Enhancing Human-AI Collaboration through a Conversational Agent for Energy Efficiency

Riccardo Campi, Mathyas Giudici, Nicolò Oreste Pinciroli Vago, Marco Brambilla, Piero Fraternali

Politecnico di Milano, Dipartimento di Elettronica, Informazione e Bioingegneria. Milano, Italy
{riccardo.campi, mathyas.giudici, nicolooreste.pinciroli, piero.fraternali, marco.brambilla} @polimi.it

Abstract

Among the many scenarios where humans and AI agents can collaborate, Energy Efficiency (EE) is one where such collaboration could most effectively contribute to the goal of net zero emissions, while also reducing costs and improving comfort. In this context, new AI solutions can support customers in making their energy consumption more efficient and aligned with renewable sources. In this work, we investigate the strengths and challenges of Human-AI Collaboration by proposing an AI-based Conversational Agent whose inspiration principles are derived from the theories of Human-Centered Artificial Intelligence (HCAI). It is specifically designed to augment users' capabilities in achieving EE by providing them with recommendations and practical tips. The Agent uses a Knowledge Graph (KG) trained on domain-specific energy-related documents, coupled with a RAG (Retrieval Augmented Generation) architecture to ensure factual accuracy, source accountability, fairness, and transparency. By tailoring responses to users' profiles and preferences, the system prioritizes human needs and values while addressing perceptions of technological usability and acceptability. The Agent is validated in a real-world application scenario with international customers, with the aim to test content accuracy and adaptation to the user context and uncertainties. The results show the effectiveness of the system in fostering Human-AI Collaboration for EE.

Introduction and Motivation

Individual energy consumption habits are crucial in achieving sustainability and mitigating climate change consequences, with particular regard to residential electricity consumption (Abbass et al. 2022), which accounts for a significant portion of the world's energy consumption and will increase since electrification as a strategy for future sustainable development (Stewart et al. 2018). As a result, the emphasis on Energy Efficiency (EE) is becoming increasingly important, aiming to achieve net zero emissions as recently outlined by European Institutions (European Parliament and Council of European Union 2021).

The increasing adoption of Conversational Agents, in particular with the rise of Large Language Models (LLMs) (Khosravi et al. 2024), offers opportunities to enhance Human-AI Collaboration by providing new digital scalable and user-tailored systems that increase understanding and optimization of energy consumption (Giudici et al. 2023). However, since these models may fall short in ensuring factual accuracy (Frieden et al. 2020), Retrieval-Augmented Generation (RAG) systems tackle this challenge by integrating LLMs with a structured knowledge base that formally encapsulates domain-specific information (Arslan et al. 2024).

Considering this scenario, our research aims to investigate the strengths and challenges of Human-AI Collaboration for EE using a Human-Centered Artificial Intelligence (HCAI) approach. We propose a Conversational Agent designed to provide users with recommendations on EE using a Knowledge Graph (KG) trained on source documents containing domain-specific energy-related information. Integrating this graph-based approach with a RAG architecture ensures the factual correctness of the answers and proper citation of sources, emphasizing system accountability, fairness, interpretability, and transparency. The Agent uses contextual data to keep track of each user's profile and preferences, including their name, country, language, and past interactions. Thus, responses are customized based on the individual's specific requirements and characteristics. The system prioritizes human needs, values, and goals rather than other aspects, and takes into consideration user perception as well as technological acceptability and usability.

Our Conversational Agent is designed to satisfy these research questions, inspired by Horvitz's concept of *Mixed-Initiative User Interfaces*:

- (RQ1) Does the Agent respond with correct content?
- (RQ2) Does the Agent adapt to the user context?

To answer these questions and validate our Conversational Agent, we apply it within *ENERGENIUS*, a H2020 European research project[1] on EE. The knowledge base is built starting from various sources and documents on energy consumption, EE, regulations, and incentives from different countries.

