UNIVERSITÀ POLITECNICADELLE MARCHE

FACOLTÀ DI ECONOMIA“GIORGIO FUÀ”

Corso di Laurea triennale in

Digital Economics and Business

An AI agent-based platform for automating sustainability audits

Relatore:

Prof. Alex Mircoli

Rapporto Finale di:

Kashish Khurana

Anno Accademico 2025/2026

A big thanks to my Professor Alex Mircoli for alway being there at every step, for motivating me always, for being a great ispiration as a professor but more importantly as a human being.

TABLE OF CONTENTS:

CHAPTER 1: INTRODUCTION

1.1 The Importance of Corporate Sustainability and the ESG Impact

In the modern global economy of our century, companies are expected to be sustainable. In fact, corporate sustainability has transitioned from being just a public relations marketing strategy to a fundamental pillar of digital economics and business management. The modern market is increasingly becoming aware of the long-term economic profitability linked to environmental and social responsibility and attentive to the issues as well. This is why ESG, which stands for Environmental, Social and Governance, is becoming so important. ESG is a way to measure how well a company is doing in these areas. It looks at how a company treats the environment, how it treats its employees and customers, and how it is governed to evaluate their ethical impact and risk profile. The “Environmental” component of ESG looks at how an enterprise performs as a steward of nature so how it affects the planet. This includes its carbon footprint,
resource management, and ecological initiatives. The “Social” part evaluates relationships with employees, customers, suppliers and the communities they operate in. Finally, the “Governance” part looks at how a company is led and managed so the internal leadership structures and audit controls. However, the current process of adaptation of Sustainability has led to one of the major market inefficiencies known as “Greenwashing". This phenomenon occurs when organizations intentionally mislead, or make excessive, outright false claims about its sustainability to create the impression and attract customers and B2B

clients. Since analyzing a company’s true environmental footprint means dealing with vast amounts of unstructured data (third party audits, reports, controversies, social media news) information, asymmetry persists. Companies can easily publish heavily curated, eco-friendly messaging on their corporate websites while simultaneously engaging in environmentally damaging practices. Places like the European Union and other regions are trying to prevent this by tightening regulations and requiring stricter sustainability reporting. But the sea of data is massive, making manual factchecking slow, expensive, and tough for any business trying to maintain a truly sustainable supply chain. In digital economics, there’s a real need for scalable, automated technology that can audit corporate sustainability claims quickly and objectively bridging the gap between what companies say and what they actually do.

1.2 Project Goal: Market B2B Networking - Social Utility - Target Audience

The main, definitive goal of this thesis project is to assist new companies in constructing and expanding their market networks, especially with verified, eco- friendly business partners while mitigating the risks of greenwashing. In today's business world working with partners is no longer a choice; it is a must if you want to stay competitive. Building a network of businesses is hard for companies because it takes a lot of time, research, and resource allocation to find truly sustainable partners and it is even harder to approach external companies blindly without prior verification of their compliance standards.

The AI Agent developed in this project directly solves this problem by acting as an intelligent bridge between firms. It uses the N8N automation platform which, through the use of various tools and programming, lowers the barriers of entering ecofriendly B2B networks. It handles all the work of finding and verifying potential partners, saving time and money on research, and removing the hesitation and friction linked to reaching out new corporate entities.

The Social Utility of this project consists of the ability to access sustainable supply chains, minimizing analytical and administrative costs of finding and approaching green partners, and creating a highly functional economic network effect. The Target Audience who benefits the most from this are corporate figures such as procurement officers, business development managers and supply chain directors who are tasked with sourcing green allies but lack the manual operational capacity of contacting candidates every day. The project is also helpful to digital economists since it shows how automated business logic can change the way market works, and how its structures generate a transparent and sustainable business cluster.

1.3 STRUCTURE OF THE THESIS

This thesis unfolds across four chapters, each addressing a different angle of the AI-driven sustainability auditor.

The present chapter, **Chapter 1** has introduced the overall picture of what the project is about, laying down the economic environment around ESG compliance,

outlining the hardships of sustainable business networking and defining the core objectives, utility, benefits and target audience.

Chapter 2 will provide the technological background behind the project's underlying architecture. Here you will find a close in depth analysis of the Large Language Models (LLMs) and the evolution of independent AI agents. It will feature how these systems process language, execute logical tasks, and output real solutions. Moreover, the chapter also offers a careful evaluation of N8N, the automation platform used to give life to the AI Agent by managing complex data and integrating APIs.

Chapter 3 is the heart of the thesis: hands-on, empirical, and thorough. This section walks through exactly how the project gets built, step-by-step with screenshots of the n8n workflow, the sample companies considered for the experiment and explanation of the nodes involved: from Firecrawl's web scraping to SerpAPI's search engine querying. So, you get to see how data flows and how the AI Agent generates outputs utilized by the system to evaluate and analyze companies and consequently automate B2B outreach.

