

On the role of artificial intelligence in feature oriented multi-criteria decision analysis

Heyuan Liu^{a,b}, Yi Zhao^{a*}, and Francois Maréchal^a

^a Industrial Process and Energy Systems Engineering (IPESE), Ecole Polytechnique Fédérale de Lausanne, Sion, Switzerland

^b MSc&T AiViC, École Polytechnique, Palaiseau, France

* Corresponding Author: yi.zhao@epfl.ch.

ABSTRACT

Balancing economic and environmental goals in industrial applications is critical amid challenges like climate change. Multi-objective optimization (MOO) and multi-criteria decision analysis (MCDA) are key tools for addressing conflicting objectives. MOO generates viable solutions, while MCDA selects the optimal option based on key performance indicators such as profitability, environmental impact, safety, and efficiency. However, large datasets pose a challenge in selecting the preferred solution during the MCDA process. This study introduces a novel machine learning-enhanced MCDA framework and applies the method to analyze decarbonization solutions for a European refinery. A stage-wise dimensionality reduction method, combining AutoEncoders and Principal Component Analysis (PCA), is applied to simplify high-dimensional datasets while preserving key spatial features. Geometric analysis techniques, including Intrinsic Shape Signatures (ISS), are employed to refine the identification of typical configurations for baseline evaluations. Once typical configurations are identified, Large Language Models (LLMs) are utilized to enhance decision-making by providing contextual problem explanations and generating weight proposals for the weighted sum method, ensuring alignment with decision criteria. This framework is designed to support stakeholders in making informed, transparent decisions in complex, uncertain environments.

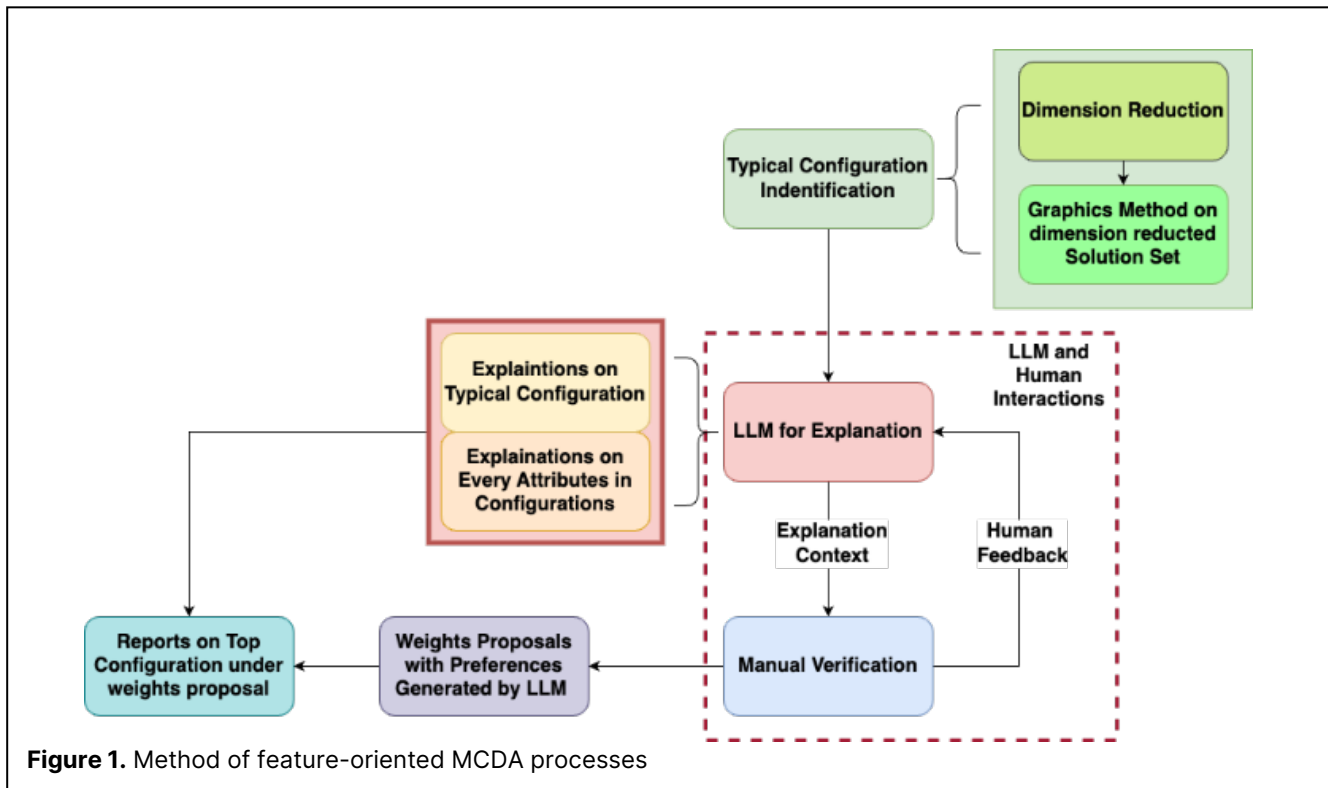
Keywords: Artificial Intelligence, Key performance indicator, Machine Learning, Multi-Criteria Decision Analysis

