STEVENS
INSTITUTE OF TECHNOLOGY
1870

School of Business

BUSINESS TRANSFORMATION USING DATA ANALYTICS AND AI EXPO

RESEARCH POSTER BOOK

MAY 16, 2025

Poster Number	Poster Title
1	Precedents Thinking Agent
2	Teaching Agent
3	Design Thinking Agent
4	Simulation of Hospital Surgery Scheduling using Genetic Algorithms
5	SDG Classification
6	Gen AI Chatbot
7	Automating Financial Report Generation with GenAI
8	Time Series Analysis for Air Quality Forecasting
9	Senior Sage - AI Powered Voice Assistant for Senior Citizens
10	Leading The Charge
11	Patient Opportunity Model
12	HIV Treatment Policy and Reinforcement Learning
13	Machine Learning for Work Task Classification
14	Predicting Macro and Microeconomy at different levels
15	Housing Analysis: New York City
16	Predicting NBA Fan Churn Risk Using In-Game Metrics and Player Availability
17	AI Writing Assistant
18	Immigration Concierge
19	AI-Powered Stock Market Prediction & Sentiment Analysis Dashboard
20	"From One Path to Another: Your Career, Your Comeback!"
21	Silverfish Safe Passage System
22	PharmaSales SlackGPT
23	Personalized Indian Cooking for Modern Kitchens Using Gen Ai
24	DealFinder AI
25	MediBuzz - Never miss a dose
26	Fraud Detection in Health Insurance Claims
27	OPTIMIZING SUPERMARKET CO-BRANDED CREDIT CARDS
28	Jobify
29	Optimising NYC Public Wifi using Clustering Model
30	Predicting the US Inflation Rate
31	Automating Fashion Tagging with GPT-4o: A Scalable AI Solution

LEVERAGING PRECEDENTS FOR BREAKTHROUGHS

Sofia Savchuk & Arun Kashyap

Instructor: Dr. Alkiviadis Vazacopoulos

STEVENS INSTITUTE of TECHNOLOGY

THE INNOVATION UNIVERSITY®

Business Intelligence & Analytics

What is Precedent Thinking? - *who says innovation has to be scary and completely out of the blue?*

*Precedents Thinking is **a repeatable, powerful approach** to innovation that involves creatively combining **existing ideas** what Albert Einstein called **"combinatory play"** to generate breakthrough ideas. It's like assembling the best discoveries of others to tackle challenges, and the cool part is it often feels **less risky** because you're building on proven success.*

THE CHALLENGE: Rethinking "New"

Is True Innovation a Myth?

- Many groundbreaking advancements, as highlighted by Harvard Business Review, are not entirely new but are creative recombination of existing ideas.
- Ford's assembly line and Netflix's streaming model exemplify this "Precedents Thinking."
- Our Approach: The Precedents Thinking Framework

  FRAME: Clearly define your core challenge.

  SEARCH: Discover proven precedents across diverse domains.

  COMBINE: Synthesize insights into novel solutions.
- HBR shows novel framing is key

STRATEGIC VALUE: WHY PRECEDENTS MATTER

Engineer Your Competitive Advantage

- **De-Risk Innovation:** Build upon established successes.
- **Accelerate Breakthroughs**: Leverage existing knowledge for faster solutions.
- **Unlock Cross-Industry Wisdom:** Find novel applications from diverse fields.
- **Informed Strategic Planning:** Make data-driven decisions grounded in proven approaches.
- **Foster a Culture of "Combinatory Play":** Empower teams to connect ideas for impactful results.

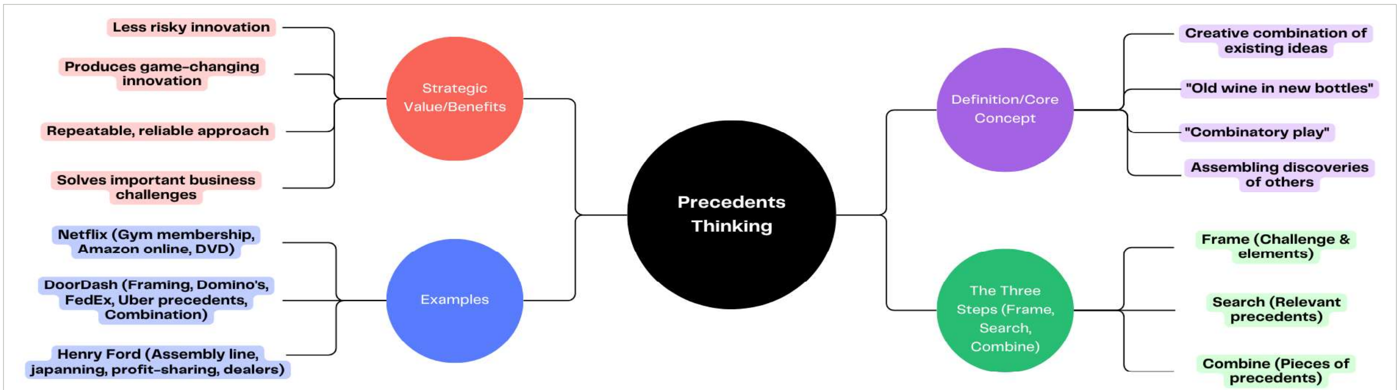

OUR SYSTEM: AI-POWERED PRECEDENT DISCOVERY & ANALYSIS

Accelerating Innovation with Intelligent Technology

Our platform operationalizes the Precedents Thinking Framework. It empowers you to:

- Conduct research-intensive analysis across vast, diverse datasets.
- **Identify and evaluate relevant precedents** from academic papers, news, industry reports, web sources, and your uploaded documents.
- Synthesize complex information into actionable insights.
- Visualize connections and patterns to spark "aha!" moments.
- Generate comprehensive reports to support strategic decision-making.

How It Works:

- **Intelligent Multi - AI Agents (CrewAI):** Autonomous agents perform specialized research and analysis tasks.
- **Advanced Language Models (LLMs):** Deep understanding and generation of nuanced information.(Gemini)
- **Custom Search & Data Tools:** Specialized tools (News, Scholar, Serper, Document Processing) ensure comprehensive data gathering.
- **Structured Output:** Delivers insights in clear, usable formats, including dynamic visualizations (networks, timelines, mind maps).

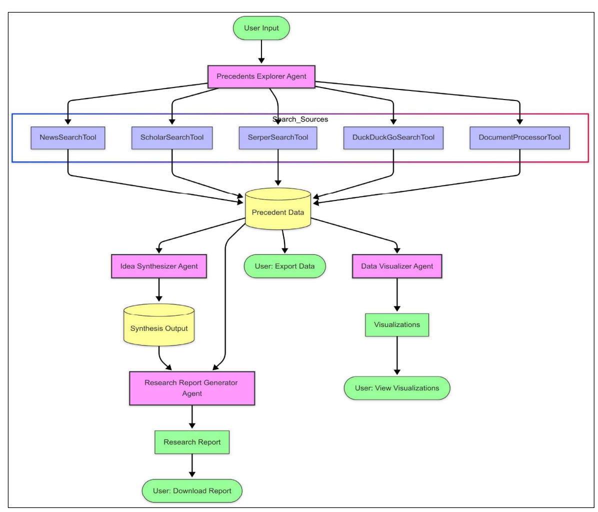

Our Precedents Thinking System employs a sophisticated workflow driven by intelligent AI agents. This diagram illustrates the journey from your initial challenge input to the generation of actionable insights, research reports, and compelling visualizations. Each step is designed to efficiently discover, analyze, and synthesize relevant precedents, supercharging your innovation process.

Generative AI Agents in Education: Automating Case Study Prep, Transcript Summaries, and Grading

Sofia Savchuk & Arun Kashyap
Instructor: Dr. Alkis Vazacopoulos

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
1870
Business Intelligence & Analytics

The Challenge: Repetitive Instructional Workflows

Instructors often face the burden of **preparing teaching materials, grading assignments, and summarizing lectures—**essential tasks that are **time-intensive, inconsistent, and hard to scale**. These repetitive workflows take valuable time away from **student engagement and strategic teaching**. Despite the rise of AI in education, few solutions are tailored to support educators directly in these core responsibilities.
We aim to address this gap with **specialized, modular AI agents** that **automate the heavy lifting**, while **keeping instructors in control of quality and outcomes**.

Our System: Modular AI Agents for Education

We built a suite of modular agents to automate three core instructional tasks: case prep, transcript analysis, and grading. Each system uses CrewAI to coordinate specialized agents and Google Gemini for high-quality content generation.

- **Architecture:** CrewAI pipelines manage writing, reviewing, and scoring tasks.
- **Format Support:** Inputs include PDF, DOCX, PPTX, TXT, and VTT; outputs are Word, PDF, and Markdown.
- **Use Cases:** Teaching note creation, transcript summarization, and rubric-based grading.

Impact: Time, Feedback, and Learning

Our agent systems streamline academic tasks that typically take hours, producing consistent, high-quality outputs in minutes.

- **Time Savings:** Reduced prep and grading from several hours/days to under 1 hour.
- **Improved Feedback:** Personalized, rubric-aligned feedback with clear strengths and improvement areas.
- **Scalable Learning:** Enables instructors to maintain quality while supporting larger classes and diverse materials.

Case Study Analysis Agent

Builds comprehensive teaching materials from raw case files using a six-agent pipeline.

- **Identifies** key themes, stakeholders, and ethical issues from the case
- **Generates** a structured 10-section teaching note, including synopsis, objectives, and teaching strategies
- **Produces** a detailed 2-hour lesson plan with timing, activities, and framing questions
- **Outputs** a visual board plan highlighting major themes and decision points
- **Streamlines** case preparation from multiple tools and hours of manual work into one unified, editable package

Assessment Agent

Grades assignments using a structured, rubric-aligned pipeline to ensure clarity, fairness, and speed.

- Assessor Agent **analyzes** submissions against reference answers and grading policies
- Feedback Agent **generates** detailed, personalized responses for each student
- Grading Agent **assigns** a final score and compiles a full report with strengths and improvement areas
- **Reduces instructor grading time** with >90% rubric alignment
- **Exports** clean results in CSV format for instructor review or LMS integration

Conversation Analysis Agent

Transforms raw transcripts into structured educational documents for use in flipped classrooms, reviews, or discussion recaps.

- Content Analyzer **extracts** speaker info and identifies structure and themes
- Quote Extractor **pulls** high-impact quotes and categorizes them by topic
- Content Writer **assembles** the final output, including a briefing doc, key ideas, FAQs, quizzes, and essay prompts
- **Automatically redacts** pharma mentions for neutrality
- Output is **formatted** for Word, PDF, and Markdown and includes 12 distinct educational sections

Agent Workflows in Action

Step-by-step diagrams showing how each agent extracts, processes, and generates educational content.

Case Study Analysis Agent

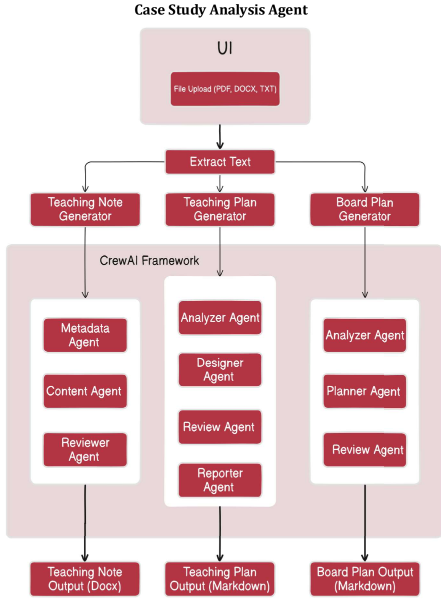

Assessment Agent

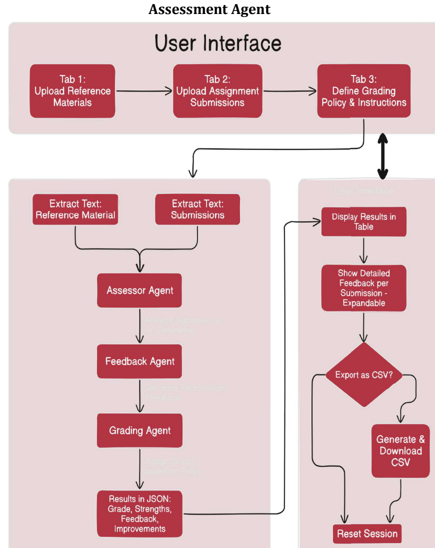

Conversation Analysis Agent

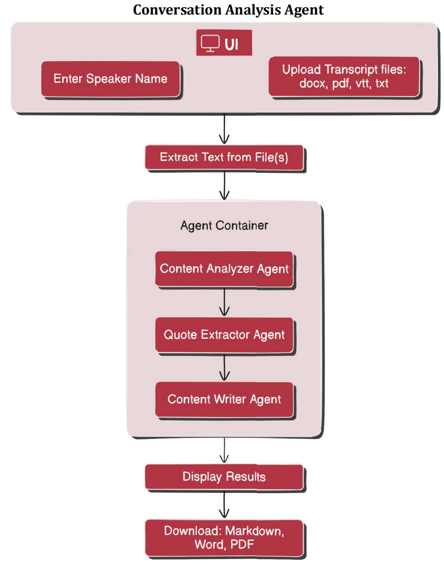

DESIGN THINKING AI AGENT CLUSTER (DTAAC)

Alexander Fromholtz & Arun Kashyap

Instructor: Dr. Alkiviadis Vazacopoulos

STEVENS INSTITUTE of TECHNOLOGY

THE INNOVATION UNIVERSITY®

1870

Business Intelligence & Analytics

INTRODUCTION

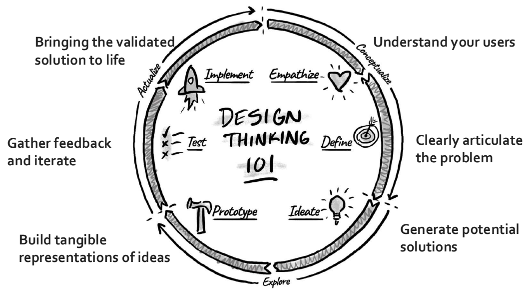

Source: voltagecontrol.com

WHAT IS DESIGN THINKING?

Design thinking is a human-centered innovation approach that tackles complex challenges through empathy, experimentation, and iteration. It combines deep user understanding with creative problem-solving techniques to develop solutions that genuinely meet human needs while remaining technically feasible and economically viable.

Key Defining Factors:

- A Human-Centered Approach to Problem Solving
- Focuses on deeply understanding user needs.
- Emphasizes iterative prototyping and testing.
- Non-linear process: stages can be revisited/refined

KEY FEATURES

- 🤖 **Intelligent Process Ecosystem:** Specialized Multi-Agent System collaborates under central coordination, automatically generating reports and tracking decisions throughout the design thinking journey.
- 🤝 **Human-AI Synergy**: Designed to augment human creativity and empathy, not replace it.
- ⚙️ **Multi-Model Flexibility:** Supports leading commercial LLMs with architecture designed for future on-premise deployment for sensitive innovation projects.
- 📄 **Dynamic Contextual Understanding:** Integrates user-uploaded PDFs and live web research into agent tasks.
- 🔄 **Built-in Iteration & Refinement:** Interactive chat and task revision capabilities for continuous improvement.
- ⚖️ **Comprehensive Evaluation Framework:** Assess outputs against defined criteria and expert baselines.

THE DTAAC ADVANTAGE: ADDRESSING DESIGN THINKING CHALLENGES

CHALLENGE	DTAAC SOLUTION
⌛ Time-Consuming Process	Accelerates research, synthesis, and reporting through AI-powered agents, significantly reducing cycle times.
Requires Specialized Expertise	Provides access to specialized AI agents (Empathy Researcher, Problem Definition Specialist, etc.) embodying best-practice methodologies for each stage.
Difficult Collaboration & Alignment	A central Process Manager agent ensures cohesive workflow, consistent context transfer, and clear documentation for team alignment.
Slow Iteration Cycles	Facilitates rapid prototyping ideation, testing plan generation, and incorporates feedback efficiently for continuous refinement.

MEET YOUR AI DESIGN TEAM WORKFLOW

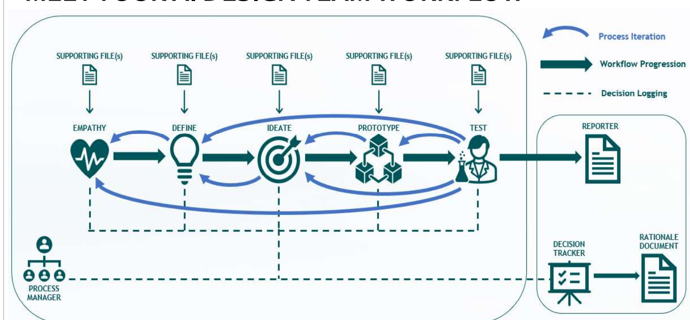

Phase-by-Phase Workflow

- **Define** your design challenge with AI assistance that formulates clear problem statements
- **Navigate** through each design thinking stage guided by specialized AI agents (Empathy Researcher, Problem Definition Specialist, etc.)
- **Enhance** understanding by uploading research documents and utilizing integrated search tools with citation tracking
- **Evaluate/Refine** solutions through feedback loops and rapid iteration based on new insights

BENEFITS & IMPACT

- 🚀 **Accelerate Innovation:** Rapidly progress from problem understanding to validated solutions with AI-powered design thinking.
- 💡 **Enhance Insights:** Generate deeper research analysis and more diverse ideation approaches.
- 📈 **Improve Solutions:** Ground innovations in robust user understanding through iterative testing.
- 🎓 **Empower Users:** Help product teams, design agencies, startups, and educators achieve better results faster with guided design thinking assistance.

EVALUATION

Uses "LLM as a Judge" to provide objective, automated analysis of agent outputs against defined criteria and baseline documents, offering an additional layer of quality assessment.

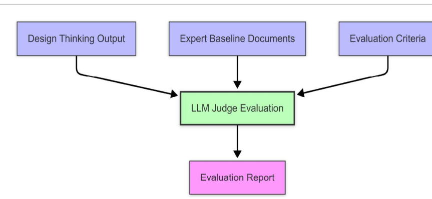

CONCLUSION & FUTURE DIRECTIONS

- DTAAC represents a significant step towards harmonizing the analytical power of AI with the empathetic core of human-centered design. It offers a structured, flexible, and powerful platform for tackling complex innovation challenges.
- **Future work includes** Expanding LLM support, enhancing the evaluation framework with more nuanced metrics, integrating local LLM capabilities for enterprise use, and exploring applications in specialized regulated domains (e.g., medical device development).

