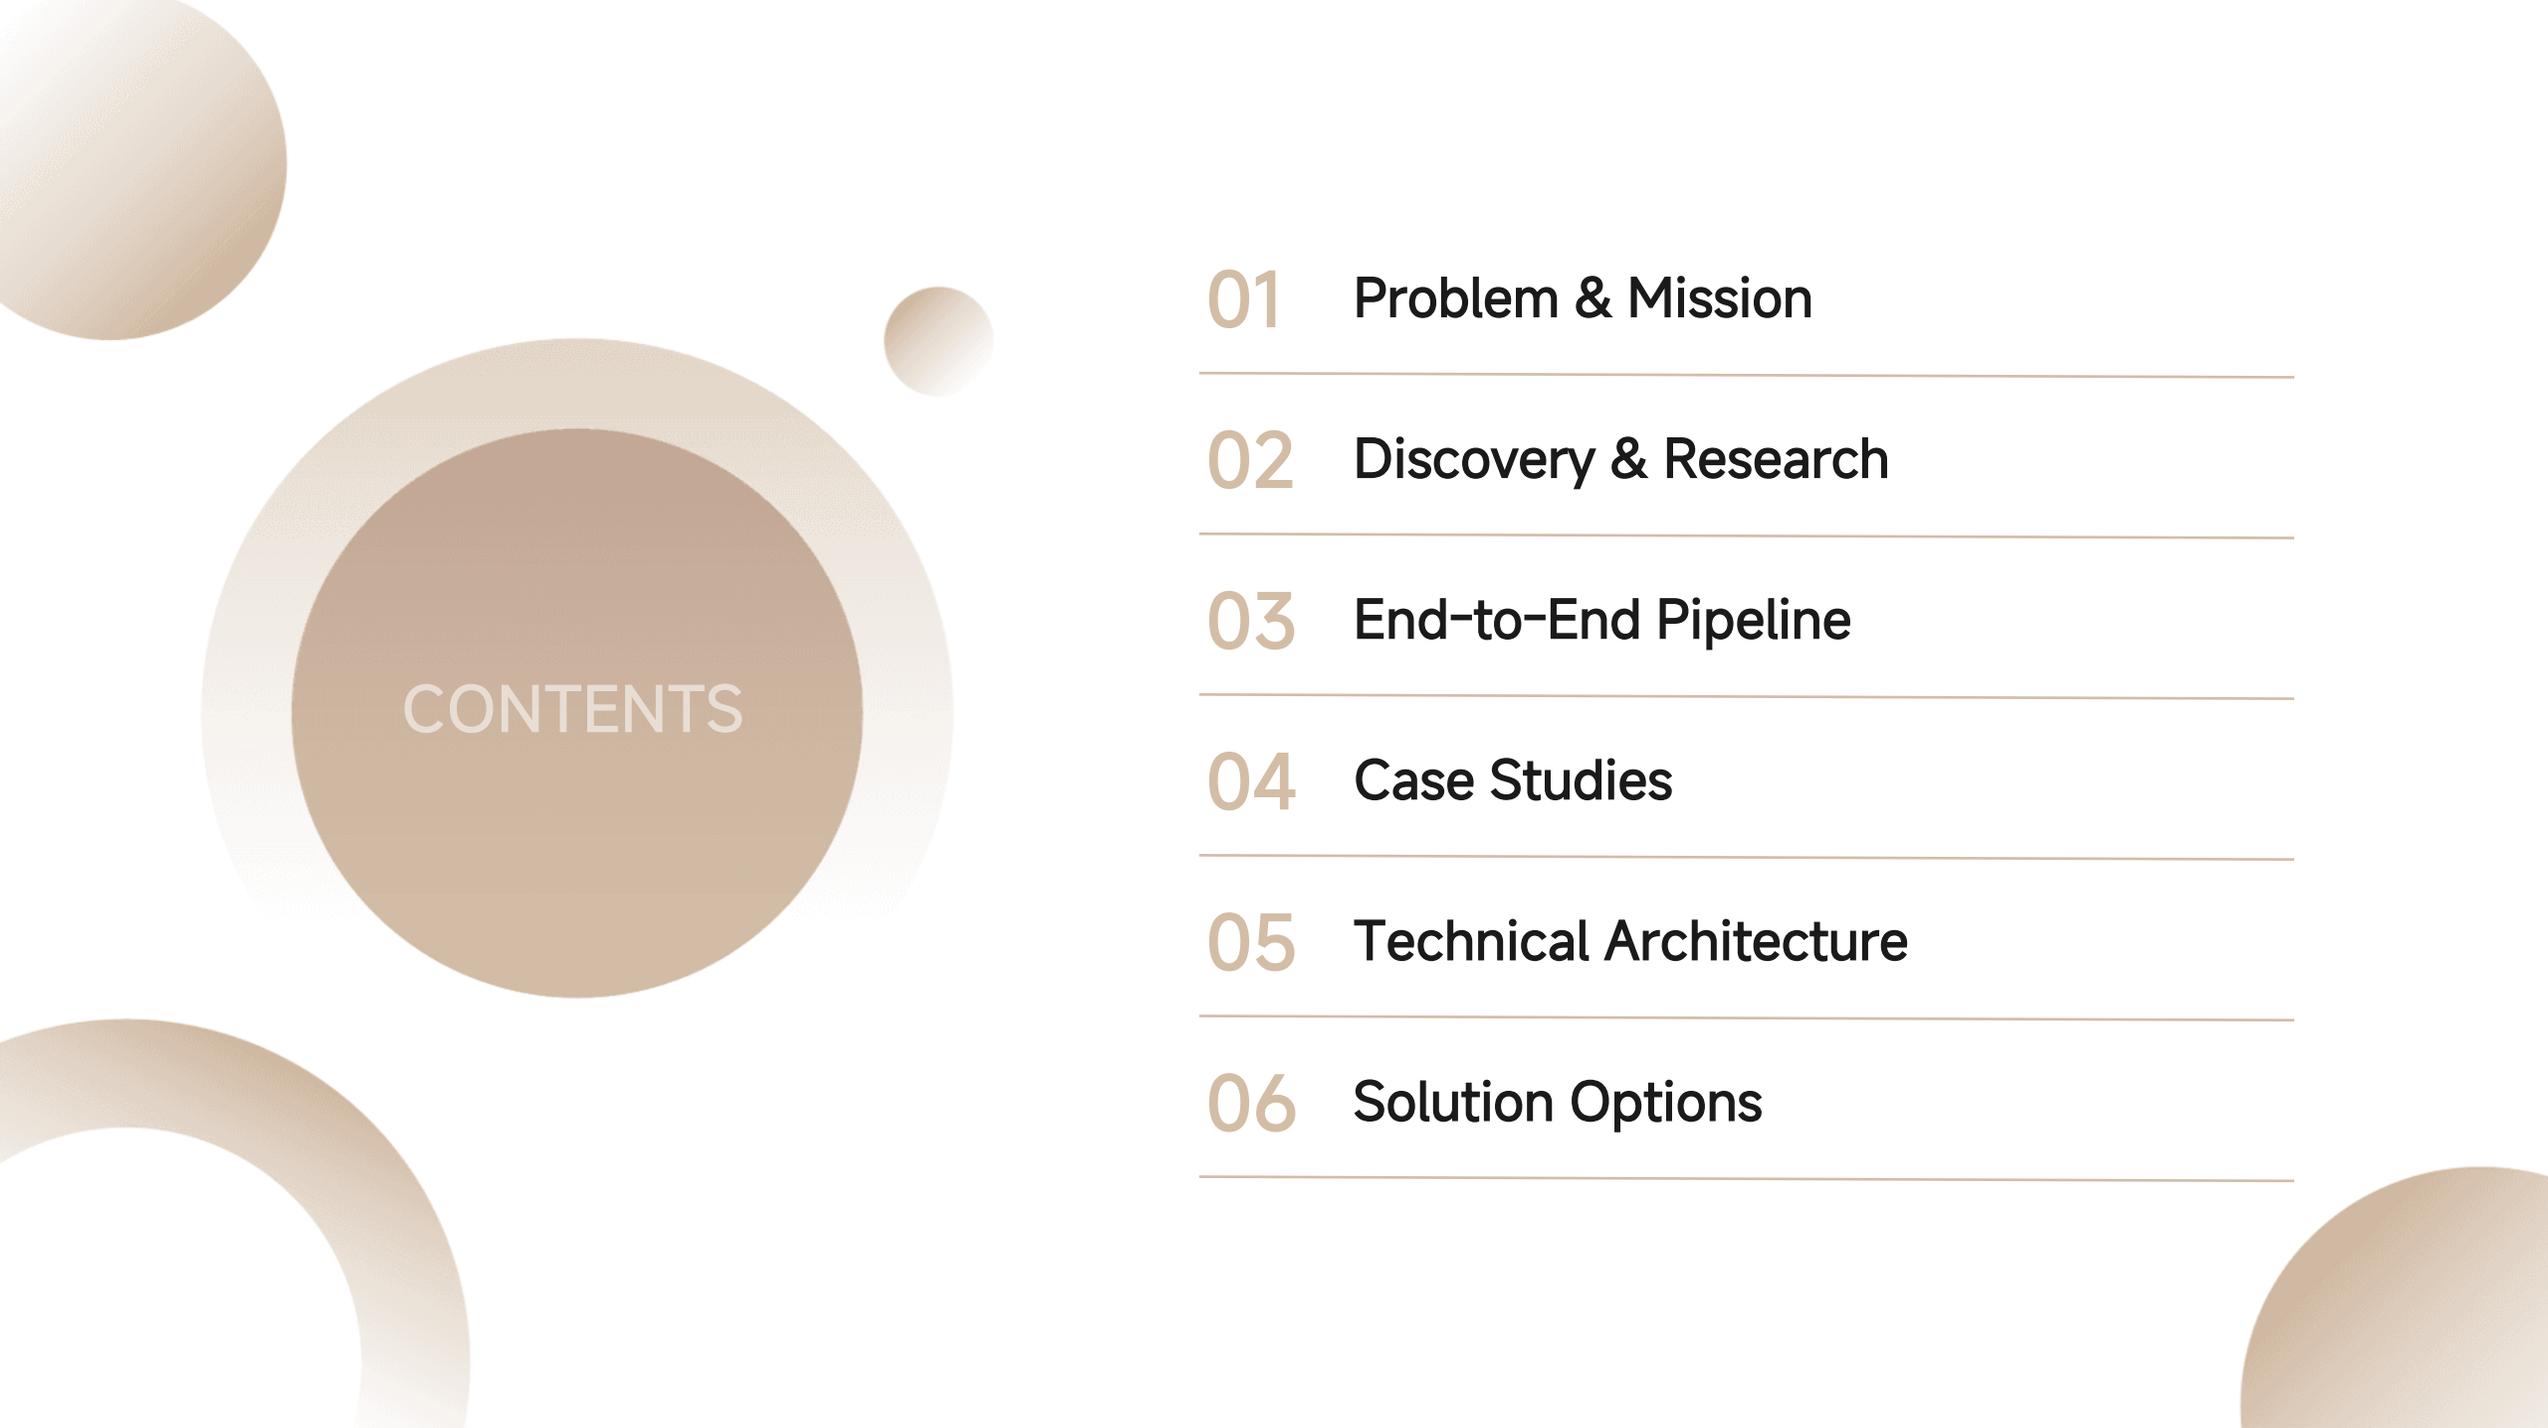


AI-Driven Patent Revival

Process Book

12/14/2025



CONTENTS

01 Problem & Mission

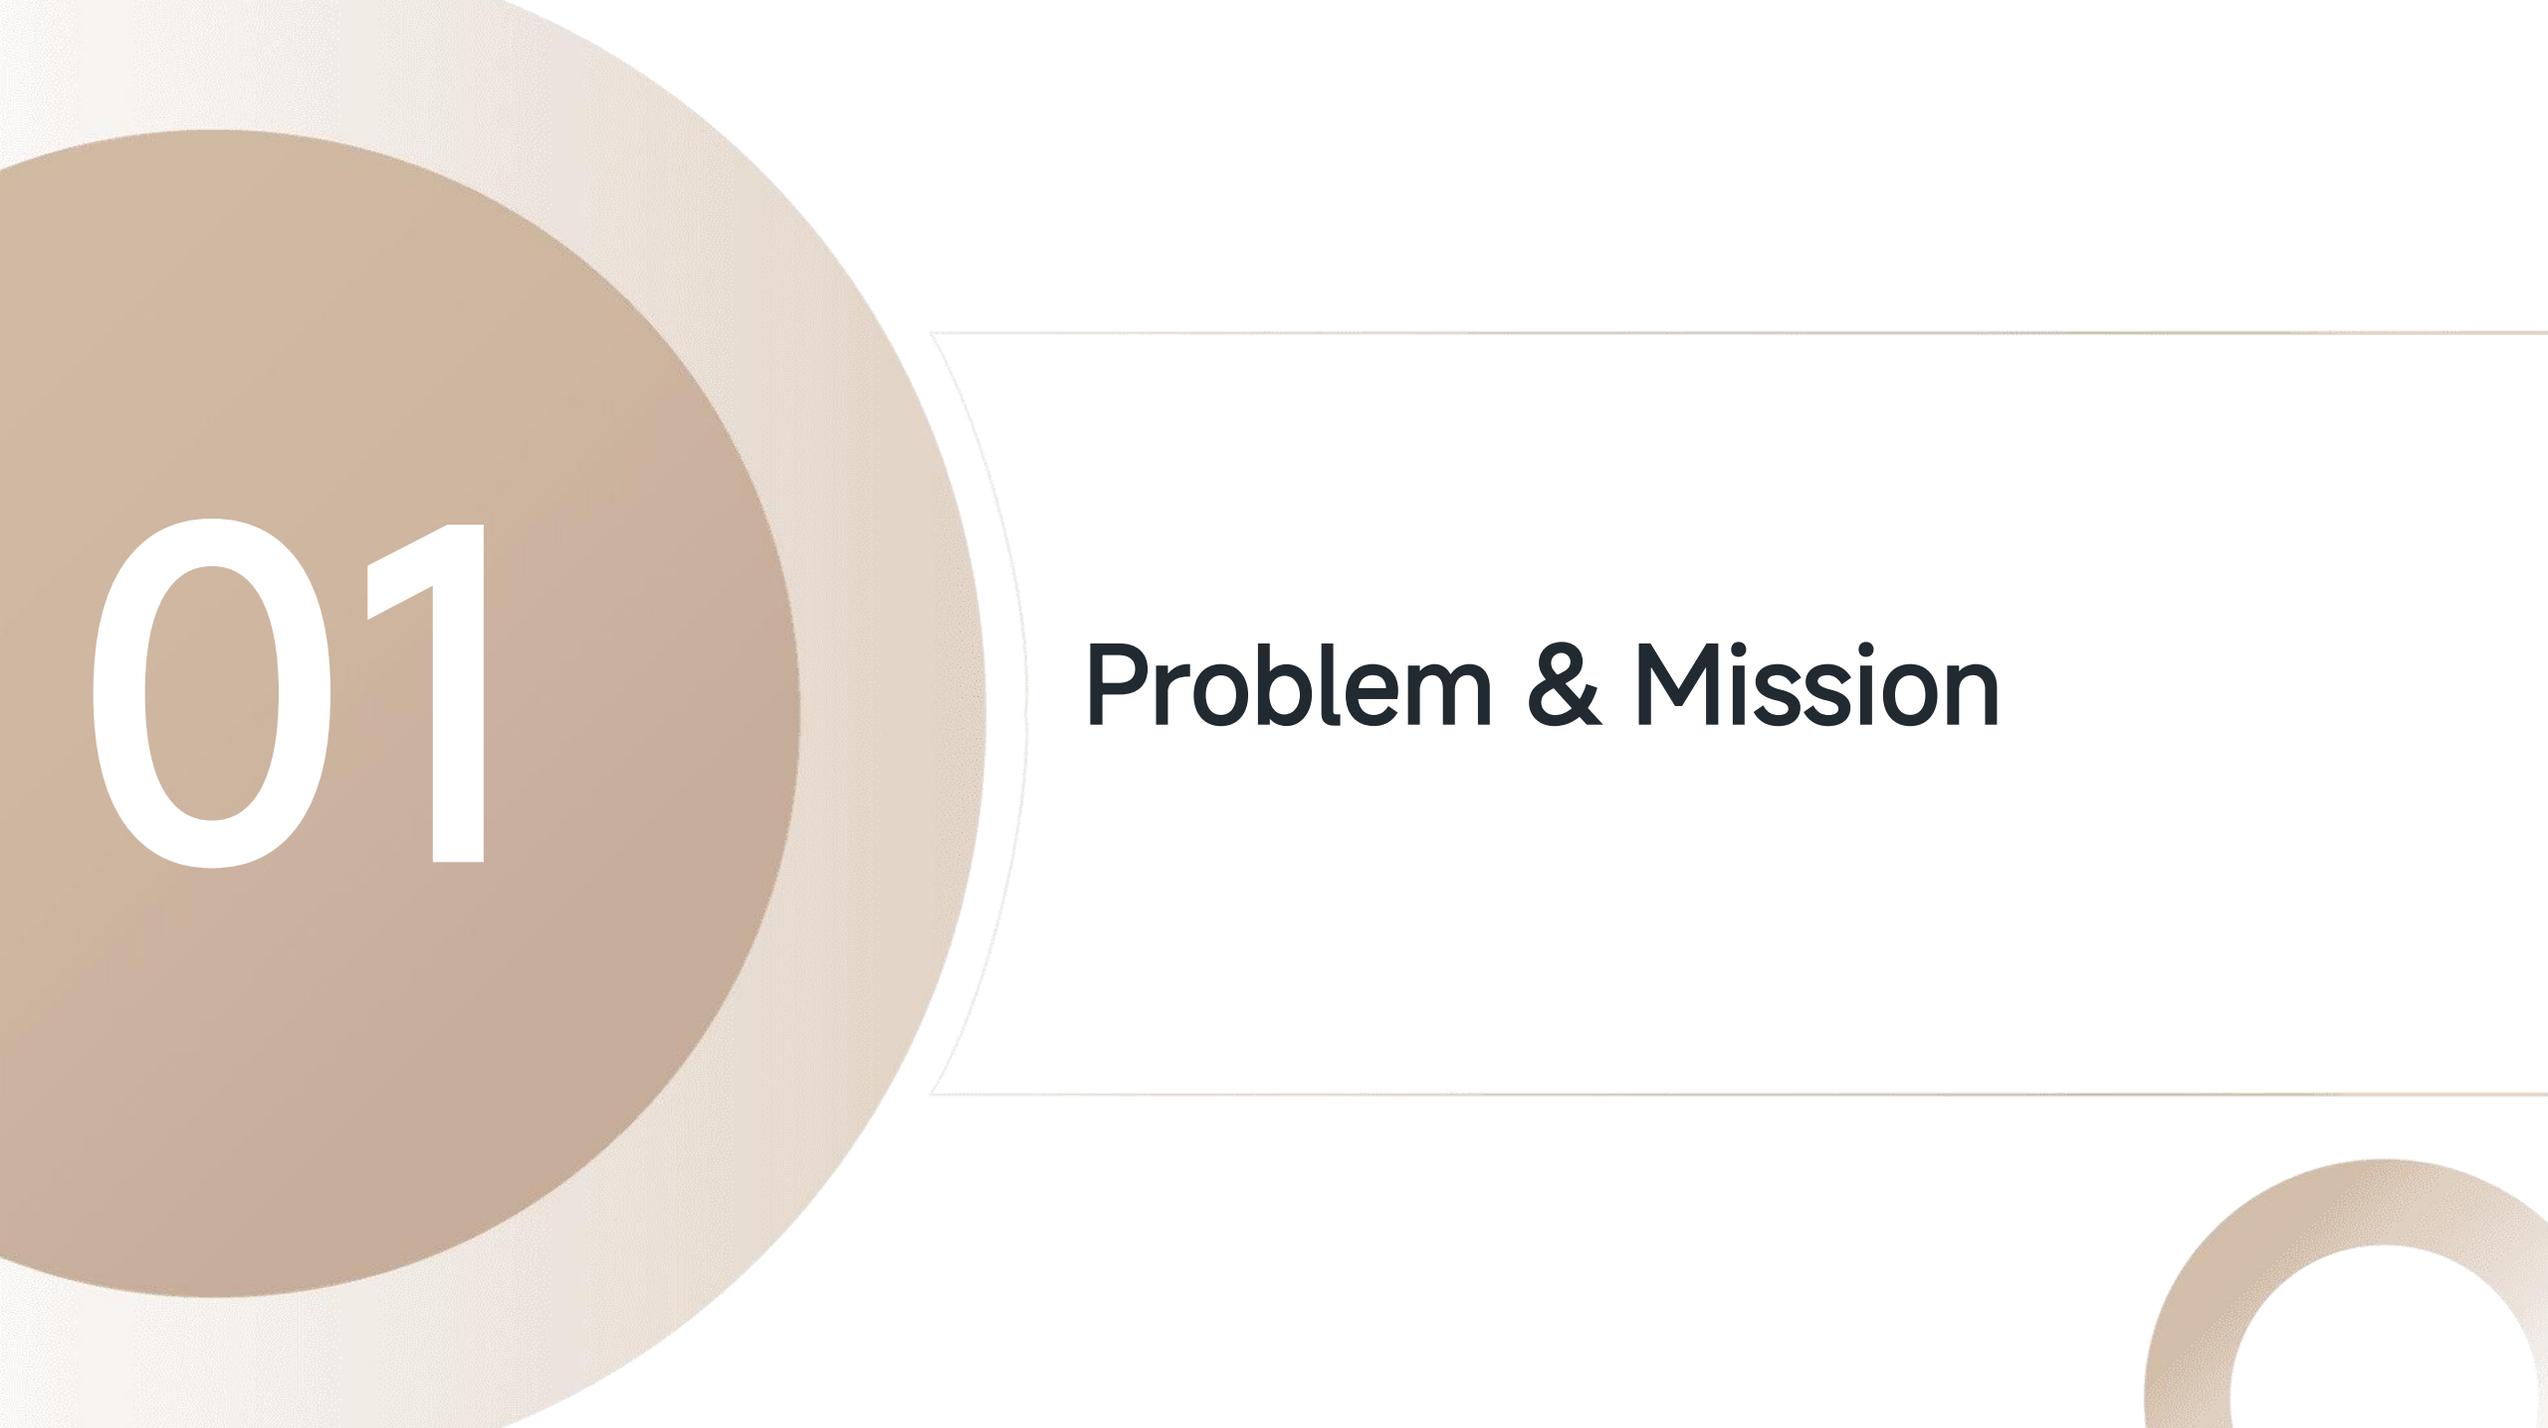
02 Discovery & Research

03 End-to-End Pipeline

04 Case Studies

05 Technical Architecture

06 Solution Options

The background features a large, light brown circle on the left side. A white callout box with a thin brown border extends from the right edge of this circle. The number '01' is written in white inside the circle. The text 'Problem & Mission' is centered within the callout box. In the bottom right corner, there is a partial view of another light brown circle.

01

Problem & Mission

AI-Driven Patent Revival

This project introduces an agentic AI framework designed to unlock hidden innovation by reviving dormant patents. It transforms unused IP into strategic business opportunities through automated analysis and cross-industry discovery.

The project is led by a dedicated team including Karan Qasba, Siddhant, and Tejas. It is sponsored by the Honda Research Institute and advised by Arun Venkatadri, ensuring strong industry support and expertise.

The objective is to demonstrate how AI can convert dormant patents into actionable insights, driving innovation and business value through an end-to-end agentic pipeline.

Project Introduction

Team and Sponsorship

Objective

Executive Summary



AI Capabilities

AI automatically analyzes patents, decomposes them into technology modules, discovers cross-industry applications, and scores opportunities based on novelty, feasibility, and market potential.

Deliverable

The final deliverable is a agentic pipeline that converts dormant patents into strategic insights, demonstrating the potential of AI to unlock hidden innovation.

Why This Problem Matters

Patent Accumulation

Large organizations accumulate thousands of patents each year, but only a fraction become products, generate IP revenue, or form partnerships.

Untapped Innovation

Many patents contain valuable ideas that remain unused because discovering new applications is slow, manual, and expert-dependent.

Dormant IP

Dormant IP represents untapped innovation and missed business value, highlighting the need for an automated solution to unlock its potential.

Current Challenges

The current process of discovering new applications for patents is slow and manual, making it difficult to fully leverage the value of existing IP.

