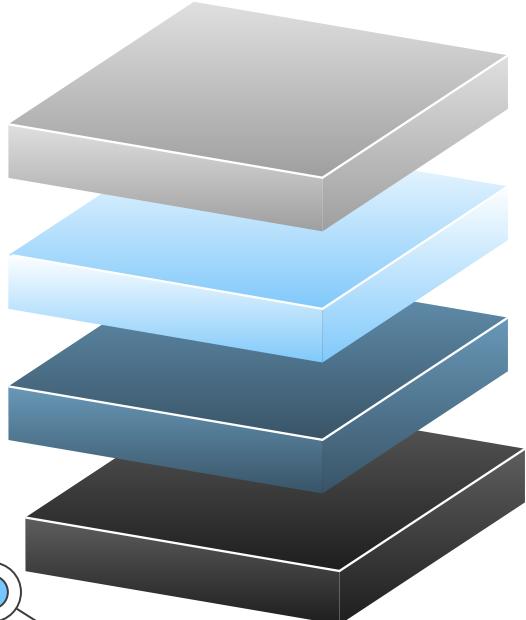
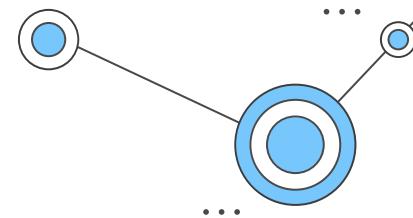
2

Integrate Internal Proprietary Data Into The Sources

- Combine public sources with Honda's internal reports, historical insights and third-party data.
- Make the system more specific to HRI's needs and context.



Recommended Next Steps (2/4)



3

Add A Verification Agent To Validate Sources and Citations

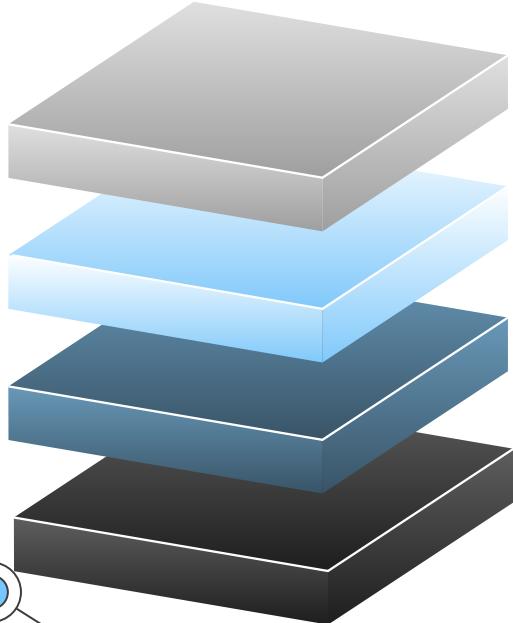
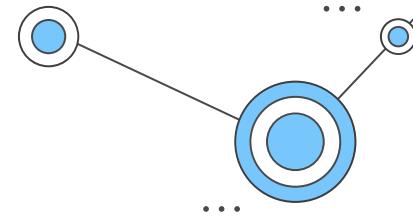
- Add an agent focused on checking citations and flagging possible inconsistencies or hallucinations.
- Reduce the manual burden of source verification.

4

Expand To Multiple User-Personas And Industries

- Reuse the workflow for additional personas and domains by adjusting scoring parameters and source weighting.

Recommended Next Steps (3/4)



5

Implement The Evaluation Workflow

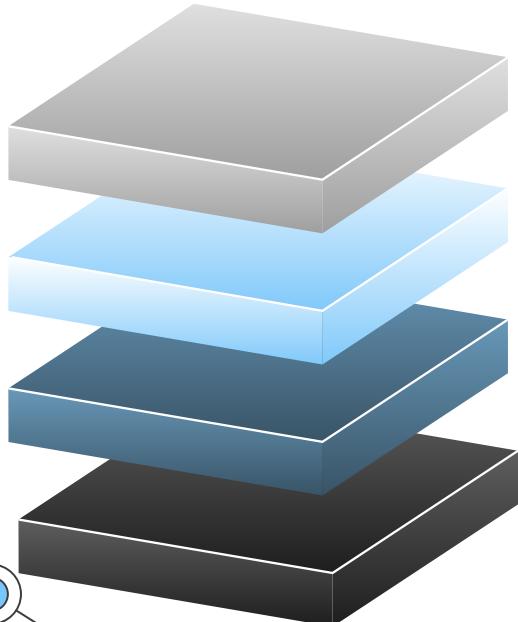
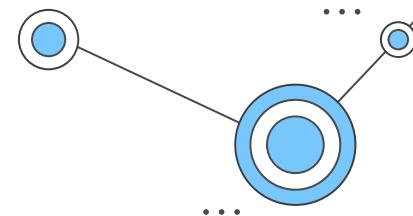
- Persist signals and scores over time to build institutional memory of how topics evolve.

6

Create A Living Database Of Trend Signals

- Use the workflow to generate articles, briefs or social content for external communication.

Recommended Next Steps (4/4)



7

Generate Shareable Outputs For Social Channels

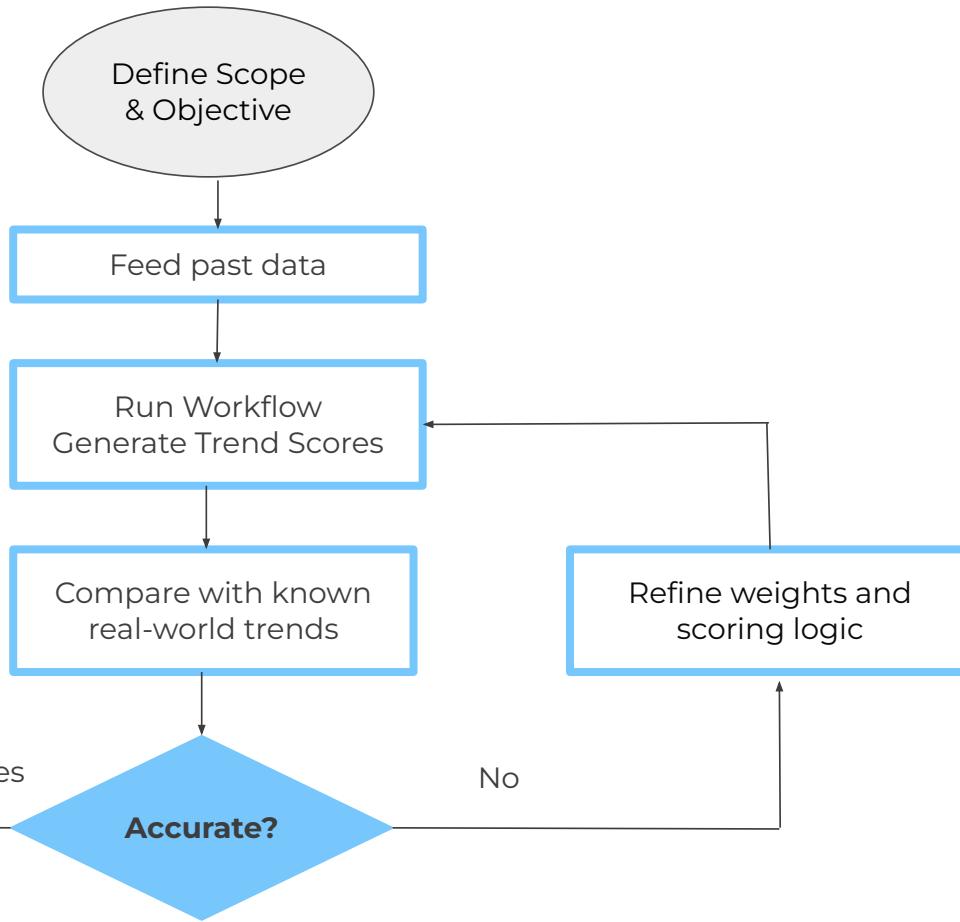
- Use the workflow to generate articles, briefs or social content for external communication.

