Self-Starting ML Agents

Mines CS Summer Field Session 2026

Building collaborative AI agents for end-to-end machine learning in enterprise environments

Company	bpx energy
Project Type	Enterprise AI / Machine Learning Systems
Academic Program	Colorado School of Mines — CS Summer Field Session 2026
Team Structure	2 teams, 4 students per team

Prepared from the provided project concept and last-year proposal template.

Executive Summary

This proposal outlines a practical, real-world AI project for the Mines CS Summer Field Session 2026. The project focuses on building a self-starting machine learning system composed of two collaborative agents: **DocIntel**, which identifies promising ML opportunities from enterprise data, and **ModelForge**, which prepares data, trains models, and explains modeling decisions. Together, the agents support the full lifecycle from discovery to insight generation in enterprise environments.

Company Background

bpx energy — bp's US onshore oil and gas business — is headquartered in Denver and manages world-class assets spanning Texas and Louisiana. bpx is making a significant step toward reducing emissions, enhancing production, and improving asset reliability.

The company operates a large-scale development program focused on drilling, completing, and producing wells efficiently and in an environmentally conscious way.

This project aligns well with that operating context because it explores how enterprise AI agents can help identify high-value machine learning use cases and accelerate data-driven decision making across complex enterprise data ecosystems.

Proposal Contents

- Project Overview
- Problem Statement
- Project Goal
- Team Structure
- DocIntel: Discovery Agent
- ModelForge: Execution Agent
- Team Integration
- Prerequisites and Resources
- Demo Requirement
- 5-Week Timeline
- Challenges and Risks
- Expected Value
- Tactical Recommendations

Project Overview

This project introduces students to a next-generation AI system composed of two collaborative agents that work together to support end-to-end machine learning in enterprise environments.

- DocIntel identifies opportunities for machine learning.
- ModelForge executes data preparation and modeling.
- Together, the agents form a self-starting ML system that extends beyond traditional AutoML.

End-to-end lifecycle coverage: Discovery → Data Preparation → Modeling → Insight

Problem Statement

In enterprise environments, valuable insights are often buried across multiple data types and systems:

- Unstructured documents such as PDFs, PowerPoint files, and reports
- Semi-structured datasets such as Excel and CSV files
- Distributed storage systems such as SharePoint

Current challenges include:

- Identifying where machine learning can add value is manual and inconsistent
- Data is messy, fragmented, and difficult to prepare
- Machine learning workflows require significant human effort

Project Goal

The goal is to build an intelligent system that can:

- Discover opportunities for machine learning modeling
- Transform messy enterprise data into usable datasets
- Automatically build and evaluate ML models
- Integrate with enterprise systems

Vision: Create a workflow in which AI can both identify high-value ML opportunities and execute the modeling process.

Team Structure

- Two teams, each with four students
- Each team owns one agent
- Teams collaborate to deliver an end-to-end integrated system

Ownership:

- Team 1: DocIntel
- Team 2: ModelForge

Success depends on a clear interface, a shared schema, and consistent coordination between teams.

DocIntel — Discovery Agent

Objective: Identify high-value machine learning opportunities from messy and unstructured enterprise data.

Core capabilities:

- Analyze unstructured and semi-structured data
- Detect prediction opportunities
- Detect anomalies
- Detect optimization problems

- Generate structured ML problem definitions

Key focus areas:

- Reasoning and interpretation
- Justification of identified opportunities
- Going beyond simple information extraction

DocIntel Output

DocIntel must produce structured JSON that can be consumed by the downstream agent.

```
{
    "problem_type": "regression",
    "target_variable": "production_rate",
    "candidate_features": ["pressure", "temperature"],
    "data_location_hint": "sharepoint://folder/path",
    "confidence": 0.85
}
```

Why this matters: A structured interface enables reliable handoff from discovery to model execution.

ModelForge — Execution Agent

Objective: Automatically transform enterprise data into machine learning models.

Core capabilities:

- Integrate with Microsoft SharePoint
- Integrate with Microsoft Entra ID
- Retrieve data from enterprise sources
- Perform data cleaning
- Perform feature engineering
- Perform model training

Illustrative modeling approaches:

- AutoML
- Advanced methods such as TabPFN

ModelForge Output

- A trained model
- Evaluation metrics
- A clear explanation of modeling decisions

Success criterion: the agent should not only build models, but also make its choices understandable and defensible.

Integration Between Teams

The two teams must define and maintain a clear interface contract.

- DocIntel produces structured ML plans
- ModelForge consumes those plans and executes them

Integration requirements:

- Weekly sync between teams
- Shared schema definition
- End-to-end demo required

Key principle: strong integration is essential for turning two separate agents into one usable system.

Prerequisites and Resources

Students will be provided with:

- Access to an enterprise Claude LLM model in the bpx environment
- bpx credentials for Microsoft SharePoint
- bpx credentials for Microsoft Entra ID

These resources support both agent development and enterprise system integration.

Demo Requirement

Each team should build a local demo application using Streamlit.

The demo should:

- Show the agent workflow end to end
- Accept user input
- Display results clearly

Goal: Create a demo that is interactive, understandable, and presentation-ready.

5-Week Timeline

Week	Theme	Planned Work
Week 1	Problem Framing and Setup	Understand project scope; define agent responsibilities; agree on interface contract; set up environment and access
Week 2	Core Capabilities	DocIntel: data parsing and initial reasoning; ModelForge: data

		ingestion and basic pipeline
Week 3	Functional Prototypes	DocIntel: ML opportunity detection; ModelForge: model training and evaluation
Week 4	Integration	Connect both agents; validate end-to-end workflow; handle edge cases
Week 5	Demo and Refinement	Build Streamlit demo; improve reliability and outputs; deliver final presentation

Challenges and Risks

DocIntel Risks

- Ambiguous problem definition
- Hallucinated or unsupported insights
- Difficulty evaluating correctness

ModelForge Risks

- Enterprise authentication complexity
- Messy and inconsistent data
- Automation versus correctness tradeoffs

Integration Risks

- Poor interface design
- Misaligned assumptions between teams

Expected Value

Students will gain practical experience in:

- End-to-end ML system design
- Working with unstructured enterprise data
- Building AI agents beyond simple chatbots
- Integrating AI with real enterprise systems

Real-world relevance: This project mirrors a modern AI architecture in which one system identifies opportunities and another executes solutions.

Tactical Recommendations

1. **1. Define a strict interface early —** Agree on the JSON schema in Week 1 to prevent integration failure.
2. **2. Limit SharePoint scope —** Restrict access to reading files from a single folder to avoid unnecessary API complexity.
3. **3. Provide fallback datasets —** Use local sample data if enterprise access fails so progress is not blocked.

4. **4. Enforce explainability —** DocIntel should justify opportunities, and ModelForge should justify model choices.
5. **5. Start integration early —** Begin integration testing in Week 2 rather than waiting until Week 4.

Closing

The goal is not just to automate modeling, but to create a system that can:

- Identify where machine learning is useful
- Prepare messy enterprise data
- Build models automatically
- Generate actionable insights

Outcome: A practical, real-world AI project that combines reasoning, engineering, and enterprise integration.

Appendix: Suggested Proposal Summary Fields

Project Name	Self-Starting ML Agents
Company	bpx energy
Preferred Team Size	8 students total (2 teams of 4)
Primary Deliverable	Integrated multi-agent ML discovery and execution demo
Demo Platform	Streamlit
Enterprise Integrations	Microsoft SharePoint, Microsoft Entra ID
Primary Outputs	Structured ML opportunity definitions, trained models, evaluation metrics, explainable decisions