INTRODUCTION

In chemical engineering, balancing economic and environmental goals is crucial amidst challenges like climate change. To address conflicting objectives, tools like multi-objective optimization (MOO) [1] and multi-criteria decision analysis (MCDA) [2] are utilized. MOO generates a range of viable solutions on Pareto fronts, while MCDA helps select the most suitable option considering key performance indicators (KPIs) like profitability, environmental impact, safety, and efficiency. These tools aid in making informed decisions amidst complex trade-offs and uncertainties. However, the vast number of solutions generated by MOO presents a significant challenge, as it is impractical for human decision-makers to compare and analyze all configurations manually. This raises a critical question: how can preferred solutions be effectively selected from such a large solution database?

The selection process is generally divided into two steps: first, identifying typical configurations from the database, and second, selecting preferred configurations from this identified set based on MCDA.

To identify typical configurations, data-driven approaches have garnered significant attention, particularly those utilizing clustering methods [3]. These methods analyze the solution space by dividing it into several clusters instead of examining the entire dataset. This approach effectively captures the global distribution of the data by segmenting it into distinct groups. However, existing clustering techniques are generally based on selected KPIs of each configuration, like operating costs, which are typically calculated under uncertain market conditions. Configurations should instead be analyzed based on their inherent features, where each feature represents an independent dimension such as the material and energy flows describing the configuration. Notably,



two different configurations may share the same KPIs despite being fundamentally distinct. Consequently, KPI-based clustering may fail to accurately represent the underlying differences among configurations. In addition, a notable drawback of the traditional clustering method is the neglect of intra-cluster distributions. The loss of local features within clusters can significantly impact the depth and accuracy of the analysis.

For preferred solution selection, the weighted sum method [4] remains the most popular approach in both research and industry. This method involves assigning a weight to each KPI, computing the score of each solution, ranking the scores, and selecting the top configurations as optimal solutions. Despite its widespread use, this approach suffers from several limitations:

Lack of Diversity: Due to the inherent properties of Multi-Objective Optimization (MOO), solution datasets often contain repeated configurations. As a result, the top-ranked solutions derived from the weighted sum method frequently share the same configuration, leading to a lack of diversity in the selected solutions.

Limited Individual-Oriented Decision Making: The subjective nature of weight assignments also poses challenges for accommodating diverse perspectives, which can hinder personalized or individual-oriented decision-making. The core of the weighted sum method lies in the assignment of weights to each KPI, which translates user preferences into quantifiable metrics. This process often requires domain expertise, making it challenging to generalize. Furthermore, the assigned weights are highly

influenced by the individual expert's subjective preferences, introducing bias into the decision-making process.

Difficulty in weight collection: Gathering accurate and reliable weights for KPIs is a significant challenge. The process requires domain expertise and often lacks sufficient data support, making it difficult to establish objective and robust weight assignments.