8

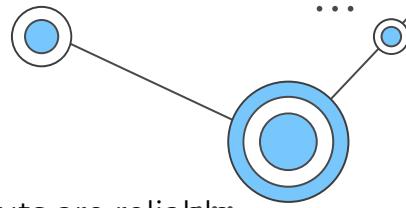
Build Real time Monitoring Dashboards

- Once the core workflow is stable, build dashboards on top of it to surface live updates and alerts.

Evaluation Framework: Flow Diagram



Evaluation in Detail

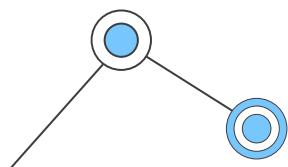


A central question for long-term adoption is how to know whether the system's outputs are reliable. Our primary future direction is therefore evaluation.

The proposed approach is based on historical back-testing:

1. **Feed Past Data** – for example, signals from a historical period such as 2018.
2. **Run Workflow** – execute the full pipeline using only information available before that year.
3. **Generate Trend Scores** – obtain the system's ranking of candidate trends.
4. **Compare with Known Real-World Trends** – check whether highly ranked trends align with developments that actually materialized.
5. **Refine Weights and Scoring Logic** – adjust scoring parameters when gaps are observed.

Over time, this creates a self-improving system that becomes more predictive and better calibrated to real-world outcomes.



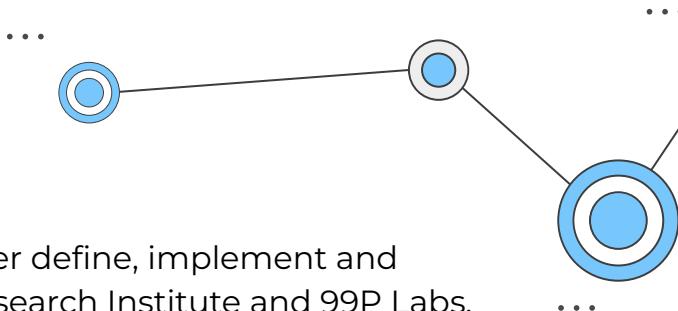
10

Project Deliverables

Final Deliverables and Artefacts

Final Deliverables

This project resulted in a set of core deliverables that together define, implement and communicate the Trend Identification System for Honda Research Institute and 99P Labs.



- **Trend Identification System & Workflow**

Four-layer architecture (scope, source, scoring, synthesis) with clear roles, inputs and outputs.

- **Prototype Implementation** ([Github Link](#))

Multi-agent system in OpenCode, plus a GitHub repository with code, scoring logic and configuration files.

- **Process Book**

Comprehensive documentation of context, research, persona, pain points, solution design, prototype and future roadmap.

- **Medium Article** ([Article Link](#))

Public-facing narrative that explains the problem, approach and Trend Identification Pyramid for a broader audience.

Intermediate Deliverables



Along the way, we created several intermediate deliverables that helped share progress, validate our direction and refine the solution with stakeholders.

- **User Research Presentation**

Early readout of interview findings, current-state analysis and initial pain points.

- **Poster Presentation Session**

Live session to walk stakeholders through the problem framing, approach and early solution concept.

- **Project Poster**

Visual summary of the problem, approach and proposed solution for internal sharing at Honda and CMU.

- **Final Presentation**

End-of-semester presentation demonstrating the full journey from problem to prototype and outlining future directions.

11

Key Learnings

What We Learned as a Team



Key Learnings



1. Applying Product Thinking Principles

- We learned to slow down and approach the problem with a product mindset rather than rushing to build something.
- Nate emphasized understanding the “why” before the “what,” which taught us to anchor every decision to user value and outcomes.
- We realized that strong products emerge from clarity of problem definition, not from the speed of solution-building.

2. Staying Focused on Users and Their Problems

- One of the biggest learnings was to consistently return to the user’s needs, pains, and motivations not our assumptions.
- Early in the project, we tended to jump ahead to features and technical implementations; Nate’s feedback pushed us to redirect our energy to understanding the user journey more deeply.
- We learned how critical it is to validate the problem space thoroughly before exploring solutions.



Key Learnings



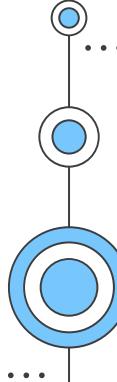
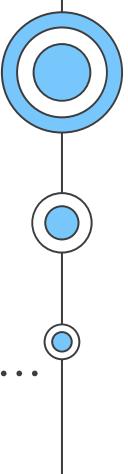
3. Attention to Detail in Delivery & Professionalism

- Nate helped us recognize that attention to detail is not superficial—presentation etiquette, clarity of slides, and even how we dress communicates respect for the work and the stakeholders.
- Small details—formatting, consistency, visual hierarchy—had a big impact on how our work was perceived.

12

Conclusion

Closing Remarks



Fragmented to a Structured System

The challenge facing Honda Research Institute was a classic paradox: drowning in information while struggling to extract meaningful, actionable insights. The previous trend research process was manual, fragmented, and broken in structure.

Our collaboration has successfully redefined this process by designing the Trend Identification Pyramid, an LLM-driven multi-agent system that institutionalizes a transparent, repeatable, and an evidence-based process for future-scouting.



Impact on HRI's Research Capability



The result is not just a tool for summarization, but a strategic thinking partner that provides three critical shifts in HRI's research capability:

- **Strategic Velocity:**

By transforming raw, unstructured data into executive-ready narratives categorized by Lifecycle Stage and Primary Driver, the system accelerates decision-making on high-stakes research investments.

- **Risk Mitigation:**

The objective Scoring Agent prioritizes signals based on Credibility, Frequency, Recency, and Relevance, ensuring Honda invests in validated trends and not just market hype.

- **Knowledge Orchestration:**

It solves the problem of scattered information by triangulating sources across Primary, Secondary, and Social buckets, establishing a robust and centralized knowledge layer

Closing the Loop: A New Way Forward

This multi-agent system marks a fundamental shift in how HRI discovers, interprets, and operationalizes external innovation trends. Instead of depending on manual, inconsistent workflows, it delivers a streamlined, transparent, and evidence-driven process that researchers can rely on for strategic decision-making.

Annexure

User Research Presentation to Clients on 10/28/25

User Research

This synthesis report aims to understand how experts currently identify emerging trends and where existing methods fail, specifically focusing on the need for **AI-based intelligence tools**.

Expert Interviews

We interviewed 3 academics and 6 researchers across multiple scientific disciplines (e.g., neuroscience, AI) and leading industry teams.

Core Focus

The primary focus was to diagnose how does each lab define a trend

CHAPTER 2: TREND DEFINITION

VARIABLES

How Different Fields Define "Emerging Trends"



Neuroscience / Biomedical

Slow-moving; trends stabilize only after **reproducibility across multiple labs** and extensive peer review to confirm maturity.



AI / Language

Technology massive post-paper explosion leads to high noise. Hard to separate **hype from fundamental impact**.



Robotics & Engineering

Requires sustained, reproducible methods. Labs pivot only when **core scientific assumptions are actively challenged**.