Finally, in the last chapter, **Chapter 4** reviews the outcomes generated and tests whether the AI agent successfully worked or not. The test consists of using a sample batch of companies and analyzing the accuracy and reliability of the agent's decisions. Besides this, also flagging up the challenges that came up, other different approaches for this experiment and sketching out ways for this technology's further improvement and broader application.

CHAPTER 2

2.1 LARGE LANGUAGE MODELS AND AUTONOMOUS AI AGENTS

2.1.1 Foundational mechanics and operational limits of LLMs

The rise of Large Language Models (LLM) has completely transformed the pace of the digital economy. These models rely on a Transformer Neural Network which uses deep self-attention mechanisms to comprehend and compute semantic relationships of words and phrases in extensive texts models. From an economics perspective, LLMs have commodified cognitive tasks, once exclusive to human labor such as classifying unstructured data, synthetizing abstract textual, and extracting semantic data. Furthermore, these models train on vast volume of text, absorbing petabytes of language and learning mathematical representations of words and phrases withing high dimensional vector spaces. Consequently, the model can predict the most statistically likely next word given a specific textual input sequence.

But in a business workflow, a language model holds various operational boundaries. Firstly, it suffers from a certain knowledge cutoff date becoming incapable of assessing real time market shifts, regulatory updates and latest public corporate news controversies. Second, LLMS have no built-in tie to real external truth: it means that they just judge based on the linguistic probability rather than the empirical facts exposing the business to a serious risk of “hallucinations” (generation of false or ungrounded data). Third, these generative

models are passive structures, so they cannot access corporate databases, activate transactional software, or extract data from the web. They are limited to the text in the immediate current session, and this proves to be a negative aspect for the companies seeking a live B2B sustainability audit. In conclusion, a simple LLM is fundamentally insufficient since it lacks verifying answers in a dynamic, evolving market ecosystem.

2.1.2 The Paradigm shift toward Autonomous AI Agents

The passive nature of generative language models, as mentioned before, created clear practical limitations leading digital economists and computer science researchers to develop a more capable alternative: Autonomous AI Agents. Unlike LLMs, this approach is innovative because it operates at the center of a larger system, enabling it to reason, plan, and take actions through a centralized processing unit (CPU). In technical terms, it coordinates a broader workflow instead of being a simple text generation tool.

Also, it is mathematically architected to mimic human problems solving cycles through an iterative process, the **ReAct** paradigm (stands for “Reasoning and Acting”). For example, when given a high-level goal as identifying supply chain partners with strong environmental records, the agent does not simply generate one response. Instead, it breaks the goal down into smaller steps, works through each one, and adjusts its approach along the way.

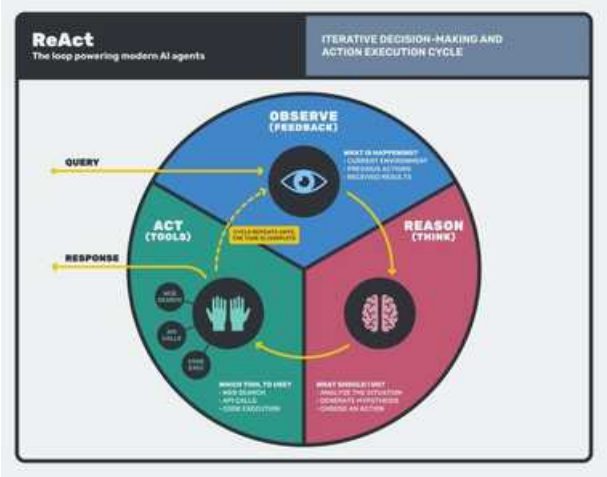

An Autonomous AI Agent becomes operationally independentthrough three key components: a reasoning framework, a memory architecture,and an interface for tool execution. Reasoning framework consists in what the aiagent thinks of its current state analyzing discrepancies between its goal and theavailable data and choosing which operational path to pursue. The memory system works on two levels: the short-term memory which remembers thought and outcomes of previous steps withing current execution loop and a long-term memory which keeps track of past interactions in previous sessions. The third peripherical system is the most crucial in automation, the tool execution interface. The latter allows the agent to connect to external sources such as search engines, web scrapers, software utilities, and databases through standardized APIs. When the agent encounters a question that cannot be answered from its internal knowledge, it pauses, retrieves information from an external source, and incorporates that result into its reasoning before continuing. This loop provides ai agents with their practical value and sets them apart from traditional language models.

2.1.3 Agentic Applications in Modern Digital Business Environments

To understand the business relevance of autonomous agents, it helps to think in terms of costs. In the past, evaluating a potential business partner for environmental compliance required a significant amount of manual work. Compliance teams had to review sustainability reports, search through legal records, monitor news coverage, and examine social media activity. All this was time-consuming and expensive so most companies limited their due diligence to a relatively small number of partners and often continued working with familiar suppliers even when better alternatives may have existed, simply because finding those alternatives was not cost-effective.

Autonomous agents change this situation in a meaningful way. Because they can process large amounts of unstructured information quickly and at low cost, the effort required to assess any individual company is greatly reduced. This makes it realistic for businesses to monitor entire market sectors on an ongoing basis, rather than conducting occasional, selective reviews. An agent can track regulatory updates, check for changes in a company's legal or compliance status, evaluate ESG ratings, and review carbon disclosure reports - all at the same time and without the delays associated with manual research.