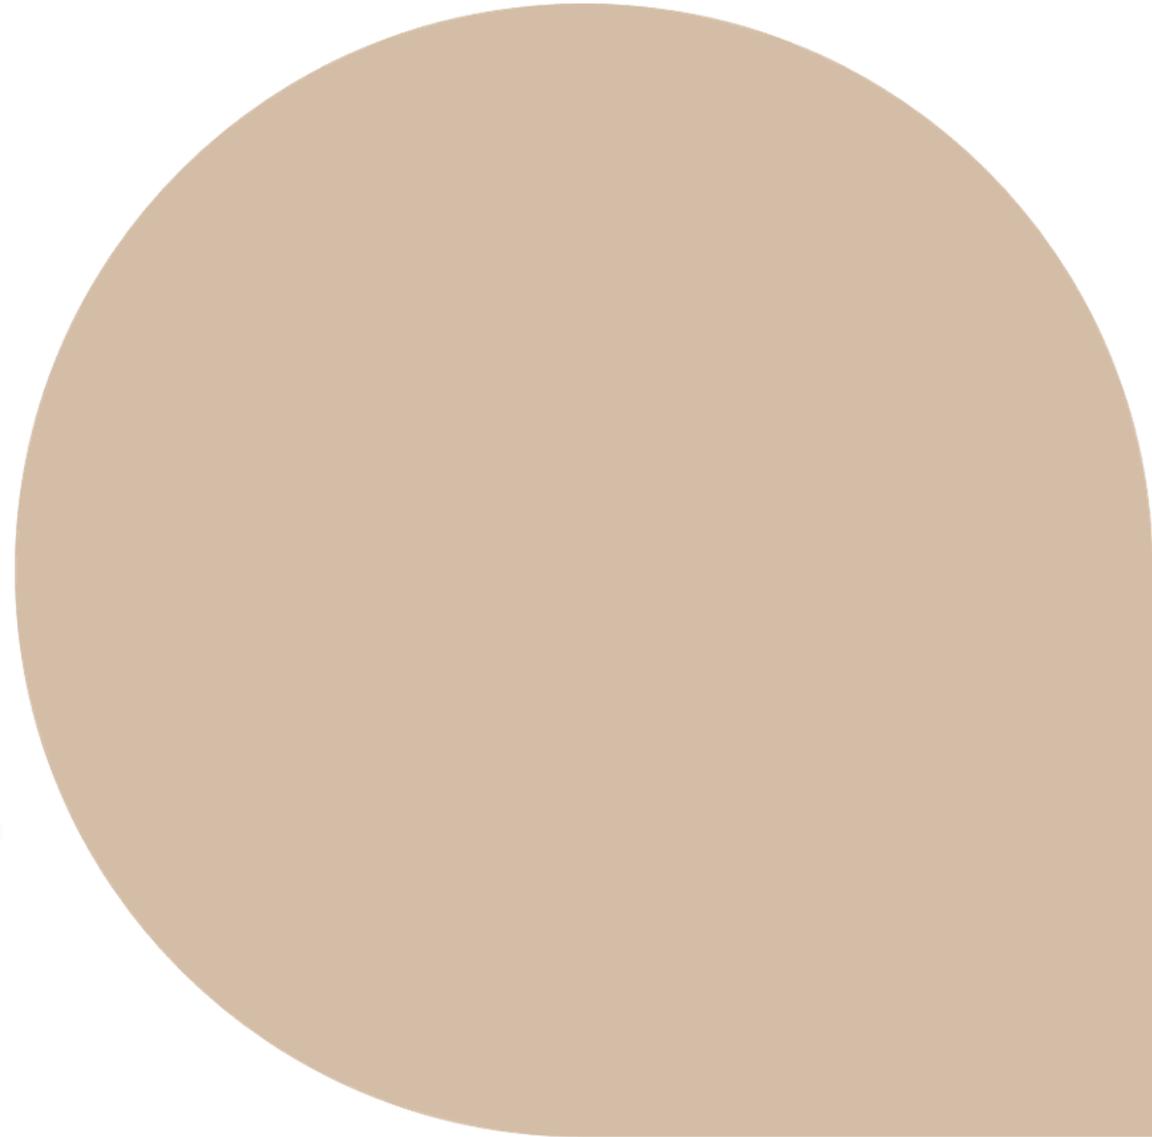


Market Reality



Industry Statistics

Industry statistics show that approximately 95% of patents never generate revenue, over 30% have transferrable potential but lack discovery mechanisms, manual review cycles take 6–18 months, and cross-industry applications are almost always missed.





Sponsor Context

High Volume of Patents

Organizations with large IP portfolios face challenges managing high volumes of technical patents and ensuring visibility across different domains.

Scalability Issues

There is no scalable way to map patents to new industries, making it difficult to identify which opportunities deserve investment.

Strategic Needs

The sponsor seeks a solution that can prioritize opportunities, enabling strategic decision-making and maximizing the value of existing patents.



Initial Sponsor Challenge



Sponsor's Question

The sponsor asked, "Can AI help us understand what else our existing patents can do?" This reframed the project from incremental search to strategic opportunity discovery.

Project Reframing

The project shifted focus from summarizing patents to discovering strategic opportunities hidden within them, aligning with the sponsor's need for actionable insights.

Our Mission

System Design

Design a scalable and explainable system that reads and understands patents, decomposes them into core technologies, and identifies cross-domain use cases.

Visualization

Visualize relationships between patents and potential applications, making it easier for stakeholders to grasp complex connections.

Opportunity Prioritization

Prioritize high-value opportunities based on strategic criteria, enabling informed decision-making and efficient resource allocation.

Business Impact

The mission is to transform dormant patents into actionable business opportunities, driving innovation and maximizing the value of existing IP assets.



Why AI + Multi-Agent Approach

AI and Multi-Agent Benefits

Traditional NLP is insufficient for patent revival, which requires reasoning, not just text extraction. Multi-agent systems offer stepwise reasoning, traceable logic, domain-agnostic generalization, transparency, and modularity, making them ideal for this challenge.

01

End-to-End Pipeline

An end-to-end agentic pipeline capable of transforming raw patent text into technology modules, cross-industry analogies, knowledge-graph connections, strategic scores, and visual dashboards.

02

Strategic Insights

The pipeline delivers strategic insights by identifying hidden opportunities within patents, enabling organizations to unlock new revenue streams and innovation pathways.

03

Business Value

The final output is designed to provide clear, actionable recommendations that guide investment decisions and maximize the value of dormant patents.



02



Discovery & Research

Research Goals

Understand Current Workflows

The research aims to understand current patent evaluation workflows, identifying bottlenecks and unmet needs in the process.

Benchmark Industry Tools

Benchmark existing industry tools and their limitations to identify gaps that the proposed solution can fill.

Explore Feasibility

Explore the feasibility of using a multi-agent AI approach to address the challenges of patent evaluation and opportunity discovery.



Inform System Requirements

The insights gathered will inform the design of the AI system, ensuring it meets the needs of stakeholders and delivers meaningful results.



What We Learned

Manual Evaluation

Patent evaluation remains deeply manual, with cross-industry discovery rarely attempted, highlighting the need for automation.

Hidden Value

Hidden value resides at the modular technology level rather than in entire inventions, indicating the importance of modular decomposition.

Explainability

Stakeholders demand explainable insights instead of black-box answers, emphasizing the need for transparency in the AI system.

Insights That Shaped Approach



Key Insights

Key insights include the need to break patents into functional building blocks, the critical role of analogical reasoning, the importance of visualization, and the necessity for scalability and transparency.

Problem Reframing



Initial Objective

01

The initial objective was to summarize patents with AI, but this was reframed to focus on discovering strategic opportunities hidden within patents.

New Focus

02

The reframed objective aligns the project with business impact, emphasizing the discovery of high-value opportunities rather than mere summarization.

Design Principles

Interpretability



Ensure that every recommendation is traceable, allowing stakeholders to understand the reasoning behind AI-generated insights.

Scalability

Design the system to handle thousands of patents effortlessly, ensuring it can grow with the organization's needs.

Domain-Agnostic Reasoning

Enable cross-industry transfer by designing the system to reason about patents in a domain-agnostic manner.

Human-in-the-Loop

Maintain human oversight by ensuring the system is compatible with human-in-the-loop decision-making, keeping experts in control.

Approach Summary

Ingestion

The Ingestion agent transforms raw patent text into structured data, preparing it for downstream processing.

Decomposition

The Decomposition agent isolates technology modules from patents, enabling granular analysis and cross-industry matching.

Evaluation

The Evaluation agent scores and ranks opportunities, providing actionable insights for strategic decision-making.



03

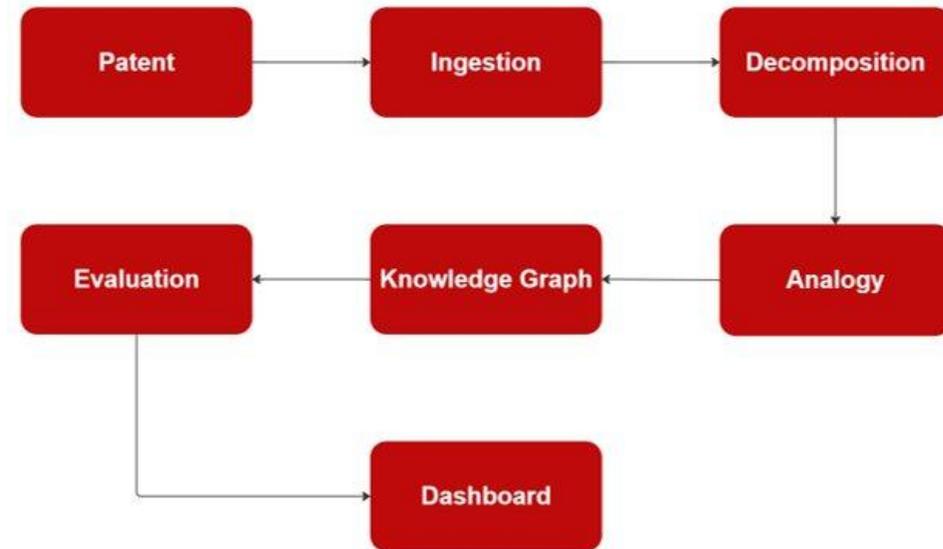
End-to-End Pipeline

Pipeline Overview



Pipeline Flow

The pipeline visualizes the flow from raw patent text through Ingestion, Decomposition, Analogy Discovery, Knowledge Graph Construction, Evaluation, and Dashboard stages, showcasing the end-to-end process.



Agent 1: Patent Ingestion



Ingestion Purpose

The Ingestion agent transforms raw patent PDFs into structured data, extracting metadata, claims, descriptions, summaries, tags, and vector embeddings.

Output

It produces a structured patent object containing a plain-English summary, domain-relevant keywords, cleaned text sections, and embedding vectors.

Importance

This standardization enables downstream agents to operate on clean, machine-readable data, ensuring consistency and reducing redundant preprocessing.

Ingestion: What It Produces



Structured Object

The Ingestion agent outputs a structured patent object with a concise summary, keywords, cleaned text sections, and embedding vectors.

Downstream Feeding

These elements feed directly into decomposition and analogy agents, ensuring a consistent data format throughout the pipeline.

Agent 2: Modular Decomposition

Granular Matching

By isolating functional building blocks, the system enables granular cross-industry matching that would be invisible at the whole-patent level.



Decomposition Purpose

The Decomposition agent parses patents into discrete technology modules such as sensors, control logic, mechanical systems, algorithms, and interfaces.



Reusable Modules

Modules become reusable innovation assets, allowing the same technology to be applied in multiple industries.



Unlocking Value

Decomposition unlocks latent value by shifting analysis from document similarity to functional complementarity, revealing hidden opportunities.

Why Decomposition Matters



Cross-Industry Matches

Cross-industry matches emerge at the module level because functions like sensing or control transcend domains. Modules enable granular matching and reuse, unlocking hidden value.



Agent 3: Analogy Discovery

Analogy Purpose

The Analogy agent asks, “Where else does this problem exist?” It combines embedding similarity with LLM reasoning to perform cross-domain matching.



Analogies Generated

It generates functional, system, component, and industry-level analogies, translating technical modules into business opportunities.



