

Generic Bills-of-Materials-Operations-Sustainability (GBOMOS) for Reconfigurable Green Manufacturing: A Cognitive Intelligent Reasoning Pilot through Retrieval Augmented Generation (RAG)

Catherine K. Jiao

School of Industrial and Systems Engineering
Georgia Institute of Technology, USA
catjiao@gatech.edu

Mulang Song and Yiyun (Cindy) Fei

School of Mechanical Engineering
Georgia Institute of Technology, USA

Shu Wang

TE Connectivity, Middletown
Pennsylvania, USA

Bening Mayanti

School of Technology and Innovations
Industrial Management, University of Vaasa, Finland

Abstract

Green manufacturing extends traditional manufacturing planning from the cost or utilization dimension to incorporating a sustainability assessment. Product variety and planning of process reconfigurations are critical in green manufacturing sustainability assessment because different manufacturing processes have varying environmental impacts. It is imperative to incorporate sustainability assessment cohesively with process planning to optimize resources and select the right processes that ensure minimal waste and efficient material utilization while meeting sustainability goals. This paper proposes a generic bills-of-materials-operations-sustainability (GBOMOS) data model for integrating product, process, and sustainability configuration into a unified framework. Adopting a cognitive intelligent reasoning approach using Large Language Models (LLMs) and knowledge graphs, the paper implements a pilot system of integrated product, process, and sustainability configuration to support cognitive intelligent reasoning for domain applications of Life Cycle Assessment (LCA).

Keywords

Green Manufacturing, Production Reconfiguration, Life Cycle Analysis, Bill of Materials, RAG

1. Introduction

Aiming at high-variety production, reconfigurable green manufacturing integrates flexible, adaptive systems with sustainable practices to efficiently handle diverse product demands while minimizing environmental impact (Milisavljevic-Syed et al., 2024). It enables manufacturers with market flexibility to adjust production lines for diverse

product variations without excessive waste or downtime. Production systems are hinged upon reconfiguration of machines and workflows in order to be adaptable to different product designs. While reducing material consumption and energy usage by tailoring processes to specific needs, it facilitates sustainability compliance to meet evolving environmental regulations and carbon reduction targets (Almusaed et al., 2024).

A generic bill-of-materials (GBOM) is fundamental to reconfigurable manufacturing because it provides a structured data model that supports production planning and control in dynamic manufacturing environments (Kombaya et al., 2021). It ensures consistency across different product configurations, enabling seamless reconfiguration while allowing manufacturers to optimize workflows by dynamically adjusting material requirements based on demand (Dahmani et al., 2022). Towards sustainability integration, green-BOM principles are widely advocated to ensure eco-friendly material selection and minimizing environmental impact (Kurniadi and Ryu, 2021).

To manage product variety and process variation in reconfigurable manufacturing, the integration of a GBOM and its generic bill-of-operations (GBOO) into a unified data structure is essential (Jiao et al., 2000). In addition to ensuring seamless coordination between product design and manufacturing processes, a unified BOMO structure enables dynamic adjustments to both material requirements and operational workflows within a coherence configuration framework (Leng et al., 2024). A generic BOMO (GBOMO) data model ensures seamless coordination between product design and manufacturing processes to achieve mass customization (Jiao et al., 2007).

On the other hand, sustainability assessment is critical to ensure that green manufacturing meets the environmental, economic, and social goals (Vrchota et al., 2020). It helps manufacturers reduce environmental impact by identifying areas for minimizing carbon emissions, energy consumption, and waste. Assessment of green manufacturing sustainability commonly involves life cycle assessment (LCA) for evaluating environmental impact from raw material extraction to product disposal, carbon footprint analysis for measuring greenhouse gas emissions across production processes, or circular economy metrics for evaluating recyclability, reuse potential, and waste reduction strategies (Ijaz et al., 2024).

In practice, product variety and planning of process reconfigurations are critical in green manufacturing sustainability assessment because different manufacturing processes have varying environmental impacts. The selection of process alternatives directly influences sustainability metrics such as energy consumption, material efficiency, and carbon footprint (Ren and Mia, 2025). Different manufacturing methods (e.g., additive vs. subtractive manufacturing) consume varying amounts of energy and materials, affecting sustainability outcomes. Carbon footprint also varies among diverse processes that generate lower or higher emissions. It is thus imperative to incorporate sustainability assessment cohesively with process planning to optimize resources and select the right processes that ensure minimal waste and efficient material utilization while meeting sustainability goals.

In line with the general gist of GBOMO, this paper proposes a general bill-of-sustainability (BOS) structure and extend GBOMO to a generic bills-of-materials-operations-sustainability (GBOMOS) data model for integrating product, process, and sustainability configuration into a unified framework. To demonstrate the potential of the GBOMOS data structure, the paper further implements a pilot system of integrated product, process, and sustainability configuration to support cognitive intelligent reasoning for domain applications of LCA. The pilot adopts a cognitive intelligent reasoning approach using Large Language Models (LLMs) and knowledge graphs. While knowledge graphs are employed for knowledge representation, generative artificial intelligence (GAI) is utilized for knowledge acquisition in the form of head-relation-tail and for case adaptation by developing a retrieval-augmented generation (RAG) app.

2. Emerging Technologies of Cognitive Intelligence, GAI, and RAG

Cognitive intelligence, a fusion of AI, neuroscience, and psychology, is advancing rapidly with neuro-symbolic AI, deep learning, and reinforcement learning (Kowalczyk and Czubenko, 2023). Cognitive intelligence allows systems to analyze complex problems, understand context, and generate solutions that are nuanced and well-informed by domain-specific knowledge. Systems with cognitive intelligence can adapt to new information and changing circumstances, providing more flexible and resilient decision-making capabilities. Advanced cognitive systems can identify patterns and trends within specific domains that might be missed by human analysts, leading to more insightful decisions (Otero et al., 2022). By leveraging historical data and domain knowledge, cognitive intelligence can make accurate predictions about future events, helping in proactive decision-making and thus, enhancing the capability to make well-informed, timely, and effective decisions, which is crucial in today's data-driven and fast-paced manufacturing environments.