To address these issues, a novel method is proposed that utilizes machine learning techniques, specifically Large Language Models (LLMs), to enhance the MCDA process under uncertain market conditions. The weight proposals for the weighted sum method are generated using the LLM model by incorporating specific preferences, eliminating the need for random identification or sole reliance on expert judgment. The primary contribution of this work lies in applying advanced machine learning methods to capture both global and local features in the solution space while leveraging LLMs for weight generation and MCDA.

METHODS AND MATERIALS

The method is outlined in Figure 1, beginning with the identification of typical configurations, followed by the application of an LLM model to generate weight proposals for different KPIs to support the selection of preferred configurations.

Typical Solution Identification

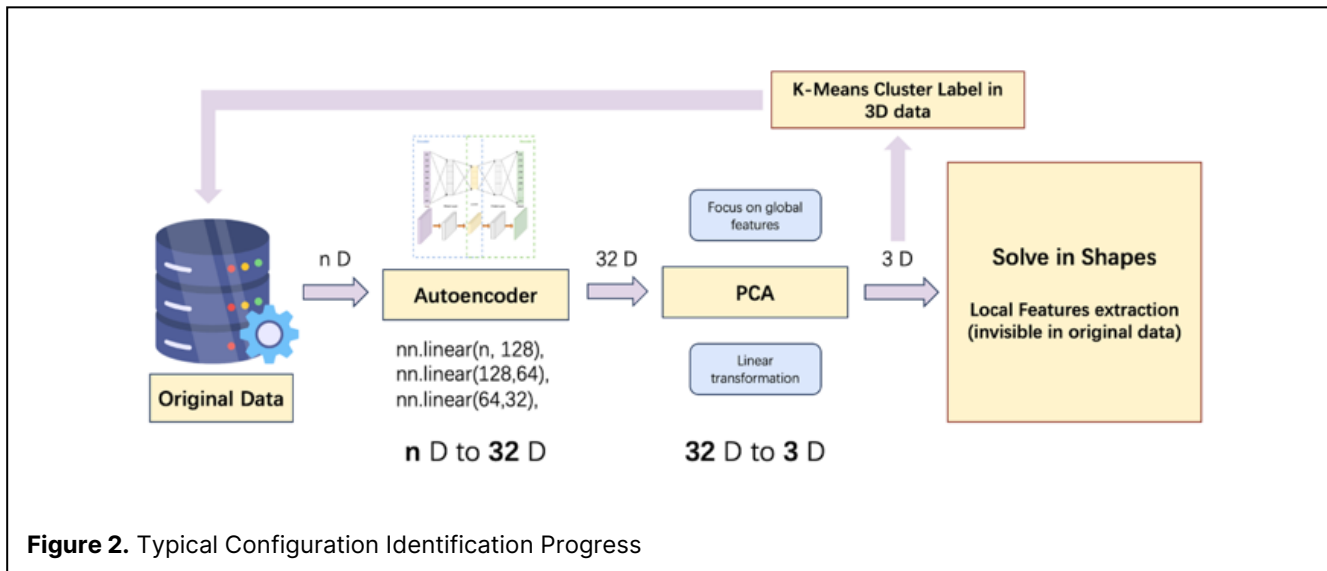


Figure 2. Typical Configuration Identification Progress

In complex energy systems, configurations often involve a vast number of features, making them a high-dimensional space impractical for manual analysis. Dimension reduction serves as an effective strategy to address this issue by simplifying the scale of the data while retaining its essential features. Clustering methods are then applied in dimension reduction datasets to identify the representative configurations, as shown in Figure 2.

Dimension Reduction

In this study, the Linear Combiner (LC) is proposed, a two-stage dimensionality reduction approach combining AutoEncoder [5] and Principal components analysis (PCA) [6] to enhance clustering performance on high-dimensional datasets. The PCA can identify the linear connections in the resulting dataset in linear space. However, the feature space exhibits non-linear relationships due to interdependencies. The Autoencoder captures the complex, nonlinear structures of the data by reducing the input dimensionality to an intermediate dimensional latent space. Subsequently, PCA is applied to further reduce this latent representation to 3 dimensions, effectively preserving the most significant linear features while eliminating redundant information.