Enhancing Surgical Efficiency with Genetic Algorithm-Optimized Team Assignments

Sanjeet Mishra
Instructor: Prof Edward Stohr

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction & Problem Statement

- Elective surgeries demand efficient team coordination to minimize procedure times and operational costs.
- Inefficient team assignments increase wait times, costs, and resource strain.
- Goal: Use Genetic Algorithms (GA) to automate optimal surgical team selection based on historical performance data.
- Team Composition: Each team includes a Surgeon, Anesthesiologist, Scrub Nurse, Circulator Nurse, First Assistant, and an optional Second Assistant.
- Fitness Evaluation: Score teams using fitness = 1 / (penalties + 100 × avg. procedure time), penalizing staff reuse (>3 for non-surgeons/anesthesiologists)

Genetic Algorithm Workflow

- Initialization: Generate 100 random team configurations (chromosomes)
- Fitness Evaluation: Score teams using fitness as function of penalties , avg. procedure time, staff reuse
- Apply tournament selection to retain high-fitness teams.
- Perform crossover (80% rate) and mutation (10% rate) to explore new team configurations.
- Evolve teams over 20 generations to converge on assignments.

Chromosome

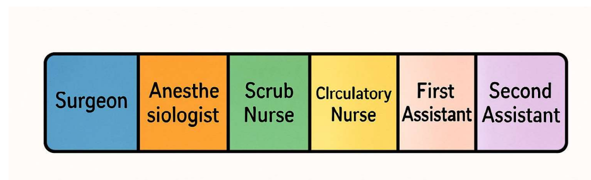

Algorithm Performance

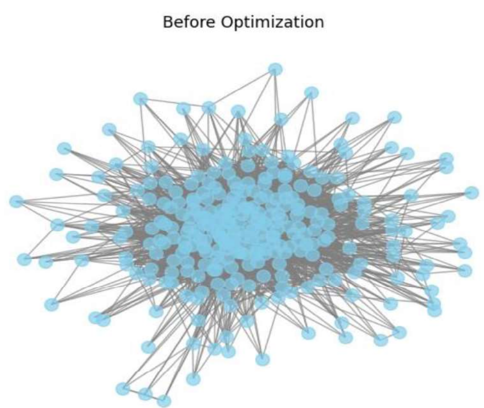

289 Nodes, 5147 Edges

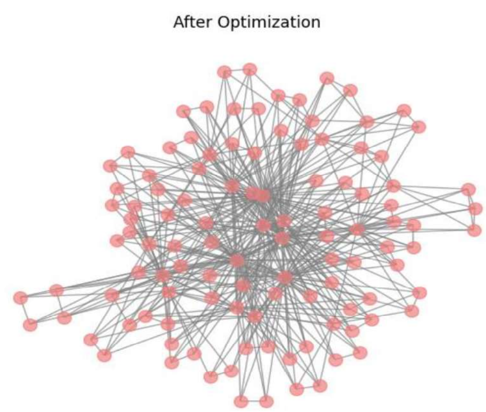

111 Nodes, 485 Edges

Optimal team assignment brings savings of ~14.4% to average surgery time

Nodes are the surgical team members and edges are their historical associations

Outcomes

- Generated optimal teams for each surgeon, reducing average procedure times
- Nodes reduced from 289 to 111 -> ~61% reduction
- Edges reduced drastically from 5181 to 485 → ~90% reduction
- In case a particular team member is not present for a particular day, the model can be re-run to get the next best team

From Disclosure PDFs to SDGs:
Novel way of SDG classification and AI-driven sustainability insights

Author Sanjeet Mishra
Instructor: Prof Emily (Rong) Liu

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Context

- UN SDGs: The 17 United Nations Sustainable Development Goals, adopted in 2015, address global challenges like poverty, inequality, and climate change under the 2030 Agenda.
- Objective1 – Identify SDG Alignment: Which UN Sustainable Development Goals (SDGs) are prioritized in various companies and how do these priorities vary across companies, esp. the IT Sector?
- Objective 2 - Sustainability Profiles: How can organizations be clustered based on their SDG contributions, revealing distinct sustainability strategies and focus areas?
- Challenge: Processing long documents (PDFs) with high accuracy and efficiency
- Solution: We finetuned pretrained RoBERTa model with Low Rank Adapation (LoRA) and processed reports

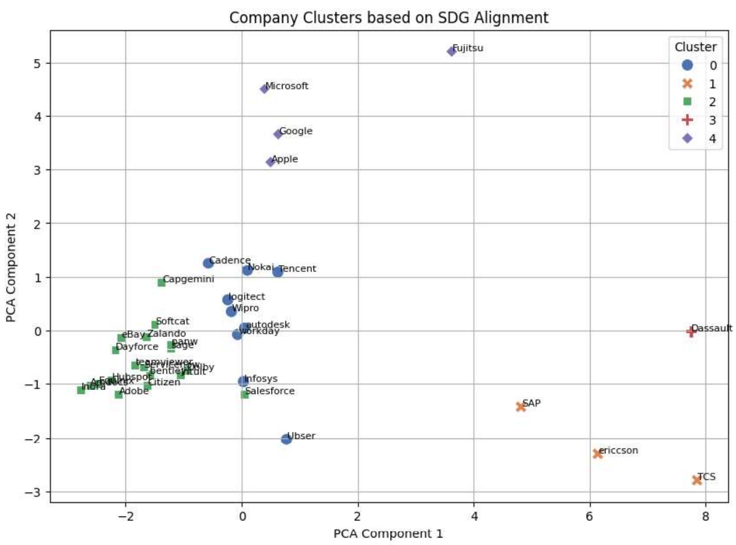

Dataset: ESG & Sustainability Reports of Time 500 IT&ITES Companies

Methodology

- Data Preparation: Trained on the OSDG dataset for single-label SDG classification, cleaning and preprocessing text data
- Training Pipeline: Fine-tuned RoBERTa with LoRA ($r=32$, $\alpha=64$, dropout=0.0758) on Apple Silicon GPU, using 12 epochs and AdamW optimizer.
- Prediction Framework: Classified chunks into SDGs (1–15) with a 10% probability threshold via softmax outputs.
- Clustering Approach: Employed K-Means with 2D PCA to identify five distinct clusters based on SDG profiles.

Insights:

Cluster 0:Balanced Contributors with Emphasis on Industry, Innovation & Infrastructure - Companies with a broad but moderate focus across infrastructure, employment, and innovation goals.

Cluster 1:Heavyweight Industry & Energy Innovators - Highly specialized firms with major emphasis on industrial innovation, equality, and energy use.

Cluster 2: Companies with minimal SDG alignment, possibly indicating comparatively limited focus wrt peers in dataset.

Cluster 3: Niche, high-impact company dominating environmental and infrastructure SDGs

Cluster 4: Prominent tech companies exhibiting strong alignment with climate, energy, and water goals.

Conclusion

- The RoBERTa model with LoRA effectively identifies SDG-related content in sustainability reports with high accuracy and efficiency.
- Companies cluster into distinct sustainability archetypes — from climate-conscious tech giants to industrial specialists and generalists — revealing varied levels of SDG integration across the corporate landscape.

Instructor: Christopher Asakiewicz

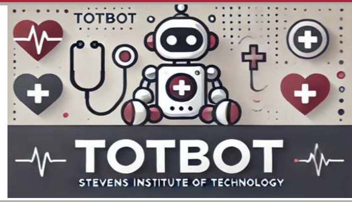

By: Brenna Bockoven, Paula Schindele, Rodrigo Polanco , Abhijit Ghosh

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

TotBot is a cognitive chatbot designed to alleviate the burden on overwhelmed healthcare systems by addressing non-urgent pediatric inquiries. Parents often struggle to access quick, reliable health advice for their children, leading to unnecessary doctor visits and anxiety. Totbot provides 24/7 evidence-based guidance on minor symptoms, nutrition, and developmental milestones using trusted sources like the American Academy of Pediatrics (AAP), CDC, and Mayo Clinic. This scalable solution improves accessibility for parents, reduces misinformation, and empowers caregivers to make informed decisions.

Experiment

Business Focus: Identifies common non-urgent parental health concerns and leverages automation for effective solutions.

Data-Driven Insights: Utilizes datasets from trusted sources (IBM Watsonx Assistant, AAP, CDC, Parenting.com) on child health, milestones, and medication dosages.

NLP-Enabled Approach: Employs Natural Language Processing to understand user intents and provide age- and symptom-specific guidance, with modular scalability for future API integrations.

Ethical Standards: Adheres to strict data privacy, medical compliance, and reliance on verified healthcare information.

Results

Insightful Engagement: The system excels at meeting the nuanced needs of its audience—ranging from dosage guidance and non-urgent medical advice to mom hacks and self-care practices—through continuous refinements driven by detailed user insights and real-world usage scenarios, ensuring it remains both empathetic and authoritative.

Qualitative Impact: Enhanced parental confidence in managing common child health concerns and improved access to credible healthcare information.

System Performance: Harnessing IBM's watsonx, the chatbot delivers seamless, context-sensitive responses while integrating trusted health resources reliably.

Healthcare Benefits: Alleviates strain on healthcare providers, optimizes resource allocation, and supports efficient management of non-urgent pediatric inquiries.

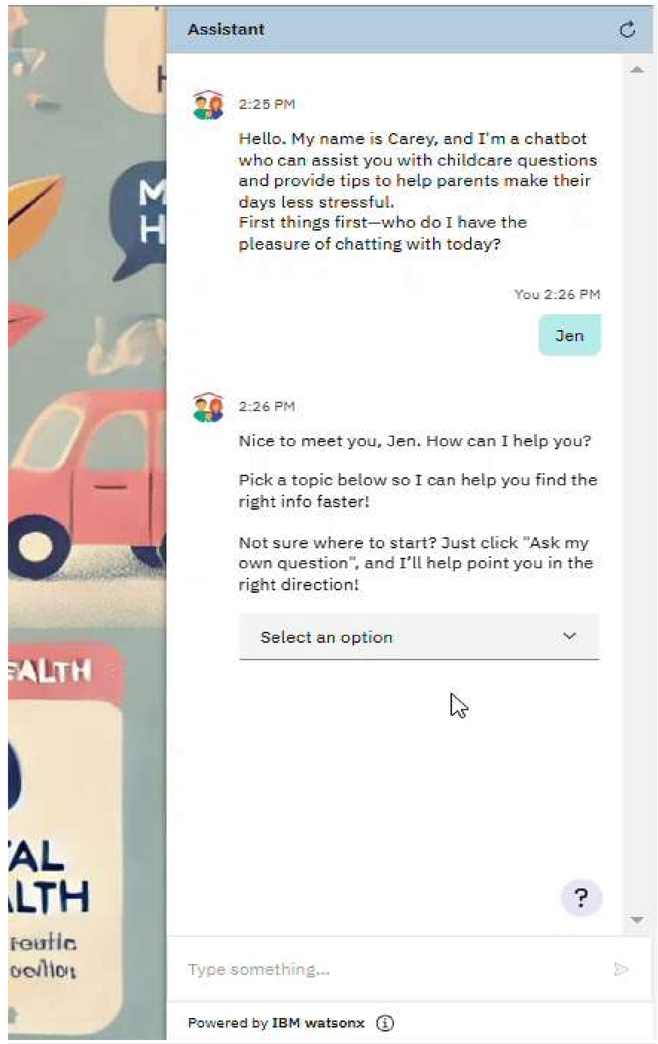

Architecture

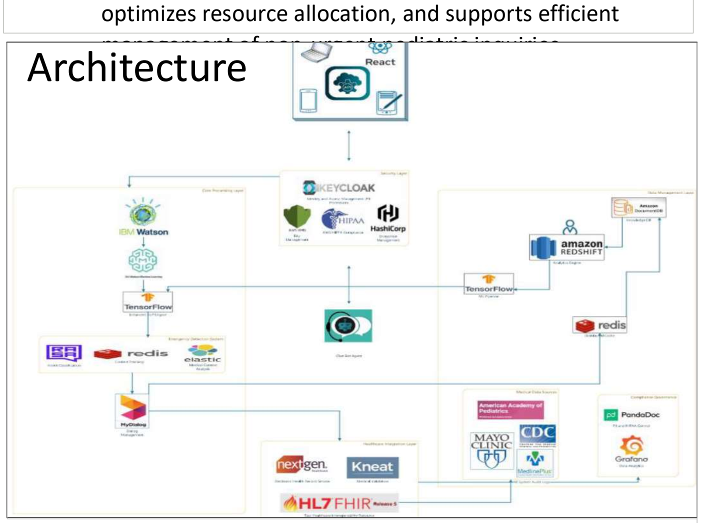

Conclusion

Healthcare Efficiency: Reduces strain on providers by addressing non-urgent pediatric inquiries, enabling better focus on critical cases.

Parent Empowerment: Delivers actionable, reliable advice to parents, boosting confidence in managing child health.

Economic Impact: By guiding parents toward efficient self-care and reducing unnecessary consultations, the solution promotes a more sustainable, resource-conscious healthcare experience.

Scalability and Equity: Offers a scalable solution to enhance pediatric healthcare access and address disparities in underserved communities.

Fin AI- Financial Report Co-Pilot

Sneha Dharne, Gunik Luthra, Parth Dharod

Instructor: Seyed Mohammad Nikoeu

STEVENS INSTITUTE of TECHNOLOGY

THE INNOVATION UNIVERSITY®

Business Intelligence & Analytics

Introduction

Financial reporting is a mission-critical yet time-intensive task for investment firms, particularly when preparing quarterly (10-Q) and annual (10-K) reports. These reports, often generated four times per year by hundreds of firms, require the manual consolidation of multi-year financial data, model building in Excel, and detailed narrative writing—demanding up to **80 hours per report**. This results in thousands of hours lost annually across the industry and introduces risks of inconsistencies, errors, and delays.

Objective

FinAI Copilot was developed to address this inefficiency by automating the report generation process. The tool leverages generative AI, model creation, and accuracy guardrails to convert structured financial data into polished, compliant reports within minutes. Its core objective is to enhance analyst productivity, improve report credibility, and enable the personalization of financial insights—all while preserving the transparency and accuracy essential for regulatory compliance.

Application

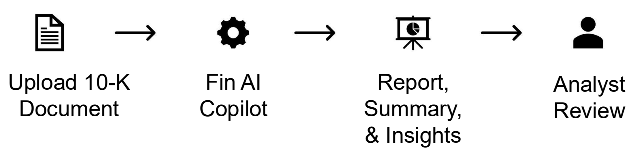

To evaluate FinAI Copilot in real-world conditions, we applied it to a variety of 10-K filings and reporting scenarios. The following summarizes how the system was implemented, tested, and adapted to mimic actual analyst workflows:

- Build on a modular architecture that can scale with your workflow and integrate multi-agent systems as needed.
- Generate accurate reports by combining AI-driven insights with logic-based validation to minimize inconsistencies.

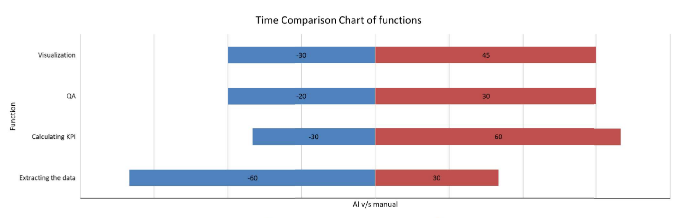

- Identify and highlight trends in cost efficiency, R&D spending, and profitability across multiple years—just like a seasoned analyst.
- Track business performance over time instead of relying on isolated metrics, enabling longitudinal analysis at scale.
- Produce executive-style summaries tailored to stakeholders, free from the legalese of standard SEC filings.
- Apply the tool to a wide range of company structures and formats, from tech giants to mid-market firms.
- Deliver investor-ready briefs in under 30 minutes—customized to your firm's tone, focus, and strategic priorities.

Results

Metric	Manual Process	FinAI Copilot
Time Taken	~40 hours	~30 mins
Errors & Typos	Frequent	Guardrails enforced
Accuracy Checks	Manual QA	Auto + human review
Compliance	Analyst effort	Prebuilt constraints
Transparency	Excel hell	Traceable lineage
Personalization	Manual rewrite	Tuned output

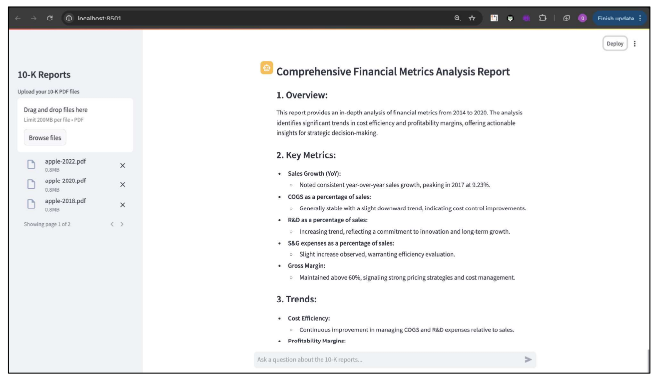

Credibility Analysis

Accuracy:

- Schema-based generation ensures that summaries are grounded in structured data — no hallucinations, no inconsistencies

Compliance:

- Every number is traceable back to its source: page number, snippet, and filing year. "Net income → Page 44, FY2023: 'Net income was $96,995"

Data Lineage:

- Outputs follow reporting logic used in analyst briefs: intros, quant KPIs, qualitative commentary — not legal structure

Reduced Manual Efforts:

- FinAI eliminates manual extraction and Excel drift by validating every field programmatically — saving hours and stress

Key Takeaways

Scan this QR for demo

- 90% reduction in reporting time
- Smart briefs, not SEC copies
- Every number is traceable
- Modular architecture → scalable
- Next up: multi-agent report generation

Future Scope

- Dashboard in minutes
- Reimagining reports with unstructured data
- Plug & Play for Enterprise- Integrated Co-Pilot
- Auditing trails with built-in compliance
- Multi-Agent Architecture

From Analysts to founders, we're making AI first reporting a reality

Time Series Analysis of PM₂.₅ Air Quality In portland

Raja Sree Vytla
Instructor: Mahmoud Daneshmand

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

- This poster presents a study of hourly $PM_{2.5}$ data from four Portland-area sensors, combining exploratory analysis with ARIMA and SARIMAX models to uncover daily pollution cycles and deliver accurate next-hour air quality forecasts.