For companies, this creates two practical advantages. On one hand, continuous monitoring helps reduce the risk of reputational or regulatory damage; on the other hand, it makes it easier to discover new suppliers who meet high

sustainability standards but have not yet come to a company's attention through traditional channels.

2.2 N8N: ADVANCED WORKFLOW AUTOMATION

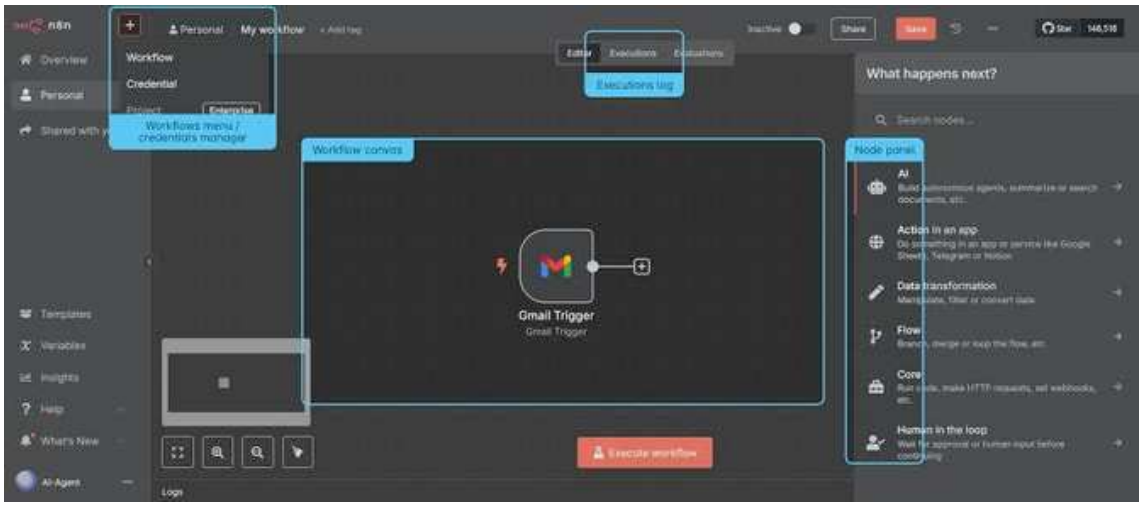

2.2.1 Node-Based Visual Programming and Deterministic Business Logic

It is challenging to develop a functional autonomous agent, but it is even harder for its real-world use within a business context which requires an environment capable of handling data flow, coordinating external service access, and catching, handling errors. For this project, n8n appeared to be the best option.

N8n is an open-source workflow automation tool capable of modeling and executing complex data pipelines via visual design by connecting blocks on a canvas instead of coding the entire process.

What distinguishes n8n from other automation tools is its design: where other tools rely on a sequential approach, n8n models' workflow on a Directed Acyclic Graph allowing data to take multiple conditional paths rather than a linear sequence.

Every n8n workflow consists of two basic concepts:

- **Nodes**: atomic units each performing a distinct function such as fetching data from an API, performing data manipulation, executing a small block of code, or branching logic flow based on an outcome.
- **Edges**: connections between nodes that dictate the sequence of execution and the flow of data.

This means we can easily map out different decision branches, run tasks in parallel, and set up fallback mechanisms making sure the AI only does exactly what we want it to do. In addition, N8n also offers two specific features that are necessary for real-world business applications:

- **Standardized JSON formatting:** Data moving between nodes is strictly formatted as JSON arrays, making it straightforward to connect to different services.
- **Complete execution logging:** N8n records the exact inputs and outputs for every step of a workflow. This creates a clear audit trail for debugging and provides transparency, which is necessary for business processes involving financial or legal data.

2.2.2 Hown8n Uses AI to Make Decisions: LangChain and Agentic Nodes

LangChain is an important concept while dealing with n8n, crucial in the evaluation of a company's sustainability; traditional automation would not be

enough. It is officially defined as "a framework for building LLM-powered applications" (LangChain, 2024). [1]

The central element to this integration is the AI Agent node. It operates on a ReAct (Reasoning and Acting) basis. According to Yao et al. (2023) [2], ReAct allows Large Language Models "to generate both reasoning traces and task-specific actions in an interleaved manner, allowing for greater synergy between the two." Additionally, authors explain, "reasoning traces help the model induce, track, and update action plans as well as handle exceptions, while actions allow it to interface with [external environments]" (Yao et al., 2023). [2] These external environments are the “Tools” in an n8n workflow context.