To validate the effectiveness of the approach, its ability to preserve the essential spatial distribution characteristics of the data after dimensionality reduction is evaluated. Clustering is performed in a three-dimensional space, and the resulting cluster labels are projected back to the original high-dimensional space, which is assessed using two metrics: the Davies-Bouldin Score (DB Score) [7] and the Calinski-Harabasz Index (CH Index) [8].

Clustering based on the graphical method

In the pursuit of identifying the most representative solutions within the dataset, cluster centroids obtained through the K-Medoids method are selected as the

reference case. However, empirical observations revealed critical limitations in this approach. The derived cluster centers proved overly dependent on intra-cluster density distributions, occasionally failing to encapsulate the essential characteristics that define typical configurations across the entire solution space. This realization motivated the development of a more robust identification strategy.

Building on the prior dimensionality reduction of the dataset into three-dimensional embeddings, graphical are used to extract latent spatial patterns that reveal structural relationships and density variations hidden in high-dimensional data. This 3D projection creates a visually interpretable landscape, where topological properties and neighborhood continuity can be algorithmically quantified, offering a more robust criterion for identifying typical configurations beyond centroid proximity. To achieve this, an advanced technique from graphics and computer vision—Intrinsic Shape Signatures (ISS) [9], is employed, providing a unique perspective to better understand the dataset's structure and identify typical configurations.

LLM for MCDA

To effectively integrate Large Language Models (LLMs) into Multi-Criteria Decision Analysis (MCDA), the first step is to enable the LLM (e.g., ChatGPT series [10]) to comprehend the dataset. This involves providing the model with a curated subset of typical configurations rather than the entire dataset, as processing large datasets directly is impractical. The LLM is tasked with generating explanations for each attribute in the dataset, enhancing its reasoning ability through structured prompts and human verification. By manually verifying the explanations provided by the LLM, one can assess whether the LLM has properly understood the dataset. This allows the LLM to interpret the data contextually and provide meaningful

insights, paving the way for subsequent weight generation and decision-making processes.

Data set interpretation

The reasoning ability of LLMs is crucial for accurate dataset interpretation. To enhance this capability, hierarchical explanation mechanisms and structured prompting are implemented to improve depth, accuracy, and reliability. Additionally, a human verification mechanism is introduced, enabling experts to review and validate AI-generated outputs at key stages. This interactive process identifies and corrects potential errors or biases while refining the model's reasoning through iterative feedback. By integrating human oversight, the framework ensures reliable, domain-aligned outputs, significantly enhancing the system's robustness and credibility.

Hierarchical Explanation:

A hierarchical explanation approach is applied, instructing the large language model (LLM) to first categorize attributes into their respective subcategories, and then provide detailed explanations for each attribute based on its assigned category. This method is beneficial because it enforces a structured organization of data, increasing the inductive bias, reducing hallucination, and enabling clear verification of each reasoning step. The overall workflow (Figure 3) proceeds as follows: (1) Identify attributes that require explanation. (2) Classify attributes into relevant subcategories. (3) Provide detailed explanations for each attribute under its subcategory. (4) Present the final output by listing each attribute's classification alongside its detailed explanation.

Meanwhile, a verified explanation ensures the stakeholders without mathematical expertise can interpret and validate decisions, which is particularly useful in a diverse disciplinary team.

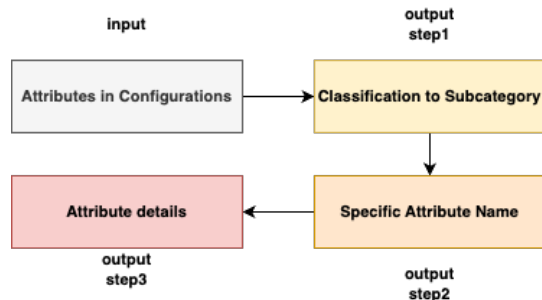


Figure 3. Hierarchical Explanation for increasing the reasoning ability of the LLM agent