Industry & Market

Action-oriented; prioritizes **continuity over time** and visible investment signals from major competitors.

Every domain defines 'trend' differently—some chase novelty, others wait for proof—but all struggle to connect evidence over time.

Quotes

“Labs don’t pivot because of a single spike on Twitter/Google News. Labs pursue core problems. We’d only reconsider if a core assumption is broken.”

Example: “If my assumption is ‘robots must work with humans’ and someone demonstrates robots that don’t need humans, that would make me rethink.”

“A self-driving crash is a news trend, not a research trend.”

CHAPTER 3: STATUS QUO

Current (Manual) Trend Extraction Process

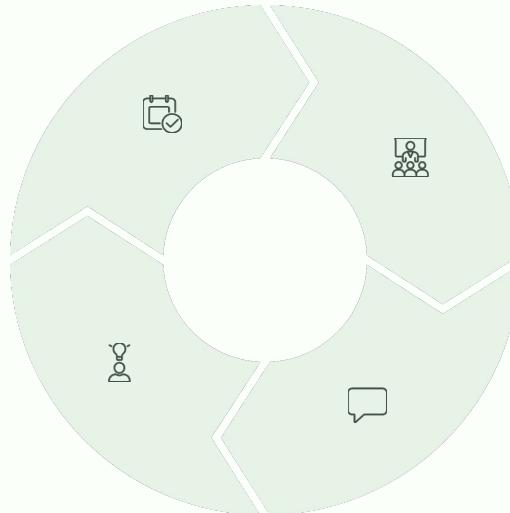
The current process relies on inconsistent, resource-intensive human effort.

Literature Review & Search

Relying heavily on standard library searches (Google, Scholar) and manual database queries.

Intuition & Internal Synthesis

No standardized process; experts rely on personal experience to validate and synthesize signals.



Conference Scanning

Following key speakers, monitoring accepted papers, and attending high-value industry events.

Peer & Competitor Discussion

Informal discussions combined with following funding announcements and competitor investment signals.

Everyone's doing their own version of signal extraction—manually and inconsistently.

CHAPTER 4: CURRENT BARRIERS

The Five Core Pain Points

1 Fragmented Sources

Journals, preprints, patents, regulatory policy, and news remain disconnected. They do not communicate or cross-reference effectively.

2 No Historical Memory

Difficulty tracing the origin and evolution of an idea, making it impossible to see when or how an idea actually emerged.

3 Manual Verification

Analysts must double-check every critical claim, leading to significant delays and resource drain.

4 Information Overload

Too many papers or reports; users are short on filtering mechanisms and actionable context.

5 Low Trust in Summaries

Skepticism towards automated outputs; insistence on traceable **provenance and clear citation trails** for all insights.

Across all domains, people aren't short on information—they're short on **connected context**.

Poster Presentation

1. INTRODUCTION

Today, while vast amounts of trend-related information already exist, the challenge lies in accessing it: the current path to gather, curate, and synthesize this information is either nonexistent or painfully manual.

With the rise of LLMs and AI-driven tools, there is now an opportunity to build an automated pipeline that can bridge this gap, turning unstructured, distributed information into curated knowledge that fuels innovation insights.

PROBLEM STATEMENT

Design a systematic process that uses LLM by building a knowledge layer to identify trends and transform them into structured, actionable knowledge outputs.

3. RESEARCH METHODOLOGY

Primary User Research

Conducted interviews with 9 experts (2 Honda R&D, 7 academics) across AI, Robotics, and Neuroscience to gather firsthand perspectives.

Current Industry Solutions

Analyzed reports from leading firms like Gartner, McKinsey, S&P Global, Accenture, WEF, and MIT Tech Review to understand existing frameworks.

Technology Research

Explored advanced tools including GPT and other deep research platforms to assess their capabilities in trend identification and analysis.

4. USER PERSONA

Monica Geller, a Research Executive at Honda USA for over five years, leads a team of 15-20 analysts and defines research direction.

Goals and Motivations:

- Provide quarterly reports on key technologies & trend shifts
- Monitor competitors and summarise findings
- Define research direction and strategy for the coming year



"I spend hours researching trends across scattered sources. It's manual, overwhelming, and I often can't connect the dots."