Objective

- Analyze hourly $PM_{2.5}$ time series to identify trend, seasonality, and stationarity
- Develop walk-forward ARIMA(2,0,10) models for next-hour forecasting
- Enhance forecasts with SARIMAX by adding daily seasonal and exogenous (neighbor-sensor) terms
- Evaluate and compare model accuracy using Mean Absolute Percentage Error (MAPE) across all sensors

Dataset

- Source : OpenAQ API CSV of hourly $PM_{2.5}$ readings from 4 Portland sensors.
- Fields : Location(sensor Id), Date_Converted,Value($PM_{2.5}$ concentration, μg/m³), latitude ,longitude
- Preprocessing:
- Parsed timestamps to datetime, , enforced hourly frequency
- Dropped invalid/missing values

Exploratory Analysis & Diagnostics

1. Raw Time Series

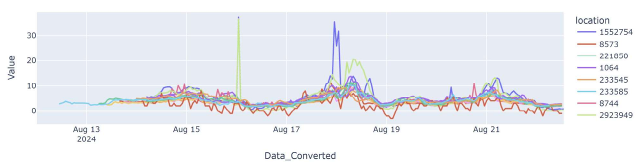

From this time-series plot, there's no obvious long-term trend, but the readings from each location indicate strong correlation across all sensors.

2. Seasonal Decomposition

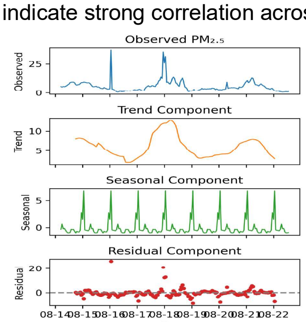

- Decomposition shows persistent upward or downward drift but a pronounced recurring spike every 24 hours confirming strong daily cycles that must be modeled.
- The residual component fluctuates randomly around zero, confirming that the trend and daily seasonal cycle capture the dominant structure in the data.

3. Stationarity Testing (ADF)

- Augmented Dickey–Fuller test on each sensor series
- Null (H_0): non-stationary (unit root)
- Results: $p < 0.05$ for all → stationary (d=0)

4.ACF and PCF Diagnostics

- **ACF → q = 10** (lags where autocorrelation remains outside confidence bands)
- **PACF → p = 2** (two significant partial-autocorrelation lags)

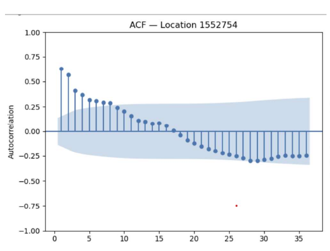

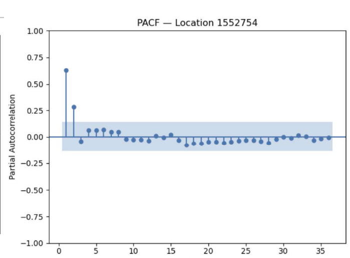

ACF for Location 1552754 → slow decay, seasonal spike at lag 24
PACF for Location 1552754 → significant cutoff after p = 2

Results

1.Forecast Accuracy

- **Model:** Walk-forward ARIMA(2, 0, 10)
- **Train/Test Split:** 80 % train (160 h) / 20 % test (41 h)

2. Sensor-wise MAPE

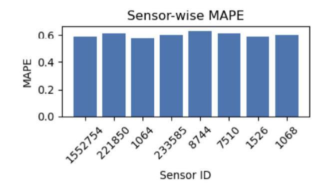

Sensor ID	MAPE
1552754	0.589
221850	0.614
1064	0.578
233585	0.601
8744	0.630
7510	0.615
1526	0.590
1068	0.602
Average	**0.600**

3. Results for ARIMA(2.0.10) Walk-Forward Forecast

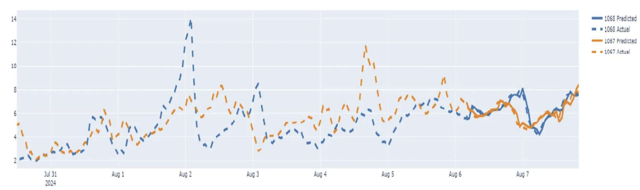

- **Sensors:** 1068 & 1067
- The ARIMA models achieve an average MAPE of **0.589**, indicating a reasonably good fit.
- They capture the overall daily cycle but tend to underestimate sharp morning/evening spikes.
- Sharp peaks on August 2 and August 5 are smoothed in the forecast—predictions peak ~20 % below the observed maximums.
- During the steepest rises and falls (e.g., early on August 3), predictions lag actual values by 1–2 hours.
- **Error (MAPE = 0.042):** Average absolute percentage error of ~4.2 % confirms a strong baseline performance for short-term forecasts.

4. Forecast vs. Actual(SARIMAX)

- Improved Peak Prediction especially around the August 2 and August 5 peaks, SARIMAX forecasts reach closer to the observed maxima—reducing underestimation seen in the ARIMA output.
- Compared to the ARIMA baseline (4.2 %), SARIMAX cuts the average absolute percentage error by over 25 %, demonstrating the benefit of modeling both seasonal and exogenous effects.

Conclusion

- Compared to the ARIMA baseline (4.2 %), SARIMAX cuts the average absolute percentage error by over 25 %, demonstrating the benefit of modeling both seasonal and exogenous effects.
- Incorporating daily seasonality (m=24) and exogenous inputs from neighboring sensors slightly reduced forecast error for correlated sites, demonstrating value of seasonal and spatial information.
- While ARIMA provided a solid univariate baseline, SARIMAX's ability to model both seasonal cycles and inter-sensor influences makes it more robust for multivariate air-quality forecasting.
- For real-world deployment, SARIMAX models leveraging both temporal (daily) and spatial (nearby stations) structure can yield more reliable $PM_{2.5}$ forecasts, supporting timely air-quality advisories.

Senior Sage: AI Voice Assistant for Senior Citizens

Aadyant Agrawal, Harsh Raghuwanshi, Senin Farheen, Prasoon Parashar, Sparsh Sarode
Instructor: Dr. Zhongyuan (Annie) Yu

STEVENS INSTITUTE of TECHNOLOGY
1870
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

Senior Sage is a personalized AI voice assistant tailored for senior users in healthcare management. It features:

Holistic Elderly Healthcare:

- Designed to enhance elderly care by merging AI-driven voice assistance with empathetic, natural interactions.

Natural Conversational Interface:

- Simulates a doctor-patient dialogue with short, concise responses, clarification handling, and end-of-conversation detection to create a comfortable interaction.

Advanced AI Integration:

- Leverages Whisper for real-time, high-fidelity speech-to-text conversion. Uses OpenAI's API for context-aware natural language processing.
- Incorporates a custom Retrieval-Augmented Generation (RAG) pipeline to dynamically fetch and contextualize latest medical data.

Personalization & Continuous Learning:

- Remembers user preferences and conversation context for tailored health guidance.
- Collects nuanced data on emotions, mood, and general well-being to support research on chronic conditions like hypertension and heart failure.

Proactive Reminders & Rewards:

- Configures critical health reminders and rewards through voice for medications, appointments, and other key health behaviors.

Results

Enhanced Health Interaction Capabilities:

- Successfully conducts efficient health questionnaires and records nuanced user responses, fostering natural and comfortable doctor-patient style conversations.

Advanced AI & Personalization Achievements:

- The integration of Whisper, OpenAI's API, and the RAG pipeline enables real-time, context-aware responses with improved accuracy by referencing current medical data.
- Personalized health reminders and context retention have elevated user engagement and improved overall guidance quality.

Robust, Scalable Backend Performance:

- The backend architecture, with container orchestration via Kubernetes, CI/CD automation, and role-based access control ensures rapid, secure deployment and efficient data management.
- Comprehensive error logging and agile iterative testing have reduced cycle times while enhancing system stability and scalability.

Measurable Impact on Healthcare Delivery:

- Improved data collection on patient emotions, mood, and general health provides actionable insights for research on chronic conditions.
- The solution supports proactive health monitoring and contributes to long-term research initiatives, aligning with broader healthcare goals.

Conclusion

- **Transformative Healthcare Solution:** Senior Sage integrates state-of-the-art AI and robust backend infrastructure to deliver personalized, real-time health guidance for the elderly.
- **Holistic & Engaging:** By combining natural, conversational interactions with proactive health reminders and a rewards system. It not only enhances patient engagement but also creates a comfortable, doctor-patient-like dialogue
- **Impact & Future Potential:** Iterative agile development has enabled measurable improvements in system performance and user experience, positioning Senior Sage as a transformative tool in elderly healthcare management with significant potential for future expansion and research integration.

Methodology & Development

- Our agile process, executed over multiple sprints, enabled rapid innovation and continuous improvement. Key technical components include:

Backend Infrastructure & Architecture:

- Containerization & Orchestration: We deploy microservices using Kubernetes, tested locally with Minikube, ensuring scalability and resilience.
- CI/CD Pipeline: GitHub Actions automates testing, integration, and deployment, reducing downtime and accelerating feature rollout.
- Secure Data Management: Utilizing MongoDB Atlas, we enforce role-based access control and HIPAA-compliant data handling, critical for sensitive healthcare information.

Advanced AI Integration:

- Speech Processing: Whisper handles real-time, high-fidelity audio transcription, forming the first step in our data pipeline.
- Natural Language Processing: OpenAI's API generates contextually rich responses, enabling natural dialogue.
- RAG Pipeline: Our custom RAG module dynamically retrieves and contextualizes medical data, ensuring that responses are both timely and clinically relevant.
- Iterative Enhancements: Agile sprints allowed us to fine-tune LLM features such as conversation context retention, personalized health reminders, and efficient error logging, all informed by continuous feedback and rigorous testing.

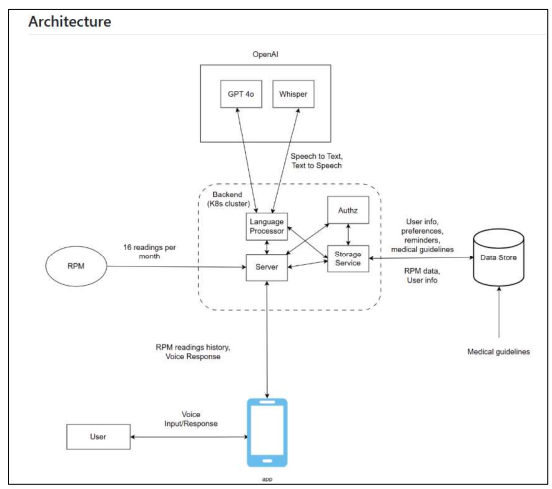

- We leverage OpenAI APIs, specifically GPT-4o mini and Whisper, for AI processing. We have opted for microservice based architecture with scalability in mind as each service (as shown in the diagram) can be horizontally scaled to meet the growing user base.
- **HIPAA Compliance:** Storage Service enforces strict authorization rules controlling access to resources and data. It also logs data access requests, including whether they were granted, for auditing purposes. Data stored in the Database is completely encrypted and will be decrypted by Storage Service on the fly while serving the requests
- **Language Processor Service:** This is the brain of the Voice Assistant. Implemented in Python and leverages OpenAI APIs with prompt engineering to achieve the desired functionalities
- **Server:** This is the server that is exposed to outside world for phones to connect to. Implemented in Golang
- **Authz Service:** This is the Authentication and Authorization service that is responsible for login, signup, authorization services. Implemented in Golang.
- **Storage Service:** Exposes internal APIs for other services to access the Data. Leveraging the Authz service, enforces strict authorization policies for data access control. Logs data access requests for auditing purposes.

Leading the Charge: Predicting the Demand for EV

Tarun Varma Mudunuri, Ravikiran Sriram
Instructor: Prof. Edward Stohr

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Motivation:

The electric vehicle (EV) adoption is rapidly rising, with the U.S. EV market projected to reach **$104.7 billion** in revenue by 2025, growing at a **10.54% CAGR** through 2029

In Washington State, a national leader in EV adoption:

EV registrations grew by **34% YoY** in 2024, reaching **223,995** vehicles

Counties like Snohomish **(+53.9%)** and King **(+44.2%)** show explosive growth

EVs now make up **20%** of new vehicle sales, more than double the U.S. average **(9.5%)**

Methodology:

- **Electric Vehicle Registrations:** Washington State API (Electric Vehicle Title and Registration Activity Data)
- **Charging Infrastructure:** National Renewable Energy Laboratory (NREL) API (Alt Fuel Stations Data)
- Aggregated monthly EV counts by county and created cumulative time series.
- Validated model forecasts with train/test splits and visual inspections.

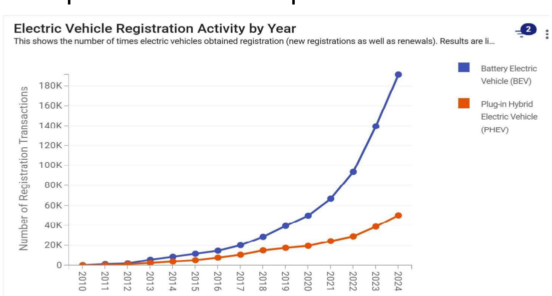

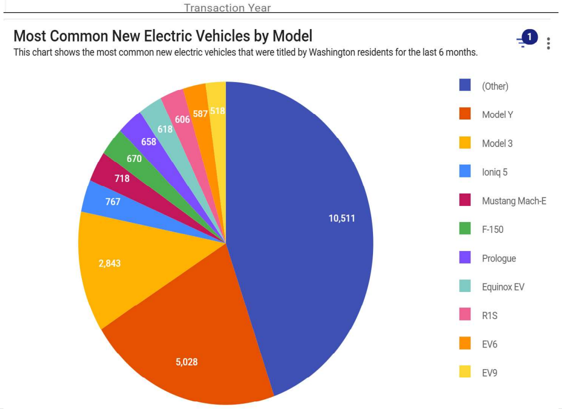

Results:

EV adoption shows exponential growth in many counties.

High-gap examples:

- *Clark County: 180 EVs per charger*
- *Whatcom County: 90 EVs per charger*

Forecast models predict strong continued demand, especially in urban and commuter-heavy areas

Heatmap of EV Registrations Across Counties

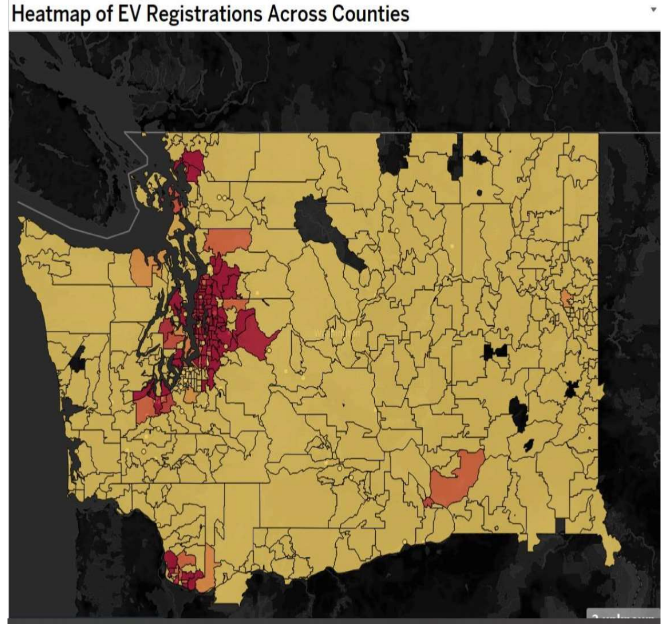

Conclusion:

✅ **Explosive Growth:** EV registrations surged by **124%** in Clark, **76%** in Snohomish, and 50% in Whatcom over two years - Yet infrastructure lags with up to **180 EVs per charger**, these counties face a critical shortage.

✅ **Investment Gap:** While Washington has invested $8M, an **additional $12–15M is needed** to keep pace

✅ **$100M+ Market Opportunity**: Charger deployment—backed by federal incentives, state grants, and user demand—can generate **$100M+ in revenue by 2029** via usage fees, partnerships, and energy resale

Lung Cancer Patient Opportunity Model

Hamid Dastgir, Harshita Hiremath, Ravikiran Sriram

Instructor: Sanjiv Koshal

STEVENS INSTITUTE of TECHNOLOGY

THE INNOVATION UNIVERSITY®

Business Intelligence & Analytics

INTRODUCTION

In lung cancer (NSCLC/mNSCLC), patients often transition between therapies due to progression, intolerance, or new targeted treatments

Steps: Build a Patient Opportunity Model to identify

1. Patients likely to remain on their current therapy (retention score)

2. Patients likely to switch to a new therapy (switching score)

Key Insight: A higher score indicates higher likelihood of switching; interventions or patient outreach can be made earlier to improve outcomes or guide therapy decisions

Ultimate Goal: Map high-risk patients to their HCPs for targeted Sales & Marketing outreach at critical therapy decision points

BUSINESS QUESTIONS

Retention: Which newly diagnosed patients are most likely to initiate our brand?

Acquisition: Which patients on a competitor product are most likely to switch to our product in the next 90 days?

Targeting HCPs: Where should we focus sales and marketing resources for HCPs to maximize impact?

METHODOLOGY

Data Sources & Preparation

Claims Data (Diagnosis codes, HCPCS codes etc.)