Prompts Structure: To further enhance the reasoning process, a structured prompting framework is designed inspired by the Chain of Thought (CoT) and its related research [11], which has demonstrated significant improvements in reasoning tasks. This framework involves 3 independent LLM agents proposing distinct strategies, each addressing specific aspects of the

problem. These strategies are then evaluated through a voting mechanism by a panel of 5 additional LLM agents, ensuring an unbiased assessment. The most effective strategy is subsequently selected and refined into an optimal solution, integrating the strengths of individual proposals while minimizing potential biases and limitations. Notably, all 8 experts involved in this process are distinct LLM agents.

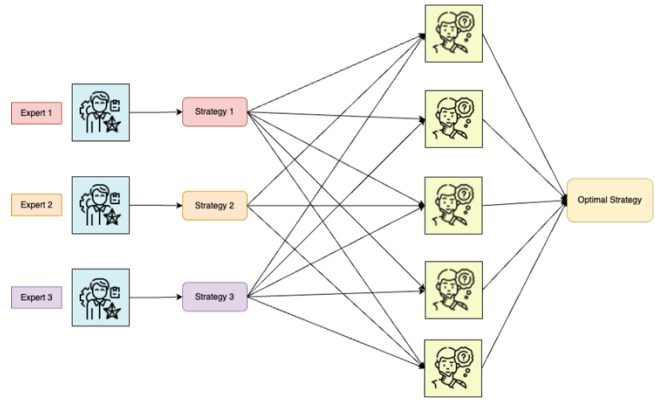


Figure 4. Prompts Structure with 3 experts for strategy generating and 5 experts for voting

Weights Proposals Generation

For weight generation, three key inputs are utilized, including user-defined preference settings, typical configurations identified in previous sections, and explanations generated by LLMs. This process leverages the structured prompting framework introduced earlier, which guides the weight assignment in a manner analogous to group-based decision-making mechanisms. Specifically, the framework facilitates the aggregation of preferences and the systematic integration of insights from the three inputs, ensuring a balanced and well-informed weight assignment process.

Traditional weight generation in MCDA often relies on expert input, which presents several challenges, including the need for domain expertise, difficulty in collecting reliable weight proposals, and potential bias from individual preferences. Moreover, integrating multiple expert opinions into a coherent weighting scheme remains a complex task. In contrast, LLMs offer a more flexible and scalable alternative. With their extensive training in diverse knowledge sources, LLMs can generate weight proposals more efficiently, accommodate various preference settings, and reduce dependence on subjective expert judgment. This approach enhances efficiency and consistency in decision-making while allowing for greater adaptability across different application scenarios.

CASE STUDY

The method is applied to rank solutions for decarbonizing petroleum refining processes. The case study is