Business Impact

This step maps technical modules to analogous challenges in distant markets, enabling the discovery of high-value opportunities.

Types of Analogies Generated



Analogies Types

The agent outputs four analogy types: functional, system-level, component-level, and industry-level, providing a taxonomy for clear opportunity identification.

Taxonomy Benefits

This taxonomy helps stakeholders quickly grasp the nature of each opportunity and select relevant go-to-market strategies.

Agent 4: Knowledge Graph

01

Graph Construction

The Knowledge Graph Architect creates visible relationships by constructing nodes for patents, modules, domains, and applications.

02

Edge Creation

Edges describe functional similarity, dependencies, and viability, enabling cluster-level insights and white-space discovery.

03

Graph Benefits

The graph provides clear visibility of transferable technology, supports interactive exploration, and accelerates strategic planning.

04

Stakeholder Value

Stakeholders can navigate from a single module to multiple industries, uncovering partnership routes and licensing targets.



Knowledge Graph Benefits



Graph Advantages

The knowledge graph delivers clear visibility of transferable technology, supports cluster-level insights, and uncovers white-space opportunities, enabling strategic decision-making.



Agent 5: Evaluation & Scoring

Evaluation Purpose

The Evaluation agent scores each opportunity across novelty, modularity, transferability, feasibility, and market potential.

Composite Rankings

Composite rankings surface high-value, low-risk applications first, enabling portfolio managers to focus resources on the most strategic opportunities.

Business Alignment

The scoring system balances innovation potential with implementation practicality, ensuring that ranked opportunities are both visionary and executable.



Score Interpretation



01

Transferability

High transferability signals wide reuse potential, indicating that a module can be applied across multiple industries.

02

Modularity and Feasibility

High modularity correlates with many possible applications, while high feasibility reflects low adoption barriers, guiding investment decisions.

The background features a large, light brown circle on the left side, partially overlapping a larger, slightly darker brown circle. A white callout box with a thin brown border extends from the right edge of the larger circle. The number '04' is centered within the larger circle, and the text 'Case Studies' is centered within the callout box. In the bottom right corner, there is a partial view of a brown ring.

04

Case Studies

Sunshade Patent Selection



Selection Criteria

The sunshade patent was selected for its intuitive mechanics and clear module boundaries, making it ideal for demonstrating the pipeline's ability to extract functional blocks and reason about cross-domain applications.

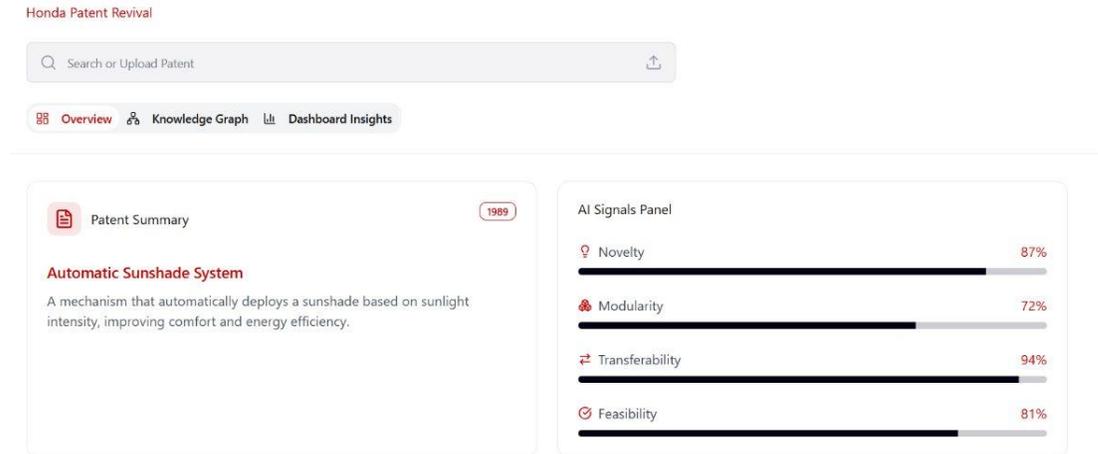
Sunshade Patent Summary

Patent Description

The patent describes an automatic sunshade mechanism featuring light sensing, motor control, adjustable barrier, and safety logic, providing a straightforward architecture for analysis.

Decomposition

Decomposition yields four modules: Light Sensor, Control Logic Unit, Motor Actuator, and Shade Barrier, which become the basis for analogy discovery.



Demonstration

This simple yet functional patent demonstrates how the pipeline can identify reusable technology modules applicable to multiple industries.

Extracted Modules



Module Identification

01

The decomposition process isolates functional modules such as Light Sensor, Control Logic Unit, Motor Actuator, and Shade Barrier from the sunshade patent.

Reusable Blocks

02

These reusable blocks become the seeds for analogy discovery, illustrating how even simple patents contain technology lego-bricks applicable to multiple industries.



Opportunity Identification

The agent proposes three high-match applications: smart automated blinds for homes, greenhouse shading systems for agriculture, and solar panel temperature protection for energy.

Functional Analogies

Each opportunity reuses the same sensing-control-actuator pattern, demonstrating functional analogies across disparate markets.

Business Potential

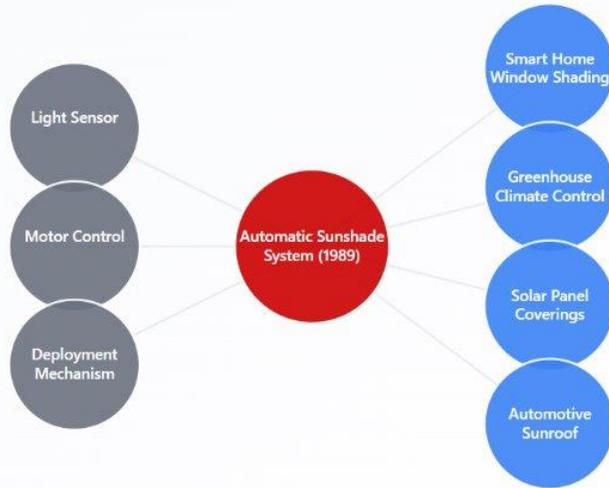
These cross-domain opportunities highlight the pipeline's ability to uncover hidden value in seemingly simple patents.

Sunshade Knowledge Graph



Honda Patent Revival

Overview Knowledge Graph Dashboard Insights



Number of connected domains

3

Top Match

Smart Homes

91% Match

AI Insight

"High transferability due to shared control and sensing logic."

Modules 3

Applications 4

Graph Snapshot

A subgraph snapshot shows sunshade modules linked to smart-home, agriculture, and energy domains via similarity edges, revealing technology communities and white-space gaps.

Opportunity Evaluation



Evaluation Results

The system ranks Smart Home Automated Blinds at 91% match score, citing high modularity, strong market demand, and low regulatory barriers.

Business Recommendation

This concrete numeric output exemplifies how the pipeline converts technical patents into prioritized business recommendations ready for executive review.

Active Vehicle Control Patent

01

Patent Overview

Introduces a 2020 Honda patent on driving assistance that combines direct and indirect indicators to predict vehicle behavior and generate stability decisions.

02

Complexity

The patent's complexity tests the pipeline's ability to handle sophisticated control algorithms and multi-sensor fusion beyond simple mechanical cases.

03

Opportunity

Its rich functional layers provide fertile ground for cross-domain analogies, especially in autonomous systems where prediction and stability are universal challenges.

Vehicle Control Summary



System Integration

The system integrates behavior prediction, sensor fusion, auto-control logic, and safety/stability models to assist drivers in real time.

Analogies Potential

These integrated functions provide a foundation for cross-domain analogies, particularly in mobility platforms where prediction and stability are critical.

Module Extraction Results

Module Isolation

Decomposition isolates Behavior Prediction, Sensor Fusion, Auto-Control Logic, and Safety/Stability Model from the vehicle control patent.

Algorithm Extraction

These modules encapsulate algorithms and data flows rather than physical parts, showcasing the pipeline's domain-agnostic capability.



Software-Centric Patents

The pipeline demonstrates its ability to handle both mechanical and software-centric patents, expanding its applicability.

Cross-Domain Reuse

The extracted modules highlight potential for cross-domain reuse in autonomous systems, smart infrastructure, and other mobility platforms.

Dashboard Snapshot



Patent Number

US 10,625,776 B2

Title

Driving Assistance Technique for Active Vehicle Control

Year Granted

2020

Assignee

Honda Research Institute Europe GmbH

Description

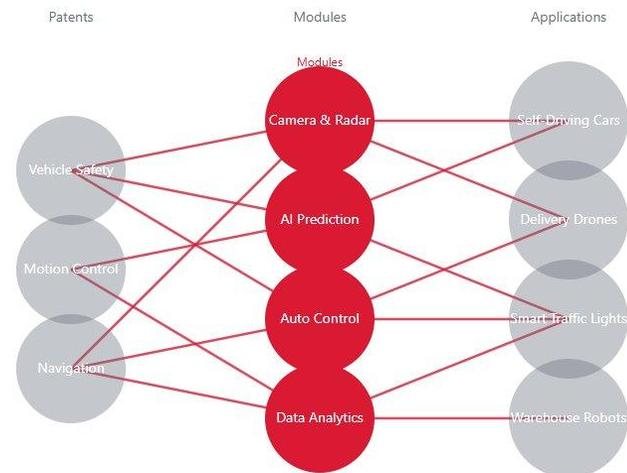
An advanced driving assistance system that combines direct and indirect indicators to predict vehicle behavior and generate safer, more stable control decisions.