The development of natural language processing (NLP) techniques has made it plausible to construct databases with LLMs. Current research often explores the relationship between SQL model-based databases and LLMs, investigating the potential of using LLMs as APIs to enhance database functionalities. BIRD, a benchmark designed to evaluate text-to-SQL models on realistic, large-scale databases, underscores the challenges posed by noisy data, external knowledge grounding, and SQL efficiency (Li et al., 2023). RAG is an advanced technique that enhances information retrieval and natural language processing capabilities, closely linked to the development of LLMs. RAG models combine sequence-to-sequence (seq2seq) transformers with a dense vector index of documents, accessed via a neural retriever, thereby improving the accuracy and factuality of generated text (Lewis et al., 2020). In the context of databases, RAG is employed to generate specific queries from text, facilitating the search process within databases by accurately translating plain text questions into the appropriate query languages.

Knowledge graphs are used to construct and describe large entity-relation networks and can represent knowledge bases in the form of graph structure (Chen et al., 2020), where entity types are represented with nodes and relations with edges. Common reasoning tasks for knowledge graphs include node classification, link prediction, graph classification, clustering, and predictive queries. Node classification and link prediction focus on a single entity and relationship in an incomplete knowledge graph while graph classification is at the graph level.

Neo4j, a graph database management system, is specifically designed to store, manage, and query data in graph structures. The Cypher query language used by Neo4j can be enhanced by LLMs, such as GPT-4, to interpret and process complex natural language queries (Neo4.com). This integration offers significant potential for researchers to construct domain-specific knowledge databases and utilize LLMs for data retrieval. The application of RAG in this context is particularly important, as it helps generate appropriate queries for effective information retrieval. For instance, the Neo4j website demonstrates a movie information retrieval example using RAG to produce suitable queries (<https://graphacademy.neo4j.com>).

3. Generic Bills-of-Materials-Operations-Sustainability (GBOMOS)

The GBOMOS model is anchored to a generic representation principle that is proved to be an effective means to describe a large number of variants with minimal data redundancy (Jiao et al., 2007). An item is generic in the sense that it represents a set of similar items (namely variants) of the same type (namely a family). The item may be a component of the product (an end product, a subassembly, an intermediate part, or raw material) or an operation of the process (a machining or assembly operation).

The generic representation of variety differentiation is illustrated as the GBOM schematics in Figure 1(b). Instead of using part numbers, the identification of individual variants from a generic item is based on variety parameters and their instances (i.e., sets of parameter values). Such an indirect identification entails a type of class-member relationships (exhibiting a meta-structure) between a family and its variants (Jiao et al., 2000). In this way, generic variety representation facilitates the specification of feasible variations of the items (products, processes, or sustainability assessments) with respect to optional and alternative values of variety parameters.

3.1 GBOMOS

Figure 1 shows the operations principle of a GBOMOS structure to support planning of green manufacturing. Product data can be represented by a BOM that is used for an end product to state raw materials and intermediate parts or subassemblies required for making the product. On the other hand, production information is concerned with how a product is produced, that is, the specification of operations sequences to be performed at corresponding work centers along with related resources such as machines, labors, tools, fixtures, and setups. Similar to describing a product structure using a BOM, an operations routing can be constructed to represent the process structure for a given product.

Corresponding to a GBOM that encompasses diverse product variants, process variety is characterized by a generic process structure that consolidates related production processes into standard routings, forming a GBOO. These standard routings serve as the foundation for managing process variations resulting from product variety, leading to distinct assessments of sustainability impacts. These impacts arise from differences in material requirements and the varying resource consumption of specific processes necessary to produce individual product variants. Similar to the GBOM and GBOO structures, differentiation in reconfigurable sustainability assessment can be represented through a generic bill-of-sustainability (GBOS). Together, these three types of variety representation can be integrated into a unified and cohesive GBOMOS framework, as illustrated in Figure 1(a).

3.2 Reconfigurable Green Manufacturing Planning

The relationships between the product structure (i.e., BOMs) and the process structure (i.e., routings) are embodied in the materials required by particular production operations. The link between BOM and routing data can be established by specifying each component material in the BOM as required by the relevant operation of the routing for making its parent product. Likewise, these variety propagation links are extended to items of sustainability, e.g., a LCA spreadsheet as shown in Figure 1(e). As conceptually described from Steps (i) to (iv) in Figure 1, the GBOMOS distinguishes the common configuration structure of green manufacturing, by which different product, process, and sustainability variants are derived coherently through instantiation of variety parameters.

While the GBOM directly associates each component material with its parent product, the GBOMOS links a component material to the relevant operation in the GBOO responsible for producing its parent component. For each manufactured end product or intermediate product, a single-level GBOO can be derived by specifying the sequence of operations required for production, along with the necessary materials and resources—categorized by work centers, cycle times, and setups—for each operation. A multi-level GBOO is then composed by linking the single-level GBOOs of lower-level intermediate parts through the operations that manufacture or assemble require them. In turn, all sustainability assessment items—relating to material consumption and manufacturing resources required for a specific customer order—can be derived, as shown in Figures 1(d) and 1(e).

