MF

CASE STUDY

FINANCE / SUSTAINABLE INVESTMENT

Agentic AI for Investment Research

A multi-agent research system combining academic and web search, RAG and cited report generation to build a shared knowledge base for investment research.

KEY METRIC

400 sources

across web and academia

Manos Fragkiadakis

AI Engineer / Consultant

FINANCE / SUSTAINABLE INVESTMENT

400 sources synthesized across web and academia.

PROJECT
Agentic AI for Investment Research

CLIENT
Baillie Gifford (through UU/DT Lab)

INDUSTRY
Finance / Sustainable Investment

ROLE
AI Engineer / Consultant

CONTEXT
Baillie Gifford was exploring expansion into the Indian market and wanted to examine how systemic climate risks could affect investment portfolios and support discussion with investors and other stakeholders about sustainable investment strategies.

Challenge

This analysis required evidence from academic literature and web sources, structured around frameworks for modeling change in complex socio-technical systems. For each region and climate-risk driver, the team needed to identify relevant sources, examine the academic literature in depth, and synthesize the findings into material for portfolio analysis and workshop discussions at Baillie Gifford. Doing this manually meant substantial time spent searching, evaluating and selecting sources before synthesis could begin.

Approach

To automate this process, I built the research pipeline for the India market as part of a consulting project through Utrecht University's Center for Global Challenges and the Deep Transitions Lab, working within a cross-disciplinary team of professionals.

The system used a multi-agent researcher → writer → polisher architecture. A researcher agent gathered evidence from academic and web sources: academic search used Semantic Scholar, retrieved full papers, and used RAG to retrieve and synthesize evidence from them, while a custom SearxNG deployment was used by the agent to gather evidence from the web. A writer agent synthesized the findings into a structured report with citations, and a polisher agent refined the final output for clarity and consistency while preserving the source trail. The resulting reports were added to a shared knowledge base used by other parts of the project. **Every research output retained citations back to its underlying sources.**

Solution Overview

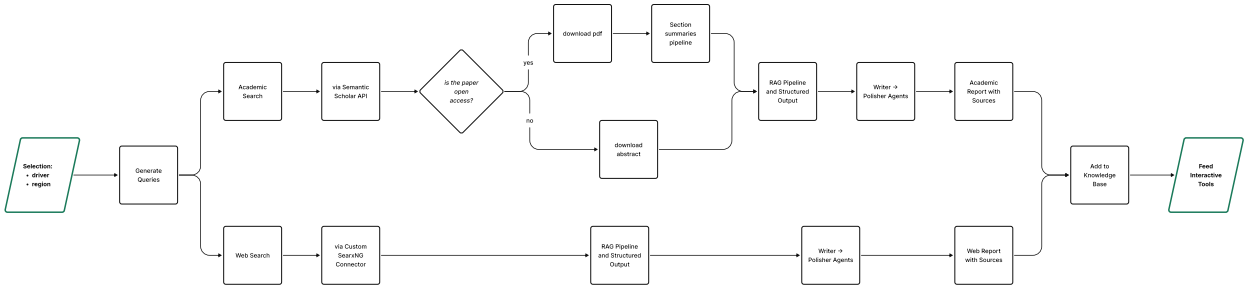

Research flow for one driver and region: academic and web search run in parallel, each through retrieval and writer and polisher agents, into a shared knowledge base that feeds the workshop tools.

How It Works

`Multi-agent research` `Academic search` `Web search` `RAG` `Citation grounding`

Tech Stack

`LangChain` `Semantic Scholar API` `SearxNG` `Qdrant` `Docker`

Results & Impact

The pipeline became the research layer behind the tools used in a live stakeholder workshop. It populated a shared knowledge base with structured, cited evidence from 400 academic and web sources across 10 Indian regions, which other members of the team used to build interactive tools for exploring systemic climate risks. The result was a quick, cited overview analysis in a fraction of the time it would have taken to fill the knowledge base by hand — with structured, source-grounded inputs that could be reused across the wider consulting project.

Project Details

LOCATION
Netherlands / UK

TIMELINE
2024

AGENT PLATFORM
Custom LangChain deployment

TEAM SIZE
Cross-disciplinary team of 5

ENGAGEMENT MODE
Consulting project through Utrecht University / Deep Transitions Lab

ABOUT MANOS

Manos Fragkiadakis is an Applied AI Engineer building document intelligence and AI knowledge systems. He has over five years of Python software development experience and has built custom ERP and CRM systems for professionals. His recent work includes cited RAG across 3,000+ page legal case files and court-monitoring automation software used by 10+ law offices. He previously worked as an AI engineer and taught Natural Language Processing at Utrecht University.

GitHub _d dunderscore Technical blog