Dashboard Features

An executive dashboard displays ranked industry matches, novelty vs. feasibility scatter plots, and module-level similarity scores.

User-Friendly Interface

The visual interface allows non-technical decision-makers to grasp opportunities instantly, supporting data-driven portfolio investment.



Signals



Novelty



Modularity



Transferability



Feasibility

AI Matchmaking Results

Top analogies include self-driving vehicles at 91%, delivery drones at 87%, and smart traffic control systems at 82%.

These matches validate that prediction and stability functions are transferable across mobility platforms, confirming the pipeline's ability to surface non-obvious reuse scenarios.

The identified analogies highlight opportunities for leveraging existing IP in emerging mobility domains, driving innovation and revenue.

Top Analogies

Transferability

Business Impact



Key Findings

Key findings highlight that safety systems exhibit the highest transferability, while modularity correlates with wider applicability across industries.

Strategic Guidance

These insights guide future patent filing strategies, suggesting that modular, safety-focused innovations offer the broadest monetization potential.

Knowledge Graph Results

Graph Expansion

The expanded graph reveals dense interconnections between vehicle control modules and emerging mobility domains such as drones and smart infrastructure.

Technology Convergence

Cluster analysis identifies technology convergence zones, informing R&D collaboration and standardization efforts before competitors enter the space.

Strategic Insights

The graph provides strategic insights by visualizing relationships and identifying white-space innovation zones for future exploration.

Final Opportunity Ranking



Ranking Results

The pipeline concludes with Self-Driving Vehicles ranked first, followed by Delivery Drones and Smart Traffic Systems, providing a clear roadmap for leveraging existing IP.

AI Matchmaking Results

Top industries where this technology can be applied



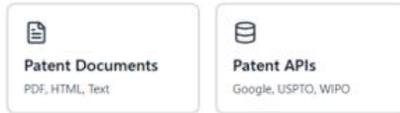
05

Technical Architecture

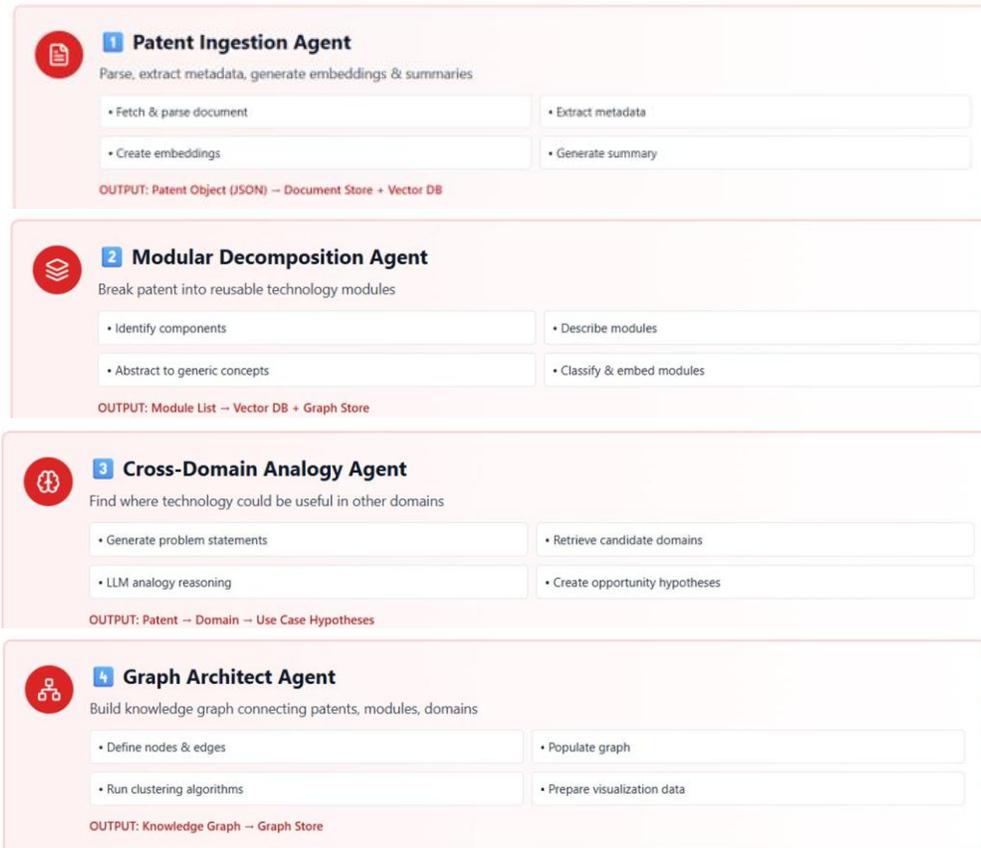
System Architecture Overview



INPUT SOURCES

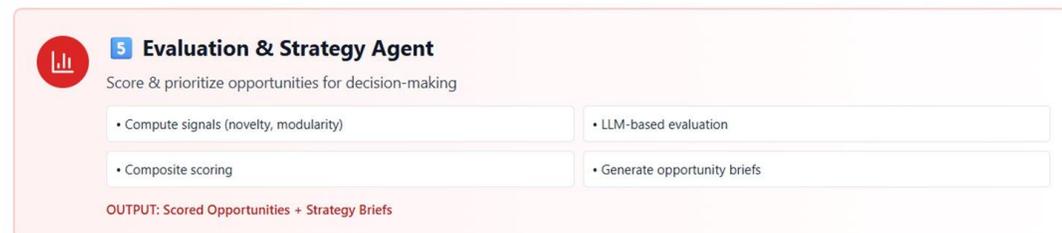


AGENT PIPELINE



Architecture Diagram

A high-level diagram illustrates data flow from patent documents through document store, vector database, knowledge graph, multi-agent orchestrator, reasoning models, and dashboard layer.



Core Components

Document Store

The Document Store holds raw patent documents, ensuring they are accessible for processing and analysis.



Vector Database

The Vector Database stores dense embeddings for similarity search, enabling efficient retrieval of relevant patents and modules.



Knowledge Graph Database

The Knowledge Graph Database captures relationships between patents, modules, domains, and applications, providing a structured view of connections.

Embeddings & Representation



Embedding Generation

Sentence-transformers generate dense vectors at patent, module, domain, and application levels, enabling hierarchical similarity search.

Semantic Coherence

Hierarchical embeddings maintain semantic coherence, ensuring reliable retrieval of cross-domain candidates for analogy discovery.



Knowledge Graph Design

01

Node Types

Nodes represent Patents, Modules, Technology Themes, Industries, and Applications, capturing the core elements of the patent ecosystem.

02

Edge Types

Edges carry similarity scores, functional dependencies, and application viability weights, providing rich context for relationship analysis.

03

Traversal Algorithms

The schema supports weighted traversal algorithms that rank potential transfers and detect white-space innovation zones programmatically.

04

Graph Utility

The knowledge graph enables clear visibility of transferable technology, cluster-level insights, and white-space opportunities, supporting strategic decision-making.

LLM Prompting Strategy



Prompt Templates

Prompt templates cover summarization, decomposition, analogical reasoning, scoring, and explanation generation, guiding agent tasks.

Chain-of-Thought Prompts

Chain-of-thought prompts enforce stepwise logic, while few-shot examples maintain consistency in reasoning chains.

Versioning and Testing

Prompt versioning and A/B testing ensure continuous improvement without breaking traceability required by stakeholders.