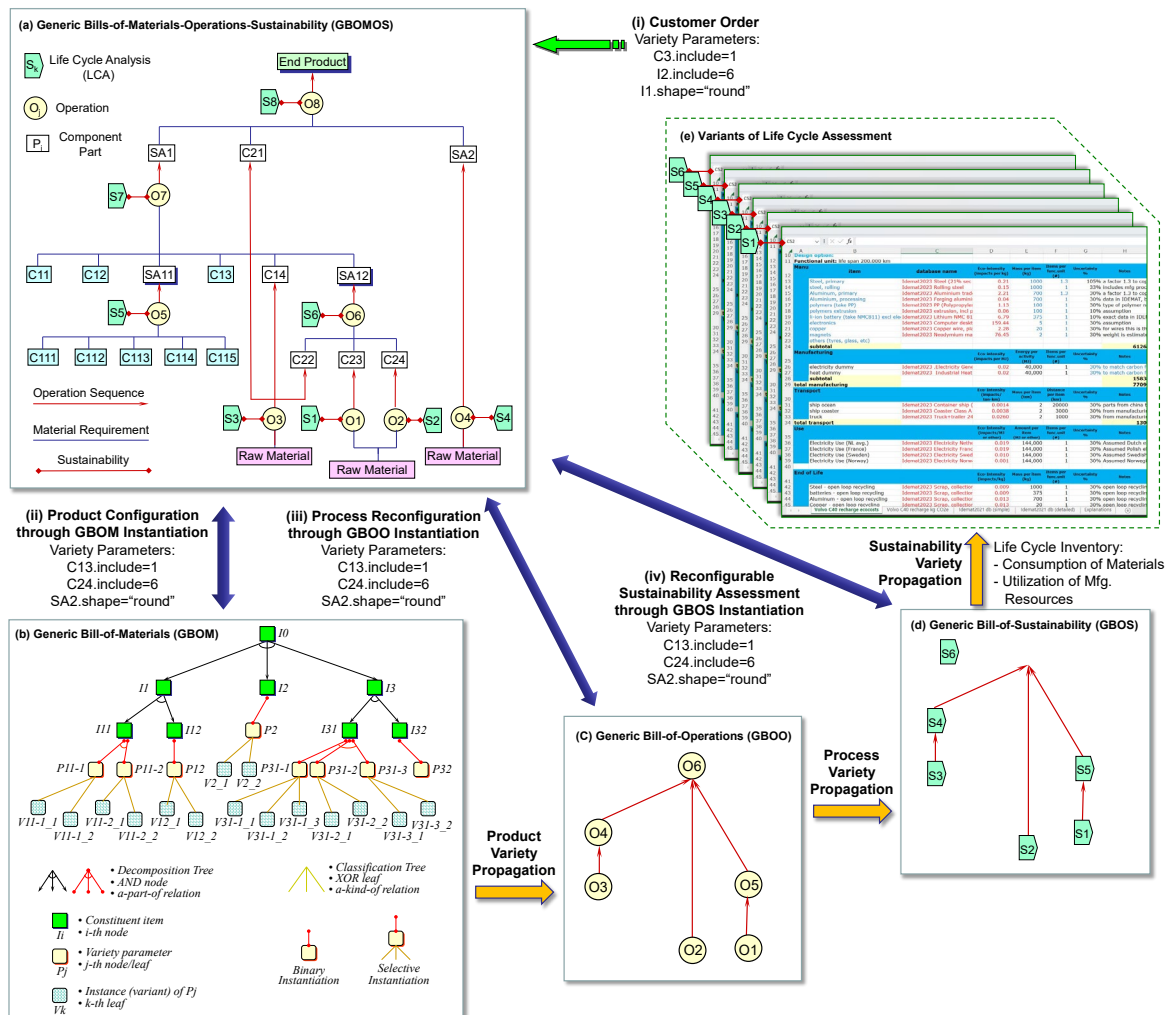


Figure 1. GBOMOS for reconfigurable sustainability assessment

4. Case Study of LED Luminaire Green Manufacturing

LED luminaire green manufacturing is crucial for sustainability due to its significant environmental and economic benefits. LED luminaires consume up to 75% less energy than incandescent bulbs and last 25 times longer, reducing electricity demand and lowering carbon emissions. Switching to LED lighting helps cut greenhouse gas emissions, supporting global climate goals (Gao et al., 2025). Sustainability planning of LED production increasingly contributes to incorporating recyclable materials and low impact manufacturing processes, further reducing environmental harm (Wang et al., 2020).

4.1 LCA Framework

LCA is a widely recognized methodology used to evaluate the environmental impacts associated with all stages of a product's life cycle from raw material extraction to disposal. By systematically assessing each stage, LCA helps identify opportunities for reducing environmental burdens such as resource consumption and emissions. This case study follows the internationally recognized ISO 14040 and ISO 14044 standards for LCA, ensuring a consistent and scientifically rigorous approach. Additionally, ISO 14025 offers principles and procedures for Type III environmental declarations (EPD), ensuring the accuracy and transparency of environmental statements.

The LCA in the case study is structured around four phases: (1) Goal and Scope Definition focusing on comparing environmental impacts across three manufacturing strategies: whole-lamp assembly, remade of substrates, and reuse of light source modules; (2) Life Cycle Inventory (LCI) involving the collection of data on material composition, energy consumption, emissions during manufacturing, and other environmental impacts across the product's lifecycle; (3) Life Cycle Impact Assessment (LCIA) using the global warming potential (GWP) as the primary impact category, with results quantified in CO₂ equivalent; and (4) Interpretation reflecting the LCA outcomes with the goal defined in the first phase.

The LCA calculations are carried out using LCAfE software (Sphera, 2025), a professional tool recognized for its ability to handle complex life cycle assessments. Integrated with the GaBi database that contains comprehensive life cycle inventory data for various materials and processes, LCAfE enables precise calculation of environmental loads associated with raw material acquisition, manufacturing processes, transportation, and end-of-life treatments. Using LCAfE in combination with GaBi and the Product Environmental Profile (PEP) framework, the case study develops a detailed and scientifically sound assessment of the environmental impacts of the luminaire's manufacturing lifecycle, as shown in Figure 2.

4.2 Product Environmental Profile and Guidelines

To ensure that the environmental impact assessment is consistent and comparable, this study follows the PEP framework, a specific type of EPD covering electrical or electronic products; in particular, the Product Specific Rules (PSR) for luminaires provided by the PEP Association (PEP, 2022). The PSR framework outlines the specific requirements for calculating and reporting the environmental impacts of lighting products, which ensures that the data used in the LCA is accurate and standardized. This adherence to PEP standards provides a robust basis for comparing the environmental impacts of different luminaire designs.

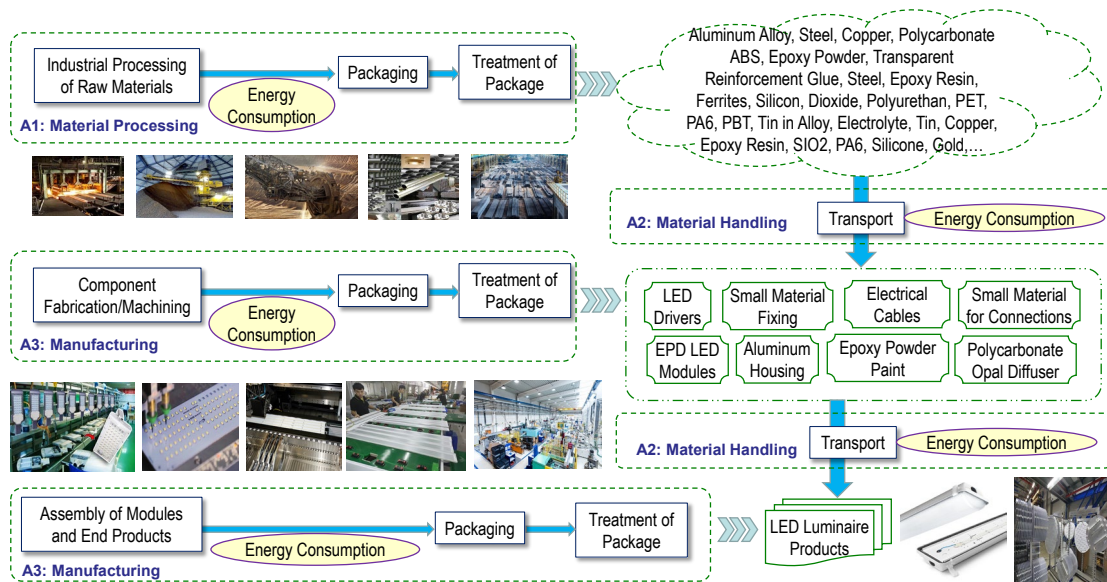


Figure 2. The life cycle of LED luminaire manufacturing

4.3 Environmental Impact Analysis of Polycarbonate Granulate Production

The assessment showed that the environmental hotspots were polycarbonate granulate production and polycarbonate component injection molding, hence further analysis focuses on these two critical production stages, which are significant contributors to the product's overall environmental footprint. By quantifying the Global Warming Potential (GWP) at each stage, the study provides actionable insights into opportunities for reducing the environmental impact of luminaire production.