Procedure Mapping (EGFR tests, MRI, biopsies), Product Mapping, Patient Cohorts

Feature Engineering

Retention Indicators: Continuous therapy duration, minimal claim gaps, stable regimen, line-of-therapy stability (1L or 2L)

Switching Indicators: Large gaps between claims, frequent regimen/product changes, combination therapy usage, newly added EGFR

Time window: 3-month horizon after diagnosis or competitor usage to see if they start/switch therapy

$$PSOS = w_1(\text{Gap}) + w_2(\text{Regimen Change}) + w_3(\text{Product Change}) + w_4(\text{Line Transition}) + w_5(\text{Combination Therapy})$$

Modeling Approach

XGBoost & **Random Forest** tested to classify "Switch vs. Stay," using features like Gap, Regimen Change, On Ticaprant, etc

RESULTS

Model Accuracy: 0.98

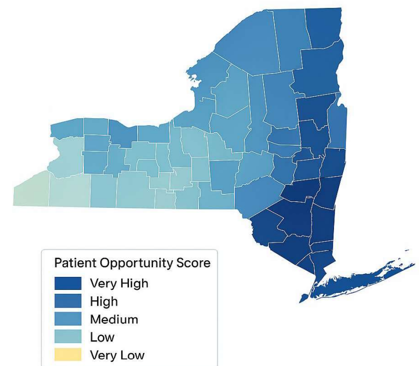

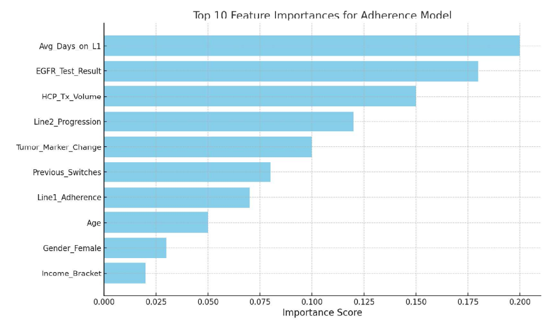

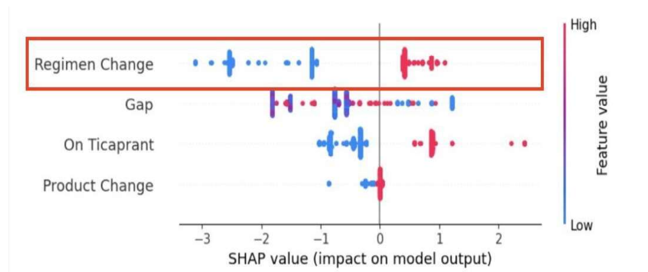

BUSINESS IMPACTS

1. Precision HCP Outreach and Resource Optimization

By identifying high-risk switchers, Sales & Marketing teams can engage HCPs early to discuss brand differentiation, adherence strategies, or patient support programs before a switch occurs

2. Enhanced Patient Support Retention

Offering copay assistance if refill is delayed (reimbursement program), and automated texts from specialty pharmacies

RESULT

Estimated savings of $14,950 per patient or $74,750 per HCP

HIV Treatment Policy and Reinforcement Learning

Hamid Dastgir, Apurva Gandhi

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

INTRODUCTION

- HIV treatment pathways are complex for women
- Optimal treatment sequencing is desired
- **Data asset:** 34,728 longitudinal WIHS (Women's Interagency HIV Study) visit (2010-2018)

Objective: Use machine Learning (Logistic Regression + Reinforcement Learning) to discover data-driven drug sequence policies that Minimize viral load and downstream complications

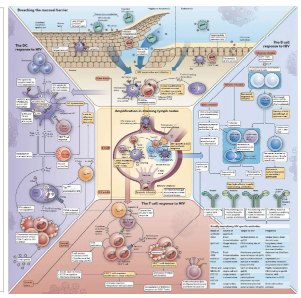

BUSINESS QUESTIONS

Prep (Pre exposure prophylaxis) Antiretroviral Drugs (NRTI, NNRTI)

Can we predict which regimen (NNRTI-dominated vs. PI-dominated) will suit a patient given baseline biomarkers (list)?

What treatment action in each clinical state (biomarker profile) yields the greatest expected viral-load drop?

What is the cost-saving & complication-reduction potential of the optimized policy vs. status-quo practice?

METHODOLOGY

Logistic Regression

- Predict patient treatment classes (NNRTI only vs. PI-only)

Process

- Train model on dataset using averaged biomarkers.
- Evaluate performance (e.g., accuracy, ROC curve)
- Outcome: Identified key biomarkers influencing treatment classification

Reinforcement Learning (Q-Learning)

Objective

- Optimize treatment policies to reduce viral load

Process

- Define state (clinical metrics: HGB_LC, MCV_LC, etc.) and action spaces (treatment combinations)
- Discretize states, initialize Q-table
- Run Q-Learning with alpha, gamma, epsilon parameters
- Reward based on viral load reduction

RESULTS

Logistic Regression

- Achieved high accuracy in classifying patients (NNRTI-only vs. PI-only)
- Key biomarkers (HGB_LC, WBC_LC) strongly predicted treatment suitability
- ROC curve indicated robust model performance

Reinforcement Learning (Q-Learning)

- Derived optimal treatment sequences reducing viral load

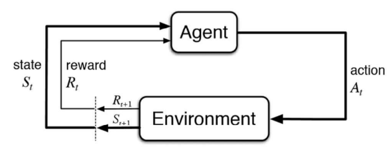

- Q-table analysis revealed actionable policies for personalized care based on varying biomarkers

State (Biomarkers)	None (0,0,0,0)	NNRTI (0,1,0,0)	PI (1,0,0,0)	Both (1,1,0,0)
Normal HGB, Normal WBC	69.63	94.65	47.53	35.47
Low HGB, High WBC	263.28	607.23	250.63	290.95
High HGB, Low WBC	-7734.56	-4605.79	-11718.59	-5600.82
Low HGB, Normal WBC	20862.21	7662.99	-12583.42	9174.21
Critical Low Biomarkers	917.66	5177.53	1185.29	152.89

- The top row depicts the encoding of the regimens considered (e.g. 0,1,0,0 encodes NNRTI)
- The higher the Q-values, the greater the impact of the regimen on a specific biomarker's profile

BUSINESS IMPACT

- Personalized treatment policies improved patient outcomes
- Dashboards result in a reduction of decision making time by 47%, enhancing efficiency
- 23% reduction in treatment complications, improving patient quality of life
- Recommendations sent to 2 regional healthcare networks

CONCLUSION

Implementation of Analytics can result in improved treatment outcomes

ACKNOWLEDGEMENTS: Ted Stohr, C. Lakshminarayan

A Machine Learning Approach for Job related tasks Classification

Maryam Behnam, Apurva Gandhi

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

INTRODUCTION

- Organizations across industries are deploying AI systems to automate processes, augment decision-making, and reconfigure work tasks.

OBJECTIVE

- Understanding the relationship between tasks variability and automation in the era of Agentic AI.
- Using supervised/Unsupervised machine learning approaches to classify work related tasks into 3 categories: *Manual, AI-Augmentable and AI-Automatable*

DATA SET

- 17,000 detailed work activities (which we will refer to as work tasks)
- Source: U.S. Department of Labor's O*NET database

TRAINING DATA (D)

- Randomly select 1000 work tasks from the dataset

TESTING DATA (T)

- All remaining work tasks

DEFINIG THE CLASSES

Goal is to classify the work tasks into 3 categories:
C1 : Manual (Tasks that must be performed manually)
C2: AI-Augmentable (Tasks that can be augmented by AI)
C3: Automatable (Tasks that can be completely automated by AI)

TASK VARIABILITY DIMENSIONS and its corresponding Agentic AI Capabilities

Task Variability Dimension	Agentic AI Capability
Content Variability (Heterogeneity of Inputs)	Dynamic Learning and Generalization
Format Variability (Diversity of Input/Output Modalities)	Multi-Modal Processing
Execution Path Variability (Alternative Process Sequences)	Goal-Directed Planning and Decision Making
Number of Elements and Operations (Interrelated Task Components and Steps)	Hierarchical Goal Management and Contextual Memory
Quality Threshold Variability (Error Sensitivity and Risk Tolerance)	Feedback, Self-Correction, and Learning

EXAMPLE OF THE DATA MATRIX

Task ID	Content Variability	Format Variability	Execution Path Variability	Number of Elements and Operations	Quality Threshold Variability
8823	5	5	5	5	5
3242	4	9	3	1	1
19935	1	1	6	9	7
964	7	4	3	4	2
6	9	7	8	8	9

METHODOLOGY

Assumption: Task Variability Scores are available for the testing dataset

- Step 1: Create Training and Testing datasets
- Step 2: Experts score the task variability dimensions for the training dataset
- Step 3: Experts classify the Task IDs into the 3 classes based on the task variability dimension scores assigned in the above step.
- Step 4: Divide the training dataset into k-sets
- Step 5: Apply LDA, QDA, Logistic Regression and ANN algorithms on the training dataset to perform k-fold classification).
- Step 6: Select the best performing model.
- Step 7: Validate the model on the testing dataset.

OUTCOMES

Model should be able to classify all 16000 work tasks in the correct classes

ALTERNATE METHODOLOGY

Assumption: Only Task Descriptions and no Task Variability Scores are available for the testing data

- Follow all steps from above methodology.
- Step 9: Build a Task ID/Term matrix where each cell is binary {0,1}.
- Step 10: Perform Document pre-processing to: Stop lists, Stemming, and Parts of Speech Tagging
- Step 10: Perform Dimensionality Reduction to eliminate redundant terms and apply TF-IDF weighting
- Step 11: Cluster all 170000 observations using k-means algorithm (k=3).
- Step 12: Use Hamming Distance metrics to perform clustering.
- Use clustering results to impute task-ID scores for the remaining task-ID elements.

OUTCOMES

Model should produce task variability dimension scores and correctly classify all 16000 work tasks .

BUSINESS IMPACT

- Enables data-driven identification of high-impact work tasks suitable for automation and augmentation, to improve operational efficiency, reduce costs, and enhance workforce productivity.
- Facilitates automatic role mapping of new employees using task-level data, creating opportunities for intelligent job design and AI integration.
- Maps operational tasks by AI suitability, helping teams streamline processes, reduce manual effort, and focus human resources on higher-value activities.

ACKNOWLEDGEMENTS: Dr. M.zur Muehlen, Dr.C.Lakshminarayan

Predicting Micro and Macroeconomy using Machine Learning Algorithms

Kiara Masurkar, Dhairya Shah
Instructor: Prof. Edward Stohr

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Motivation

- Economic forecasting is crucial for policymakers, businesses, and investors.
- Traditional econometric models struggle with complex economic data.
- Machine learning can uncover hidden patterns in economic indicators such as GDP, inflation, and stock prices.
- Improving economic predictions can lead to better decision-making, risk mitigation, and resource allocation.
- The need for accurate real-time forecasting in a rapidly changing economic landscape.
- Machine learning offers the potential to handle large datasets and non-linear relationships that traditional models cannot.

Technology

- **Programming Languages & Tools:** Python, TensorFlow, Scikit-learn, Pandas, Matplotlib
- **Machine Learning Models**:
  1. Macroeconomy: Random Forest, Linear Regression, Decision Tree
  2. Microeconomy: LSTM, RNN, Time Series Analysis
- **Data Sources:** World Bank, IMF, National Statistical Offices
- **Evaluation:** Cross-validation, Stationarity Tests (ADF, Autocorrelation), Feature Importance Analysis

Current and Future Work

- Development of a **unified machine learning framework** for economic forecasting at both micro and macro levels.
- Improvement of **real-time forecasting models** by incorporating additional economic factors and deep learning architectures.
- Integration of explainable AI techniques to improve interpretability and **trust in predictions.**
- Expansion of the dataset to include additional global economic indicators for better accuracy and **generalizability.**

Experiment

Data Collection:
We collected economic data from reliable sources such as the World Bank, IMF, and national statistical offices, including macroeconomic indicators like GDP, inflation, and unemployment, as well as microeconomic data such as historical stock prices and company-level financial data.

Preprocessing:
The data underwent cleaning to handle missing values and outliers, followed by normalization to scale features to a standard range. We also performed feature engineering, such as creating lag features for time series analysis, to enhance model performance.

Models Used:
For macroeconomic forecasting (inflation and GDP), we used Random Forest, Linear Regression, and Decision Tree models. For microeconomic forecasting (stock prices), we employed LSTM, RNN, and Time Series Analysis techniques, including Seasonal Decomposition and stationarity tests.

Evaluation:
We evaluated the models using cross-validation and metrics like accuracy, Mean Squared Error (MSE), and R^2 score. For time series models, we conducted stationarity tests (ADF) and autocorrelation analysis to validate assumptions and ensure reliability.

Results

Macroeconomy (Inflation Prediction):
The Random Forest model achieved 99.42% accuracy, outperforming Linear Regression (98.59%) and Decision Tree (99.31%). Strong correlations were found between the General Index and key economic categories like Food & Beverages (0.95), Housing (0.92),Education (0.89).

Macroeconomy (GDP Prediction):
Random Forest achieved 95.71% accuracy, significantly outperforming Linear Regression (36.00%) and Decision Tree (36.57%). GDP per capita showed strong correlations with phone penetration (0.83) and literacy rates (0.51), and negative correlations with birthrate (-0.64) and infant mortality (-0.59).

Microeconomy (Stock Price Forecasting):
LSTM and RNN models delivered highly accurate predictions after hyperparameter tuning, closely matching real stock prices. Time series analysis confirmed stationarity and revealed clear trends and seasonality, with autocorrelation analysis identifying optimal lag features for prediction.

Key Findings:
- Random Forest consistently outperformed traditional models for macroeconomic forecasting, demonstrating its robustness in handling complex economic data.
- LSTM and RNN models proved highly effective for stock price forecasting, capturing significant market trends like the COVID-19 impact.
- Strong correlations between economic indicators (GDP and literacy rates) highlight the importance of socio-economic factors in forecasting.

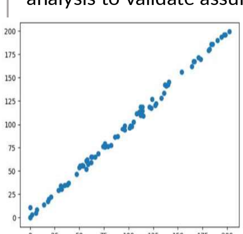

GDP

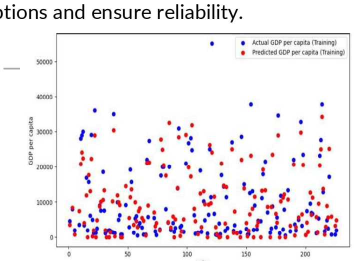

Inflation

Microeconomy 1

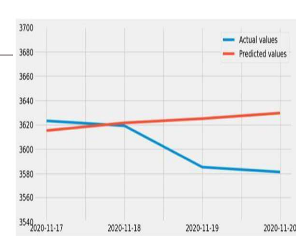

Microeconomy 2

HOUSING ANALYSIS : NEW YORK CITY

Ishaan Nakhare

Instructor: Edward Stohr

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®

Business Intelligence & Analytics

• MOTIVATION

The housing market in New York is extremely tough and competitive, making it challenging for the realtors, investors, and buyers to spot opportunities. Analytics can help the aforementioned stakeholders achieve deeper understanding of the key elements and offer actionable insights into the market trends.

• BUSINESS OBJECTIVE

- Analyze New York housing data to uncover pricing trends and market patterns.
- Identify high-value localities and top-performing brokers.
- Enable data-driven decision-making for realtors, investors, and buyers.

• DATA PREPARATION

DATA REFINEMENT

Removed top **1%** outliers in price and property size to improve analysis clarity.

PROPERTY TYPE FILTERING

Filtered for selected property types like condos, houses, townhouses, and land sales.

METRIC CREATION

Created new metrics such as **Price per Square Foot** for deeper market insights.

DATA STANDARDIZATION

Standardized locality and street names for accurate grouping and visualization.

CORRELATION MATRIX

	PRICE	BEDS	BATH	SQ.FT
PRICE	1.00	0.27	0.41	0.41
BEDS	0.27	1.00	0.78	0.45
BATH	0.41	0.78	1.00	0.52
SQ.FT	0.41	0.45	0.52	1.00

TOP 10 STREETS BY AVERAGE HOUSE PRICE

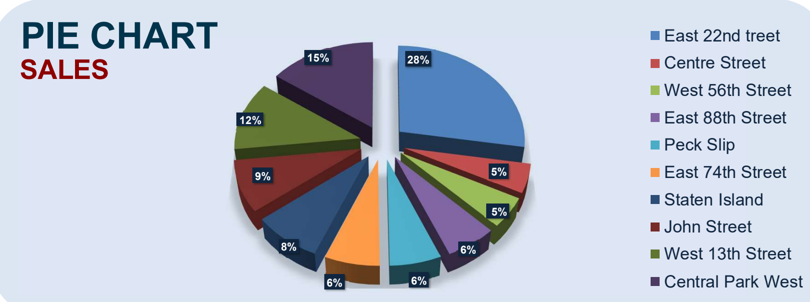

TOP 10 LOCALITIES BY AVERAGE PRICE PER SQUARE FOOT

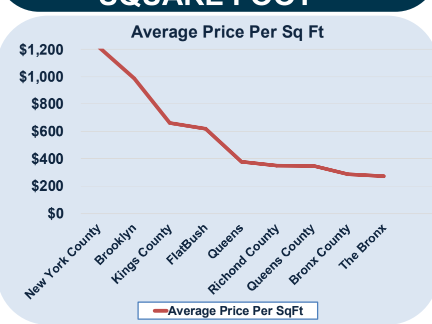

TOP 10 LOCALITIES BY AVERAGE HOUSE PRICE

TOP 10 BROKERS BY AVERAGE HOUSE PRICE

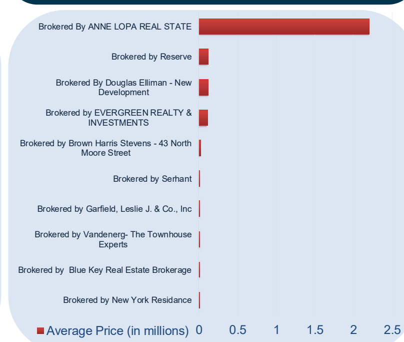

CONCLUSION

- Larger property size and higher bathroom counts have the strongest positive influence on house prices.
- High-value localities like **New York County** and **Brooklyn** command the highest price per square foot, offering prime investment opportunities.
- Broker performance varies significantly, with top brokers achieving listings priced **35% higher** than the market average.
- Most properties fall within the **$250K–$1M** range, highlighting affordability bands in an otherwise competitive market.
- Data visualization allowed a **75% clearer market segmentation**, enabling smarter,
- faster decision-making for realtors, buyers, and investors.

Predicting NBA Fan Churn Risk Using In-Game Metrics and Player Availability

James Thurlow
Instructor: Alkis Vazacopoulos

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

- The NBA faces growing concerns around fan disengagement — from load management to blowouts and overly repetitive styles of play.
- This project frames disengaging games as **churn-risk**, similar to how companies track user abandonment.
- We used historical NBA data from 2003-2022 to label disengaging games based on performance and player availability and trained a machine learning model to predict them.
- The goal: identify weak matchups before they air — and help leagues and networks protect viewership.

