

# The Carbon Footprint Wizard: A Knowledge-Augmented AI Interface for Streamlining Food Carbon Footprint Analysis

Mustafa Kaan Aslan<sup>1</sup>, Reinout Heijungs<sup>2</sup>, Filip Ilievski<sup>1</sup>

<sup>1</sup>Department of Computer Science, Vrije Universiteit, Amsterdam, The Netherlands

<sup>2</sup>Department of Operations Analytics, Vrije Universiteit, Amsterdam, The Netherlands

m.k.aslan@student.vu.nl, {r.heijungs,f.ilievski}@vu.nl

## Abstract

Environmental sustainability, particularly in relation to climate change, is a key concern for consumers, producers, and policymakers. The carbon footprint, based on greenhouse gas emissions, is a standard metric for quantifying the contribution to climate change of activities and is often assessed using life cycle assessment (LCA). However, conducting LCA is complex due to opaque and global supply chains, as well as fragmented data. This paper presents a methodology that combines advances in LCA and publicly available databases with knowledge-augmented AI techniques, including retrieval-augmented generation, to estimate cradle-to-gate carbon footprints of food products. Our methodology is implemented as a chatbot interface that allows users to interactively explore the carbon impact of composite meals and relate the results to familiar activities. A web demonstration showcases our proof-of-concept system with user recipes and follow-up questions, highlighting both the potential and limitations —such as database uncertainties and AI misinterpretations—of delivering LCA insights in an accessible format.

**Demo** — <https://carbonfootprintwizard.labs.vu.nl>

**Code** —

<https://github.com/mkaanaslan/carbon-footprint-wizard>

## Introduction

As sustainability is high on the agenda for consumers, producers, and policymakers, indicators of sustainable production and consumption are key to present-day research and consulting (Krajnc and Glavić 2003). While sustainability is a multidimensional concept (Purvis, Mao, and Robinson 2019), in the context of climate change, environmental considerations take precedence over social and economic factors (Yañez, Sinha, and Vásquez 2019). The most popular indicator for quantifying the environmental impact to climate change is the carbon footprint (CF), which aggregates greenhouse gas (GHGs) emissions using global warming potential (GWP) values (Čuček, Klemeš, and Kravanja 2012).<sup>1</sup>

Copyright © 2026, Association for the Advancement of Artificial Intelligence ([www.aaai.org](http://www.aaai.org)). All rights reserved.

<sup>1</sup>In this work, we focus on the most recent GWP values from the Intergovernmental Panel on Climate Change (IPCC), with a 100-year time horizon.

Carbon footprints are typically measured across three scopes (Hertwich and Wood 2018). Scope 1 covers the emissions of GHGs directly from the activity itself, and Scope 2 includes the GHG emissions from the energy production required to run the activity. Scope 3 includes non-energy-related inputs from the whole supply chain, such as materials and processes. Scope 3 is vital for food processes, where it includes issues as diverse as methane emissions from rice cultivation, CO<sub>2</sub> emissions from fertilizer production, and methane emissions from cattle breeding for the meat industry. If an analysis includes all three scopes, it takes the form of a life cycle assessment (LCA) (Finkbeiner 2016). Most food LCAs are cradle-to-gate (or, with a popular term, farm-to-fork) (Schebesta, Bernaz, and Macchi 2020), stopping short of including the end-of-life phase.

Conducting LCAs is complex due to the long, often global supply chains. To facilitate this, databases such as ecoinvent (Wernet et al. 2016) and the U.S. Life Cycle Inventory Database (Deru 2009) have emerged from early manual studies (PET 2003; Frischknecht et al. 1994). These databases offer data on either individual unit processes (e.g., operating a tractor), covering materials, transportation, energy, and products, or aggregated data (e.g., the CF of 1 kg of wheat). Meanwhile, it is typically more informative for consumers to know the carbon footprint of a compound meal, such as pasta with tomato sauce and parmesan cheese, cooked for 10 minutes. As such data are typically not readily available, producing them from the data for the meal components requires the analyst to fill the gaps, including the materials they are composed of, their quantities, and how the preparation or assembly process works. An illustrative example is an LCA study on pizzas (Cortesi et al. 2023). Moreover, CO<sub>2</sub>-equivalent values are not always intuitive. For consumers, framing a carbon footprint in relatable terms (e.g., “like driving 10 km”) is often more meaningful.