The production of polycarbonate granulates is an essential step in manufacturing luminaire components. This stage involves chemical synthesis using raw materials such as Bisphenol A and phosgene through the phosgenation process. Despite its economic efficiency, the process is energy-intensive and contributes significantly to carbon emissions due to high energy consumption and the use of fossil fuels in many production facilities. The system boundary for this stage includes (1) Energy consumption - Electricity from both renewable and non-renewable sources is modeled to reflect regional energy mixes and transmission losses; (2) Transportation emissions - Emissions associated with transporting raw materials, including maritime and inland shipping, trucking, and pipelines; and (3) Process emissions - Direct emissions from phosgenation and related manufacturing steps.

The granulates are then processed into luminaire components such as housings and diffusers through injection molding. This stage primarily consumes electricity and water while generating plastic waste. Excess plastic is either recycled within the system or sent for energy recovery through incineration. Key considerations within the system boundary include (1) Electricity use - Modeled based on production conditions, including regional energy grid mixes; (2) Water consumption - Utilized for cooling and operational purposes, with recycling measures factored into the analysis; and (3) Plastic waste management - Excess material from the injection molding process is carefully managed, with non-recyclable fractions incinerated to recover energy and reduce net waste. The production of polycarbonate components injection molding represents a moderate contribution to the total GWP, primarily driven by electricity consumption and plastic waste management. Adopting energy-efficient molding technologies and improving material recovery rates could further reduce the environmental impacts.

Table 1. LCAfE software output of GWP assessment for LED luminaire throughout its manufacturing life cycle

LED Luminaire LCA	GWP Total (kg CO ₂ eq.)
1. Luminaire Fixture Components	1063.56
1.1. Control gear	1.89
1.2. Luminaire structure	7.71
1.2.1 HDPE Extrusion	0.01
1.2.2 Polyamide 6 injection	0.06
1.2.3 Polycarbonate Injection part	3.88
1.2.4 Polypropylene Injection	0.03
1.2.5 POM Injection	0.00
1.2.6 Silicone Rubber Extrusion	0.35
1.2.7 Stainless steel sheet part	2.11
1.2.8 Transport (PEP PCR)	0.04
1.2.9 Transport (PEP PCR)	0.06
1.2.10 Transport (PEP PCR)	0.02
1.2.11 Wires & Cables	0.73
1.2.12 GLO: Connector Steck Klemme Leiste (3 g, 2 pins) based on parametric plan model	0.12
1.2.13 GLO: EMS Shielding Sphera	0.25
1.2.14 RER: Fixing material screws galvanized (EN15804 A1-A3)	0.04
1.3. Packaging	-0.26
1.4. Light Source	4.64
1.5. Upstream packaging (PEP PSR)	-0.08
2. Component Manufacturing	0.88
3. Transport	0.61
4. Downstream Packaging	0.51
4.1. Module Packaging	0.13
4.2. Treatment of Packages	-0.54
5. Assembly	1047.90
5.1. Sub-assembly with energy recovery	0.11
5.2. Final Assembly with energy recovery	0.09
6. Product Shipment	0.00
Carbon Correction - Landfill <u-so>	0.00
RER: Electricity grid mix	-0.02
RER: Thermal energy from natural gas	-0.03

4.4 Assessment of Global Warming Potential in Production of Entire Lighting Fixtures

With focus on the polycarbonate components, the case study applies LCAfE software to comprehensively assess the GWP of the entire lighting fixture production. This assessment covers the full manufacturing lifecycle of all components of the lighting fixture, including raw material acquisition, manufacturing, assembly, and transportation, to ensure the comprehensiveness and accuracy of the sustainability assessment results. GWP is measured in kilogram of CO₂ equivalent to quantify the potential of a greenhouse gas that contributes to global warming compared to carbon dioxide over a specified timescale, as shown in Table 1.

4.5 Assessment of Total Carbon Emission in Green Manufacturing

A thorough analysis of the carbon footprint resulting from the three green manufacturing strategies is further conducted. As shown in Table 2, the first strategy, whole-lamp assembly, generates carbon footprint including control gear (1.1), luminaire structure (1.2), light source (1.4), and the manufacturing process (2) for each LED luminaire product. The second strategy involves remade of mounting substrate yields carbon footprint from control gear (1.1), the stainless-steel sheet part of the luminaire structure (1.2.7), transportation for the stainless-steel sheet part (1.2.8), and the light source (1.4). The strategy of reusing light source modules produces carbon footprint solely from the light source (1.4).

Table 2. Carbon footprint comparison of reconfigurable green manufacturing

Green Mfg. Strategy	#1: Whole-Lamp Assembly		#2: Remade of Mounting Substrates			#3: Reuse of Light Source Modules		
	1.1. Control gear	1.89	1.1. Control gear	1.89	1.1. Control gear	1.89	1.1. Control gear	1.89
Components and their Carbon Emissions (kg CO2 eq.)	1.2. Luminaire structure	7.71	1.2. Luminaire structure	7.71	1.2.7 Luminaire structure- Stainless steel sheet part <LC>	2.11	1.2. Luminaire structure	7.71
	1.4. Light Source	4.64	1.4. Light Source	4.64	1.2.8 Luminaire structure- Transport (PEP PCR) <LC>	0.04	1.4. Light Source	4.64
	2. Manufacturing	0.88	2. Manufacturing	0.88	1.4. Light Source	4.64	2. Manufacturing	0.88
	Total Carbon Emission (kg CO2 eq.)	15.12		15.12		8.68		15.12
								4.64

5. Generative GBOMOS Reconfiguration Using LLMs

Akin to most engineering practices, planning of reconfigurable green manufacturing is a knowledge-intensive creative process that involves tremendous domain-specific contexts and decisions. The ability to discover meaningful knowledge from large amounts of LCA information is important for process reconfiguration while minimizing overall environmental impacts. We approach domain problem context awareness from a perspective of knowledge acquisition, representation, and reasoning, which is the typical problem-solving mindset of AI and intelligent systems. Specifically, we propose to implement the GBOMOS-based green manufacturing planning as a cognitive intelligent configurator by orchestrating efficient domain context-aware symbolic reasoning over knowledge graph embeddings using LLMs. Through automated knowledge graph construction, domain context-aware intelligent reasoning is achieved as a unified RAG application via differentiable messages passing over the graph structure and embeddings. This symbiosis enables generative intelligent configuration to leverage both statistical learning and symbolic logic, i.e., combining the strengths of neural networks and structured knowledge representation.