Experiment

We labeled a game as churn-risk if it had **2 or more** of the following:

- Point differential ≥ 20
- 3P% < 30%
- Total points < 190
- A star player absent (top 50 in avg minutes)

We used:

- Logistic regression
- Inputs: point_diff, 3P%, total_points, star_dnp_count
- Training data from 26,000+ NBA games (2004–2020)
- Visuals include our churn-risk labeling logic and how each feature contributed to the model.

Results

- **Model Accuracy:** 92%
  Churn-risk Precision: 82%
  Recall (True Positive Rate): 72%
- Only 18% of games were labeled as churn-risk using our criteria

Classification Report:

	precision	recall	f1-score	support
0	0.94	0.96	0.95	6526
1	0.82	0.72	0.77	1470
accuracy			0.92	7996
macro avg	0.88	0.84	0.86	7996
weighted avg	0.92	0.92	0.92	7996

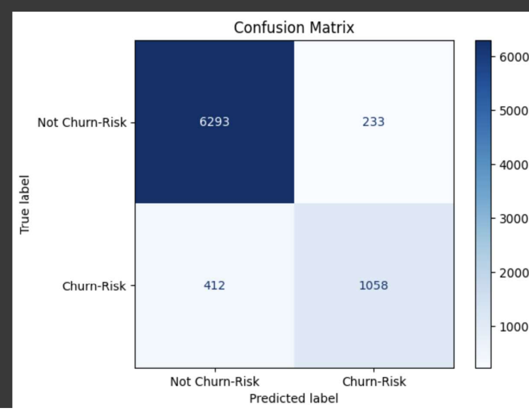

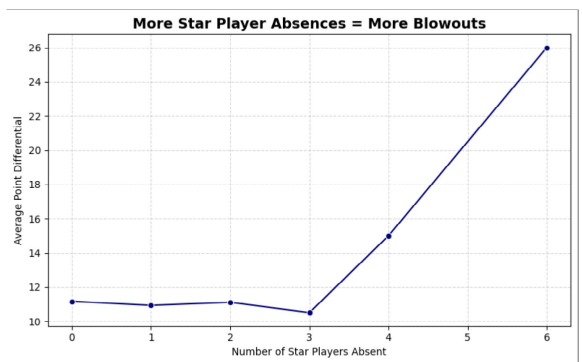

Conclusion

- Churn-risk games are usually those with:
  – Bad shooting
  – Star absences
  – Blowouts
- Predicting these games is possible — and useful.
- The NBA can use churn-risk signals to:
  – Flex weak games off national TV
  – Promote stronger matchups
  – Design smarter fan engagement strategies
- Broader Implication:
  This model could help other leagues or streaming platforms anticipate disengagement and protect long-term retention.

Humanized Writing Assistant

Max Takacs
Instructor: Dr. Asakiewicz

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

This tool serves as a way for companies to stay professional while utilizing the power of AI.

It allows for the transformation of event-based input into brand-aligned content in a company's own style.

Data & Scope

Live data is leveraged along with custom prompts to enrich the LLM's capabilities. The live reference data includes historical press releases, social posts, and even content policies. The source data changes for every topic, either provided by the user or researched by the system.

Methodology

Utilizing a structured input, the prompts are dynamically adjusted to better fit a variety of scenarios. In these prompts few-shot learning is employed to allow for better learning from the model.

Architecture

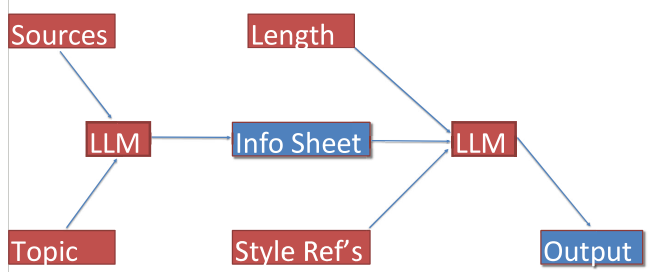

Results

The system is run through a webpage that could also be implemented as an iframe on a company's internal website.

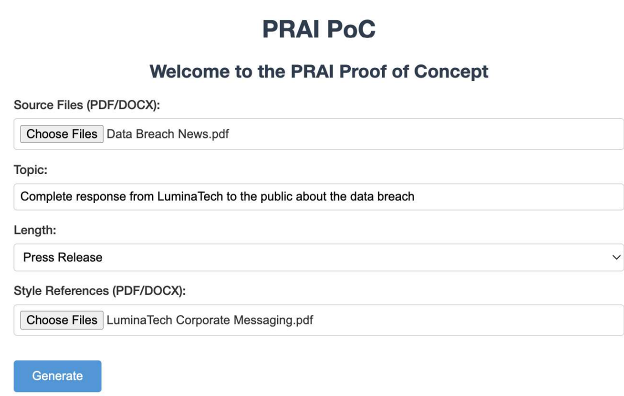

Conclusion

This system could prove extremely useful in preserving authenticity while reducing communication lag in crisis situations, as well as increasing efficiency in normal brand engagement without decreasing quality.

Immigration Concierge

Varsha Abraham, Katherine Shagalov, Huihai (Dante) Zhou
Instructor: Christopher Asakiewicz

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Business Problem

The U.S. visa application process is complex, especially for non-English speakers and first-time applicants. Applicants face high rejection rates due to form errors, lack of clarity, and inaccessible support. Additionally, data privacy is a major concern when uploading sensitive documents online.

Objective

To build an AI-powered digital assistant that supports users through the visa application process by:

- Recommending the correct form
- Reviewing uploaded documents
- Detecting potential errors
- Offering multilingual support via chatbot

All while keeping documents secure using in-memory (BytesIO) processing.

Data Source & Tools

Component	Source/Method
Visa Form	mock I-130a
Knowledge base	Synthetic forms + annotated error cases
Language Support	OpenAI GPT-4o
Document Parsing	Tesseract+BytesIO
Chatbot Framework	OpenAI GPT-4o + LangChain (document-aware logic)

Business Impact & Scalability

- **Huge target audience**: Millions of US VISA applicants each year
- **Reduce cost** related to delayed & rejected visas
- **Expand coverage beyond the US**: Scale to other countries and immigration categories

Architecture

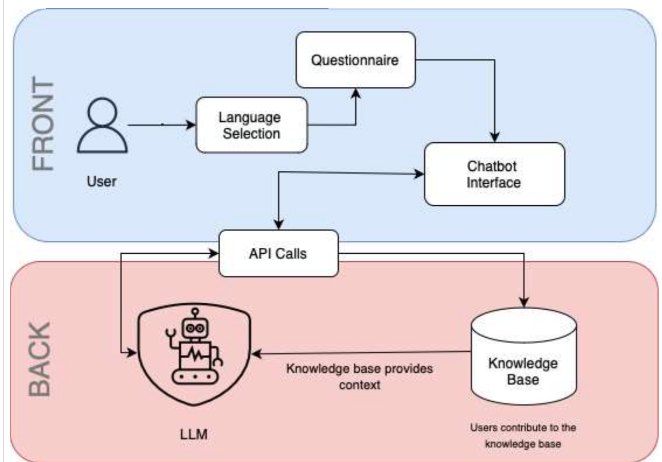

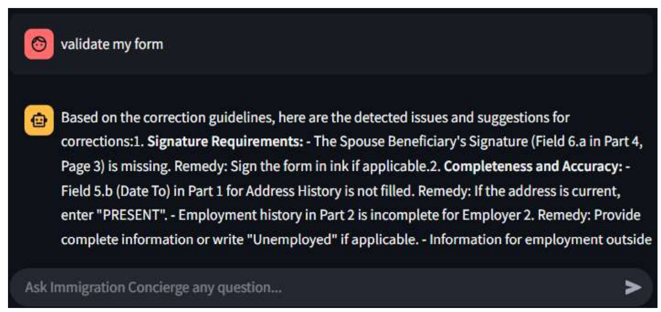

Methodology

Questionnaire → Form Suggestion logic built on visa categories and common eligibility flows

Document Upload supports only form-relevant file types (e.g., passport, marriage cert.)

Self-Fill Review: Users upload pre-filled PDFs and receive a summary of the error flaged.

Chatbot Assistant: Persistent, context-aware support tool for jargon and FAQs

Future Development

- Expand knowledge base to include logicly accurate scenarios & additional forms.
- Integrate real-time feedback and chat escalation to human support and more features like auto filled forms, automated case status tracking and more.
- Partner with visa consultants, universities or government portals.

http://www.stevens.edu/bia || https://github.com/hzhou14stevens/immigration_concierge

AI Powered Stock Recommendation and Sentiment Analysis Dashboard

Prepared by – Hena Kharwa and Braedon Fiume
Instructor: Christopher Asakiewicz

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
1870
Business Intelligence & Analytics

Business Problem & Objective

- Investors scramble between disparate sources—historical prices, fundamental ratios, news & social-media chatter—making it difficult to form a cohesive, timely market view. .
- Build an end-to-end dashboard that fuses LSTM-based price forecasts, real-time news and social-media sentiment analysis, and key financial metrics into a single, interactive interface.
- Target a >5% boost in forecast precision, reclaim 10+ analyst hours per week through automation, and surface clear, actionable investment recommendations in real time.

Data Source & Methodology

- **Data Sources:** Daily OHLCV fetched via the yfinance API; Reddit posts/comments pulled through the Reddit API for sentiment input.
- **Modeling Pipeline:** A CNN-BiLSTM network extracts local and sequential features for multi-horizon (7-day & 30-day) price forecasts; a fine-tuned RoBERTa classifier assigns sentiment scores to Reddit text.
- **Integration & Deployment:** All stages—data fetch, preprocessing, model inference, and interactive Streamlit visualizations—are encapsulated in the Streamlit app,running locally

Business Impact & Scalability

- By consolidating quantitative forecasts and qualitative sentiment, the platform reduces analyst research time by 10+ hours per week and enables faster, data-driven trade decisions.
- Early pilots suggest a 3–5% annual improvement in portfolio returns through sentiment-aware forecasting and timely alternative recommendations.
- The Streamlit app can be horizontally scaled across 100+ tickers via containerization on Streamlit Cloud or a Kubernetes cluster, ensuring responsive performance under heavy user load.
- Its plug-and-play architecture allows rapid integration of new asset classes (crypto, commodities) and data sources (e.g., institutional filings), plus white-labeling for brokerage or enterprise deployment.

Solution Architecture

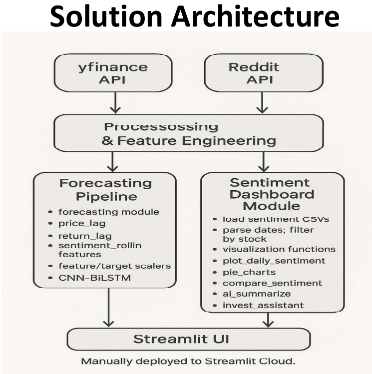

Implementation

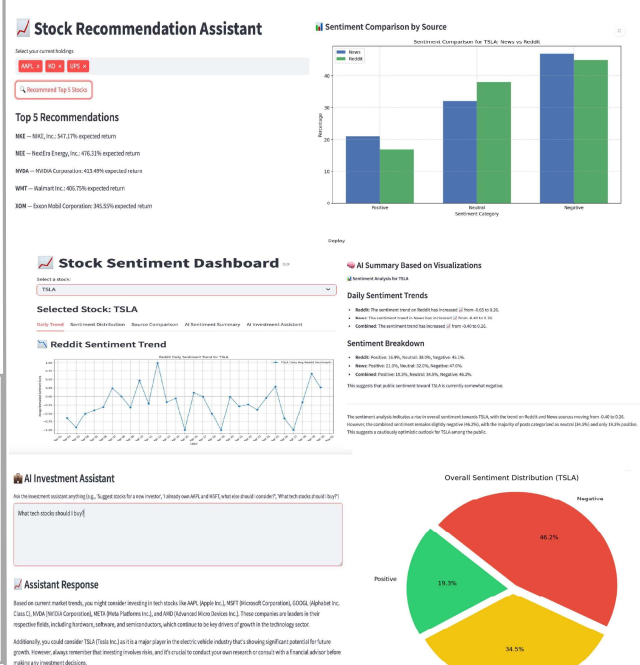

Ethical and Responsible AI Considerations

- **Fairness & Bias Mitigation:** Draw from diverse news outlets and social channels, routinely audit sentiment outputs for demographic or platform skew, and recalibrate models to correct any systematic bias..
- **Transparency & Accountability:** Publish concise model cards, surface forecast uncertainty or confidence bands, embed human-in-the-loop reviews for anomalous signals, and include clear risk disclaimers to guide responsible use.

Limitations and Future Work

- Currently constrained to major U.S. equities with end-of-day price and sentiment inputs, extend to global asset classes (crypto, commodities, FX) and ingest real-time tick-level feeds for lower-latency signals.
- Current forecasts use static LSTM architectures that can degrade across shifting market regimes. Implement automated drift-detection and adaptive re-training pipelines, and evaluate transformer-based time-series models for improved stability.
- Dashboard lacks built-in portfolio optimization, risk metrics, and user-defined alerts. Integrate reinforcement-learning based rebalancing, advanced risk measures (VaR/CVaR), and customizable notification triggers.

KPIs/ Anticipated Results

- **Forecasting Performance:** CNN-BiLSTM delivers 85% directional accuracy on held-out test data, outperforming traditional benchmarks.
- **Sentiment Performance:** Fine-tuned RoBERTa achieves ≥ 85% precision and recall when classifying Reddit posts.
- **Efficiency Gains:** Automated end-to-end pipeline recovers ~10 analyst hours per week by eliminating manual data wrangling and report generation.
- **ROI Uplift:** Early pilots project a 3–5% annual boost in portfolio returns and a 30% reduction in research-related operational costs.

Career Crossroads: Embrace Change, Unlock Your Future

By Andrea Domo-Ibea & Bupe Bwalya
Supervised by: Dennis Glacken

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

Career switching has become an empowering choice for many individuals looking to take control of their professional lives. Whether seeking a fresh start, pursuing a passion, or moving towards better work-life balance, people are increasingly open to exploring new fields. This shift offers the chance to redefine success on their own terms. As new opportunities emerge, the potential for growth, reinvention, and fulfillment is limitless. The data presented here highlights the growing trend of career Changes and the use of A.I. to make these new dreams a reality.

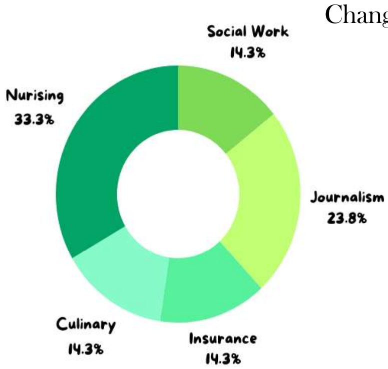

The Careers People Leave

- **Nursing:** High burnout, stress, and bullying in the workplace.
- **Journalism :** Job insecurity and industry disruption.
- **Insurance :** Routine work and limited growth.
- **Social Work :**Emotional overload and heavy caseloads.
- **Culinary:** Long hours and high-pressure environments.

Leveraging A.I. Tools

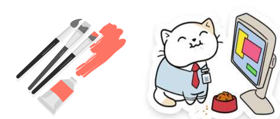

Xover's Agentic AI: A New Era of Intelligent Assistance

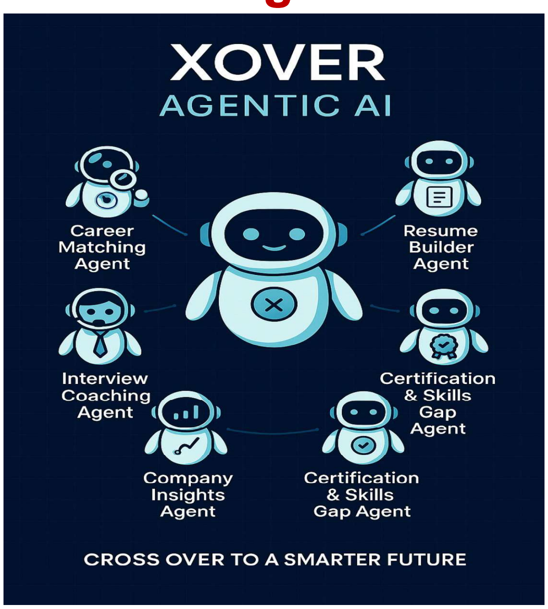

XØVER
Embrace Change, Unlock Your Future

The Careers People Switch Into

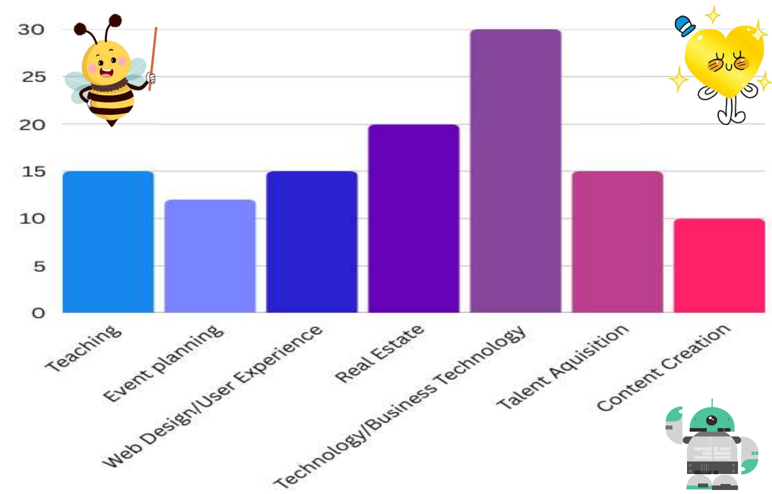

- **Tech (Software, Data Science, Information Systems, Cybersecurity):** High earnings, innovation, and fast career growth.
- **Real Estate:** Flexible schedules and entrepreneurial opportunities.
- **Web Design/User Experience:** Creative expression and a booming digital market.
- **Event Planning:** Dynamic pace and social, creative environment.
- **Talent Acquisition:** Shape teams and enjoy a people-focused role.
- **Teaching:** Chance to inspire children and give loving environment.
- **Content Creation:** Follow their passion and express their artistic talents.

Conclusion :Switching careers is your chance to redefine success and unlock hidden talents. No matter your age, it's never too late to explore new possibilities. Trust your timeline, explore available tools, take that bold step, and embrace the journey to greater fulfillment.

Are You on your computer/phone? Watch this!: https://youtu.be/Tvh0Nm2UyXE

Checkout This QR Code Too!