To make this data more accessible and user-friendly, we explore using artificial intelligence (AI), particularly large language models (LLMs), augmented with domain knowledge. LLMs, such as the GPT and Claude series, provide natural language interfaces that can process queries in natural language and be used in an interactive, dialogue form. Since LLMs often lack information and may “hallucinate” an answer, they can be enhanced with domain-specific documents (Lewis et al. 2020) or structured data (Edge et al.

2024) through retrieval-augmented generation (RAG) techniques and scripted dialogues.

Recent studies have proposed integrating LLMs with LCA (Cornago, Ramakrishna, and Low 2023; Preuss, Al-shehri, and You 2024; Tu et al. 2024), often at a conceptual level. Applications are beginning to emerge, e.g., in hydrogen production (Chen et al. 2024), construction materials (Turhan 2023), and bio-based materials (Turhan-Haskara and Demirbilek 2024). In parallel, procedures for data retrieval in CF calculations (Luo et al. 2024; Wang et al. 2024), fine-tuning LLMs with CF data (Li et al. 2025) have been proposed, estimating food’s water footprint (Joshi, Ilievski, and Pujara 2024), or emission factor recommendation (Balaji et al. 2025). Still, we note a lack of methodologies that effectively model the CF of full meals, reconcile data from multiple overlapping databases, or provide analogs to intuitive processes.

In this paper, we investigate *how the results of LCA studies and existing LCA databases can improve the performance of LLMs in modeling the CF of cradle-to-gate food processes*. We develop a methodology that utilizes three freely accessible databases with granular carbon footprint information on food products and other items (De Rosa and Weidema 2019; Colomb et al. 2015; CONCITO 2021). Furthermore, we develop a chatbot interface that enables users to explore the carbon footprint of meals interactively and contextualize emission values with familiar activities (e.g., sending an email).

Our solution is implemented as a public demo, supported by open-source code. While still a proof of concept, our method and its implementation demonstrate how AI can help translate complex LCA data into user-friendly, actionable insights, acknowledging that uncertainties remain in both the underlying data and AI interpretation.

## Methodology

Our proposed methodology implements an integrated pipeline to generate environmental impact assessments from free-form recipe inputs, designed to overcome key challenges in food sustainability analysis. As illustrated in Figure 1, the methodology consists of four stages:

1. **Ingredient processing**, which uses an LLM to extract structured ingredient data from textual descriptions;
2. **Product matching**, which uses vector embeddings to match ingredients with database entries for products despite naming variations, followed by user verification;
3. **CF querying**, which retrieves CF values from the corresponding databases;
4. **Interactive exploration**, which synthesizes data from multiple LCA databases into comprehensive impact reports and visualizations, and enables the user to continue to chat with the large language model.

Existing LCA approaches often fail to serve non-expert users as they lack readily available data for compound meals and require specialized skills to assemble carbon footprints from individual ingredients. Our system addresses these limitations by leveraging AI techniques to interpret natural lan-

Database	BONSAI	Agribalyse	Big Climate DB
# Products	411	2,616	540
# Categories	-	11	13
Geo Focus	Global	France	Europe (5)

Table 1: Number of food products and categories, and geographical focus of the three data sources in our framework.

guage queries and match them with LCA data, while maintaining human oversight at critical decision points. This hybrid approach combines AI and humans, making complex LCA data accessible to laypeople while ensuring high quality through human validation of crucial mapping decisions. We describe our procedure in this section, whereas the exact prompts for the LLMs are provided in the Appendix.

## LCA Databases

Our system leverages environmental impact data from three key LCA databases: BONSAI (De Rosa and Weidema 2019), Agribalyse (Colomb et al. 2015), and Big Climate Database (CONCITO 2021). These databases were chosen primarily for their open-source availability, which ensures transparency and accessibility, while also providing robust, high-quality data critical for evaluating the carbon footprints of food products. Each brings distinct advantages that enhance the accuracy, granularity, and contextual relevance of our environmental impact assessments.