2.2.3 Processing Data for Sustainability Analysis

Automated ESG due diligence requires dirty web data management. In fact, the latest academic research highlights "ESG data exists in varied formats, e.g., tables, narratives, scanned documents, and multimedia, making consistent extraction difficult" (Zhuang et al., 2026). [3] To solve this, n8n utilizes HTTP requests to extract HTML, converting it into standardized JSON. Visalli et al. (2023) note that "the evolution of LLM has greatly optimized text processing methods, making the transformation of unstructured documents in ESG reports into structured data a feasible reality." [4]

Structuring data before providing it to our agent requires a certain accuracy. As demonstrated by Moodaley and Telukdarie (2023) [5], processing clean,

structured text significantly "enhanced the ability to identify green claims and detect 'greenwashing'."

2.2.4 Keeping the System Reliable: The Human-in-the-Loop

The introduction of autonomous agents in digital economics brings also risks, primarily due to LLM errors. Ji et al. (2023) [6] define this phenomenon as: "Hallucination in Large Language Models (LLMs) refers to outputs that appear fluent and coherent but are factually incorrect, logically inconsistent, or entirely fabricated." The authors warn that "hallucinations in LLMs affect the reliability and efficiency of AI systems, particularly in high-impact domains" (Ji et al., 2023). [6]

N8n requires a Human in the Loop (HITL) architecture to ensure that the agent does not approve of a risky partnership with a corporation. The HITL according to the official documentation "combines AI speed with human judgment, adding review checkpoints at critical moments to prevent costly errors, compliance issues, and brand damage" (n8n, 2024). [7]

CHAPTER 3

3.1 GOAL AND SOCIAL UTILITY OF THE AI AGENT

The practical focus of the thesis is to translate concepts like Environmental Sustainability and ESG metrics into a real everyday automated tool that helps in business operations. Whereas the main goal of this project is to build an intelligent, autonomous AI Agent that can optimize how firms create and expand their corporate networks. It surely saves time and money, making the approach simpler, and most importantly this sustainability-friendly approach. The AI Agent can crosscheck through social media, external web data, and media reports if there are controversies or false sustainability claims from a green perspective. This really helps before making a business connection.

From an economic standpoint, this tackles the problem of information asymmetry. Big companies spend a lot of money on polished sustainability reports, while smaller firms do not have the resources to check if those claims are actually true. To solve this, the workflow uses the AI to compare what a company says about itself on its website with what the news and search engines say about it. The practical value of this tool is that it gives businesses a scalable way to find genuinely eco-friendly partners, avoiding reputational risks and respecting environmental regulations

This project is a practical example of the exact skills needed in the digital economy today. Businesses need professionals who can combine IT capabilities

with real business strategy. Building an automated system that processes unstructured web data to help with B2B networking is a direct application of data-driven business analysis. Instead of looking at corporate compliance through a purely theoretical lens, this AI agent forces a strict, data-first approach. It shows that automation can realistically help managers make better decisions, allowing them to build supply chains based on verified facts rather than just trusting corporate marketing.

3.2 Workflow Implementation and Node-by-Node Analysis

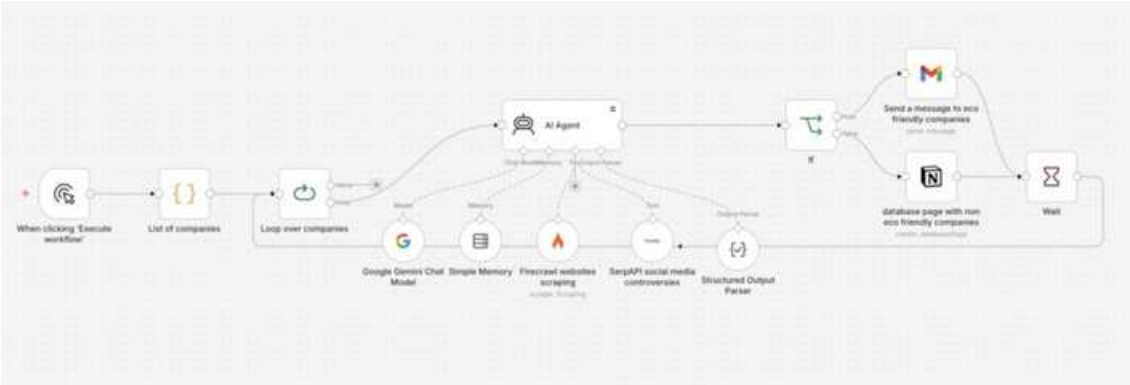

System Architecture and Workflow Implementation

To make the project work in reality, the automated auditing system was built using the n8n visual programming platform. This platform uses a visual canvas where data moves through a series of connected blocks, called nodes instead of writing traditional lines of code. Each node performs a specific task. By linking them together, the project creates an automated pipeline that can search the web, read text using artificial intelligence (Gemini, Chatgpt etc..), and send the results to external apps. Below is a detailed breakdown of every node used in the project,

explaining what it does, how the data flows, and the practical problems I had to fix during the setup.

1. Workflow Trigger Node

Purpose and how it works: This is the starting point of the workflow; it tells the system to start running. This is a manual trigger which allows me to click a button and watch the data move step-by-step, necessary for testing and debugging in the experiment.

Inputs and Outputs: It takes zero input. When clicked, it outputs a basic starting signal to activate the next node in line.