Module Consistency

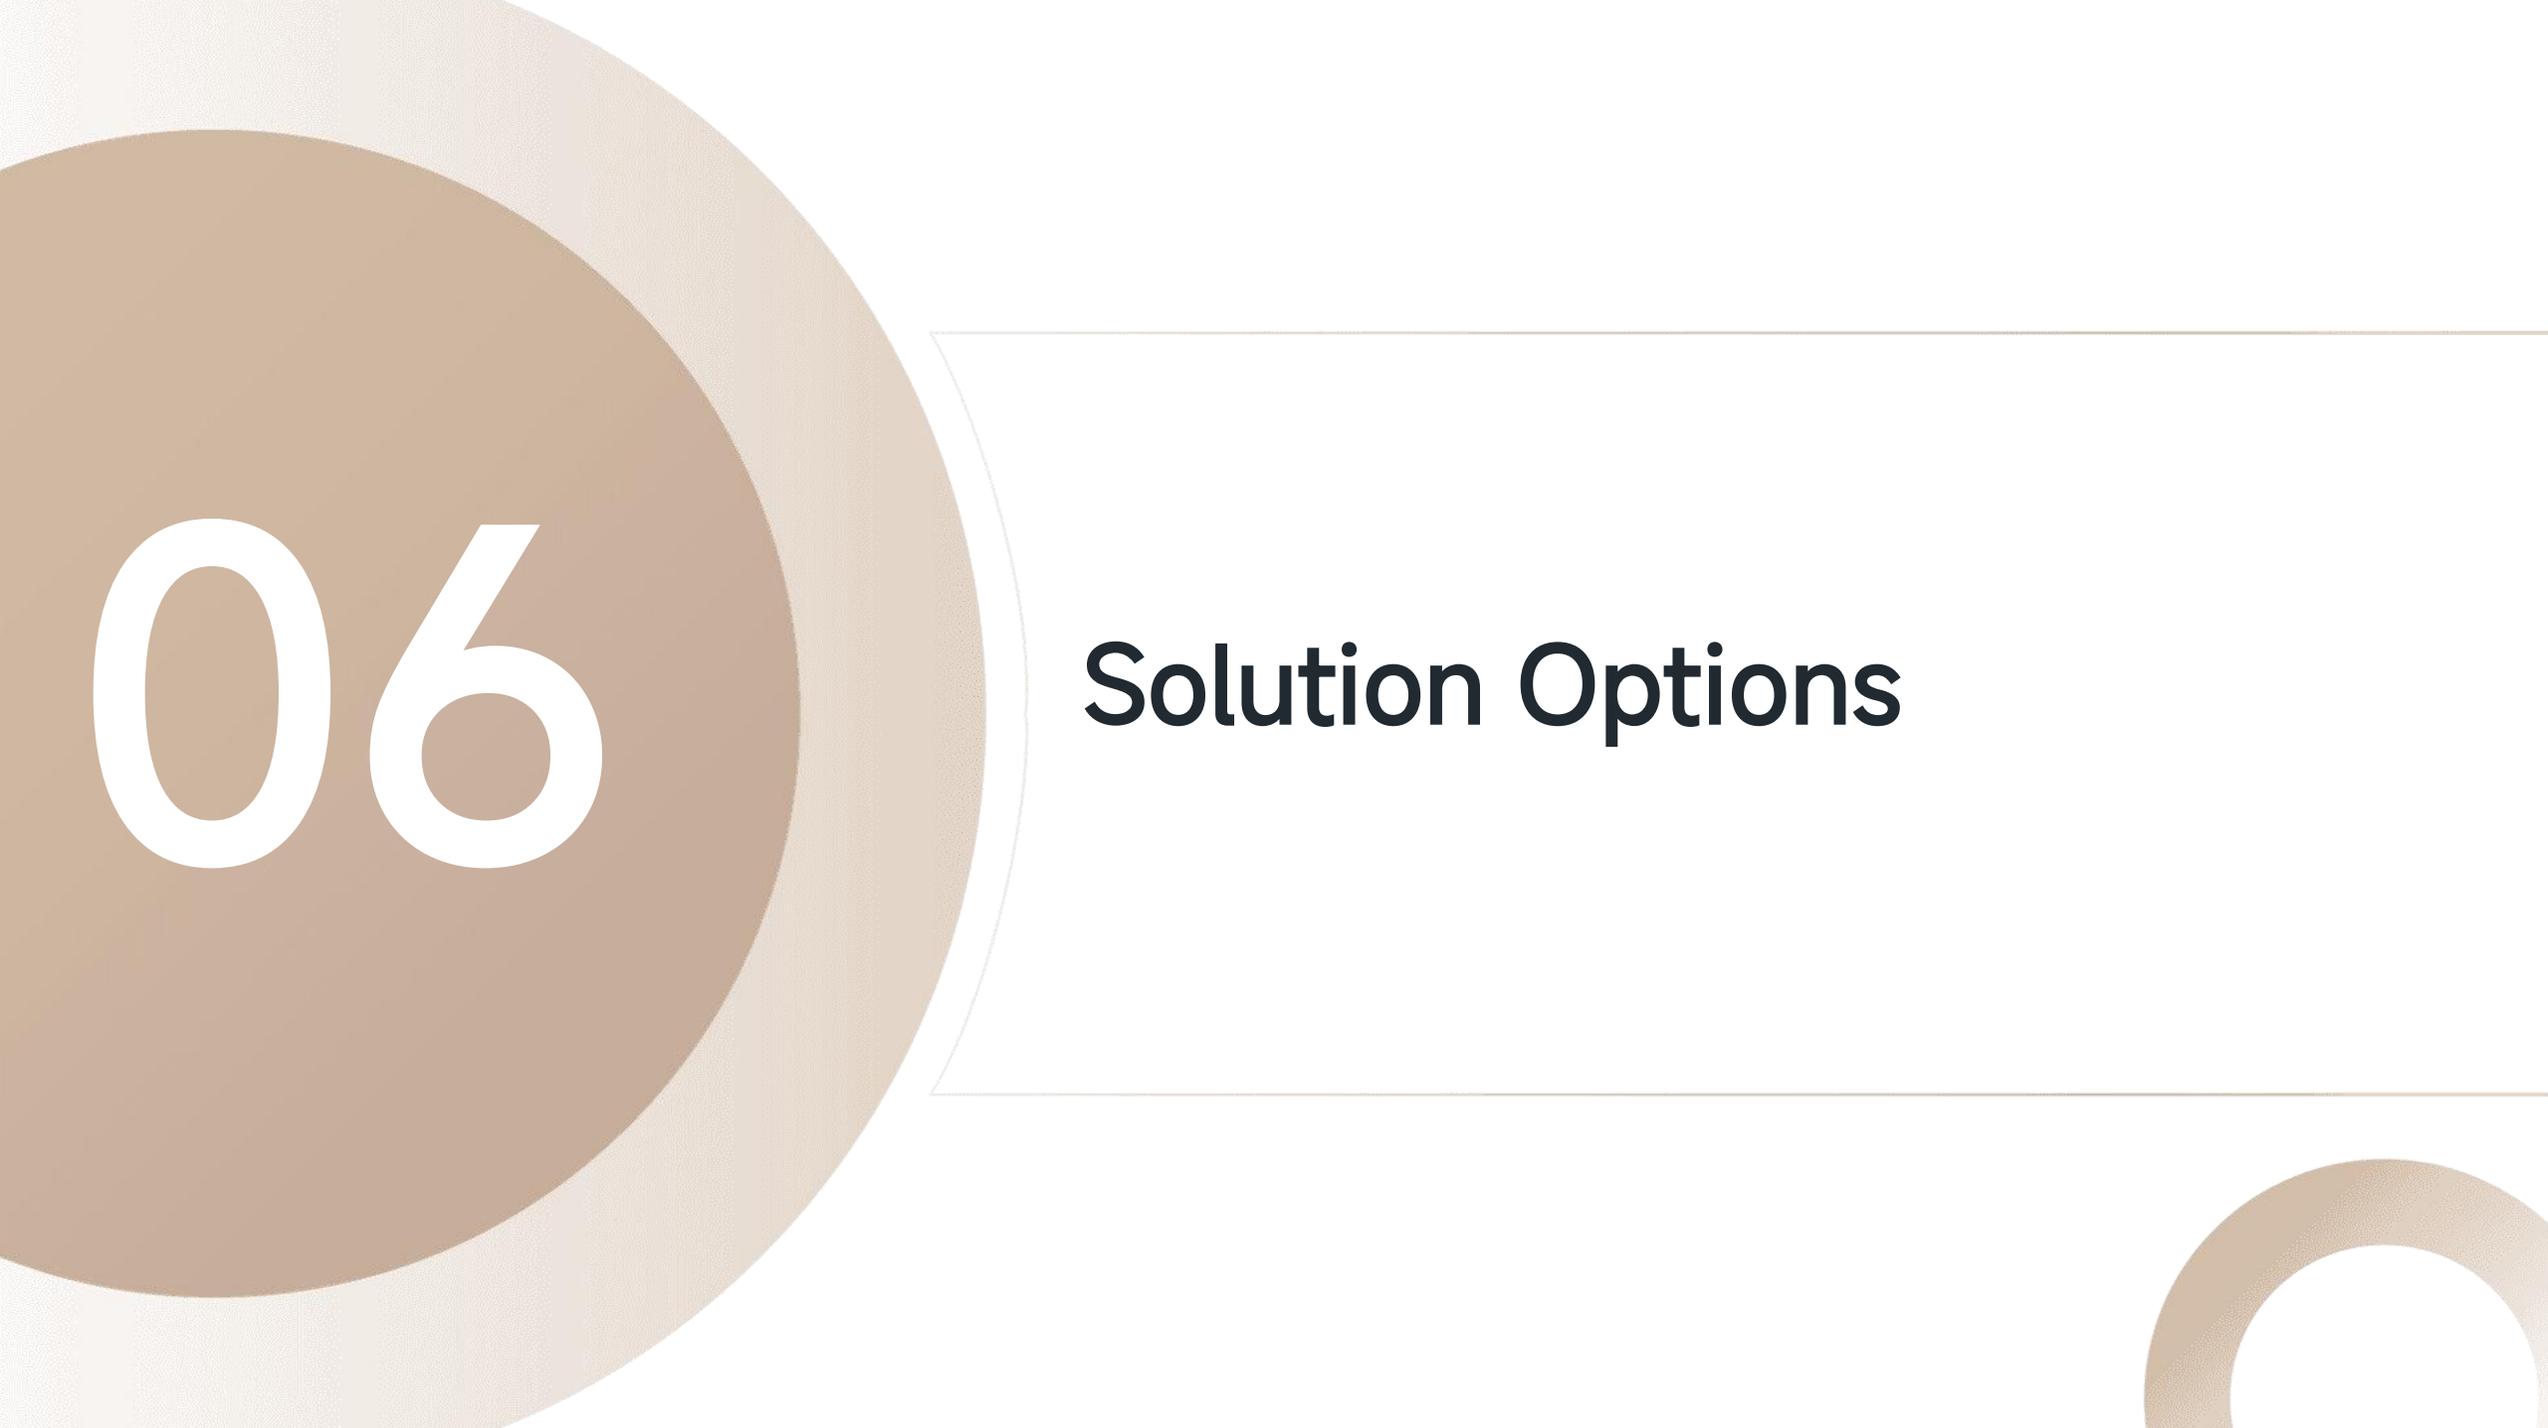
Internal validation checks module consistency via expert review to ensure accurate decomposition.