5. PAIN POINTS



Fragmented Sources



Information Overload



No Historical Memory

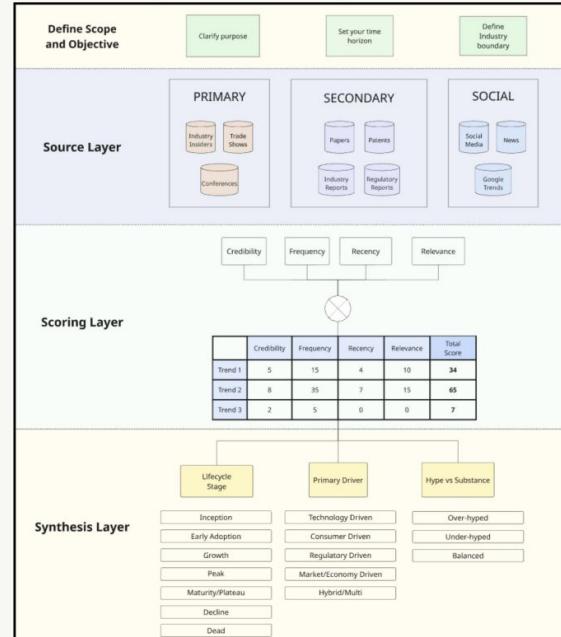


Low Trust in Automation



Manual Verification

6. SOLUTION

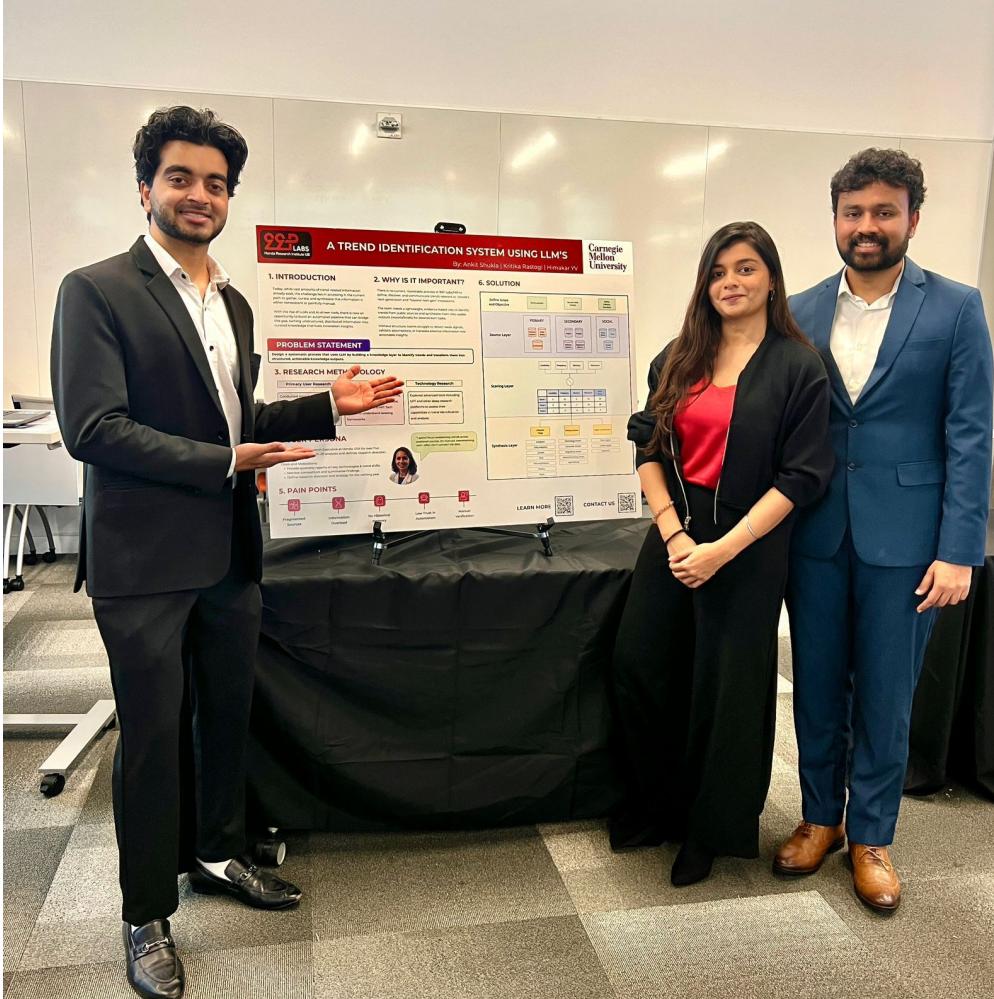


LEARN MORE



CONTACT US





Final Presentation



**How many of us truly know
where the industry is
heading?**

You can't steer the future if you can't see it.



Turning **Chaos** into **Clarity**

A systematic approach to identify trends using AI



Ankit Shukla



Kritika Rastogi



Himakar YV

Agenda

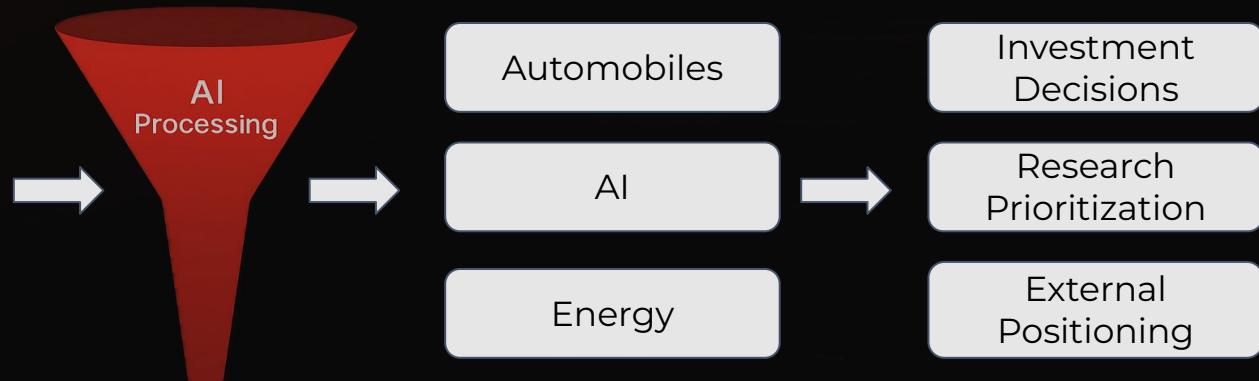
- 1 Problem Statement
- 2 Solution Overview
- 3 Process to Building the Solution
- 4 Deeper Solution Understanding and Demo
- 5 Current Limitations
- 6 Future Scope
- 7 Next Steps

PROBLEM STATEMENT



How might we define a process for using AI to research trends and generate high quality content?

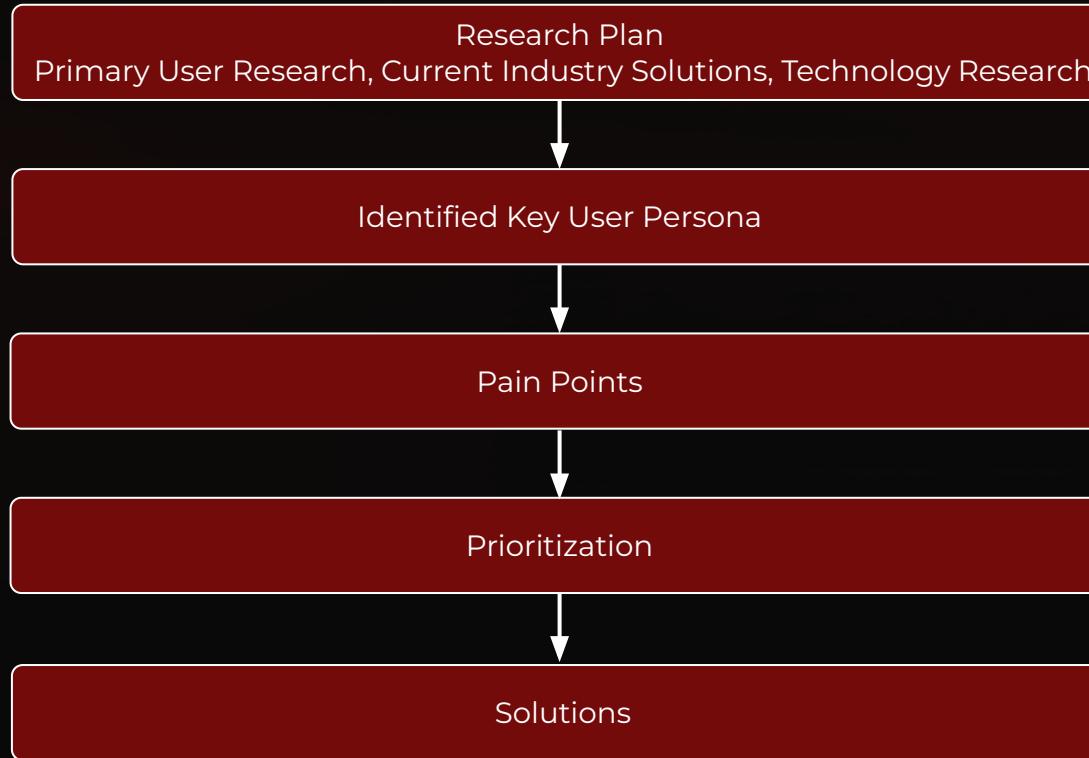
Lot of Information



Solution



Process to Building the Solution



Research Plan: Identified Current State

Primary User Research

Conducted interviews with 9 experts across AI, Robotics, and Neuroscience to understand their pain points and customer journey.

Current Industry Solutions

Analyzed reports from leading firms like Gartner, McKinsey etc, review to understand existing frameworks.

Technology Research

Explored tools like GPT and other deep research platforms to assess their capabilities in trend identification and analysis.

Key Insights: Manual & Fragmented | Inconsistent Time Horizons | Inconsistent Trend Ranking

Meet Monica Geller from Honda



"I spend hours researching trends across scattered sources. It's manual, overwhelming, and I often can't connect the dots."

Goals, Motivation & Tasks

- Provide **quarterly reports** on key technologies & trend shifts
- Monitor **competitors and summarise findings**
- Define **research direction** and strategy for the coming year
- **Summarize trends & shifts** manually and using LLM's
- Identify **the “next big thing”**

Pain Points



Fragmented Sources

78% of the respondents said they regularly miss important sources due to fragmentation.

Information Overload

80% *felt overwhelmed by the volume of emerging research in AI, robotics, and materials.*

No Historical Memory

Analysts spend **10+ hours/week** retracing old insights to connect the dots.

Low Trust in Automation

Only 2 in 9 respondents trust AI outputs without additional validation due to lack in transparency.

Manual Verification

There is no well-defined way to verify the legitimacy of a trend.