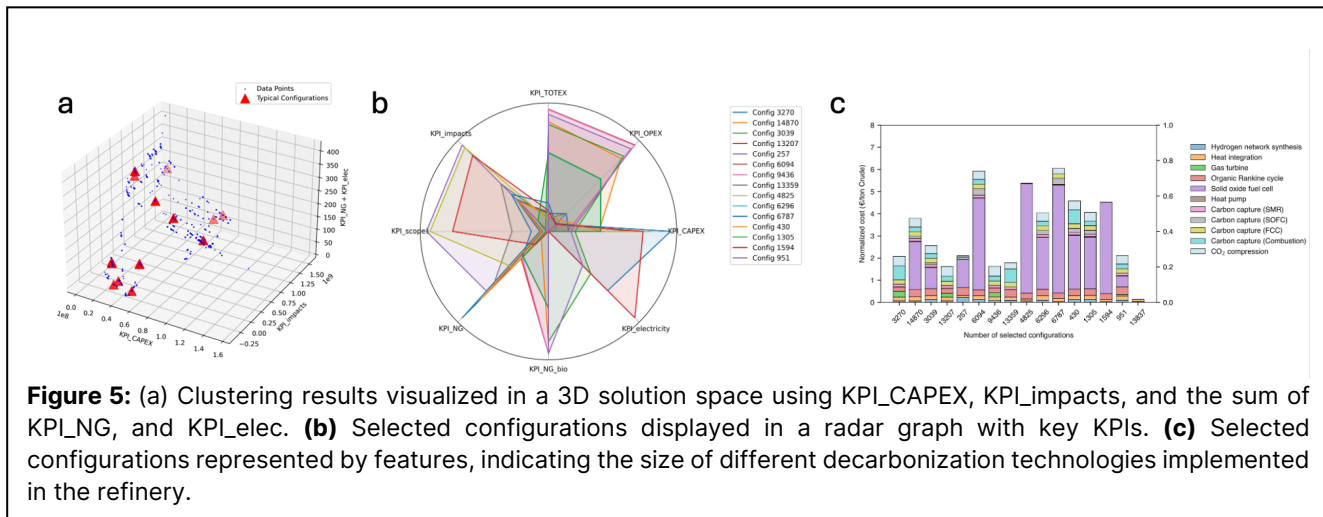


Figure 5: (a) Clustering results visualized in a 3D solution space using KPI_CAPEX, KPI_impacts, and the sum of KPI_NG, and KPI_elec. (b) Selected configurations displayed in a radar graph with key KPIs. (c) Selected configurations represented by features, indicating the size of different decarbonization technologies implemented in the refinery.

based on a superstructure of refining processes integrating various decarbonization technologies, including energy efficiency measures, carbon capture and storage, low-carbon hydrogen, and renewable energy resources. A total of 16,000 decarbonization configurations are generated using a MILP formula to balance the trade-offs between investment, electricity and natural gas consumption, and CO₂ emissions (Scope 1 and Scope 2). Solutions are derived under a wide range of market and other conditions, considering factors such as natural gas and electricity prices, carbon pricing policies, and electricity carbon footprint. Each of the 16,000 configurations represents an optimal solution for a specific set of input conditions.

Each configuration is represented by two data sets: KPIs and features. The KPI set includes total cost (KPI_TOTEX), operating costs (KPI_OPEX), annualized investment (KPI_CAPEX), and resource consumption metrics such as natural gas, biogas, and electricity usage (KPI_NG, KPI_NG_bio, and KPI_electricity), along with scope 1 and scope 2 emissions (KPI_scope1 and KPI_scope2). These KPIs are derived from the material and energy flows as well as the installed units, which constitute the configuration's features. Unlike KPIs, which vary based on different conditions, features remain constant for a given configuration.

RESULTS

Typical configuration identification

Figure 5(a) visualizes the dimension-reduced dataset, with 16 clustering centers highlighted by red triangles, identified using the graphical methods described in the method section. Clustering performance is quantified using established metrics. The selected configurations are further presented by their KPIs (Figure 5(b)) and features, such as the installed capacity of decarbonization technologies (Figure 5(c)). This provides a clearer

representation, demonstrating that the selected configurations are distinct in both KPIs and features.

Preferred configuration selection

Dataset interpretations were first generated independently by LLM, incorporating hierarchical explanation and structured prompting to enhance its reasoning ability. However, as shown in Figure 6, three key misinterpretations were identified in the initial LLM explanation: (1) *KPI_impacts* should only include scope 1 and 2 emissions, (2) *KPI_CAPEX* should be annualized, and (3) Heat Exchange Variables should be treated solely as intermediate computational values. Notably, despite lacking predefined conditions, the LLM correctly avoided assigning weights to Heat Exchange Variables. After manual feedback alignment, the second iteration showed a significantly improved understanding of the dataset and problem, leading to a more accurate interpretation.

Following the correct understanding achieved through human alignment with the LLM-generated context, Figure 7 presents an example of the complete weights-generation process. From the outset, the priority was set to "Focusing on the Environmental Impact." The deduction progress in weights generation follows the structure of the prompts we designed in the previous section, prompting experts to vote on various KPIs. After consolidating the voting results, higher weights were allocated to KPIs most strongly tied to environmental impact, while moderate values were assigned to other attributes like CAPEX and OPEX. This iterative, feedback-driven approach demonstrates how the LLM can integrate expert judgment and maintain alignment with an overarching priority—namely, reducing environmental impact—within its weights-generation process.