**BONSAI.** BONSAI is a global open-source database that provides multiregional input-output data, allowing detailed assessments of food product emissions in diverse geographical regions. Its hierarchical structure links products, activities, and recipes, offering a fine-grained view of supply chain emissions, from raw material extraction to market delivery. Inclusion of recipe-specific data further allows for precise tracking of inputs and CO<sub>2</sub> equivalent outputs in agricultural and manufacturing processes, making it an invaluable tool for complex analysis. Still, its coverage of regions and products remains incomplete, with full development expected by June 2025.

**Agribalyse.** Agribalyse is an open-source database tailored to the French food sector, providing comprehensive life cycle inventory (LCI) and impact data for more than 2,600 food products. It excels in covering the entire lifecycle—agriculture, processing, packaging, transportation, retail, and consumption—using a standardized LCA methodology that ensures consistency and reliability. A notable feature is its data quality rating for each entry, which gives confidence in the results and supports informed decision-making. However, its focus on French consumption patterns may limit its applicability elsewhere without adaptation.

**Big Climate Database.** It provides open-source, region-specific environmental impact data for 540 food products in five European countries: Denmark, the United Kingdom, France, Spain, and the Netherlands. Its standout feature is the incorporation of indirect impacts of land use change (iLUC), offering a broader perspective on the environmental

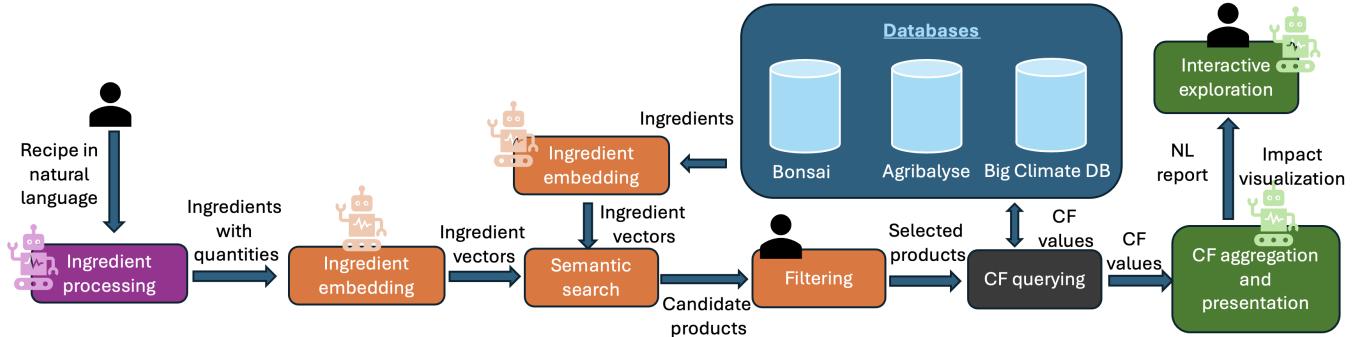


Figure 1: Overview of our four-step methodology. Each step is represented by one color: ingredient processing in purple, product matching in orange, CF querying in black, and iterative exploration in green. The blue container shows the three databases. Arrows indicate data flow. The robot icons indicate LLMs. The person icons indicate input from humans.

consequences of food production. Detailed emission breakdowns across the lifecycle stages—agriculture, processing, packaging, transport, and retail—further enhance their utility for targeted analysis. Although its product scope is narrower than Agribalyse’s, its regional precision compensates for this limitation.

## Ingredient Processing

Given a user input, the first step is converting unstructured recipe inputs into a standardized format suitable for database matching and impact calculation. The accuracy of this initial extraction step is crucial, as errors or omissions at this stage can spread throughout the entire pipeline, potentially leading to significant inaccuracies in the final impact assessment. Natural language recipe descriptions present several challenges. Ingredients may be described with varying levels of detail (e.g., “minced beef” versus “beef”), quantities can be specified in various units (cups, tablespoons, pieces), and key information might be embedded within cooking instructions or casual language (e.g., “a handful of spinach”).

