Generative AI for Life Cycle Assessment

Lotte A. Thomsen, Ning An and **Massimo Pizzol**
Aalborg University
lotteat@plan.aau.dk, ningan@plan.aau.dk,, massimo@plan.aau.dk

Abstract

Life Cycle Assessment (LCA) is a systematic method for evaluating the environmental impacts of products and services across their entire life cycle, from raw material extraction to end-of-life disposal. A major challenge in current LCA practice is the significant manual effort required to search, match, and interpret information from large, complex, and nested databases.

This project introduces a modular AI-assisted workflow that integrates Large Language Models (LLMs) with Graph Retrieval-Augmented Generation (GraphRAG) to streamline data integration and enable more advanced querying. By representing LCA databases as graphs, the system allows for flexible and semantically rich interactions that go beyond traditional search. The modular design ensures transparency and enables independent validation of each component, supporting trustworthy and adaptable AI integration in LCA workflows.

A first prototype highlights the feasibility of this vision. Initial trials with generative AI show encouraging outcomes, particularly in aligning entities, extracting contextual information, and linking knowledge across datasets. Retrieval performance was evaluated using synthetic queries from 88 Ecoinvent activities, showing strong accuracy for both exact search and top-10 results. Hit@10 exceeded 88% across all test sets, demonstrating robust retrieval performance. These results establish a strong baseline for future LLM-guided applications. Continued development is aimed at testing the boundaries of the approach, strengthening robustness, and ensuring scalability to broader applications.

The envisioned outcome is an LCA AI Assistant capable of guiding practitioners through assessments while also supporting the generation of new data entries needed to fill gaps in existing databases. This approach lays the groundwork for scalable, intelligent, and more automated sustainability analysis. It reduces manual workload, improves consistency, and expands the capabilities of LCA tools.

Keywords: Generative AI, Large Language Models (LLMs), GraphRAG, AI agent, Life Cycle Assessment (LCA), Environmental Impact

1 Introduction

Balancing societal development with environmental sustainability is one of the most urgent challenges of our time, especially in the context of climate change and biodiversity loss. Life Cycle Assessment (LCA) is a widely adopted methodology for evaluating the environmental impacts of products, services, and technologies across their entire life cycle. [International Organization for Standardization [2006]] However, the effectiveness of LCA modelling depends heavily on large, nested databases that vary in scope, format, time coverage, and completeness. These inconsistencies make data integration difficult and the modelling process manual, time-consuming, and prone to error. Missing data and the tedious task of searching and comparing inventory datasets further reduce the efficiency and reliability of LCA workflows.

To address these challenges, this project explores the use of Generative AI (GenAI) to automate and enhance key steps in the LCA process. Specifically, we propose integrating Graph Retrieval-Augmented Generation (GraphRAG) with the Ecoinvent database to improve the inventory search process.

While traditional Retrieval-Augmented Generation (RAG) enhances LLMs by incorporating external knowledge during inference, it typically relies on flat document retrieval. This limits its ability to capture the complex relationships found in structured LCA datasets. GraphRAG overcomes this by combining graph-based retrieval with LLMs, enabling more semantically rich querying and deeper contextual understanding. [Darren Edge [2025]] This approach aims to improve data reliability, reduce modelling effort, and support more robust and scalable sustainability assessments.

Recent advances in Generative AI (GenAI), particularly through Large Language Models (LLMs), offer promising opportunities to automate and enhance key steps in the LCA process. Existing attempts to apply GenAI in this domain remain largely conceptual or rely on naïve Retrieval-Augmented Generation (RAG). [Zhang et al. [2024]] for querying LCA databases [Saad et al. [2023] Wang et al. [2024] Ghose et al. [2019]]. The most advanced prototype to date, developed by Amazon, applies RAG to search LCA datasets with notable accuracy, mimicking human search behavior.[Balaji et al. [2025]]

2 Method

Figure1 illustrates the conceptual method behind an LCA AI assistant. The assistant interacts with structured LCA knowledge, external data sources, and Python-based tools to support key tasks in the assessment process.

The LCA AI Assistant is built around two core innovations: a modular architecture for orchestrating task-specific components, and a graph-based retrieval system that enables deeper semantic understanding of complex inventory data.

Modular AI Architecture
We propose a modular system in which generative models coordinate the flow of information across specialized components. Each module addresses a distinct LCA subtask, allowing flexible orchestration. Some modules are designed for automation, while others support iterative refinement or human-in-the-loop processes. This modular design enables independent validation of each component,

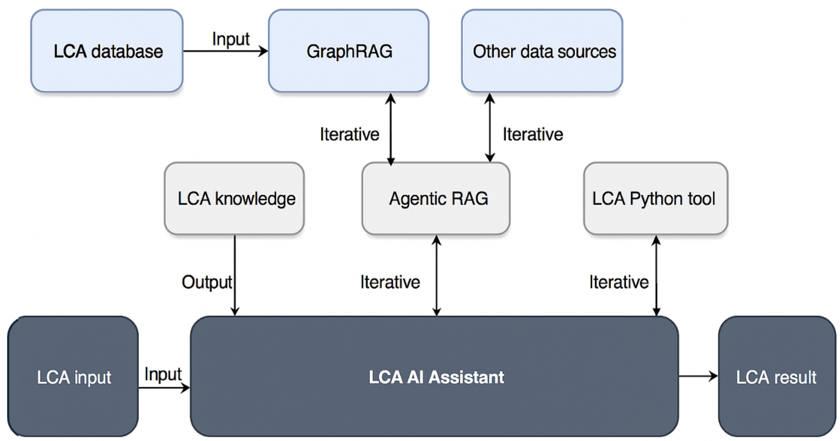

Figure 1: Overview of the LCA AI Assistant architecture.

enhancing transparency and trust in AI-assisted LCA workflows. The current focus lies in structuring data and facilitating user interaction.