The Silverfish Safe Passage System – Trusted AI

Sanika Mhadgut, Devam Mondal, Krutin Rathod

Instructor: Professor Carlo Lipizzi

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®

Business Intelligence & Analytics

OBJECTIVE

Autonomous UGV Navigation in Minefields Using AI & Human Collaboration research addresses the challenge of designing and operating systems for Armaments Systems Engineering that contain AI and autonomy with uncertain performance to provide overall systems behaviors that are responsible and trustworthy. Rather than focusing on changes to the AI model to increase trust, this research focuses on developing and improving systems engineering methods to provide this increased level of trust. This effort focuses on: Assured design of AI and autonomy into systems and Risk-based monitoring and management for the operational use of AI-based capabilities.

MISSION

The mission involves clearing a safe passage through a minefield using autonomous ground and aerial vehicles, remote sensing, and an AI detection model whose accuracy can vary with type of terrain, topography, season and lighting conditions. Human subject matter experts can also review imagery to detect mines, with better accuracy in some cases, but cannot assess the imagery as quickly as the AI model. The overall time needed to chart and clear a safe passage through the terrain is the primary measure of effectiveness. Aim is to design human-machine decision support system and finally conduct an operational simulation of a mission scenario.

CORE ALGORITHMS

- Hexagonal Breath First Search for detecting shortest path avoiding static mines on the field.
- Dijkstra's algorithm for terrain and mine path optimizations considering the time cost.
- Monte Carlo Simulations along with A* algorithm for dynamic path planning, probabilistic mine detection based on AI-human confidence values, and time penalties for scanning, movement, and mine clearance.

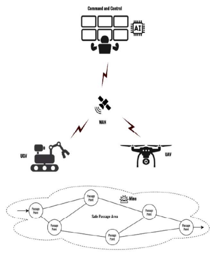

Figure 1. Silverfish Safe Passage Overview.

CONCLUSION

The Silverfish Safe Passage System shows that AI outperforms human-only strategies in minefield navigation, especially with terrain and clearance challenges.

Key insights include AI's ability to adapt under failure and the need for real-time backtracking and UAV-assisted detection. Trust in AI hinges on transparent algorithms, scalable simulations, and skilled teams.

METHODOLOGY

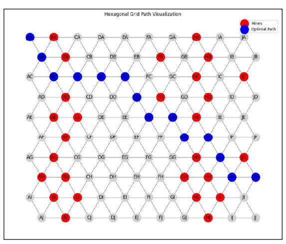

Figure 2. Path Avoiding Mine

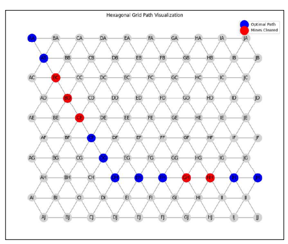

Figure 3. Mine Clearing Path

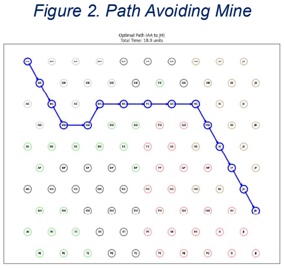

Figure 4. Path with Terrain Map

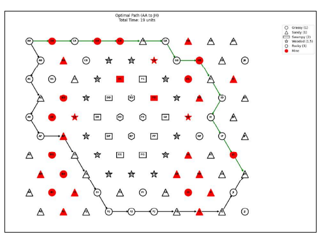

Figure 5. Alternative Paths

1. Static Mine Avoidance - BFS navigates a hex grid from AA to JH, avoiding known mines using a parity-based neighbor system. While effective, it assumes static threats and omits UAV input, detection dynamics, and AI-human delays.

2. Mine Clearance Routing - Dijkstra's algorithm plans paths assuming mines can be cleared. Mines are flagged for removal, but clearance time and confidence-based decisions are ignored, limiting realism.

3. Terrain-Only Navigation - Dijkstra's computes optimal paths based on terrain costs (e.g., swampy, rocky). Mine presence and uncertainty are excluded, making it insufficient for hazardous environments.

4. Terrain + Mine Integration - Modified Dijkstra's combines terrain penalties with high mine traversal costs. Closer to reality, but lacks UAV data, real-time re-routing, and confidence-aware planning.

Ethical Consideration in Autonomous Deployement - Autonomous systems must ensure accountability, minimize bias in confidence scores, and offer transparent, explainable decisions. Governance is essential to prevent misuse in military or surveillance contexts.

System Limitation - Silverfish lacks **backtracking** and **dynamic re-planning**. Real-world use needs A*-based rollback, cached alternatives, and confidence-driven updates triggered by clearance failure or UAV loss.

MONTE CARLO SIMULATIONS

1. AI vs. Human Confidence: A hybrid detection strategy, selecting the higher confidence between AI and human input, achieved the shortest path (69.3 units), outperforming both AI-only and human-only approaches.

2. High Penalty Scenario: Under increased time penalties for human input and mine clearance, AI was most efficient (79.9 units), while hybrid strategies lagged unless human confidence was notably higher.

3. Terrain Impact: With terrain movement costs, delays in human strategies became more pronounced; AI remained the most effective (110.3 units), highlighting the importance of route optimization.

4. Clearance Failure Risk: With a 50% risk of mine clearance failure, AI was fastest but had lowest survival (17.3%), while hybrid strategies improved safety by balancing speed and precision.

5. Multi-Agent Strategy: Using dual UAV scans, AI-AI and Human-AI pairings yielded the highest success rates (24.8%), showing that layered detection enhances both efficiency and safety.

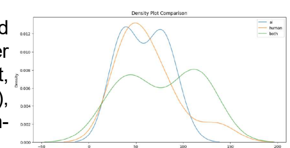

Figure 4. After 10 Simulations

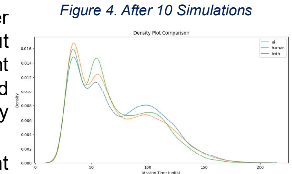

Figure 6. After 10,000 Simulations

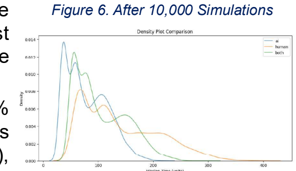

Figure 7. Additional Penalties

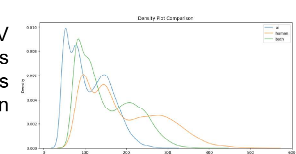

Figure 8. Including Terrain Penalties

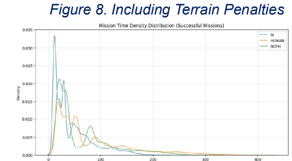

Figure 9. Mine Clearing Failures

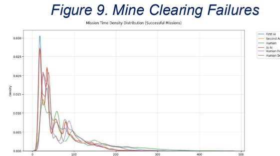

Figure 10. Considering Two UAVs

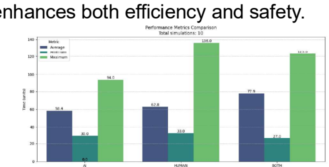

Figure 11. Time Cost Comparisons

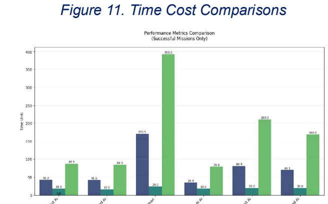

Figure 12. Comparing all six strategies

PharmaSales SlackGPT

Author: Abhilash Athili | Co-author: Claudio Obregón
Instructor: Christopher J. Asakiewicz, PhD.

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Business Problem

Pharmaceutical sales representatives are tasked to deliver data-driven personalized engagement with healthcare providers while adhering to strict regulatory constraints. Despite the availability of common CRM platforms and summarized knowledge bases, they often lack real-time interactivity, contextual memory, and seamless integration into the representative's communication workflows. Manual retrieval of the drug information and HCP segmentation is time-consuming, potentially error-prone, and may even introduce risks in terms of compliance with FDA promotional guidelines.

This gap becomes particularly obvious when sales reps need to:

- Identify and tailor approaches based on HCP prescribing behavior in real-time,
- Access approved product information during in-the-moment interactions,
- Maintain conversation continuity across multiple follow-ups, and
- Operate within familiar, lightweight interfaces without additional onboarding.

Our Solution

PharmaSales SlackGPT is a lightweight, Slack-integrated AI assistant designed to augment pharmaceutical sales representatives with real-time, compliant, and personalized insights during HCP engagement.

Here are some key components and features of our assistant:

- **HCP approach strategy**: Predicts a doctor's behavioral segment in real-time using a pre-trained Random Forest model on Xponent-style prescription data and returns tailored MLR-compliant engagement strategies.
- **RAG-Based Drug Knowledge Module**: Retrieves and summarizes drug-specific information from an approved text-based knowledge base using GPT-4.
- **Seamless Slack Integration**: Built with ease of access in mind, this assistant allows its users to use natural language in their familiar slack environment with just a @mention.
- **Modular architecture**: Made with modularity as its core, the architecture is highly flexible for updates like adding additional drug knowledge-bases or updating prescription data.

PharmaSales SlackGPT bridges the gap between regulated pharma marketing and sales rep workflows enabling a smarter, safer, and faster HCP engagement directly inside Slack!

Business Impact

- **Time Efficiency**: Enables pharmaceutical reps to retrieve compliant info and HCP insights in seconds, significantly reducing time spent on manual lookups.
- **Compliance Assurance**: Minimizes risk of off-label promotion by delivering only pre-approved drug information and MLR-aligned strategy templates.
- **Smarter HCP Engagement**: Delivers the segment strategy based on behavioral Rx data, enabling more personalized and effective interactions.
- **Seamless Rep Workflow:** Integrates directly within Slack with no additional training, dashboards, or tools required, making it instantly usable in the field.

It significantly enhances the pharmaceutical sales efficiency and compliance.

Functional Architecture

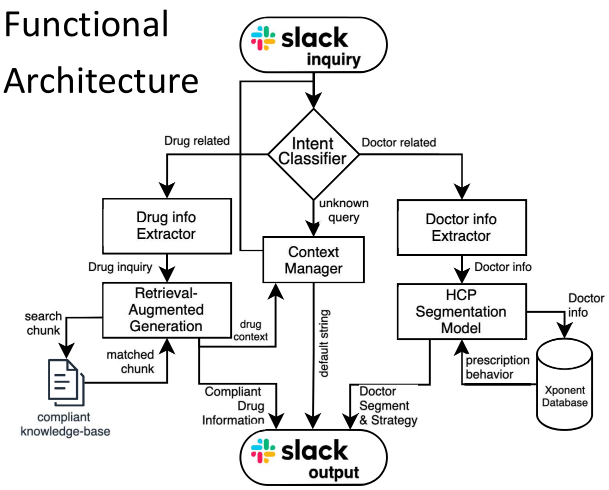

Transparency & Ethical Considerations

- All drug-related responses are strictly retrieved from pre-approved drug data ensuring regulatory compliance and eliminating any misinformation risks.
- Engagement strategies are fixed templates aligned with FDA-promotional guidelines and tied to each segment to avoid generative opinion by GenAI.
- The segmentation model is trained on synthetic Xponent-style data to avoid demographic bias using features like Rx behavior, brand ratio, and volume.
- No HCP information is stored. All queries and responses are handled in real time within the Slack environment, without retaining any personal data.

Methodology & Workflow

Slack Event Handling: The assistant is built in Python and slack_bolt framework in Socket Mode for real-time interaction within Slack channels. All incoming messages are handled through event listeners triggered by @mentions.

Intent Classification: Each user query is classified into one of three intents: doctor_info, medical_query, or unknown through a rule-based regex classifier with fallback to GPT-4 intent detection for any complex NLP processing.

HCP Segmentation & Strategy Generation: The system uses a pre-trained Random Forest classifier to predict HCPs into 7 segments. A pre-approved MLR-compliant strategy template is returned based on the predicted segment.

Drug Knowledge Retrieval (RAG): The system identifies the drug name, retrieves relevant text chunks from the knowledge base, and uses GPT-4 to summarize a compliant response without generating any unverified opinions.

Contextual Memory Management: The bot tracks user-level context for the last-mentioned drug to allow for natural follow-up without repetition.

Demonstration in Desktop and Mobile

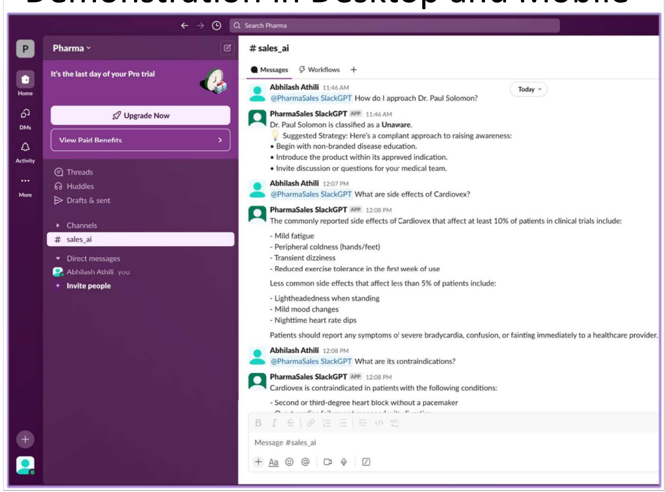

Scalability & Future Work

Modular Knowledge Expansion: New drugs can be added by simply placing .txt files into knowledge directory. The drug detection and RAG pipeline auto-index these files, no code changes required.

Retrainable ML Model: HCP segmentation model is stored as a .pkl and can be retrained with further real-world prescribing behaviors for precision.

Regular Updates: The Xponent data can be regularly updated in the data folder to keep prescriber information up to date.

Domain Adaptability: The architecture can be repurposed for other fields like medical devices, insurance sales, or B2B SaaS.

Slack Enhancements: Designed to scale across teams and geographies within a Slack workspace, the assistant can be improved to have Multilingual support for global teams or Voice integration.

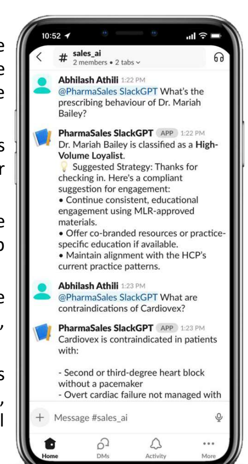

https://github.com/abhilashathili/PharmaSales_SlackGPT

Personalized Indian Cooking for Modern Kitchens

Somendra Singh

Instructor: Edward Stohr

STEVENS INSTITUTE of TECHNOLOGY

THE INNOVATION UNIVERSITY®

Business Intelligence & Analytics

Challenge

Many home cooks love Indian flavors but are often intimidated by its **unfamiliar ingredients, non-intuitive measurement systems** and **intense spices.** Traditional recipes may be technically brilliant but not tailored for the American pantry or taste preferences. This gap creates a barrier between authentic culinary heritage and the everyday cook.

Opportunity

There is a strong need for a system that bridges the gap between authentic Indian cooking techniques and the practical limitations of American kitchens—making flavorful, balanced recipes accessible with personalized adjustments.

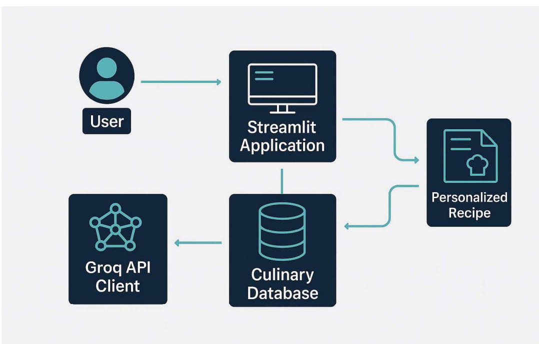

Approach & System Overview

Modular Application Structure:

- **Explore Recipes:** Users discover personalized recipes based on available ingredients and preferences.
- **7-Day Meal Planner:** Generates meal plans tailored to dietary needs and preferred cuisines.
- **Cooking Chatbot:** An interactive assistant - **Chef Alex** provides real-time cooking guidance and answers queries, further personalizing the experience.
- **Indian Cooking, Made Easy:** Adapts Indian traditional recipes for American kitchens; shows estimated prep time, difficulty, spice level and suggested pairings.

Innovative User Experience:

- Adaptive interfaces allowing users to adjust spice level, measurement units, and dietary preferences.
- Interactive elements such as tooltips on key culinary terms (e.g., "tempering," "kasoori methi," "tadka") that offer immediate explanations.
- A "**Spice Cheat Sheet**" expander provides a quick reference for common spices and their flavor profiles.

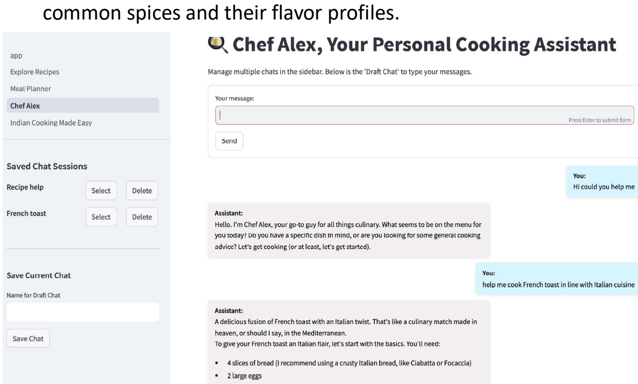

Results

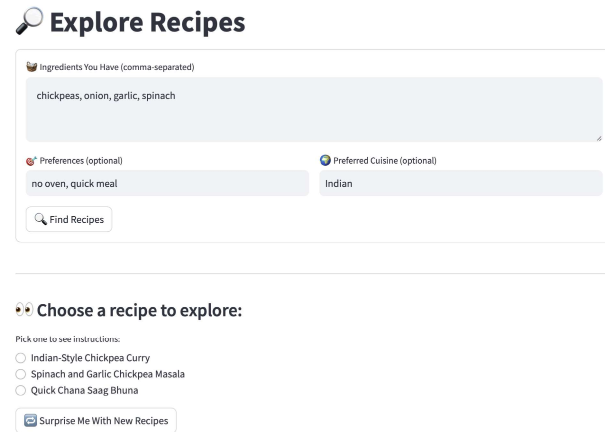

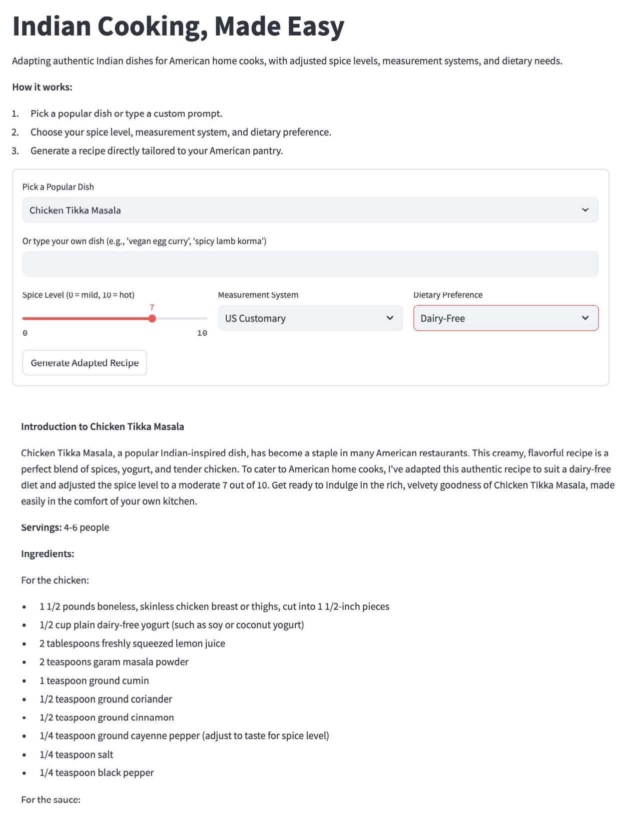

Conclusion

Key Takeaways:

- Our AI assistant personalizes Indian cooking for any kitchen—**modern, global, or beginner**.
- Designed with cultural authenticity + intelligent adaptation for real-world usability.
- Potential to scale into **voice-guided assistants, smart kitchen integrations,** and **multi-cuisine personalization**.