Key Learnings

We need to build an LLM-driven workflow that can:



- **Collect** signals across multiple sources
- **Score** them objectively
- **Synthesize** them into digestible insights

Scoping the Problem

Sources

Secondary and Social Sources

Industry

Automobiles

User Persona

Researcher at HRI

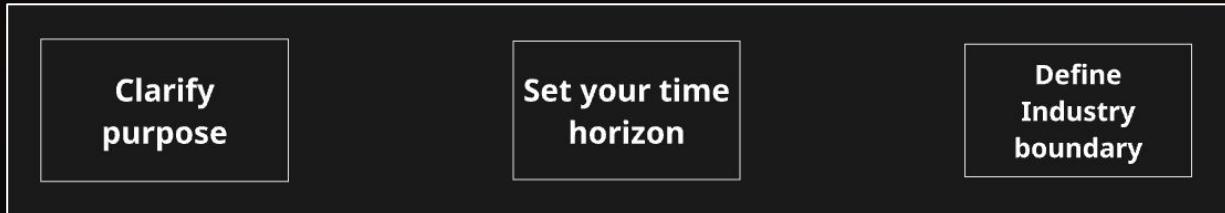
Downstream Task

Research Prioritization

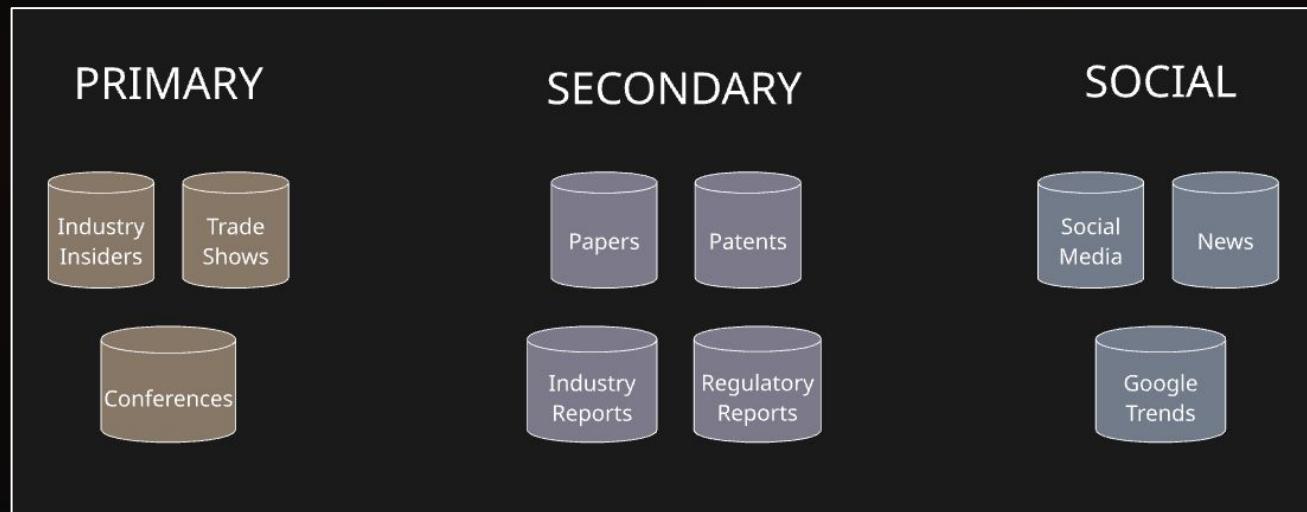
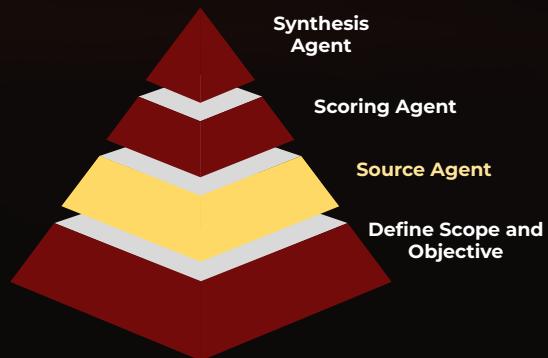
Prioritization - MoSCoW Framework

Must Have	Should Have	Could Have	Won't Have
<ul style="list-style-type: none">• Handle Fragmented Sources• Score Signals Objectively• Produce Consistent & Explainable Insights	<ul style="list-style-type: none">• Identify lifecycle stage• Identify what is driving the trend (Primary Driver)• Executive-Ready Summary	<ul style="list-style-type: none">• Visual dashboards	<ul style="list-style-type: none">• Real-time monitoring or alerts• Expansion beyond automobiles (Future Scope)

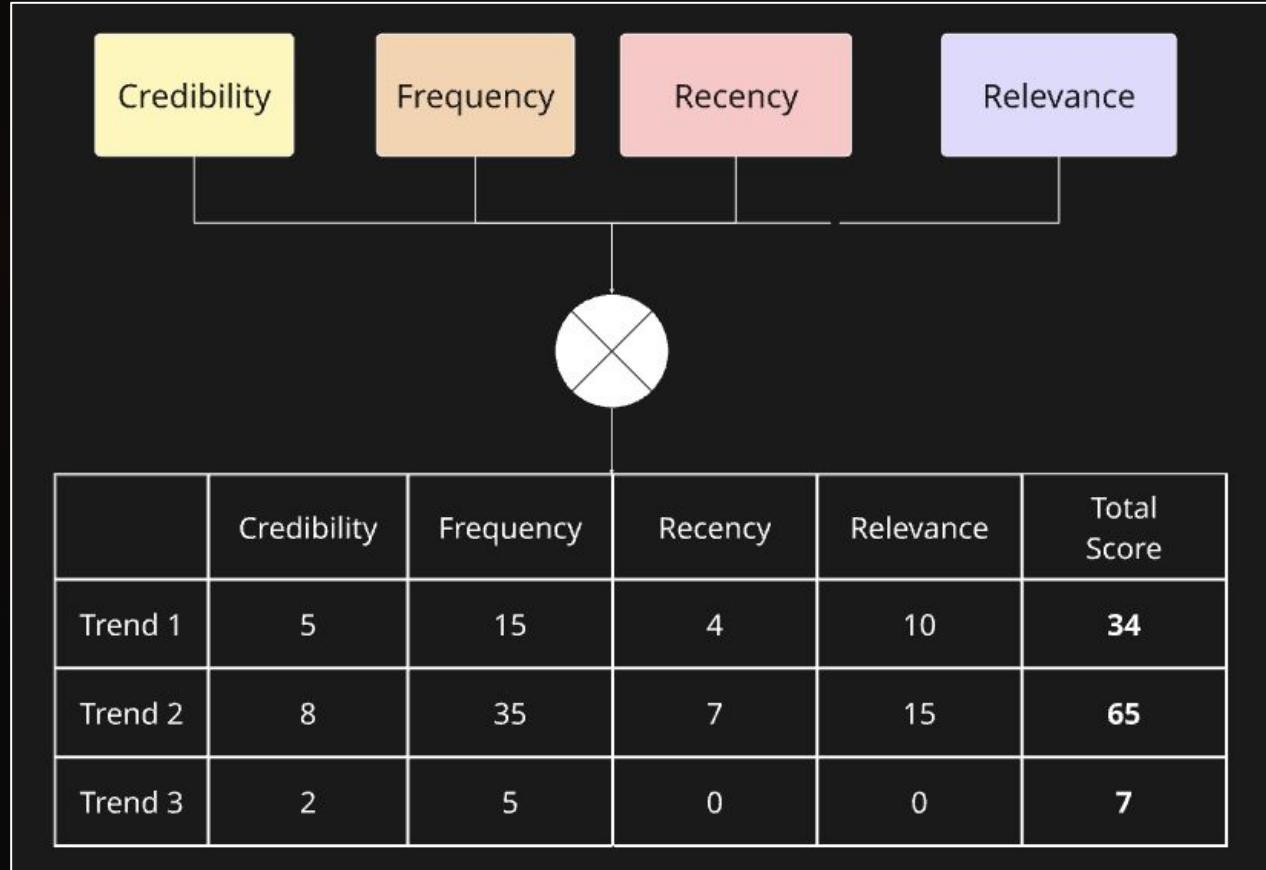
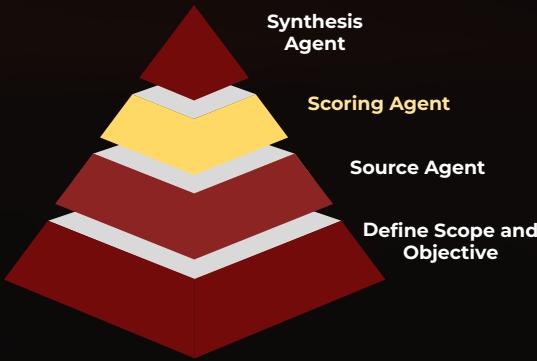
Solution Layer 1: Define Scope and Objective