Graph-Based Knowledge Integration
To manage the complexity of LCI data, we adopt a graph-based approach that improves interpretability and navigation of nested structures. This is combined with a hybrid retrieval-generation mechanism that integrates external knowledge into the generative process. By preserving relational context, the method enables deeper reasoning and mitigates common limitations of large language models, such as hallucinations and limited domain depth.

3 Validation

As a first step, we evaluated retrieval performance using synthetic queries generated from a sample of 88 Ecoinvent activities. These queries were designed to test semantic matching beyond exact string similarity. The first test set, *Prompt 1*, represents simpler rephrasings, while the second test set, *Prompt 2*, introduces more variation while remaining close to the original meaning. Accuracy was assessed using the Hit@10 metric, which captures how often a correct result appears among the top 10 retrieved items. Table 1 presents the results, showing strong performance for both exact search and Hit@10. Prompt 2 shows lower Top-1 accuracy since its queries are more abstract and linguistically varied, making exact matches harder to rank first even though semantically correct results still appear within the top 10.

Table 1: Retrieval Performance

Query Type	**Top 1 (%)**	**Top 10 (%)**
Exact Search	88%	97%
Prompt 1	83%	89%
Prompt 2	72%	90%

Table 2 includes a comparison to Ecoinvent's native search functionality. Notably, the retrieval system also provides access to biosphere exchanges, which are not available through the standard Ecoinvent search interface. The results for exact search are worse than Ecoinvent because our system is optimized for semantic matching rather than strict text similarity. However, as queries become more abstract (*Prompt 2*), performance improves compared to Ecoinvent, showing the advantage of semantic search for complex or varied phrasing
These results demonstrate that the system performs reliably for both exact and top-10 retrieval, establishing a trustworthy baseline from which more advanced, LLM-guided use cases can be explored.

Table 2: Performance Difference compared to Ecoinvent basic search

Query Type	**Top 1 Difference (%)**	**Top 10 Difference (%)**
Exact Search	-13%	-3%
Prompt 1	8%	-3%
Prompt 2	25%	2%

While the system performs well on both exact and semantic queries, its focus on semantic matching may reduce precision in dataset selection. This trade-off should be monitored, especially in critical assessments. Future work will explore user feedback and confidence scoring, and incorporate real-world case studies to validate practical utility across domains.

4 Conclusions

This project envisions a transformative AI framework that redefines how Life Cycle Assessment (LCA) workflows are conducted. By integrating modular orchestration with graph-based retrieval, the system unlocks seamless access to complex inventory data, minimizes modeling overhead, and fosters intuitive user engagement. It marks a shift from static tools to dynamic, intelligent systems that evolve with user needs.

The current prototype showcases how generative AI can serve as a guide through structured LCA datasets. This is demonstrated using Ecoinvent, but the architecture is inherently flexible, designed to adapt across databases, sectors, and applications. This adaptability opens the door to broader experimentation and cross-domain integration.

At its core, this initiative is part of a larger mission to harness AI for socially beneficial outcomes. It aims to democratize sustainability assessments, making them more inclusive, scalable, and responsive to real-world complexities.

References

B. Balaji, F. Ebrahimi, N. G. G Domingo, V. S. G. Vunnava, A.-Z. Faridee, S. Ramalingam, S. Gupta, A. Wang, H. Gupta, D. Belcastro, K. Axten, J. Hakian, J. Kramer, A. Srinivasan, and Q. Tu. Emission Factor Recommendation for Life Cycle Assessments with Generative AI. *Environ. Sci. Technol.*, page acs.est.4c12667, Mar. 2025. ISSN 0013-936X, 1520-5851. doi: 10.1021/acs.est.4c12667. URL https://pubs.acs.org/doi/10.1021/acs.est.4c12667.

N. C. J. B. A. C. A. M. S. T. D. M. R. O. N. J. L. Darren Edge, Ha Trinh. From local to global: A graph rag approach to query-focused summarization. https://arxiv.org/abs/2404.16130, 2025.

A. Ghose, M. Lissandrini, B. Weidema, and K. Hose. An open source dataset and ontology for product footprinting (awarded best poster - eswc 2019), 06 2019.

International Organization for Standardization. ISO 14040:2006 - environmental management – life cycle assessment – principles and framework. https://www.iso.org/standard/37456.html, 2006.

M. Saad, Y. Zhang, J. Tian, and J. Jia. A graph database for life cycle inventory using Neo4j. *Journal of Cleaner Production*, 393:136344, Mar. 2023. ISSN 09596526. doi: 10.1016/j.jclepro.2023.136344. URL https://linkinghub.elsevier.com/retrieve/pii/S0959652623005024.

H. Wang, M. Zhang, Z. Chen, N. Shang, S. Yao, F. Wen, and J. Zhao. Carbon Footprint Accounting Driven by Large Language Models and Retrieval-augmented Generation, Aug. 2024. URL http://arxiv.org/abs/2408.09713. arXiv:2408.09713.

Y. Zhang, X. Chen, B. Jin, S. Wang, S. Ji, W. Wang, and J. Han. A comprehensive survey of scientific large language models and their applications in scientific discovery, 2024. URL https://arxiv.org/abs/2406.10833.