To unify LLMs and knowledge graphs within a coherent framework of domain context-aware cognitive intelligent reasoning, we propose to create a GenAI agent for automated knowledge graph construction. The general gist is to perform our RAG application on knowledge graph embeddings as a chatbot. For testing the feasibility and potential of such an GenAI agent, we use OpenAI's GPT-3.5-turbo (OpenAI.com) as a local LLM and configure a Neo4j environment of graph database and analytics (Neo4j.com). For domain-specific knowledge, we utilize product descriptions and process routing data, along with their LCA information in the LED luminaire case study. There are a lot of ontologies used on the Web to describe products. The most popular ones include the Good Relations Ontology (2024) and the Product Types Ontology (productontology.org). Both ontologies extend the Schema Ontology (Schema.org), which is a collaborative community activity with a mission to create, maintain and promote schemas for structured data on the Internet. For this task, we use the Schema.org definitions for products and related process and LCA information, including their relations, to extract triplets from product descriptions.

A GenAI agent system design for GBOMOS reconfiguration is proposed as shown in Figure 3. We implement a pilot system in Python, for which we first install and import the required libraries, i.e., to read the LED luminaire case dataset as a pandas dataframe, representing the specifications of the product and process LCA that we are going to prompt GPT-3.5-turbo to extract entities and relations from. Next is information extraction, i.e., to instruct GPT-3.5-turbo to extract entities and relations from the provided product/process/LCA specifications and return the result as an array of JSON objects. The JSON objects must contain the following keys: 'head', 'head_type', 'relation', 'tail', and 'tail_type'. Using these entity types and relation types to prompt GPT-3.5-turbo for entity-relation extraction, we then map these entities and relations to the corresponding entities and relations from the Schema.org ontology. The keys in the mapping represent the entity and relation types provided to GPT-3.5-turbo, and the values represent the URLs of the objects and properties from Schema.org.

In order to perform information extraction using GPT-3.5-turbo, we create an OpenAI client, and, using the chat completions API, we generate the output array of JSON objects for each identified relation from the raw product specification. Then we proceed to prompt engineering. The system_prompt variable contains the instructions guiding GPT-3.5-turbo to extract entities and relations from the raw text, and return the result in the form of arrays of JSON objects, each having the keys: 'head', 'head_type', 'relation', 'tail', and 'tail_type'. The extract_information function for each specification in the dataset creates a list of all extracted triplets, representing the generated knowledge graph.

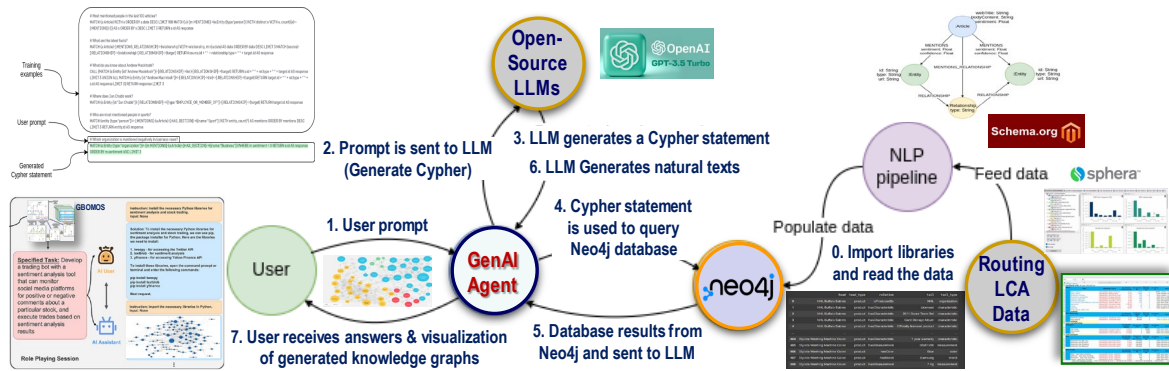


Figure 3. System architecture of cognitive intelligent reasoning through a domain knowledge graph GenAI agent for GBOMOS reconfiguration

6. Domain Context Modeling as Knowledge Graphs Using Neo4j

LLMs rely on generative algorithms to produce data based on the knowledge they have acquired. The quality of the database from which these models learn directly impacts the relevance and accuracy of their responses. Thus, creating a high-quality knowledge database is crucial. Our pilot implementation applies the Neo4j database tool for creating a domain knowledge database. Its ability of expressing relationships between domain nodes is vital for analyzing and constructing embeddings, which are necessary for representing connections between different knowledge problem contexts. Neo4j is also capable to integrate with various LLMs to simplify the management and building of functions within the database for efficient retrieval and representation. The built-in functions provided by Neo4j further facilitate the modification and application of algorithms such as embeddings and retrieval, enhancing their ease of use.