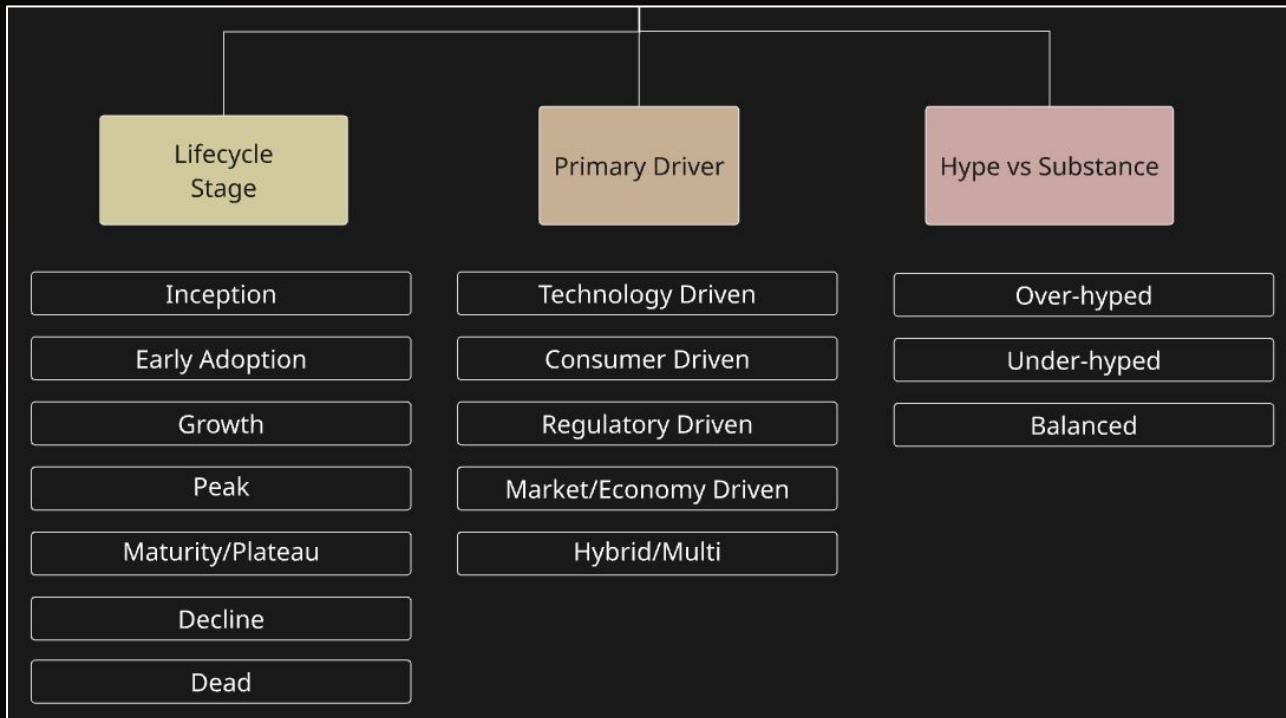
Layer 2: Source Agent



Layer 3: Scoring Agent



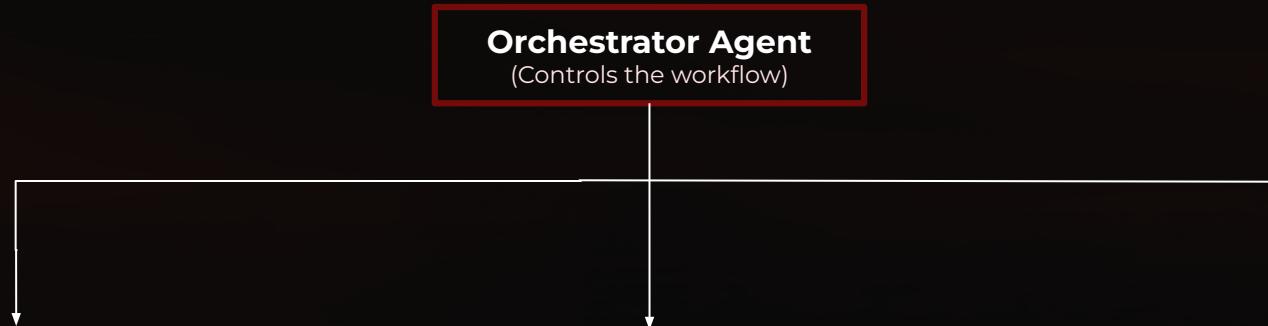
Layer 4: Synthesis Agent



System Overview - Multi-Agent System



[GitHub Link](#)



Source Agent

Collects signals from multiple sources - **papers, patents, reports, news & social media**