To address these challenges, we opt for an LLM agent with a carefully engineered prompt that guides the extraction process. Specifically, we utilize an LLM through a function-calling interface that enforces structured output. The LLM is prompted with explicit instructions to identify ingredient names and quantities, with the requirement to convert all measurements to grams using standard culinary conversions (e.g., 1 tablespoon of oil  $\approx$  15g). The model returns a JSON object containing an array of ingredient objects, each with name and quantity fields. The resulting ingredient names and their corresponding amounts in grams are then stored and serve as input for the subsequent semantic search phase.

## Product Matching

Establishing accurate correspondences between user-provided ingredients and database entries presents a significant challenge due to variations in naming conventions across different LCA databases. Direct string matching fails to capture critical semantic relationships - for instance, “minced beef” and “ground beef” refer to the same product but would be considered distinct under

exact matching. Semantic search techniques address this limitation by matching ingredients based on their meaning rather than their exact textual representation. Semantic search leverages vector space representations that capture the contextual relationships between different food items, enabling the system to identify relevant products even when their names differ significantly from the user’s input.

We use an embedding model to generate vector embeddings for both ingredient names and database entries. These embeddings capture semantic relationships between food items based on their contextual usage patterns in large text corpora. We maintain a pre-computed index for each product name in the database (BONSAI, Agribalyse, and Big Climate Database), storing normalized embeddings for all product names. The ingredient names are transformed into vectors when processing a user query using the same model and normalization process. The system then performs efficient similarity searches using these indices to retrieve the top 3 most similar products from each database, employing cosine similarity as the distance metric. This approach ensures that semantically related products are matched appropriately while maintaining computational efficiency through optimized nearest neighbor search algorithms.

Although semantic search provides a robust method for identifying potential product matches, the inherent complexity of food products and variations in naming conventions require human oversight in the final selection process. Our method includes an interactive product selection interface, where users are presented with up to three candidate products from each database for each ingredient in their recipe. For each ingredient, the system displays the ingredient name, quantity, and a list of semantically similar products identified from BONSAI, Agribalyse, and the Big Climate Database, using checkboxes. For instance, both “ground beef” from BONSAI and “minced beef” from the Big Climate Database might be displayed when searching for a “beef” ingredient. Users can select multiple products for each ingredient to account for cases where different databases use different names for the same product or when a combination of products better represents the intended ingredient. If the user cannot find a suitable match, no re-

sults are provided for that ingredient. The selected products are assigned to their corresponding database identifiers and passed to the next stage for impact calculation. If a user selects a product for which data is not available for the target country, we calculate an average impact value using the available data from other countries. This human-in-the-loop approach strikes a balance between automation and accuracy, ensuring that the final environmental impact assessment is based on correctly matched products while maintaining an efficient and user-friendly workflow.

## Carbon Footprint Querying

From BONSAI, we extract both production-level emissions data and market-based emissions that account for regional production shares. Specifically, when querying BONSAI for a selected product, the system retrieves direct production emissions, where available, followed by a comprehensive breakdown of market composition that shows percentage contributions from different geographic origins. Each market share entry includes the region name, percentage value, and associated emissions data, preserving the multi-tiered data model essential for understanding global supply chains. From Agribalyse, we gather total impact values expressed in kg CO<sub>2</sub>-equivalent, along with their data quality ratings and detailed lifecycle stage breakdowns that partition impacts into six distinct phases: agriculture, processing, packaging, transportation, retail, and consumption. The Agribalyse query process preserves both the absolute values and the proportional contribution of each lifecycle stage. From the Big Climate Database, we retrieve region-specific total impact values and their detailed breakdowns across production phases, with a particular focus on their unique components, such as indirect land-use change impacts, which are not captured in other databases. The Big Climate Database query also includes identifying regional variations when data for the target country is unavailable.