Analogy Coherence

Analogy coherence is verified through similarity metrics, ensuring that generated analogies are semantically meaningful.

Scoring Reliability

Opportunity scoring reliability is validated using hold-out test sets, providing confidence intervals for business recommendations.



06

Solution Options

Options Considered

Manual Taxonomies

Option B: Manual Taxonomies requiring expert maintenance, which is labor-intensive and not scalable.



Multi-Agent Framework

Option D: Multi-Agent Framework, which offers modular development, auditable reasoning, scalability, transparency, and domain-agnostic generalization.

Similarity Search

Option A: Similarity Search using embeddings alone, which is fast but lacks reasoning capabilities.

End-to-End Black Box LLM

Option C: End-to-End Black Box LLM lacking transparency, which is powerful but not interpretable.

Why Multi-Agent Wins



01

Modular Development

Multi-agent architecture enables modular development, allowing individual agents to be improved without retraining the entire pipeline.

02

Auditable Reasoning

It provides auditable reasoning, ensuring transparency and trustworthiness of the generated insights and recommendations.

07

Recommendations



Recommendation 1



Scale Dataset

Scale the pipeline to a larger dataset of 500–1,000 patents to validate generalizability and stress-test computational performance.

Recommendation 2



Automate Ingestion

Automate ingestion from USPTO APIs and internal patent archives to ensure continuous data refresh and immediate identification of revival opportunities.

Live Feeds

Live feeds enable monitoring of newly granted patents, maintaining competitive advantage and maximizing IP monetization windows.



Recommendation 3

Deploy Explorer

Deploy a lightweight knowledge-graph explorer that allows IP managers to interactively traverse technology clusters and filter by industry.

User-Friendly Interface

An intuitive interface accelerates stakeholder adoption and transforms technical graphs into strategic tools for rapid scenario planning.

Interactive Exploration

Interactive exploration enables stakeholders to simulate module substitutions and identify new opportunities without technical expertise.



Recommendation 4



Integrate Market Data

Integrate real market datasets such as Total Addressable Market estimates and competitor activity feeds to enrich scoring algorithms.

Align with Business

Contextualizing technical opportunities with commercial data ensures that high-ranking recommendations align with revenue potential and competitive landscape.

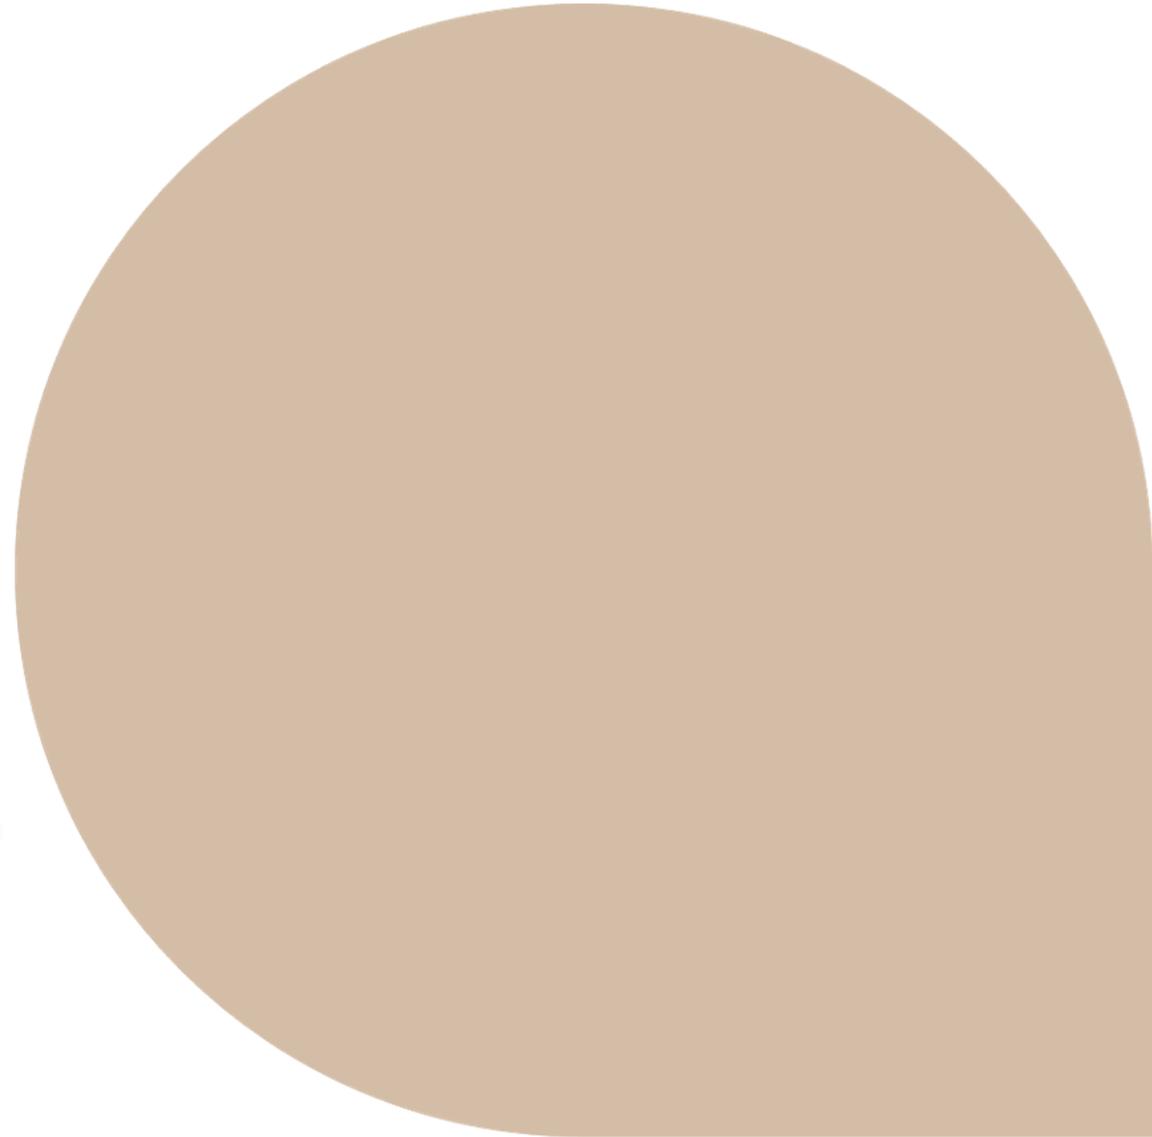


Recommendation 5



Pilot with Teams

Pilot the system with internal IP and R&D teams for usability testing, capturing feedback on dashboard clarity, ranking credibility, and workflow integration.



08

Reflections & Learnings

What Worked Well

Modular Design

Modular agent design facilitated parallel development and debugging, ensuring efficient progress and scalability.

Interpretability

Interpretability of reasoning chains built trust, while visual storytelling translated complex AI outputs into compelling business narratives.

Stakeholder Feedback

Strong stakeholder feedback loops kept the system aligned with real workflows, ensuring it met actual business needs.

Clear Communication

Clear communication of results through visualizations and narratives ensured stakeholders could understand and act on the insights provided.

Challenges Faced

01

Complex Patent Language

Complex patent language occasionally misled embedding models, requiring iterative prompt tuning to improve accuracy.

02

Analogy Consistency

Maintaining analogy consistency across diverse domains demanded iterative LLM calibration to ensure reliable cross-domain matching.

03

Balancing Creativity and Feasibility

Balancing creative novelty with practical feasibility forced refinement of scoring weights to avoid overly speculative recommendations.

Key Learnings

Clear Problem Framing

Clear problem framing drives meaningful AI solutions by ensuring the system addresses actual business challenges.

Explainability

Explainability matters more than raw model complexity, as stakeholders require understandable insights to make informed decisions.

Multi-Agent Superiority

Multi-agent pipelines outperform single-shot LLM reasoning by enforcing structure and enabling modular improvement.

Stakeholder Involvement

Early stakeholder involvement prevents technical overengineering, ensuring that AI deliverables address real business pain points.

If We Continued

Next steps would include portfolio-level clustering to detect technology themes and identify overarching trends.



Portfolio Clustering

Develop technology evolution timelines to forecast obsolescence and inform strategic R&D and investment decisions.



Technology Evolution

Create patent-to-market opportunity maps that overlay regulatory milestones, providing a comprehensive view of commercialization pathways.



Market Mapping

Closing Slide



Project Summary

AI-Driven Patent Revival unlocks hidden value of innovation by converting dormant IP into strategic business opportunities. The project demonstrates the potential of multi-agent AI to transparently reason across domains, prioritize investments, and accelerate monetization.



**THANK
YOU**