Generates trend signals from diverse sources

Scoring Agent

Scores signals on **Credibility, Frequency, Recency & Relevance**

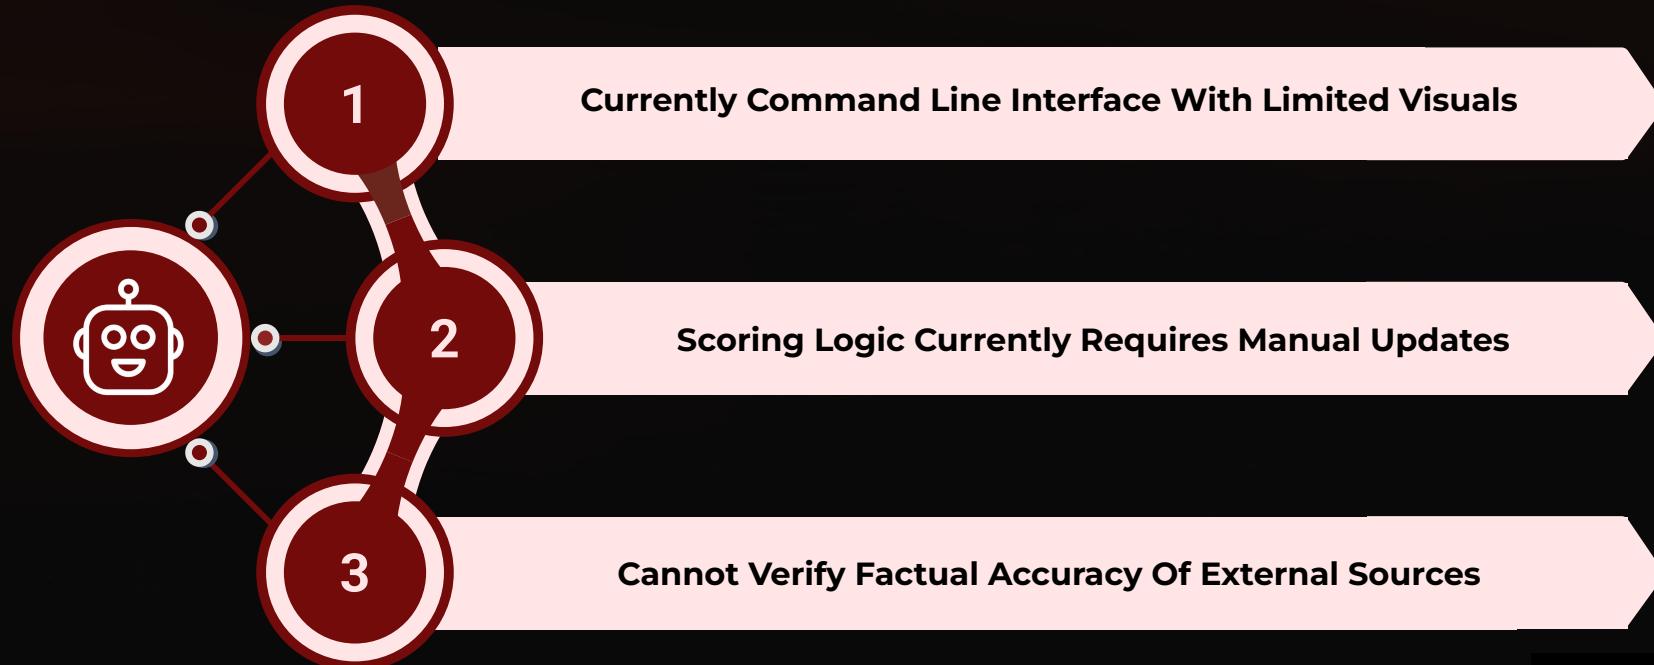
Generates score table

Synthesis Agent

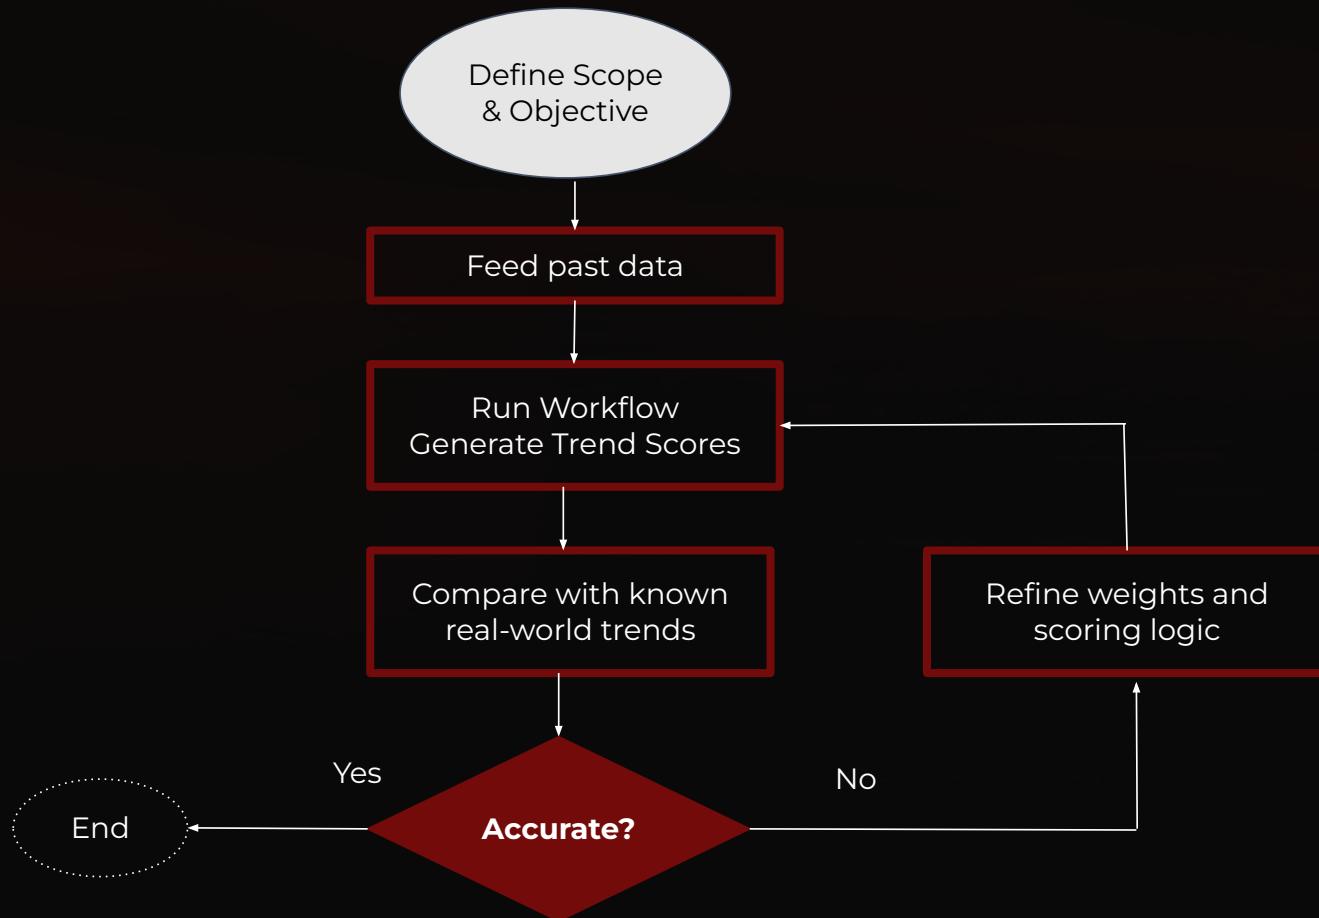
Synthesises scoring table information into trend profiles
Assigns lifecycle stage & primary drivers