Related Work

Recent research on LLMs has enabled multilingual capabilities and summarization tasks, as well as language com-

[1] https://energenius-project.eu/

prehension and content creation (Arslan et al. 2024). However, assessing the factual correctness of their answers often is difficult, because of the provenance problem (Frieden et al. 2020). Indeed, LLMs often struggle to provide or accurately cite the sources from which their answers are derived. Also, when posed with questions they are unfamiliar with (i.e., they were not trained on, domain-specific, or ambiguous), LLMs frequently provide illogical or completely inaccurate responses (Giudici et al. 2023; Sanguinetti et al. 2024). These are referred to as hallucinations and are mainly caused by a lack of context awareness or domain-specific expertise. Since hallucinations are difficult to distinguish from factual information, they can have serious negative consequences, including the spread of false information and a decline in user trust (Huang et al. 2025).

By combining the advantages of information retrieval with the creative powers of LLMs, RAG systems prove to be a potent remedy for these constraints. This makes it possible for RAG systems to handle knowledge-intensive activities like answering domain-specific queries or referencing the sources of the information they have received, producing outputs that are more precise and relevant to the context (Arslan et al. 2024). In addition, new graph-based architectures are emerging, enabling RAG to provide more precise and pertinent answers by relying on KGs (Edge et al. 2024), especially for complex questions that require synthesizing information from multiple sources.

In the field of energy efficiency, Giudici et al. (Giudici et al. 2023) investigated strategies to improve comprehension and maximize energy usage to satisfy EE by utilizing LLMs. Their chatbots answered broad queries with ease and coherence, but they were inaccurate when asked domain-specific questions, emphasizing the necessity of coupling LLMs with other retrieval methods. Arslan et al. (Arslan, Mahdjoubi, and Munawar 2024) proposed Energy Chatbot, a vector-based multi-source RAG with the aim of enhancing decision-making for SMEs, by providing comprehensive Energy Sector insights through a Question Answering system. Important conclusions highlight how the system's capacity to provide precise, pertinent, and consistent information is greatly improved by the integration of a RAG, particularly when using the Llama3.1:8B model. However, a Graph-based solution is still missing. Similarly, Bruzzone et al. (Bruzzone et al. 2023) improved urban planning simulations by integrating a RAG system with a chatbot running with GPT4. This technology provides accurate information to urban planners and encourages sustainability activities by dynamically simulating various urban environments. Finally, using multiple publicly accessible electricity usage KGs, another study (Fortuna, Hanžel, and Bertalanič 2024) examined Graph-based RAG techniques to answer complex questions on electricity. The combination of RAG and LLMs utilizing publicly available electricity consumption KGs yields encouraging outcomes, according to key findings.

To the best of our knowledge, there is no previous work leveraging a Graph-based RAG architecture that investigates the strengths and challenges of Human-AI Collaboration for EE using an HCAI approach.

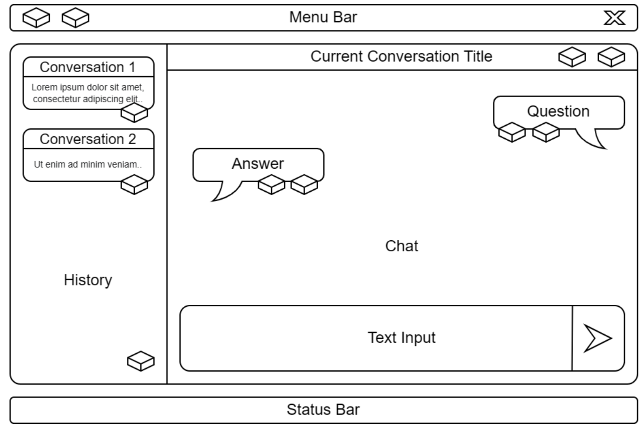

Figure 1: Proposed UI/UX Wireframe. According to *Prometeus principles*, (1, 2) user intentions are conveyed through interactive components like the text input field and multiple contextual action buttons (box icons). (3) Actions are designed to be fast, incremental, and reversible, with users receiving (4) clear feedback, (5) progress indicators, and (6) completion reports for confirmation.