2. List of companies (Target companies)

Purpose and how it works: the workflow needs companies to audit, so we give a list of companies instead of connecting external data.

I chose five well-known brands: Unilever, Patagonia, LEGO, Coca Cola and Shell and configured them in the node with a java code.

```
const companies = [
  { "company_name": "Unilever", "website_url":
"https://www.unilever.com" },
  { "company_name": "Patagonia", "website_url":
"https://www.patagonia.com" },
  { "company_name": "LEGO", "website_url": "https://www.lego.com" },
  { "company_name": "Shell", "website_url": "https://www.shell.com" },
  { "company_name": "Coca-Cola", "website_url": "https://www.coca-
colacompany.com" }
];
return companies.map(company => { return { json: company }; });
```

I**nputs and Outputs:** It uses the code input given and outputs the five companies as a packet text containing each company's name and official website address.

Implementation Hardships: In this step, I faced two main problems. So, at the beginning, I struggled a lot with the official websites of the companies I first chose to deal with. In that list, I had big tech companies such as Apple, Microsoft, etc., but I could not access their websites since their anti-bot firewalls constantly blocked me from allowing me to scrape their pages. At this point, I changed my approach and tried to make my workflow work by using the Wikipedia websites of these companies, but something went wrong again. At last, I changed the list of companies to the actual one, but now I had another problem; the node originally grouped them into one massive block of text. The web scraper crashed because it tried to search all five websites at the exact same time. I had to reformat the data, so the node output five distinct, separate items.

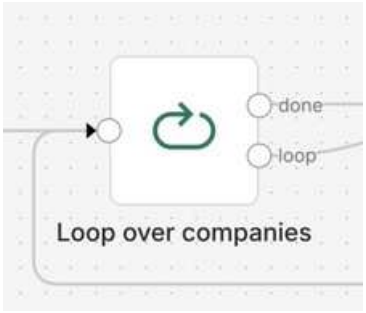

3. Loop over companies

Purpose and how it works: this acts like a “traffic controller” for the whole workflow, it is a structural flow block which regulates systematically how data items pass through the system

Inputs and Outputs: Large Language models and external web APIs have strict limits, and they receive too many requests at the same time. This node collects the list of all target companies and then provides them one by one to the Ai Reasoning Engine. So, this is a serial execution pattern which extracts one company profile at a time though the loop port, preventing data overlap or server crashes. Once downstream operations complete and return through the loop back, it sends the next index. When no items remain, it dumps the collective processing to log out of its done port to shut down the workflow.

Implementation Hardships: Setting up this one was hard since many times the workflow would freeze or loop the same first company repeatedly without proceeding forward. Even though the batch size was fixed to 1, the loop kept having the bug. Through the help of the n8n assistant, this was resolved by enabling a cleanup and aligning the loop ports.

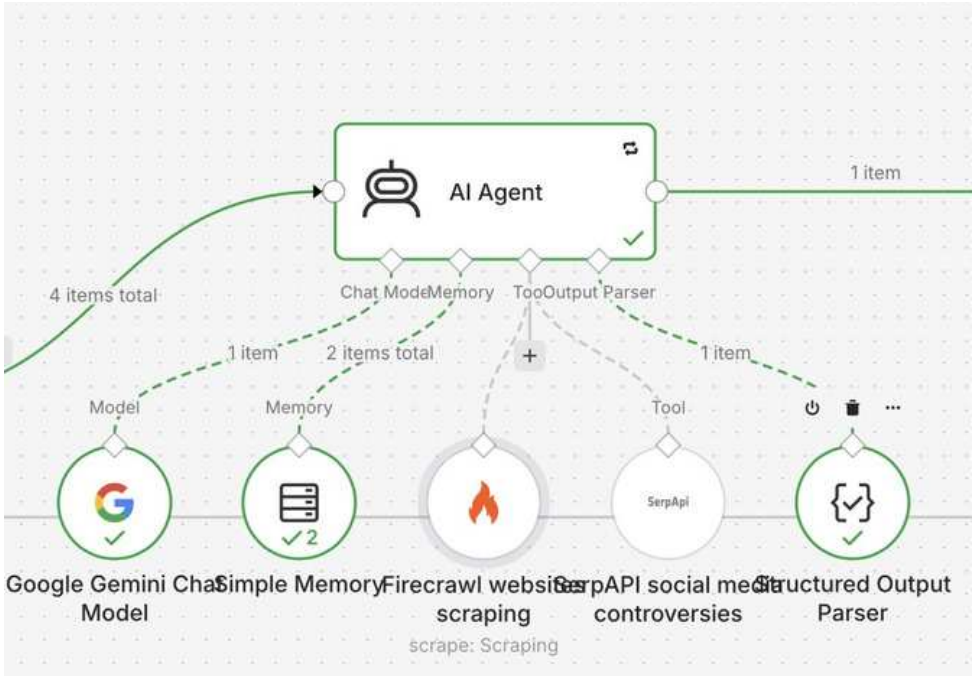

4. Core Intelligence Module (Ai Agent) & connected Tools [8]

1.AI AGENT

2.GOOGLE GEMINI CHAT MODEL

3.SIMPLE MEMORY

4.FIRECRAWL WEBSITE SCRAPING

5.SERPAPI SOCIAL MEDIA SCRAPER

6.STRUCTURED OUTPUT PARSER

• AI AGENT

Use and Purpose: This node is the central brain of the project. Connected with the different tools, it performs a deep analytical evaluation of the data provided. This is where all the main research, scraping, and execution of our workflow takes place, so it is the crucial, decisive step. It also proves if our algorithm was

revealed to be a success or not and if it was accurate enough in the research process and the outcomes provided.