Generates concise leadership summaries

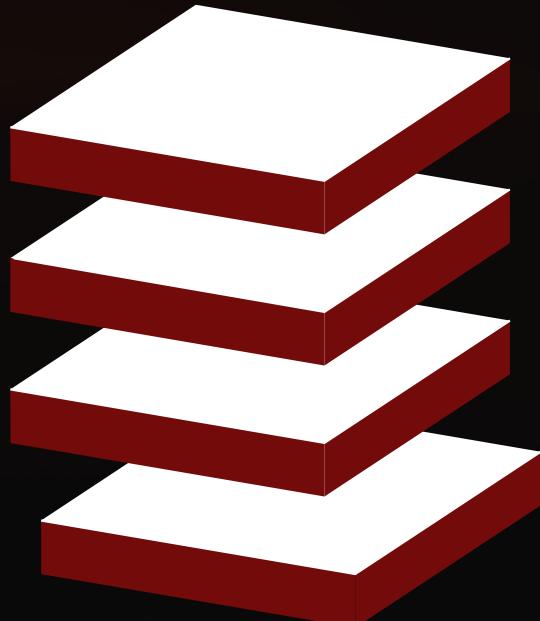
Current Limitations of the Solution



Future Scope: Evaluation



Recommended Next Steps



1

Roll out the tool internally

2

Integrate Internal Proprietary Data Into The Sources

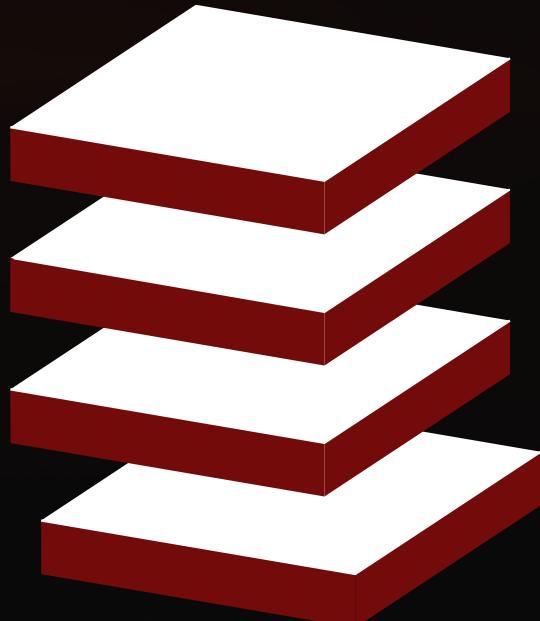
3

Add A Verification Agent To Validate Sources and Citations

4

Expand To Multiple User-Personas And Industries

Recommended Next Steps



5

Implement The Evaluation Workflow

6

Create A Living Database Of Trend Signals

7

Generate Shareable Outputs For Social Channels

8

Build Real time Monitoring Dashboards

Thank you!

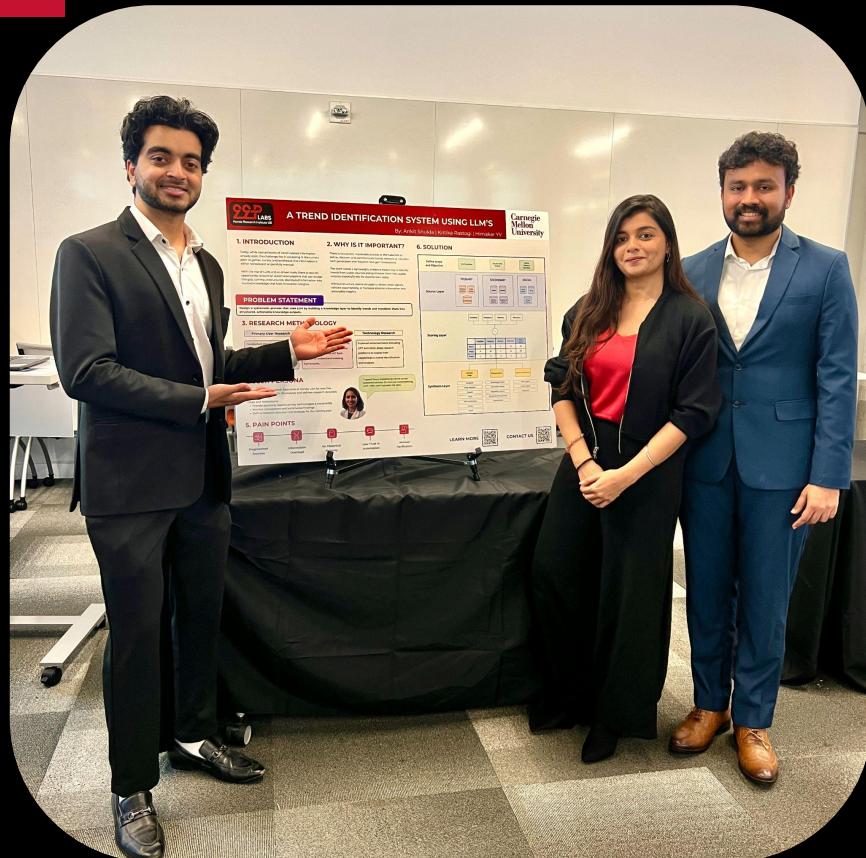
Team HRI & 99P labs:
Ryan, Rajeev, Marissa & Ito



Our Advisor:



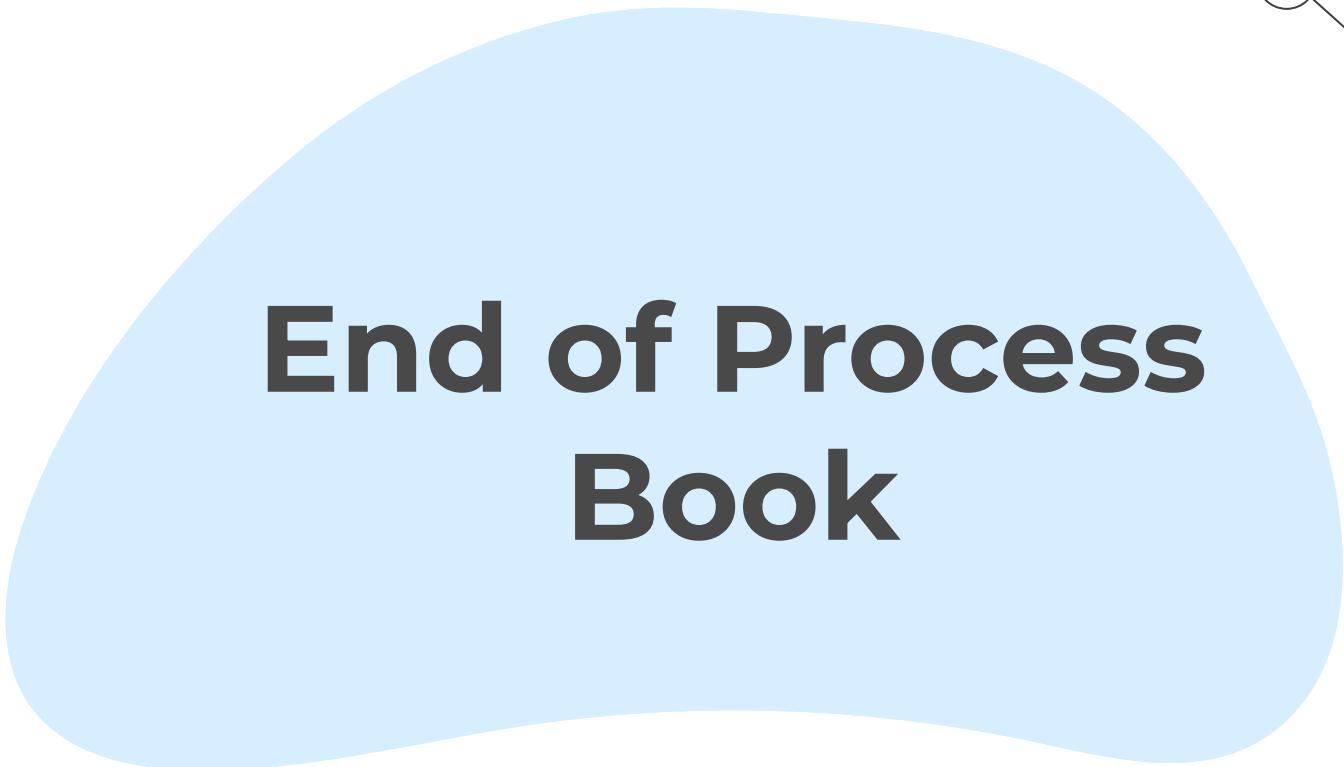
Nathan Mancine



Thank
You

CONTACT US





End of Process Book