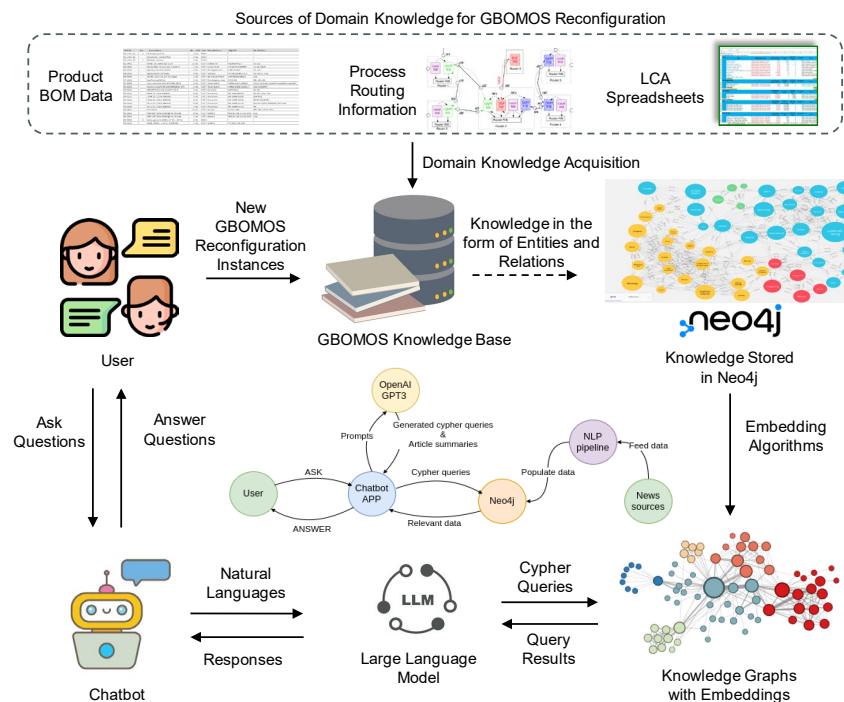


Figure 4. Graph database construction of GBOMOS reconfiguration domain knowledge using Neo4j

Figure 4 depicts the basic process flow of building a domain knowledge database. The initial data is generated from the input of domain experts with specific expertise in an industrial domain. This information is organized as nodes (representing knowledge regions) and links (representing the connections between these nodes). The graph database

empowers two primary functions: (a) visualizing the nodes and links of domain knowledge, and (b) embedding all the necessary information (such as vectors or Cypher queries) for these nodes and links using embedding algorithms. Generative algorithms can then access this data to train models for future predictions or to generate query responses to questions posed by users. Therefore, the graph database is crucial for decision support of domain experts, users, and generative algorithm agents. The agents not only retrieve data from the database but also update it with newly generated data based on expert inquiries. This step ensures that the database continuously evolves and becomes more specialized, maintaining its relevance and accuracy over time.

A knowledge graph database of GBOMOS is constructed using Neo4j, as shown in Figure 5. A knowledge graph is used to represent relationships between different knowledge bases. In a knowledge graph, nodes represent the knowledge bases while edges represent the relationships between them. These relationships can be one-to-one, one-to-many, or many-to-one, effectively visualizing the connections and logic between them. Traditional methods of constructing knowledge graphs often face challenges with data management. Neo4j stands out by addressing these challenges, providing robust database management and representing relationships between nodes in a graphical format.

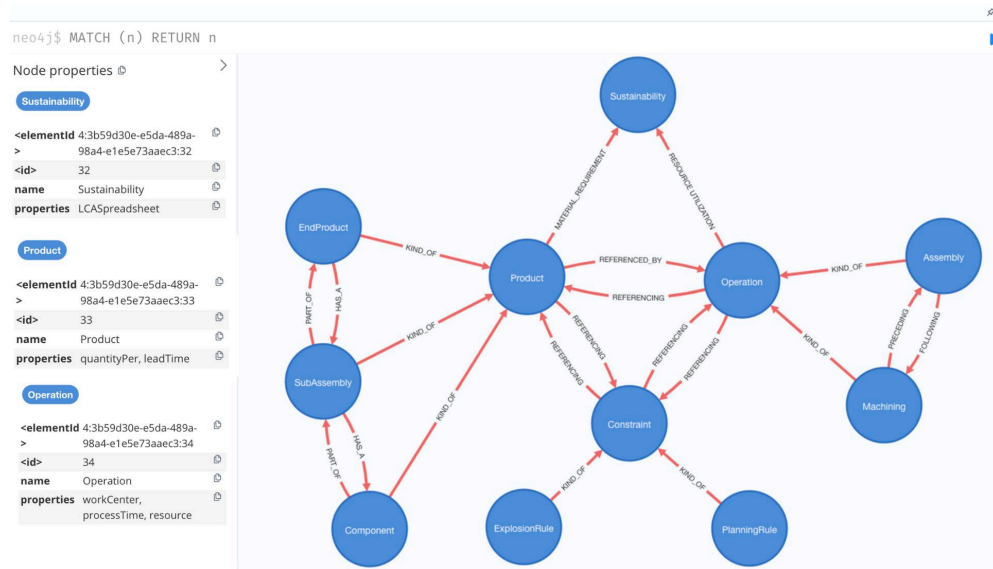


Figure 5. Graph database construction of GBOMOS reconfiguration domain knowledge using Neo4j

7. Cognitive Intelligent Reasoning through Retrieval Augmented Generation

The RAG system plays a crucial role in generating appropriate responses directly to users, ensuring that the answers generated are neither arbitrary nor fake. The construction of a RAG system typically consists of three main processes: understanding queries, retrieving information, and generating human-understandable responses. In this paper, the first step of understanding the query is addressed by using Chat-GPT APIs (e.g., GPT-3.5-instruct) to convert user queries into a format that the LLM can comprehend. The second step involves retrieving information from the database, which is where generative algorithms come into play. Langchain is a widely used tool that functions as a transformation chain to retrieve text representations for nodes. It provides flexibility to switch between different LLMs with minimal code changes. Using Langchain, users can build their own RAG systems specifically designed to answer domain-specific questions.

The first step in constructing a RAG system is to define the chain. This chain connects LLMs with various data sources and third-party APIs. The second step is to create a chat model designed for conversation. This chat model includes system messages, human messages, and AI responses, and it also requires grounding and memory history storage to refine the interactions. The third step is to create an agent that encapsulates the model and provides several tools to generate responses. These tools have access to additional data sources, enabling the RAG system to focus more on answering domain-specific questions. The final step is to define these tools. Several options can be chosen, such as vector embeddings and Cypher queries. For the vector embedding tool, text is embedded using an embedding model, and the database stores these embedding vectors as attributes of the nodes. For the Cypher query tool, a query is

generated so that the database can read, which is then used to search for the answer and wrap it in human-understandable language.