Inputs and Outputs: This is configured as an Advanced AI Agent, connected to Google's Gemini 1.5 Flash model. It allows the model to read and process massive web pages faster without slowing down. Then there is a simple memory node which keeps track of its thoughts during the audit process and an output parser that translates the final decision into reliable, clean data formats. The latter means more explicitly providing readable outputs to the next nodes in the pipeline. The input is a data packet containing the name and the website links of the companies; now our AI Agent runs its internal research loops and then outputs a single structure object containing the final evaluation results with a reasoning text summary.

Implementation Hardships: The biggest problem initially was the Gemini's output; it kept giving long detailed text results making it hard for the next nodes to read the output. This required the addition of the "Structured Output Parser" forcing the AI to rewrite the prompts and output only the required fields. Then another hardship in this step was the API keys since it was my first time using them, so it took me time to learn where and how to create and use them.

- **CONNECTED WEB RESEARCH TOOL (Firecraw [9] scraping)**

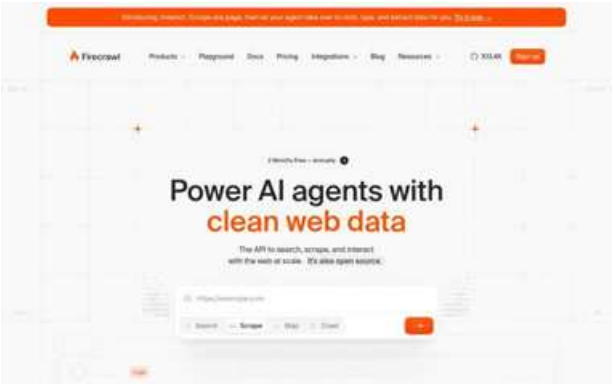

Purpose and Outputs: This tool visits thecompanies' websites, researching their data and bypassing their backgroundsto see what they claim about their sustainability and their environmentally friendly initiatives.

Issues: This step was exactly the cause of my workflow failure. I changed the choice of companies a few times since the anti bot security of some firms blocked Firecrawl from entering. After realizing what was happening, I tried with wikipedia pages of those companies, and it worked instantly. Lastly, I chose companies wisely preventing the automated scraper from being blocked.

- **CONNECTED MEDIA SEARCH TOOL (SerpAPI [10] controversy)**

Purpose: This Tool visits the companies' social media links to ensure their environmental reliability. This happens through accurate research of controversies, news, scandals, lawsuits, pollution issues, or press statements.

Issues: No issues encountered in this node.

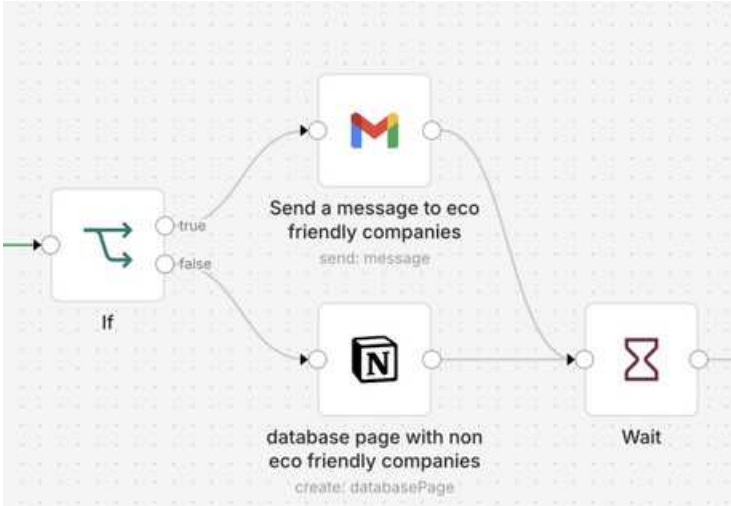

5. IF NODE SPLITTER

Purpose: This node splits the workflow into two different output paths based on the Ai's final decision.
It generates true/false for ecofriendly or not

Inputs and outputs: If the firm turns out to be ecofriendly, then the workflow executes the top right Gmail node; in the other scenario the company data is routed to Notion and here added to a database with reasoning.