Methodology

To achieve EE, it is important to engage users by offering energy-saving recommendations and motivating them to adopt more efficient habits (Stewart et al. 2018). In certain cases, achieving this goal can be challenging, especially when users are asked to change behaviors they have followed for a long time. Also, habits developed within the home can be difficult to change, as users may perceive the recommendations as intrusive. To persuade them of the soundness of the recommendations, the Agent must leverage the dual foundation of human reasoning, made up of both rationality and emotions (Rolls 2019). In this way, by *adapting* to users, the Agent significantly increases the likelihood of success. Otherwise, there is a considerable risk of forcing users to adapt to the Agent, which could result in them rejecting the recommendations and failing with EE (Jensen et al. 2018). At the same time, the Agent needs to take into account the *uncertainty* of human communication, which arises from the inherently ambiguous nature of natural language, and the fuzziness of human reasoning, which is caused by the graded human cognition (Voskoglou 2014). Thus, inferring the real intent of the user is required, and acting accordingly ultimately minimizes the damage caused by incorrect interpretations. For these reasons, our Agent is designed to prioritize the human experience, and the underlying methodology focuses more on Human-AI Collaboration and HCAI aspects over technical requirements.

We imagine it as a tool to *augment* human capabilities and improve their performances in doing tasks like gathering information and understanding complex regulations, all aimed at achieving the goal of EE. Fig. 1 shows the UI/UX Wireframe of the designed application. The UI/UX follow some design rules, formalized by Shneiderman as the *Prometeus principles*, aimed at increasing the system comprehensibil-

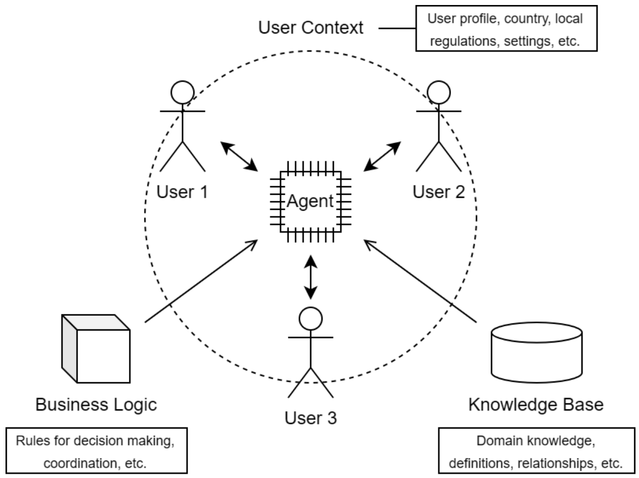

Figure 2: A high-level view of the proposed *Interaction Scheme* between the humans and the Agent. Agent interacts with the users keeping track of their *User Context*. It uses a custom *Business Logic* for decision-making and coordination, while a *Knowledge Base* contains all the domain knowledge required to provide the recommendations.

ity, predictability, and controllability (Shneiderman 2020). In particular, our agent must have (1) a consistent interface to enable users to set their intentions, (2) a visual display of the objects and actions of interest, (3) rapid, incremental, and reversible actions, (4) informative feedback after each action, (5) progress indicators to show the current status, and (6) completion reports to confirm the success or the failure.

We propose leveraging AI to develop a Conversational Agent by putting into practice the above-mentioned collaborative goals, and designed to provide accurate recommendations based on domain-specific energy-related information retrieved from official documents. It acts as an intermediary between the users and a Knowledge Base, using a custom rule-based Business Logic for decisions and coordination, and taking into consideration a User Context to provide tailored answers based on each user's profile and preferences. Fig. 2 shows a high-level view of the proposed *Interaction Scheme*. While communication between users and the Agent is made with natural language, this latter uses function calling to interact with the Business Logic. It does not allow direct communication between users.

Diving deeper into technical details, the Agent relies on KG-based RAG architecture that couples LLMs' exceptional natural language processing and understanding capabilities with the KG's retrieving and reasoning abilities. The KG is trained one-off on domain-specific EE source documents, and is then queried by the retriever using the *SPARQL Query Language* to obtain relevant information. The use of a KG-enhanced RAG guarantees reliable citation of sources, prioritizing system accountability, fairness, interpretability, and transparency, while enabling local reasoning processes to handle complex questions and provide the users with accurate and factual answers.