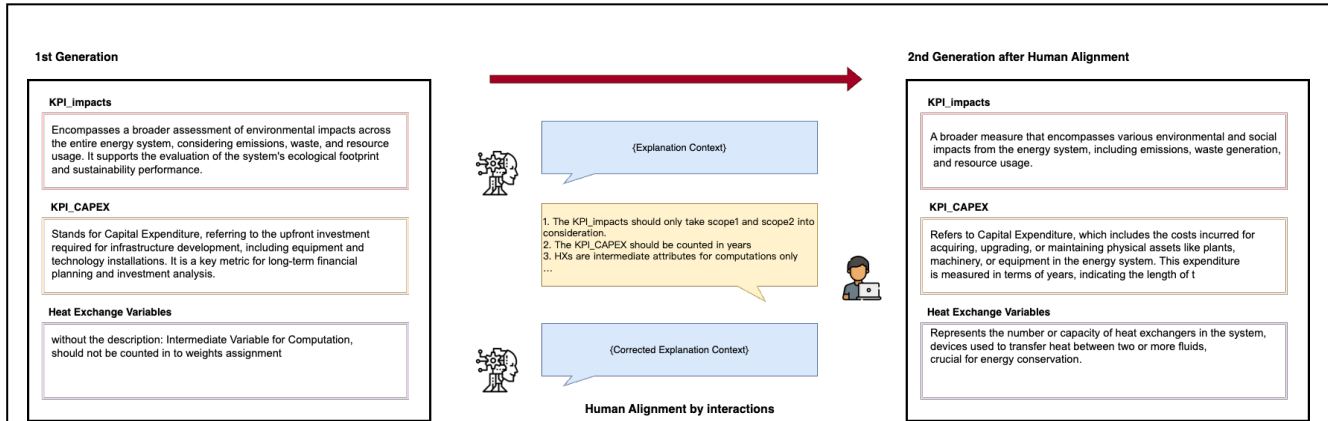


Figure 6: The Human Alignment progress in case study to ensure the AI-generated context trustworthy

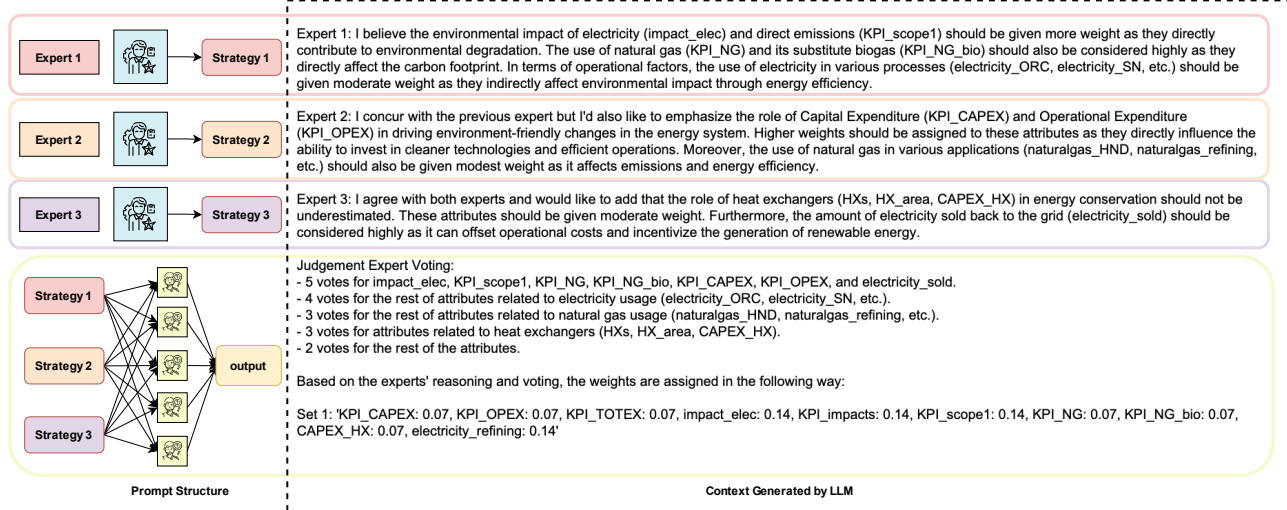


Figure 7: The Deduction Progress of generating the weights proposals with the preference is set as the "focus on Environmental Impact" and the weights proposal is listed in Set 1.