Figure 6 illustrates the principle and procedure for building a RAG system tailored to specific domain knowledge. It depicts how these components interact to provide effective communication and information retrieval for users. Construction of the knowledge graph begins at the data layer, where all relationships and domain knowledge require initial human input to establish a foundational structure. At the information layer, the database connects with LLMs to enrich each node with specific information and support model understanding and training. Finally, at the knowledge layer, the database gains the capability to generate new information based on the trained domain knowledge database, facilitated by generative algorithms. Figure 7 shows a Neo4j knowledge graph database constructed for GBOMOS

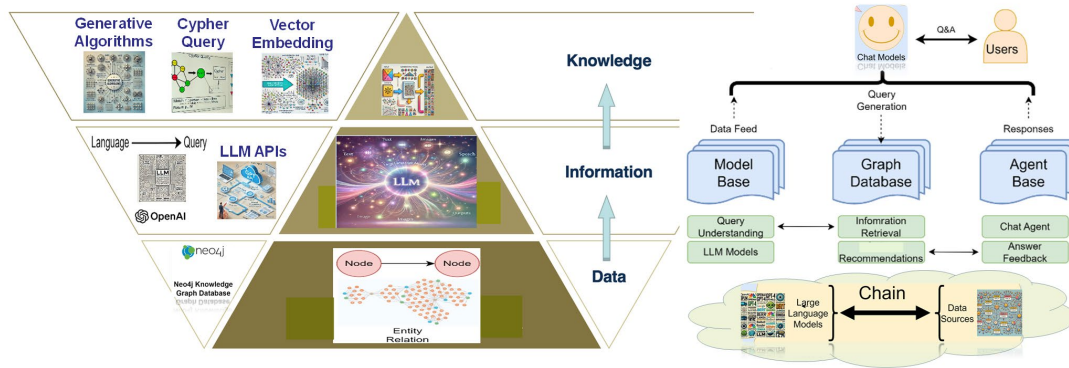


Figure 6. RAG system architecture and workflow

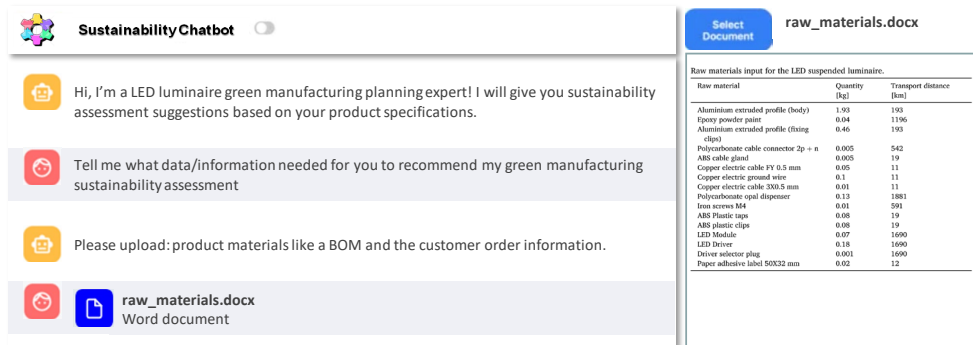


Figure 7. User interface of the chatbot for prompting input of material requirements for GBOMOS reconfiguration

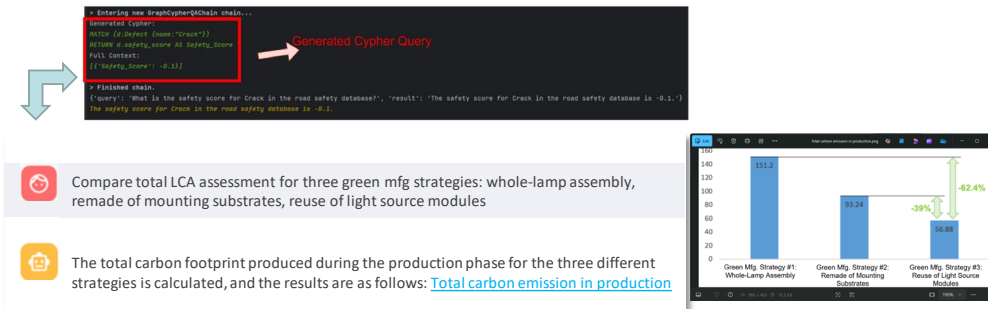


Figure 8. Graph database retrieval and reasoning by creating cypher query

After setting up the database and RAG system, a chatbot interface is essential for providing user feedback. The chatbot retrieves product and process configurations and reasons about sustainability assessment through the knowledge graph database. It offers recommendations on users' inquiry about sustainability assessment or justifications on green manufacturing planning. Figure 7 shows an example of user prompts to provide initial input of product BOM data in relation to specific customer orders. Then, the chatbot automatically generates recommendations on green manufacturing planning and presents it to the users. The Cypher query tool is used to generate queries that retrieve data from the database. Figure 8 shows an example of the query generation process using the cypher query above.

8. Conclusions

The case study and pilot implementation indicate that the feasibility and potential of the proposed GBOMOS data model for integrating product, process and sustainability configuration into a unified framework. It offers several benefits for reconfigurable green manufacturing, including (1) Holistic sustainability integration by aligning material selection, operational processes, and environmental impact within a single structured model, (2) Enhanced flexibility for reconfigurable manufacturing with support of dynamic product and process adaptation, enabling rapid reconfiguration for high-variety production, (3) Improved optimized resource utilization with reduction of material waste and energy consumption by ensuring sustainability-driven decision-making, and (4) reconfigurable LCA to facilitate real-time sustainability tracking, helping manufacturers meet carbon footprint reduction goals. Moreover, enabling cognitive intelligence for domain knowledge-intensive decision making is crucial for enhanced problem solving in reconfigurable green manufacturing. Cognitive intelligence allows systems to analyze complex problems, understand context, and generate solutions that are nuanced and well-informed by domain-specific knowledge.