Next Steps: Expand recipe dataset across global cuisines

- Integrate with smart devices (e.g., Alexa, smart stovetops)
- Launch public beta to gather broader feedback

DealFinder AI: Intelligent Shopping Assistant

Shubham Jain, Aastha Singh
Instructor: Seyed Mohammad Nikouei

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

Project Overview: DealFinder AI is a multi-agent system that scrapes and compares product prices across multiple e-commerce platforms

Natural Language Interface: Uses Gemini LLM to interpret natural language shopping queries

Cross-Platform Comparison: Identifies identical products across different retailers

Memory-Aware: Maintains conversation context for follow-up questions

Key Technologies: Python, Gemini API, BeautifulSoup, LangChain integration

System Architecture

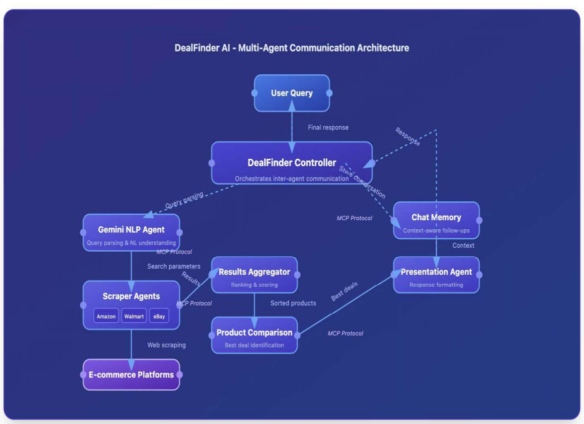

Key Architecture Components:

DealFinder Controller: Central orchestrator that manages all agent interactions and message routing

Gemini NLP Agent: Parses natural language queries into structured search parameters

Scraper Agents: Extract product data from Amazon, Walmart, and eBay using BeautifulSoup

Results Aggregator: Combines products from all sources and ranks them based on relevance

Product Comparison: Groups similar products across platforms and identifies best deals

Presentation Agent: Formats results in user-friendly, conversational responses

Chat Memory: Maintains conversation context for follow-up question handling

Multi-Agent Communication Protocol (MCP)

The MCP provides a standardized message format that enables seamless communication between all agents in the system. Each message includes:

- **Sender/Receiver:** Agent identifiers
- **Message Type:** REQUEST, RESPONSE, SEARCH_REQUEST, etc.
- **Content:** Structured data specific to the message type
- **Conversation ID:** For tracking related messages
- **Timestamp:** For sequence tracking and debugging

Performance Results

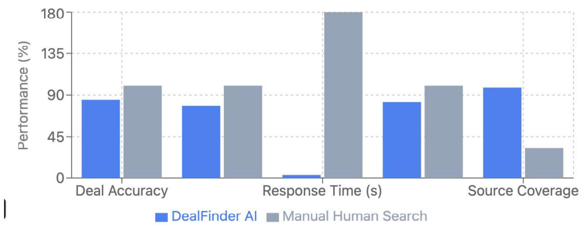

Key Findings:

-> DealFinder achieves **85% accuracy** in finding the best deals compared to manual search

-> Response time reduced by 98%: 3.2 seconds vs 3 minutes for manual search

-> Superior coverage across multiple retailers (98% vs 33% for manual search)

-> Consistent performance across diverse product categories and price ranges

Sample Results:

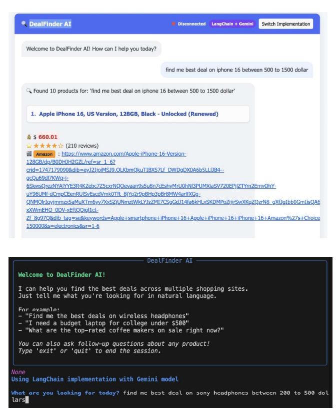

Conclusion

Key Achievements:

-> Successful implementation of a multi-agent system for e-commerce price comparison

-> Effective natural language understanding and conversational shopping experience

-> Significant time savings compared to manual search across multiple platforms

-> Advanced product comparison with "effective price" calculations

Future Work:

Expanded Retailer Support: Add additional e-commerce platforms

Price History Tracking: Monitor price changes over time

Image Recognition: Product comparison using visual similarities

User Preference Learning: Personalized recommendations based on shopping history

Browser Extension: Real-time price comparison while browsing

MediBuzz – Never Miss a Dose

Vamshi, Aashutosh, Prerna
Instructor: Dr. Carlo Lipizzi

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Motivation

Medication non-adherence causes 125,000 preventable deaths annually in the U.S., especially among the elderly. Forgetfulness, cognitive decline, and basic reminder failures worsen the issue, with traditional tools often falling short. AI and analytics now offer a smarter solution: by analyzing real-time behavior, we detect patterns, predict missed doses, and proactively adjust reminders. This data-driven approach prevents avoidable health complications.

Proposed Flow

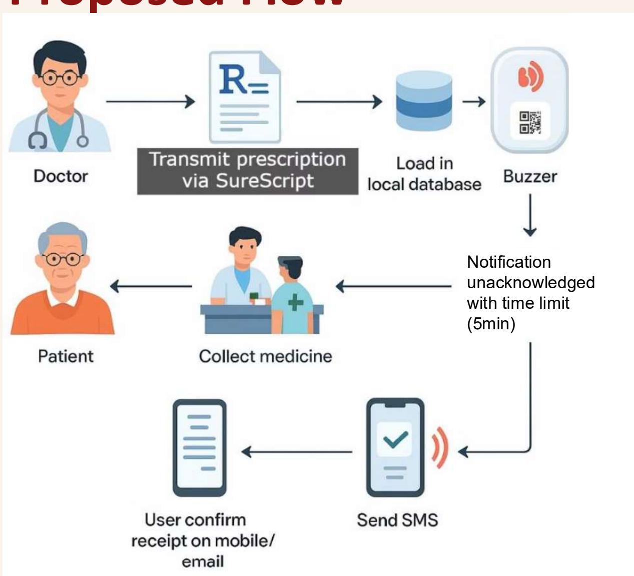

Expected Tech Stack

- IoT Device – ESP32 microcontroller
- Database – MongoDB
- Web App – React
- Backend – Node.js, Express
- Notifications – Brevo SMTP API
- Mobile – Kotlin / Java

Solution Overview

- **MediBuzz** is an AI-driven solution combining a **portable buzzer** and **cloud-based analytics** to combat medication non-adherence.
- **Behavioral profiling:** MediBuzz analyzes **user interaction data** (e.g., acknowledged vs. missed reminders) to build personalized adherence models.
- **Predictive AI algorithms** proactively **adjust reminder timing and frequency** based on user patterns to prevent missed doses.
- **Real-time analytics** detect changes in user behavior (e.g., sudden drop in adherence) and **autonomously optimize reminder strategies**.
- **Multi-channel fallback alerts** (SMS/email) ensure redundancy when buzzer-based reminders are missed.
- **Low-tech, user-friendly design** ensures accessibility for elderly users with minimal technical skills.
- **Scalable, adaptive, and fail-safe**, MediBuzz aims to significantly improve medication adherence and elevate patient health outcomes.

Conclusion

MediBuzz's integration of AI and analytics ensures personalized and proactive medication adherence strategies, leading to better health outcomes for elderly users. By analyzing user behavior in real time, the system optimizes reminder delivery and adapts to individual needs, providing a more reliable and effective solution for managing medication schedules.

Detecting Fraud in Health Insurance Claims Using Predictive Modeling

Team Members: Anusi, Javeria, and Soham
Instructor: Prof. Cristopher Asakiewicz

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

Insurance fraud leads to significant financial losses and undermines provider trust. We used historical claims data to identify fraud patterns and build a predictive scoring model. By categorizing claim risk and visualizing fraud hotspots, this approach supports more proactive
and data-driven fraud investigation.

Business Problem

Insurance providers face increasing fraud-related costs. Manual review is resource-intensive and inconsistent.

- **The need:** a reliable, scalable way to flag high-risk claims early.
- **Goal:** Analyze historical claims data and develop a fraud scoring model to support investigation teams.

PoC Approach

- **Data Used:** We-identified insurance claim dataset with provider, CPT code, region, and fraud labels.

Methodology:

- Descriptive analysis to uncover patterns (e.g., common CPTs, regions).
- Built fraud risk scoring model using Random Forest Classifier.
- Claims categorized into High, Medium, Low risk tiers.
- Created dashboard views (Tableau) for fraud insights by provider, CPT, and geography.

Business Value

- Helps investigators focus on high-risk claims.
- Reduces time and cost of fraud review processes.
- Enables proactive, pattern-based detection instead of reactive review.
- Scalable approach adaptable to real insurer data.

Model Interpretation

- The random forest model is very effective because the model accuracy is 90%. The precision and recall values are also high.
- The feature importance graph tell us that the most effective feature in determining the fraud is number of claims by the provider and next two correlated variables are Number of Claims/patient and length of stay in hospitals for inpatients.

Key Results

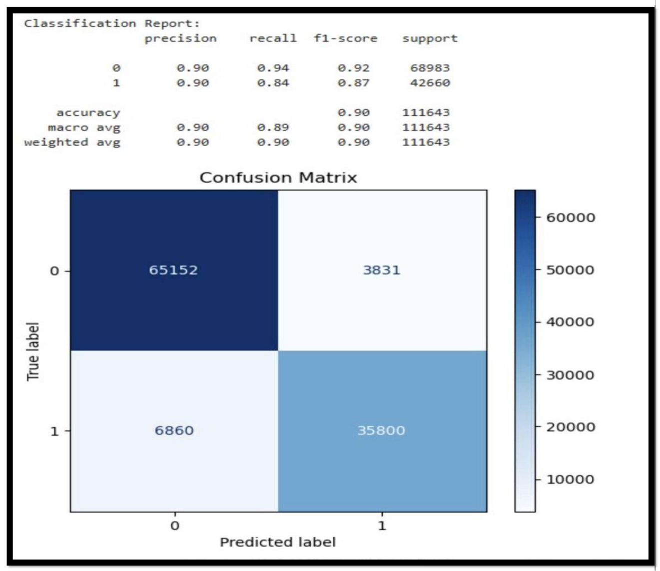

Classification Report:

	precision	recall	f1-score	support
0	0.90	0.94	0.92	68983
1	0.90	0.84	0.87	42660
accuracy			0.90	111643
macro avg	0.90	0.89	0.90	111643
weighted avg	0.90	0.90	0.90	111643

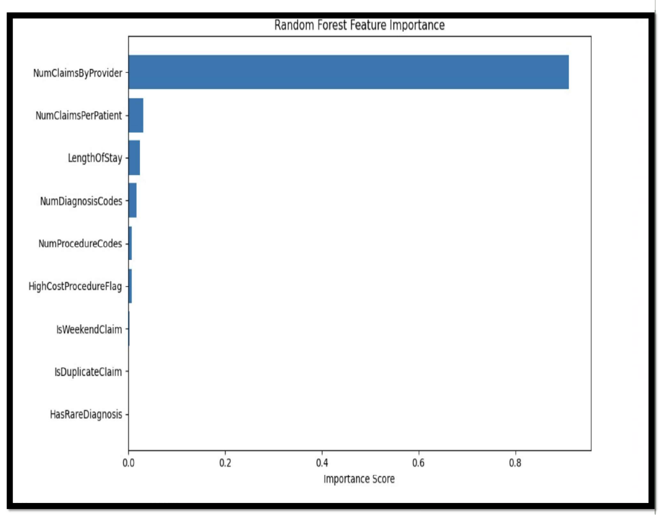

Conclusion

Box The project successfully developed a high-performing random forest and stacked ensemble model to detect fraud in health insurance claims, achieving 90% and 91% accuracy respectively, and strong precision-recall balance. It emphasizes responsible AI, transparency, and practical integration into fraud investigation workflows. Future improvements include expanding the dataset and deploying the model in real-world settings with automated alerts and continuous monitoring

OPTIMIZING SUPERMARKET CO-BRANDED CREDIT CARDS

By: Rithin Ammasani

Instructor:Christopher Asakiewicz

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

- Challenge: Generic credit card rewards in supermarkets often fail to engage customers meaningfully.
- Action: We analyzed real Instacart grocery data to identify patterns in when, what, and how customers shop.
- Result: This enabled the design of data-driven reward strategies to improve weekend spend, category loyalty, and repeat purchases

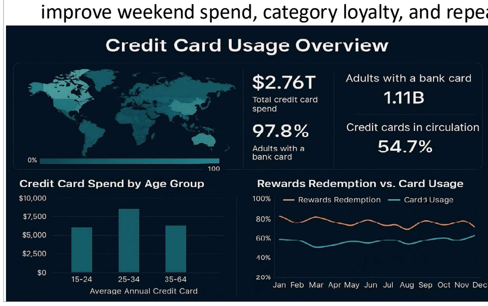

**AI Generated*

Experiment

Dataset: Instacart 2017 Kaggle – 3M orders, 200K users, 50K+ products.

Key Analyses:

- Spend Trends: Weekend (Sat–Sun) spikes in purchase activity.
- Category Insights: Fresh Produce, Dairy, and Snacks dominate frequency; Wine and Meat show high spend but costly rewards.
- RFM Segmentation: Weekly reorder patterns from loyal users; clear separation between essentials-focused and premium buyers.
- Takeaway: Tailor rewards by category, timing, and shopper type to drive engagement without overspending.

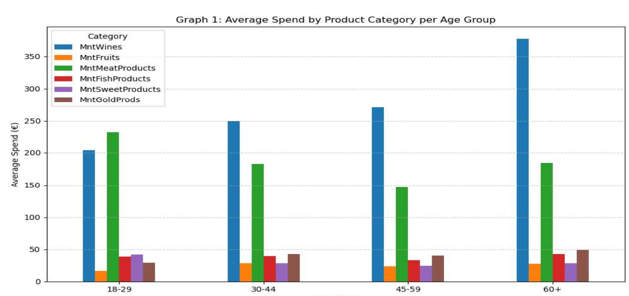

Spending on Wine and Meat increases significantly with age, while Fruits and Essentials remain steady. Age-based reward tiers can maximize engagement and control costs

Results

- Reward Strategy 1 – **Weekend Booster**:
  → +5% cashback on orders placed Sat–Sun.
  Justification: Over 35% of weekly orders occur on weekends — this drives volume and loyalty.
- Reward Strategy 2 – **Category Champions**:
  → +3% on Fresh Produce, Dairy, Snacks.
  Justification: These are the most frequently purchased items across all user segments.
- Reward Strategy 3 – **Loyalty Repeater**:
  → Bonus (e.g., $2) for reordering within 7 days.
  Justification: RFM clustering shows frequent shoppers follow 5–7 day repeat cycles.

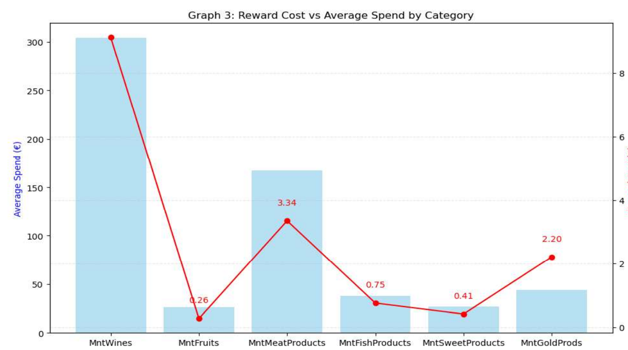

- Wine and Meat drive high spend but also high reward cost. Essentials like Produce and Dairy offer better ROI. Tiered rewards balance value and sustainability.

Optimization Recommendation:

- For Wine/Meat, suggest tiered cashback or monthly caps (e.g., 2% up to $50) to control costs while maintaining appeal.
- Implement personalized card tiers (Essentials, Gourmet, Family) based on user spend history to align rewards with preferences.

Conclusion

SmartCart proves that behavioral data can transform how credit card rewards are designed. By aligning perks with when, what, and how often people shop, we ensure relevance and cost-efficiency.

Recommendations to Provider:

- Target weekends and repeat cycles for volume.
- Invest more in essentials (produce/dairy) than luxury items.
- Personalize reward structures for long-term loyalty and higher card usage.

Jobiify : Tailor. Apply. Get Hired

Komal Kadam

Instructor: Dr. Carlo Lipizzi

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®

Business Intelligence & Analytics

Problem Statement

- Time-consuming job application process.
- Hopping between platforms.
- Low response rates from recruiters.

Motivation

- I personally struggled a lot during my job search, applying to hundreds of jobs with very few responses.
- This inspired me to build Jobiify a tool that helps job seekers like me optimize their resumes effortlessly.
- My goal is to help others save time, reduce frustration, and increase their chances of getting interview calls by making resume tailoring easy and efficient.