Results and Discussion

We assess the effectiveness of our Conversational Agent in (RQ1) delivering correct and relevant EE answers and in (RQ2) adapting to the user context and profiling information. With this aim, we carry out a preliminary evaluation by first training the Agent on source documents in Italian language on energy consumption, EE, regulations, and incentives from Italy and Switzerland. Then, we create a ground-truth dataset of 101 question-answer pairs on the same topics of the source documents. Finally, we develop a prototype of our Agent and include into the *ENERGENIUS Guru*, a Decision Support System (DSS) on EE aimed at providing users with useful recommendations, with a particular focus on the answer accuracy and the proper citation of the source documents to increase accountability and transparency.

Our preliminary results are obtained by tasking four energy-domain experts to use the Agent, taking our 101 question-answer pairs and comparing the ground-truth answers with the obtained ones. They are then tasked with the evaluation of the answers using the RAGAs framework (Es et al. 2024). In particular, an answer is considered valid if it satisfies all three RAGAs properties: *faithfulness*, the answer must be grounded in the given context, *answer relevance*, the generated answer must directly address the provided question, and *context relevance*, the retrieved context must be concise and focused, containing minimal irrelevant information. If the answer does not meet one or more of these required properties, its score will be proportionally reduced until it reaches zero.

Aggregated results show that (RQ1) answers are correct in terms of the accuracy and correctness of the delivered content in the $75.2 \pm 2.7\%$ of the cases.

Results also show that (RQ2) the Agent is effective in adapting to the user context, in terms of language and locally relevant information. Indeed, in the $77.4 \pm 2.9\%$ of the cases it correctly answers to questions in Italian, and in the $73.0 \pm 2.5\%$ of the cases for questions in English. It correctly answers questions on Italy's laws and regulations in the $73.4 \pm 0.9\%$ of the times, $71.2 \pm 2.1\%$ for Switzerland ones, and $78.4 \pm 3.0\%$ for generic questions (i.e., not related to any specific country).

Users have reported that "the collaboration is strengthened by the presence of the source documents citations, which enhance the perception of the Agent's faithfulness" and that "while some tips on energy efficiency are too general, others are useful and easy to follow". Finally, "the system does not always find the requested information, but when this happens, it displays an error message, enhancing the perception of transparency".

Conclusion

This work investigates the strengths and challenges of Human-AI Collaboration for sustainability and climate change consequences mitigation, highlighting the potential of an AI-based Conversational Agent aimed at achieving EE. The used HCAI design principles and methodologies coupled with advanced technical solutions effectively enable the Agent to offer recommendations on EE by integrating domain-specific knowledge on energy and local regulatory documents. It is meant to augment human capabilities in information gathering to reach the goal of EE in a collaborative way.

Various elements enhance the Human-AI Collaboration factor. While the adoption of the *Prometeus Principles* in the design of the UI/UX increases the comprehensibility, predictability, and controllability of the Agent "thereby increasing the users' self-efficacy", the choice to design the Agent with a KG-enhanced RAG optimizes the accuracy of responses and enhances contextual relevance. The system provides language-agnostic responses, thus decoupling the content from its language, and uses contextual data such as profile information or favorite country and language to tailor answers to the users.

The preliminary results, obtained by integrating the prototype of the Agent into the *ENERGENIUS Guru*, show the system's performance in responding with correct content and context, and in adapting to users' uncertainty and fuzziness.

Limitations and Future Work

This work serves as an initial exploration of the collaboration between humans and AI to achieve EE. Our future efforts will focus on expanding this concept by enhancing the Agent's capabilities and performance. Additionally, we plan to conduct comprehensive experiments involving a larger group of users and utilizing more advanced evaluation frameworks.

Acknowledgements

This research was funded by the European Commission's Horizon Europe project ENERGENIUS, Project ID 101160720.

References

Abbass, K.; Qasim, M. Z.; Song, H.; Murshed, M.; Mahmood, H.; and Younis, I. 2022. A review of the global climate change impacts, adaptation, and sustainable mitigation measures. *Environ. Sci. Pollut. Res. Int.*, 29(28): 42539–42559.