MESSAGING NODE GMAIL

In this node,it connects directly to my personal email account. When an eco-friendly company comes through, it automatically drafts an email given:

SUBJECT: University Experiment- Personal Approval

MESSAGE: Hello, this is an automated test for a Digital Economics and Business project. Our AI agent, utilizing Firecrawl and SerpAPI, has

analyzed your digital footprint and identified as a highly sustainable partner. Best regards, Kashish Khurana

NOTION NODE

This creates a permanent record of companies that fail the environmental audit. It connects to a Notion Database Page called ” Sustainability Audit: Rejected”. In this page, it logs the non-approved companies’ data such as company name, the website and the AI Reasoning proving and explaining the decision with some examples of controversies found.

3.3 EMPIRICAL TESTING AND CASE STUDIES

I ran the workflow to see whether my project works or not. The safe side was that, if it worked, it would send emails to my personal account since this was an experiment to test the AI Agent and not a real time approach to these companies. The main goal was to see if the AI Agent could spot the difference between real sustainability and corporate greenwashing. Eventually it worked, I received emails on my account for the approved ecofriendly companies, and the other companies were added to “Sustainability Audit: Rejected”.

3.3.1 APPROVED COMPANIES (GMAIL WORKFLOW)

These are the two examples of companies approved by theAI Agent for networking; no major controversies were found:

- **LEGO:** The reasoning of approval given by the agent is the following as seen in the image: “Our AI agent, utilizing Firecrawl and SerpAPI, has analyzed your digital footprint and identified LEGO as a highly sustainable partner. AI Audit Reasoning: Despite some public debate regarding the transition to sustainable materials, LEGO demonstrates a transparent sustainability strategy. The company has publicly acknowledged challenges in replacing fossil-fuel-based plastics and maintains ongoing,

audited reporting on its environmental initiatives, which indicates a genuine effort rather than deceptive greenwashing”

- **PATAGONIA:** In this case, the agent states: “Our AI agent, utilizing Firecrawl and SerpAPI, has analyzed your digital footprint and identified Patagonia as a highly sustainable partner. AI Audit Reasoning: Patagonia is widely recognized for its industry-leading commitment to environmental sustainability, activism, and ethical supply chain practices. While no company is immune to criticism, their transparent reporting and ownership structure indicate a high level of accountability that warrants an approach.”

3.3.2 REJECTED COMPANIES (NOTION WORKFLOW)

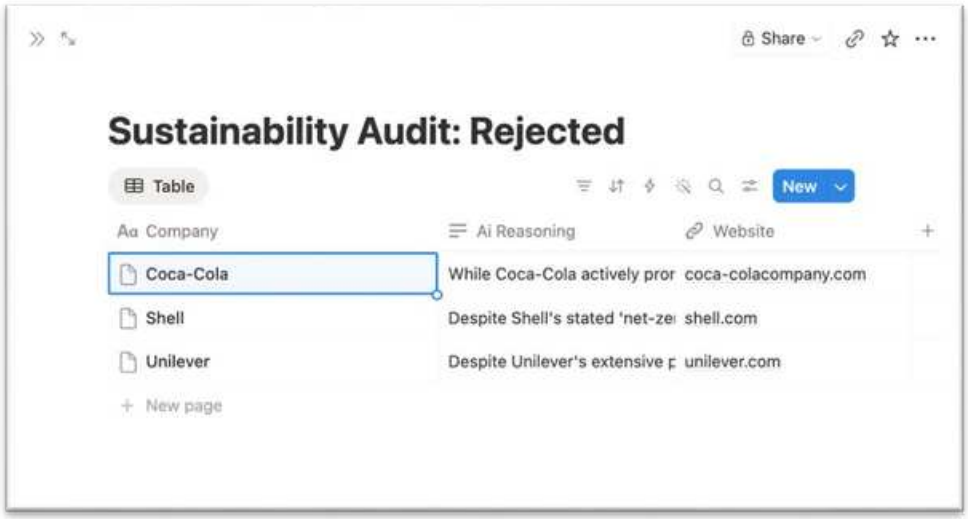

When the Ai Agent caught companies lying, hiding scandals, involved in controversies, it pushed them to the Notion database page with reasoning. These are the three rejected companies with the following reasonings:

- **COCA COLA:** AI REASONING - “While Coca-Cola actively promotes its local economic impact and brand initiatives, it remains a frequent target for criticism regarding environmental sustainability, specifically being labeled as a major global plastic polluter. Recent reports indicate potential internal restructuring impacting sustainability staff, raising concerns about the company's commitment to long-term sustainability goals.”

- **SHELL:** AI REASONING - “Despite Shell's stated 'net-zero by 2050' goal, the company continues to face significant global climate-related litigation and remains deeply invested in fossil fuel extraction (e.g., recent acquisition of ARC Resources), creating a misalignment between their climate rhetoric and core business expansion strategy.”

- **UNILEVER:** AI REASONING - “Despite Unilever's extensive public sustainability messaging regarding climate, plastics, and supply chains, the company has faced repeated regulatory and industry challenges related to its environmental and sustainability claims. Recent rulings (including the September 2024 RCC decision) regarding misleading packaging claims, combined with ongoing expert criticism concerning the efficacy of its plastic reduction strategies and sustainability marketing, indicate significant risks regarding transparency and the substance of its environmental commitments. These findings are incompatible with the rigorous accountability standards required for approval.”