Figure 8 presents the ranking results obtained from the weighted sum method, with weights generated by the LLM agent as illustrated in Figure 8. The top-ranked configurations (e.g., Configurations 257, 4825, and 1594) exhibit varying ranking weighted sum scores under the focus on "Environmental Impact," as they are identified as typical configurations. Figure 5(b-c) highlights the corresponding KPIs and features of these configurations, revealing distinct characteristics and orientations that contribute to their differentiation within the solution space.




Figure 8. Typical Configuration weighted sum scores are

computed with weights generated by LLM.

DISCUSSION

This study presents a feature-oriented MCDA framework with machine learning method that addresses key challenges in selecting preferred solutions from large industrial datasets. First, by integrating dimensionality reduction and geometric analysis, our approach effectively captures both global and local structural features, enhancing decision-making transparency and robustness. Second, through hierarchical explanation, prompt engineering, and human verification, the LLM-assisted explanation and weight generation process reduce reliance on subjective expert input, ensuring consistency, efficiency, and adaptability in preference-based decision-making across diverse industrial applications.

Applied to refinery decarbonization, our framework successfully identifies representative solutions while ensuring diversity and interpretability in high-dimensional solution spaces. Additionally, the AI-driven methodology leverages Large Language Models (LLMs) with enhanced

reasoning ability to provide context-aware explanations for decision attributes and generate weight proposals aligned with user preferences. This approach reduces reliance on subjective expert input and improves decision-making efficiency. This scalable, interpretable, and adaptable solution demonstrates its applicability to complex multi-criteria decision-making problems.

ACKNOWLEDGEMENTS

The author thanks TotalEnergies for their financial support and valuable discussion about this research.

REFERENCES

1. Zhao, Y., Hagi, H., Delahaye, B., & Maréchal, F. (2024). A holistic approach to refinery decarbonization based on atomic, energy and exergy flow analysis. *Energy*, 296, 131117.
2. Wang, J. J., Jing, Y. Y., Zhang, C. F., & Zhao, J. H. (2009). Review on multi-criteria decision analysis aid in sustainable energy decision-making. *Renewable and sustainable energy reviews*, 13(9), 2263-2278.
3. Nolzen, N., Lademann, A., Roskosch, D., Moret, S., Seele, H., Baader, F. J., & Bardow, A. (2024). Low-regret decisions for the steam supply in the chemical industry. In *ESCAPE-34 PSE-2024*.
4. Dawes, R. M., & Corrigan, B. (1974). Linear models in decision making. *Psychological bulletin*, 81(2), 95.
5. Wang, Y., Yao, H., & Zhao, S. (2016). Auto-encoder based dimensionality reduction. *Neurocomputing*, 184, 232-242.
6. Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: a review and recent developments. *Philosophical transactions of the royal society A: Mathematical, Physical and Engineering Sciences*, 374(2065), 20150202.
7. Davies, D. L., & Bouldin, D. W. (1979). A cluster separation measure. *IEEE transactions on pattern analysis and machine intelligence*, (2), 224-227.
8. Caliński, T., & Harabasz, J. (1974). A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1), 1-27.
9. Zhong, Y. (2009, September). Intrinsic shape signatures: A shape descriptor for 3D object recognition. In *2009 IEEE 12th international conference on computer vision workshops, ICCV Workshops* (pp. 689-696). IEEE.
10. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877-1901.
11. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., ... & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35, 24824-24837.

© 2025 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See <https://creativecommons.org/licenses/by-sa/4.0/>