Arslan, M.; Ghanem, H.; Munawar, S.; and Cruz, C. 2024. A Survey on RAG with LLMs. *Procedia Computer Science*, 246: 3781–3790. 28th Intl. Conf. on Knowledge Based and Intelligent information and Engineering Systems (KES 2024).

Arslan, M.; Mahdjoubi, L.; and Munawar, S. 2024. Driving sustainable energy transitions with a multi-source RAG-LLM system. *Energy and Buildings*, 324: 114827.

Bruzzone, A.; Giovannetti, A.; Genta, G.; and Cefaliello, D. 2023. Generative AI and Retrieval-Augmented Generation (RAG) in an Agent-Based Simulation Framework for Urban Planning. *23rd Int. Conf. on Modelling and Applied Simulation MAS*.

Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; and Larson, J. 2024. From Local to Global: A Graph RAG Approach to Query-Focused Summarization.

Es, S.; James, J.; Espinosa Anke, L.; and Schockaert, S. 2024. RAGAs: Automated Evaluation of Retrieval Augmented Generation. In Aletras, N.; and De Clercq, O., eds., *18th Conf. of the European Chapter of the ACL: System Demonstrations*, 150–158. St. Julians, Malta: Association for Computational Linguistics.

European Parliament and Council of European Union. 2021. Regulation (EU) 2021/1119 of the European Parliament and of the Council of 30 June 2021 establishing the framework for achieving climate neutrality and amending Regulations (EC) No 401/2009 and (EU) 2018/1999 ("European Climate Law").
http://data.europa.eu/eli/reg/2021/1119/oj.

Fortuna, C.; Hanžel, V.; and Bertalanič, B. 2024. Natural Language Interaction with a Household Electricity Knowledge-based Digital Twin. arXiv:2406.06566.

Frieden, D.; Tuerk, A.; Neumann, C.; d'Herbemont, S.; and Roberts, J. 2020. Collective self-consumption and energy communities: Trends and challenges in the transposition of the EU framework. *COMPILE, Graz, Austria*.

Giudici, M.; Abbo, G. A.; Belotti, O.; Braccini, A.; Dubini, F.; Izzo, R. A.; Crovari, P.; and Garzotto, F. 2023. Assessing LLMs Responses in the Field of Domestic Sustainability: An Exploratory Study. In *2023 Third International Conference on Digital Data Processing (DDP)*, 42–48. IEEE.

Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; and Liu, T. 2025. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Trans. Inf. Syst.*, 43(2).

Jensen, R. H.; Raptis, D.; Kjeldskov, J.; and Skov, M. B. 2018. Washing with the Wind: A Study of Scripting towards Sustainability. In *Proceedings of the 2018 Designing Interactive Systems Conference*, 1387–1400. ACM.

Khosravi, H.; Shafie, M. R.; Hajiabadi, M.; Raihan, A. S.; and Ahmed, I. 2024. Chatbots and ChatGPT: A bibliometric analysis and systematic review of publications in Web of Science and Scopus databases. *International Journal of Data Mining, Modelling and Management*.

Rolls, E. T. 2019. Emotion and reasoning in human decision-making. *Economics*, 13(1): 20190039.

Sanguinetti, M.; Pani, A.; Perniciano, A.; Zedda, L.; Loddo, A.; and Atzori, M. 2024. Assessing Italian Large Language Models on Energy Feedback Generation: A Human Evaluation Study. *10th It. Conf. on Computational Linguistics (CLiC-it)*.

Shneiderman, B. 2020. Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. *International Journal of Human–Computer Interaction*, 36(6): 495–504.

Stewart, I. D.; Kennedy, C. A.; Facchini, A.; and Mele, R. 2018. The electric city as a solution to sustainable urban development. *Journal of Urban Technology*, 25(1): 3–20.

Voskoglou, M. G. 2014. Measuring the Uncertainty of Human Reasoning. *American Journal of Applied Mathematics and Statistics*, 2(1): 1–6.