The system structures this information into a standardized results text as input for the LLM's analysis. It creates a section for each ingredient, starting with its name and quantity, followed by results from each available database. When BONSAI data are available, we first list any direct production emissions, followed by a detailed breakdown of market shares and their associated emissions for different regions. For Agribalyse matches, we include the total impact value, data quality rating, and percentage contributions from each stage of the lifecycle. Similarly, for Big Climate Database matches, we present the total impact and phase-specific contributions, including unique indirect land-use change impacts. All impact values are consistently expressed in kilograms of CO<sub>2</sub> equivalent per specified ingredient quantity, enabling straightforward comparison and aggregation during the analysis step.

## Interactive Exploration

The last phase of the system uses an LLM to integrate structured environmental impact data from various databases into a cohesive evaluation and presentation of a recipe's CF. This process is systematically executed through a series of well-defined steps that ensure accuracy, consistency, and trans-

parency in the analysis. The LLM interprets structured data extracted from multiple environmental databases using a carefully crafted prompt that provides explicit instructions for parsing data formats unique to each source. For example, the prompt enables the LLM to calculate market-based total impact values from BONSAI, lifecycle-based aggregated values from Agribalyse, and region-specific values from the Big Climate Database, using the extracted data. The prompt also specifies protocols for reconciling data from multiple sources, such as calculating minimum and maximum impact ranges and derived average values. In addition, it includes guidelines for estimating cooking-related impacts when relevant, ensuring that LLM can unify heterogeneous data into a consistent analytical output.

The recipe's environmental impact computation follows predefined rules in the LLM's prompt. Per ingredient, the LLM determines a range of impact values (minimum to maximum) when multiple values are available, or adopts a single value when data are sourced from a single database. When the recipe implies cooking processes (e.g., baking, boiling), the LLM estimates supplementary impacts based on standard energy consumption metrics associated with the cooking method, duration, and temperature. The total environmental impact is then calculated by aggregating the impacts of all ingredients and any associated cooking processes, resulting in both a range and an average value. To enhance transparency, the system flags deviations from the user-specified target country, ensuring that the analysis remains contextually relevant.

The LLM generates a structured output comprising a natural language summary and accompanying visualization data. The summary ranks the recipe's ingredients by carbon footprint and presents the total impact range and average. Moreover, it contextualizes these findings with relatable comparisons to everyday activities (e.g., number of emails that cause a similar footprint), calibrated against reference values manually extracted from the book "How Bad Are Bananas?" (Berners-Lee 2020). The visualization data supports the creation of graphical representations, including a horizontal bar chart and a pie chart. These visuals delineate the proportional contributions of each ingredient and cooking process to the overall carbon footprint, ensuring that the results are presented in a clear and standardized manner that supports user comprehension.

In its initial response, the system presents only the most critical information: ingredient impacts ordered by magnitude, cooking considerations, total recipe impact, and relatable comparisons. However, the LLM maintains access to all detailed data, including regional market shares, lifecycle stage breakdowns, and granular emissions data across different stages of production and distribution. This selective presentation approach prevents information overload while allowing for iterative, in-depth exploration through follow-up questions. Users can ask for specific aspects, such as market share distributions, regional variations, or lifecycle patterns, and the LLM draws upon the complete dataset to provide detailed responses. This creates an interactive experience where users can progressively explore the complex aspects of their recipe's environmental impact through natu-

ral dialogue.