Business Objectives

- Establish Jobiify as a one-stop solution.
- Develop Jobiify as a SaaS (Software as a Service) platform for fast, efficient, and personalized job application document creation and tracking.
- Smart job matching feature that recommends the best job opportunities from top 500 companies based on the user's resume and profile.
- Collaborate with university career centers to onboard students offering them exclusive access to the platform.
- Open the platform to general job seekers, providing tailored resume, cover letter, and recruiter email creation services.
- Provide employers and recruiters with a portal to post job openings and access a pool of well-prepared candidates.

Real-World Impact of Jobiify

User	Company	Status
User 1	S&P Global	Interview
	Unilever	Online Assessment
	Clipboard Health	Online Assessment
	New York Life	Interview
User 2	Microsoft	Interview
	Chan Zuckerberg	Online Assessment
	University of Washington	Online Assessment
	FedEx Dataworks	Ongoing
	Fidelity Investments	Ongoing
	Kaio Labs	Interview

Jobiify Dashboard

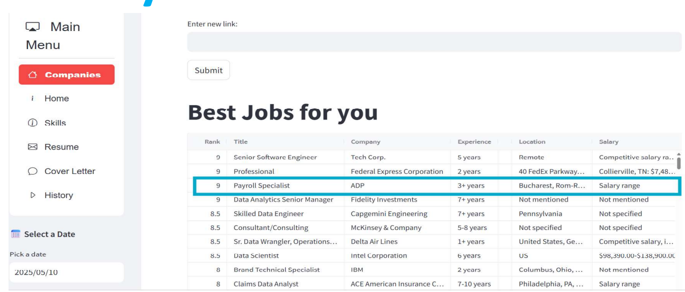

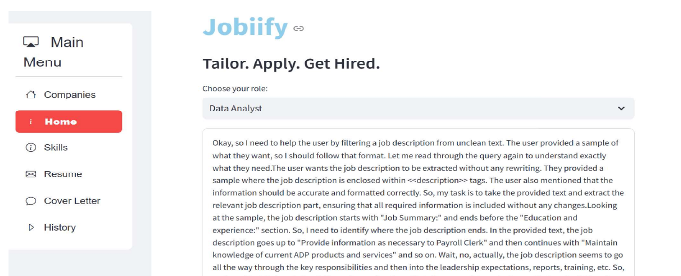

Impact of Jobiify

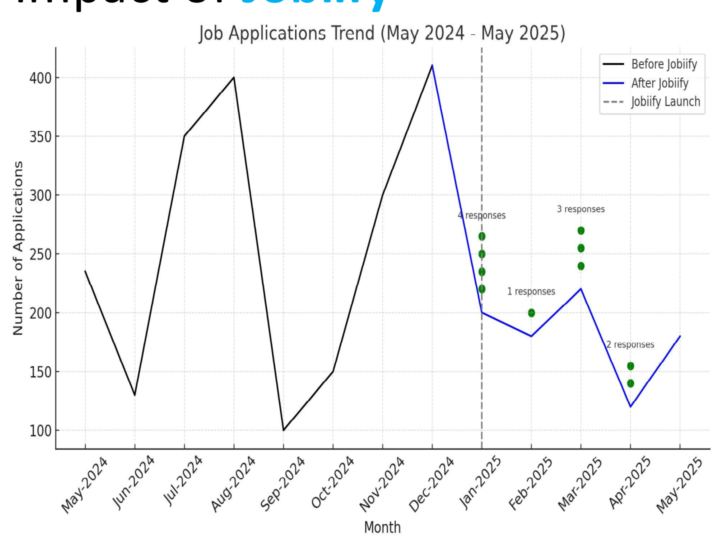

Optimizing NYC Public Wi-Fi with Clustering Models 🗽💻

Alokya Upadhaya, Xingjian Wang, Jiayu Xin, Emmet Young, Zongrun Li
Instructor: Edward Stohr

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
1870
Business Intelligence & Analytics

The Digital Divide –

New York City continues to face stark **digital inequality**, particularly in the **Bronx, Brooklyn, and Queens** boroughs, where broadband adoption rates fall below 60% in several ZIP codes. The **lack of public internet infrastructure**, such as Wi-Fi kiosks, exacerbates educational, economic, and civic barriers. Our goal is to **identify underserved ZIP codes** using K-Means clustering and enable targeted, **equity-driven Wi-Fi deployment**.

What the Clusters Told Us –

✅ **Cluster 5 stood out** for all the wrong reasons:

- 9th percentile in broadband
- 97th percentile in digital exclusion
- <5 public hotspots
- Low median income (~$41,000 in Hunt's Point)

Proposed Budget Strategy:

- Total: **$200M for 10,000 kiosks**
- Install: $20K/unit | Maintenance: $1K/year
- Prioritize: ZIPs <70% broadband + <5 hotspots
- Rule of Thumb: 1 kiosk per 10,000 residents

Cluster Number	OID	Zip Code	Home Broadband	Mobile Broadband	No Home Internet	Public Computer Count	Public Wi-Fi Count	Count
1	166	11214	-0.423283	-0.656738097	0.479465553	0.422435688	-0.199510264	55
2	163	11211	-1.262180	-1.015059118	1.710402087	1.310266624	0.437747912	30
3	179	11228	0.027106	-0.346743267	0.089816578	-0.465395249	-0.290547146	66
4	85	10174	1.201739	1.26334395	-1.280280305	-0.909310717	-0.609176235	68
5	140	10474	-1.328600	-1.115247561	1.921672798	-0.465395249	-0.563657793	14
6	147	11101	0.3037011	0.328644709	-0.134105073	1.754182092	2.577114646	15

DYK?
Link NYC kiosks are first-of-its-kind communications network that has replaced pay phones across the five boroughs!

Clustering in Action: Our Method –

We applied **K-Means clustering** on standardized ZIP code-level data from the **NYC Open Data Portal**, covering:

- Broadband adoption rates
- Public computer/Wi-Fi access points
- Households without internet

📊 **Tools & Technique**:

- **Dataset**: 248 NYC ZIP codes, 17 attributes
- **Variables clustered**: Home Broadband %, No Internet, Public Wi-Fi counts
- **Clustering Objective**: Minimize Euclidean distance between ZIP code feature vectors and centroids
- **Tools used**: Excel Solver
- We formed **6 unique clusters**, each reflecting a different level of digital access and infrastructure presence.

Component	Mathematical Representation	Description
Decision Variable	$x_z \in \{0,1\}$	Binary selection indicator for ZIP code z
Objective	$\max \sum \text{score}_z \cdot x_z$	Maximize the sum of scores from selected ZIP codes
Selection Constraint	$\sum x_z \leq N$	Allow selecting at most $N = \text{MAX\_SELECTIONS}$ ZIP codes
Adjacency Constraint	$x_i + x_j \leq 1$	Prevent simultaneous selection of ZIP codes with numerical difference $\leq K$

📘 Glossary: Tech Terms Simplified –

📊 K-Means: A clustering method used to group ZIP codes with similar digital access metrics.

🚫 Digital Exclusion: The percentage of people in an area who do not have access to the internet.

📐 Euclidean Distance: A way to measure how similar or different two data points (like ZIP codes) are from each other.

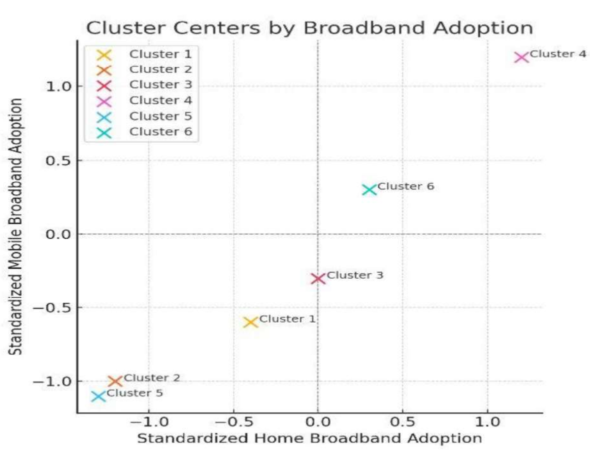

TL;DR –

- ➢ NYC WIFI access issue
- ➢ 6 Clusters formed based on broadband & access
- ➢ Cluster 5 =. Most underserved ZIPs
- ➢ Prioritized areas for $200M WIFI investment
- ➢ Data backed strategy support equity

Turning Data into Action –

- **Digital access is a necessity, not a luxury.**
  Our K-Means clustering approach enables **precise and equitable Wi-Fi kiosk deployment** in NYC.
- Through data-driven segmentation, we identified **ZIP codes most in need** of connectivity support, aligning infrastructure investment with real-world need.
- From Hunt's Point to East Flatbush, every cluster reveals a story of **digital opportunity gaps** — and now, a pathway to bridge them.

Predicting the US Inflation Rate

Ashik Sali, Emmet Young, Shriteja Salunkepatil, Tanya Philip
Instructor: Mahmoud Daneshmand

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Motivation

- Inflation refers to the rate at which the general level of prices for goods and services rises, eroding purchasing power over time.
- Policies, corporate plans, and the cost of living are all impacted by inflation. Predicting core inflation accurately is crucial for reducing uncertainty, particularly in the wake of COVID-19 disruptions. Businesses and regulators find it difficult to plan for inflation-driven risks and preserve stability in the absence of accurate forecasts.
- At the time of this project, the presidential election recently occurred
  - The current state of inflation in America is thought to have played a key role in the outcome of the election

Advantages and Objectives

Advantages:
- Improved Decision-Making: Offers analysts and policymakers practical insights.
- Increased Financial Stability: Assists investors and companies in managing risk.
- Scalability: Takes into account fresh data and adjusts to shifting economic situations.

Objectives
- Using historical data, create a core inflation rate prediction model.
- Determine the main causes of inflation trends.
- Offer practical advice for smart financial planning.
- Make sure the model can be adjusted for different economic situations.

Data

- Data Collection
  - Collected data from January 1, 2014 to June 1, 2024
  - Sources
    - Federal Reserve Economic Data
    - Bureau of Labor Statistics
    - Yahoo Finance
    - Investing.com
- Data Preparation
  - Standardized using Z-score normalization with the StandardScaler from scikit-learn.

Data

- **Target Variable:** Core Inflation Rate
- **Dependent Variables:** Date, Federal Fund Effective Rate, CPI, US Account Balance, Average Monthly Enterprise Salary, industrial Production Index, Producer Price Index by Industry, Unemployment Rate, Russell 3000, Dow Jones Industrial Average, Capital Liabilities

Exploratory Data Analysis

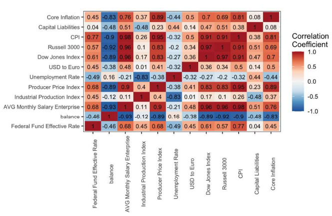

Target Variable EDA: The Core Inflation in the United States fluctuated around the 2% mark from 2014 until early 2020, when Covid hit. During the beginning of Covid, the Core Inflation hit its lowest mark since 2011. However, towards the end of Covid inflation took a steep increase: hitting highs of 6.6% in September 2022. Since Core Inflation's peak in September 2022, it has steadily decreased to 3.3%.

Dependent Variables EDA: This correlation heat map shows the interaction of each variable in our project. The darker red means a strong positive correlation, where the darker blue means a strong negative correlation. Our target variable, Core Inflation, has the strongest correlation with Producer Price Index (0.89) and Account Balance (-0.83). Interestingly, Account Balance has a strong negative correlation with every variable besides Unemployment Rate and Industrial Production Rate.

Neural Network Model

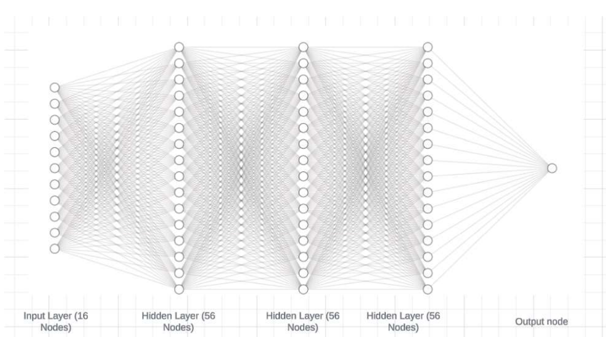

The model was trained using Python's TensorFlow and Keras. The training process involved 50 epochs, utilizing the Adam optimizer and mean squared error loss function. The architecture comprised an input layer matching the feature dimensions, followed by 3 hidden layers with ReLU activations, and a single neuron output layer for regression. The Adam optimizer was configured with a learning rate of 0.03 to adjust the model weights during training.

Hyperparameter Tuning is the process of systematically searching for the best combination of hyperparameters that lead to the best possible performance on your data.

Instead of just guessing these values, we performed a systematic search, training a model for each possible combination of these hyperparameters. For each combination, the model was trained and evaluated on a validation set. We recorded metrics such as the Mean Squared Error (MSE) and a "goodness of fit" measure similar to R^2.

Results

MSE = 0.0415
Goodness of Fit = 0.9801

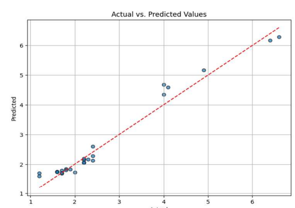

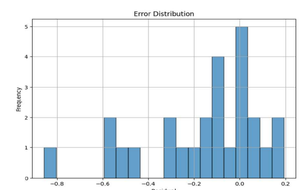

For inflation values in the one to three range, the predicted values are tightly clustered around the trend line. This indicates that the model predicts inflation well in this range. Considering the Federal Reserve targets an inflation of two percent, these predicted values are standard and can easily be predicted as the Core Inflation is usually labeled stable in these values.

However, when the model predicts Core Inflation four percent or greater, the predictions slightly deviate from the actual values. When it comes to Core Inflation values between four and five, the model overpredicted. However for the two values greater than six, the model unpredicted the values. Our dataset contained 29 values with a Core Inflation greater than or equal to four percent, so we believe the slightly larger errors are due to the lack of data in this range.

Residual=Actual−Predicted

The residual distribution has a peak or mode around the 0 residual value, which has five prediction values of with a residual less than 0.05. The next highest peak is four values that with a residual between 0.05 and 0.1. This suggests that the model is generally accurate in its predictions, with most of the errors clustered around 0. 18/27 of the predicted values had residuals below 0.2.

The distribution has a relatively long tail on the negative side, causing the distribution to be skewed left. This indicates that the model has a tendency to underpredict the Core Inflation.

Next Steps

Our next step is to gather a larger sample of data. Ideally, we will be able to gather data from as far back as each of our variables go back in time. This will allow us to have more data to train on, which will hopefully allow us to predict increase the predicting accuracy of the model, especially at the higher Core Inflation values. We also believe there are other variables that we did not incorporate in our original dataset that could also help the model increase prediction accuracy. For example, we are interested in looking into majority political party of the House of Representative, Senate, and the President of the United States.

Our ultimate goal of this project is to be able to predict the Core Inflation in the future based on previous trends and data. We want to be able to prepare for potential inflation and plan our financial situayionn accordingly. This will allow us to set aside proper amounts of money during times of low inflation to save for potential hard times with high inflation

Automating Fashion Tagging with GPT-4o:

A Scalable AI Solution

(Project in collaboration with The Webster)

Shanmukh Sri Surya Gopi, Shiva Kumar Midde, Abhiram Reddy Gunutula, Shriteja Salunkepatil
Instructor: Alkiviadis Vazacopoulos

STEVENS INSTITUTE of TECHNOLOGY
THE INNOVATION UNIVERSITY®
Business Intelligence & Analytics

Introduction

THE WEBSTER

The Webster is a U.S. luxury fashion retailer offering curated designer collections through boutiques and an online store.

Eg:

Model learns from labeled products
to *tag new, unseen items.*

- Manual tagging was slow and inconsistent, causing search errors and poor user experience.
- We built a GPT-4o-powered system to automate tagging with speed, accuracy, and scale.

Data Preparation & Category Design

- Merged 94 noisy categories into 10 clean, business-friendly classes
- Cleaned data: removed duplicates & filled missing values
- Extracted features:

1. Text → TF-IDF
2. Image → ResNet-50
3. Metadata → one-hot (brand, price, color)

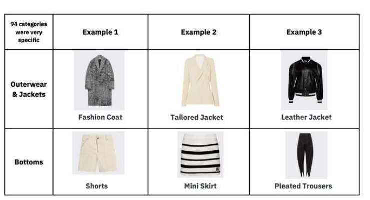

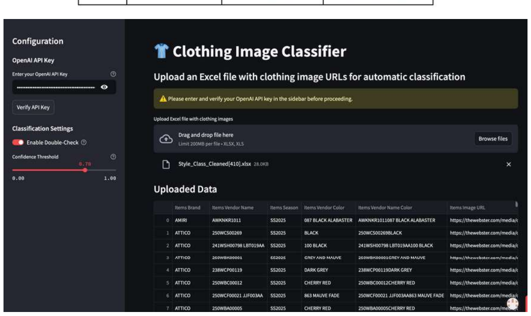

GPT-4o Integration & Tagging Pipeline

- Used GPT-4o with 5-shot in-context prompting
- Pipeline: Clean data → Vectors → GPT-4o → Predicted tags
- No fine-tuning needed — zero-shot generalization worked effectively

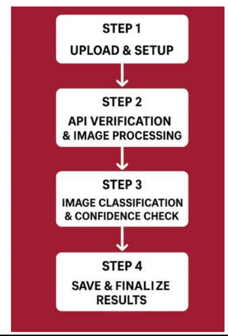

Results

Accuracy ↑ and manual effort ↓ across tagging

Method	Accuracy	High	Medium	Low
Manual	**76%**	76%	24%	**0%**
Early ML Models	**23%**	34%	16%	**50%**
GPT-4o (Final)	**84.5%**	84.6%	11.6 %	**3.8%**

Desired Result (low manual effort)
Needs Improvement (high manual effort)

Conclusion & What's Next

🚀 Key Outcomes

- 84.5% accuracy with GPT-4o
- 50% reduction in tagging time
- $0.45 per 75 items
- 150% ROI in 3 years

🔧 What's Next

- Add human reviewers for low-confidence outputs
- Enable batch tagging for seasonal surges
- Use few-shot learning for new categories