CHAPTER 4

4.1 Project Conclusion and Goal Realization

The goal of the project was to address the issue of creating business networks in an eco-friendly context. It is already a fact that the first step where companies

struggle is approaching other companies, especially more and more companies want to build business networks with eco-friendly partners. Checking whether a business is actually sustainable takes a lot of time, costs too much money, but it is even harder with the information asymmetry surrounding the corporate sustainability world.

The main objective was to build an automated tool using AI that could quickly look past a company's marketing and find out the truth about its real environmental impact. By setting up the automation workflow inside n8n, the project achieved exactly that. The testing proved that an automated system can successfully handle this kind of background research instead of a human doing it manually. Also, it used Firecrawl to scrape what a company says on its official website and SerpAPI to search the internet for news and public scandals.

The AI agent managed to find clear examples of corporate greenwashing.
As the case studies showed, the system correctly flagged companies whose ongoing investments in fossil fuels or heavy history of plastic pollution completely contradicted their green marketing. In the end, this project delivered a working, low-cost tool that helps companies grow their business networks safely, protecting them from the reputational risks of partnering with non-compliant businesses.

4.2 Future Enhancements and Technological Choices

The current workflow proved to be fully operational, but still, it is just a starting prototype. There are several ways for future upgrades which could potentially turn this simple university experiment into a real helpful Tool for companies and an enterprise-ready software solution.

Some of the ideas/advice that could lead to this enhancement are:

- **EXPANDING THE CORPORATE DATABASE**

Themost immediate improvement would be adding more companies to the system. For this experiment, the workflow only used a list of five companies to keep testing simple and stable. But in real business, you wouldn't type them in by hand. Instead, you could connect the first node directly to a company’s CRM system, like Salesforce or HubSpot. That way, the second a new supplier or partner gets added to the system, the workflow would automatically run the background check. This would let a company scan hundreds of businesses at the sametime without any extra work.

- **TESTING ALTERNATIVE AI MODELS**

Right now, the system uses Google Gemini 1.5 Flash because it is fast and can read a massive amount of text without slowing down. In the future, it would be really interesting to send the exact same company data to different AI models at the same time to see how their answers compare.

For example, trying OpenAI’s GPT-4o or Anthropic’s Claude 3.5 Sonnet would help show which AI is stricter with environmental rules or makes fewer mistakes when reading the news. Another cool update would be trying out an open-source model that is specially trained just on international environmental laws.

- **ALTERNATIVE AUTOMATION PLATFORMS**

While it was chosen for its simple visual layout and ease of connecting APIs,

The automation market has plenty of options. Future versions of this project could rebuild the exact same logic on other platforms like Make.com or Zapier Central. For davvero massive companies, relying on a visual canvas with nodes might end up being troppo limiting. The next logical step to scale this up for grandi enterprises would be to drop visual platforms completely and write the whole agent logic from scratch using pure Python code and the LangChain framework. Then, it could be hosted directly on private company servers to keep everything completely secure.

- **INTEGRATING HARD FINANCIAL DATA**

Right now, the AI agent checks website marketing against news articles and social media sentiment. To make the audit much tougher, future workflows could connect directly to financial APIs like Bloomberg or Yahoo Finance. By pulling out a company's actual financial records, the AI could compare how much money the firm spends on sustainable research and development versus how much it spends on advertising. Tracking actual project spending would create a very strict, mathematical test to catch greenwashing at the financial level.

References / Bibliography

[1] LangChain. (2024). LangChain Official Documentation:
A framework for building LLM-powered applications. Retrieved from https://python.langchain.com/

[2] Yao, S., Zhao, J., Yu, D., Du, N., Narasimhan, I., ... & Chen, K. (2023). ReAct: Synergizing Reasoning and Acting in Language Models. International Conference on Learning Representations (ICLR).

[3] Zhuang, Y., et al. (2026). Challenges in ESG Data: Varied Formats and Consistent Extraction Methodologies. Journal of Environmental Analytics.

[4] Visalli, M., et al. (2023). The Evolution of LLMs and Data Extraction in Unstructured ESG Reports. Technology and Sustainability Review.

[5] Moodaley, W., & Telukdarie, A. (2023). Greenwashing Detection: Processing Clean, Structured Text for Environmental Claims. Journal of Digital Business Economics.

[6] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., ... & Fung, P. (2023). Survey of Hallucination in Natural Language Generation. ACM Computing Surveys, 55(12), 1-38.

[7] n8n. (2024). n8n Official Documentation: Advanced Workflow Automation and Human-in-the-Loop (HITL). Retrieved from https://docs.n8n.io/

[8] Google Cloud. (2026). Gemini 1.5 Flash Architecture and API Documentation. Google AI Studio.

[9] Firecrawl. (2026). Firecrawl API Documentation and Web Scraping Guidelines. Retrieved from https://www.firecrawl.dev/

[10] SerpApi. (2026). SerpApi Google Search API Documentation. Retrieved from https://serpapi.com/