## Implementation Details

The methodological pipeline is implemented as our **Carbon Footprint Wizard** tool, which utilizes a combination of AI tools to enable efficient processing, matching, and analysis of recipe data. Ingredient processing leverages Open AI's GPT-4o-mini (Hurst et al. 2024), configured with a temperature parameter of zero for deterministic outputs, to extract structured ingredient lists from natural language inputs via a function-calling interface that enforces JSON formatting. For semantic search, the all-MiniLM-L6-v2 Sentence Transformer model generates vector embeddings for ingredient names and database entries, capturing semantic relationships.<sup>2</sup> The resulting embeddings are then indexed and queried efficiently using an FAISS index (Douze et al. 2024) to retrieve the most similar products based on cosine similarity. Database querying and result synthesis are facilitated by custom functions that interface with BONSAI, Agribalyse, and Big Climate Database. The chatbot interface uses GPT-4o-mini to process structured impact data and generate natural language responses and visualization data with zero temperature setting, integrated into a Gradio-based (Abid et al. 2019) web interface for user interaction. For real-world equivalence scaling, we extracted eight entries from daily life activities in "How bad are bananas?" (Berners-Lee 2020), such as sending an email (0.004 kg CO<sub>2</sub>-eq), driving 1 mile in a Fiat 500 (0.35 kg CO<sub>2</sub>-eq), and hand washing dishes (8.0 kg CO<sub>2</sub>-eq), specifically chosen to help users better understand and contextualize the CF of the recipe through relatable comparisons. Visualization of the CF, presented as a horizontal bar chart and a pie chart, is implemented using the Matplotlib library (Hunter 2007), which allows for intuitive and standardized graphical representations of ingredient impacts.

## Case Study Analysis

We show the effectiveness of our methodology and the corresponding Carbon Footprint Wizard tool through a step-by-step analysis using a vegetarian pizza recipe. We do not perform a quantitative evaluation because there are no ground-truth values for the CF of the meals. We outline the step-by-step workflow with the outputs of our system, highlighting its ability to provide detailed environmental impact assessments from natural language recipe descriptions.

Figure 2 illustrates the initial **ingredient processing** stage of our interface, where the user submits a recipe description in natural language: *"Could you estimate the environmental impact of my veggie pizza? Ingredients: 200g of pizza dough, 100g of tomato sauce, a handful of shredded mozzarella, half a red onion, a few olives, and a sprinkle of oregano."* The user selects the Netherlands as the target country for the analysis. The Carbon Footprint Wizard's LLM successfully extracts and standardizes the ingredients, converting qualitative descriptions into quantified amounts: *pizza dough (200g), tomato sauce (100g), shredded mozzarella (75g), red*

*onion (70g), olives (30g), and oregano (5g)*. This demonstrates the system's ability to interpret ambiguous quantities (e.g., "a handful," "half," "a few," "a sprinkle") and convert them into standardized weight measurements in grams.

After ingredient extraction, the system performs **ingredient matching**. For this purpose, the tool's semantic search retrieves and presents multiple potential product matches for each ingredient to the user through the product selection interface, as shown in Figure 4 in the Appendix. For example, *"pizza dough"* returns seven distinct matches including "bread," "mixes and doughs," "pastry," "pizza base, cooked," "pizza base, raw," "pizza dough", and "pizza with tomato and cheese, ready meals". Similarly, for *"shredded mozzarella"*, options include "cheese from skimmed cow milk", "cheese, semihard, mozzarella, 30% fdm", "mozzarella cheese, from cow's milk", and "processed cheese, in slices". The interface reveals the complexity of ingredient classification across databases: *"red onion (70g)"* yields numerous options including "garlic", "onion, cooked", "onion, dried", "onion, raw", "onions, dry", and "red onion", with only the exact match *"red onion"* selected by the user as the most adequate one. For *"olives"*, the system provides specialized options like "olives, black, without stones, in brine" alongside generic *"olives"*, with both selected to capture comprehensive data. The interface also demonstrates the breadth of alternative product classes that could be semantically related but conceptually distinct, such as *"tapenade"* as an option for *"olives"* and *"cream sauce with herbs"* as an option for *"oregano"*. This deliberate selection process addresses the challenge of varying terminology across different databases while allowing users to match ingredients based on their specific recipe knowledge. The asterisk notation indicates products without country-specific data, informing users that estimates from other countries will be used in such cases. This human-in-the-loop approach ensures accuracy in product matching while maintaining an intuitive user experience. Using the selection from the user, the **CF querying** step is performed in the background.