References

- Almusaed, A., Yitmen, I., Myhren, J. A. and Almssad, A., Assessing the impact of recycled building materials on environmental sustainability and energy efficiency: a comprehensive framework for reducing greenhouse gas emissions, *Buildings*, vol. 14, no. 6, pp. 1566, 2024.
- Chen, X., Jia, S. and Xiang, Y., A review: Knowledge reasoning over knowledge graph, *Expert Systems with Applications*, vol. 141, pp. 112948, 2020.
- Dahmani, A., Benyoucef, L. and Mercantini, J. M., Toward sustainable reconfigurable manufacturing systems (SRMS): past, present, and future, *Procedia Computer Science*, vol. 200, pp. 1605-1614, 2022.
- Gao, W., Zeng, X., Yang, H., Chen, F., Pan, A., Yang, F., Wang, Y., Li, H., Ren, Z., Zhang, G. and Sun, Z., Green manufacturing evaluation of light-emitting diodes lamps production: Resource and environmental assessment, *Renewable and Sustainable Energy Reviews*, vol. 211, pp. 115317, 2025.
- Good Relations Ontology, Available: <http://wiki.goodrelations-vocabulary.org/Cookbook/Schema.org>, Accessed on June 18, 2025.
- Ijaz, T., Qazalbash, A. H., Razzaq, A. A., Rafique, M. Z., Khan, M. A. and Jamshaid, S. H., Development of green manufacturing implementation framework based on life-cycle assessment, *Sustainable Environment Research*, vol. 34, no. 1, pp. 23, 2024.
- Jiao, J., Tseng, M. M., Ma, Q. and Zou, Y., Generic bill-of-materials-and-operations for high-variety production management, *Concurrent Engineering*, vol. 8, no. 4, pp. 297-321, 2000.
- Jiao, J., Zhang, L. and Pokharel, S., Process platform planning for variety coordination from design to production in mass customization manufacturing, *IEEE Transactions on Engineering Management*, vol. 54, no. 1, pp. 112-129, 2007.
- Kombaya, J. V., Hamani, N. and Kermad, L., Modeling and Configuration Management of Reconfigurable Manufacturing Systems (RMS), *Proceedings of the 11th Annual International Conference on Industrial Engineering and Operations Management*, pp. 2348, Singapore, March 7, 2021.
- Kowalczyk, Z. and Czubenko, M., Cognitive motivations and foundations for building intelligent decision-making systems, *Artificial Intelligence Review*, vol. 56, no. 4, pp. 3445-3472, 2023.
- Kurniadi, K. A. and Ryu, K., Development of multi-disciplinary green-BOM to maintain sustainability in reconfigurable manufacturing systems, *Sustainability*, vol. 13, no. 17, pp. 9533, 2021.
- Leng, J., Guo, J., Xie, J., Zhou, X., Liu, A., Gu, X., Mourtzis, D., Qi, Q., Liu, Q., Shen, W. and Wang, L., Review of manufacturing system design in the interplay of Industry 4.0 and Industry 5.0 (Part I): Design thinking and modeling methods, *Journal of Manufacturing Systems*, vol. 76, pp. 158-187, 2024.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W. T., Rocktäschel, T., Riedel, S., and Kiela, D., Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, *Advances in Neural Information Processing Systems*, pp. 9459-9474, 2020.

- Li, J., Hui, B., Qu, G., Yang, J., Li, B., Li, B., Wang, B., Qin, B., Geng, R., Huo, N., Zhou, X., Chenhao, M., Li, G., Chang, K., Huang, F., Cheng, R., and Li, Y., Can LLM Already Serve as A Database Interface? A Big Bench for Large-Scale Database Grounded Text-to-SQLs, *Advances in Neural Information Processing Systems*, pp. 42330-42357, 2023.
- Milisavljevic-Syed, J., Li, J. and Xia, H., Realisation of responsive and sustainable reconfigurable manufacturing systems, *International Journal of Production Research*, vol. 62, no. 8, pp. 2725-2746, 2024.
- Otero, I., Salgado, J. F. and Moscoso, S., Cognitive reflection, cognitive intelligence, and cognitive abilities: A meta-analysis, *Intelligence*, vol. 90, pp. 101614, 2022.
- PEP, Product Specific Rules for Luminaires (PSR0014), PEP Association, Available: <https://www.pep-ecopassport.org>, Accessed on June 18, 2025.
- Ren, X. and Mia, M. A., The determinants of green innovations in manufacturing industries: a systematic, *Future Business Journal*, vol. 11, pp. 42, 2025.
- Sphera, LCAfE software: A professional tool for life cycle assessment, Available: <https://sphera.com/solutions/product-stewardship/life-cycle-assessment-software-and-data/lca-for-experts/>, Accessed on June 18, 2025.
- Vrchota, J., Pech, M., Rolinek, L. and Bednář, J., Sustainability outcomes of green processes in relation to industry 4.0 in manufacturing: Systematic review, *Sustainability*, vol. 12, no. 15, pp. 5968, 2020.
- Wang, S., Su, D., Wu, Y. and Chai, Z., Application of life-cycle assessment to the eco-design of LED lighting products, *Euro-Mediterranean Journal for Environmental Integration*, vol. 5, no. 2, pp. 41, 2020.

Biographies

Catherine K. Jiao is an undergraduate student pursuing a B.S. in Industrial and Systems Engineering with a minor in Psychology at the Georgia Institute of Technology. She has served as both a teaching assistant and research assistant, as well as being involved in a variety of student organizations. Her research interests lie at the intersection of human-computer interaction and decision making. More information about her: <https://www.linkedin.com/in/catherinejiao>

Mulang Song is a Ph.D. candidate in the School of Mechanical Engineering at the Georgia Institute of Technology (USA). He is graduating in Mechanical Engineering, with research that spans both mechanical and industrial engineering domains. He received both his B.S. and M.S. degrees in Mechanical Engineering with a focus on Computer Science, also from Georgia Tech. He has led and contributed to engineering projects through direct collaboration with industry, applying research in real-world production environments. His research interests include advanced manufacturing systems, intelligent production systems, and industrial AI applications. More information about his research: <https://scholar.google.com/citations?user=jLL4nycAAAAJ>

Yiyun (Cindy) Fei is a Ph.D. candidate in Mechanical Engineering at Georgia Institute of Technology (USA), where she also earned her B.S. and M.S. degrees in the same field. Her research focuses on the modeling, analysis, and optimization of complex operational systems to support data-driven design and decision-making. Yiyun worked as a Manufacturing and Process Development Engineer at TE Connectivity, where she gained four years of hands-on industry experience. During her time there, she led and contributed to multiple projects involving AI-based machine vision, process automation, and discrete event simulation of manufacturing systems. More information about her: www.linkedin.com/in/yiyun-fei-a134a21a3

Shu Wang is a smart manufacturing engineer at TE Connectivity. His expertise spans in industrial AI and data analytics, engineering design decision-making, human-automation interaction, and operations optimization. He holds a Ph.D. in Mechanical Engineering from Georgia Tech, along with dual M.S. degrees in Mechanical Engineering and Electrical and Computer Engineering. He earned his B.S. degree in Electrical Engineering from Wuhan University.

Bening Mayanti is a Postdoctoral Researcher in Industrial Management at the University of Vaasa, Finland. Her research interests include environmental system analysis, with a focus on quantifying and reducing environmental impacts using life cycle assessment for various systems, such as manufacturing, waste management, and energy systems. She holds D.Sc (Tech) in Industrial Management from University of Vaasa and M.Sc degree in Climate Studies from Wageningen University.