Following user selection and database querying for CF values, the LLM performs a comprehensive CF analysis and visualization in the **interactive exploration** phase. The natural language report shown in Figure 3 presents the ingredients ranked by their environmental impact, with mozzarella cheese identified as the highest contributor, followed by pizza base, olives, red onion, tomato sauce, and oregano in descending order. For each ingredient, the system notes when the numbers are estimated based on data from countries other than the user-selected target country (the Netherlands), enhancing transparency in data provenance. While the system should typically estimate cooking impacts even when not explicitly mentioned in the recipe, in this specific case, it incorrectly reports no cooking impact (0 kg CO<sub>2</sub>-eq) for the pizza preparation, thus violating common sense. The analysis provides both a total recipe impact range (0.14-0.84 kg CO<sub>2</sub>-eq) and calculates an average impact value of 0.49 kg CO<sub>2</sub>-eq as the final carbon footprint estimate, which represents only the ingredients without cooking emissions. To contextualize these values, the system translates the carbon footprint into equivalent everyday activities: sending ap-

<sup>2</sup><https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>. Accessed: 2025-04-17.

The Carbon Footprint Wizard 

Recipe Input   Product Selection   Chat with Assistant

1) Enter Your Recipe

Enter your recipe with ingredients and their quantities. You can specify quantities in grams, tablespoons, or other measurements.

Could you estimate the environmental impact of my veggie pizza? Ingredients: 200g of pizza dough, 100g of tomato sauce, a handful of shredded mozzarella, half a red onion, a few olives, and a sprinkle of oregano.

Select Target Country

Netherlands

Submit Recipe

Ingredients successfully extracted!

Ingredient	Amount (grams)
Pizza dough	200
Tomato sauce	100
Shredded mozzarella	50
Red onion	75
Olives	30
Oregano	1

Continue to Product Selection

Figure 2: Recipe input interface with ingredient processing from natural language (shown on the left) to standardized quantities (organized as a two-column table on the right).

proximately 122 emails, watching TV for about 2 hours, or driving 1.4 miles in a Fiat 500. This comparison helps users conceptualize the environmental impact in relatable terms. Additionally, the system suggests potential follow-up questions to promote further exploration of the results.

The visual representation of the carbon footprint analysis in Figure 5 features a horizontal bar chart and a pie chart. These visualizations provide a clear breakdown of each ingredient's contribution to the total environmental impact, allowing users to quickly identify which recipe components have the largest carbon footprints. The charts reveal that mozzarella cheese (49%) and pizza base (26%) are the dominant contributors to the carbon footprint. At the same time, tomato sauce contributes only 3% despite making up a substantial portion of the recipe by weight.

Figure 6 demonstrates the interactive capabilities of the system, showing a follow-up inquiry on the market shares of ingredients. In response, the system provides detailed information on the geographical origins of each ingredient. Although market share data are unavailable for most ingredients, olive data reveal market share from various countries, including Italy, Germany, and Portugal. This interaction showcases the system's ability to provide granular details from the underlying databases not included in the initial summary, demonstrating the depth of information available through the conversational interface. On the flip side, it shows the importance of having high-quality LCA databases that fuel the tool in the background.

## Discussion

The Carbon Footprint Wizard effectively bridges the gap between complex environmental impact databases and user-friendly insights. By combining natural language processing for ingredient processing, user-assisted semantic search for product matching, and a conversational interface for interactive result exploration, the system makes LCA data accessible to non-expert users. At the same time, our analysis reveals three key challenges with missing data, inconsistencies between the databases, and violations of common sense by the LLM.

First, the Carbon Footprint Wizard reveals challenges regarding **data availability**. The three databases used (BON-SAI, Agribalyse, and Big Climate Database) exhibit notable gaps in country-specific product coverage, which necessitates relying on data from other regions in many cases. This geographical mismatch is evident in the user interface, where asterisks denote products that lack data for the selected country. In our pizza example, several ingredients showed data from France instead of the requested data from the Netherlands, potentially reducing the accuracy of region-specific assessments.

Second, **inconsistencies** between the footprints reported in the databases result in considerable impact ranges rather than precise values, as evidenced by the wide min-max intervals (0.13-0.84 kg CO<sub>2</sub>-eq) presented in the results. This variability is partly due to inherent differences in how each database handles lifecycle phases and system boundaries. Another aspect of data consistency concerns the lack of standardized vocabulary used in the different databases. The semantic matching process, while generally effective, occa-

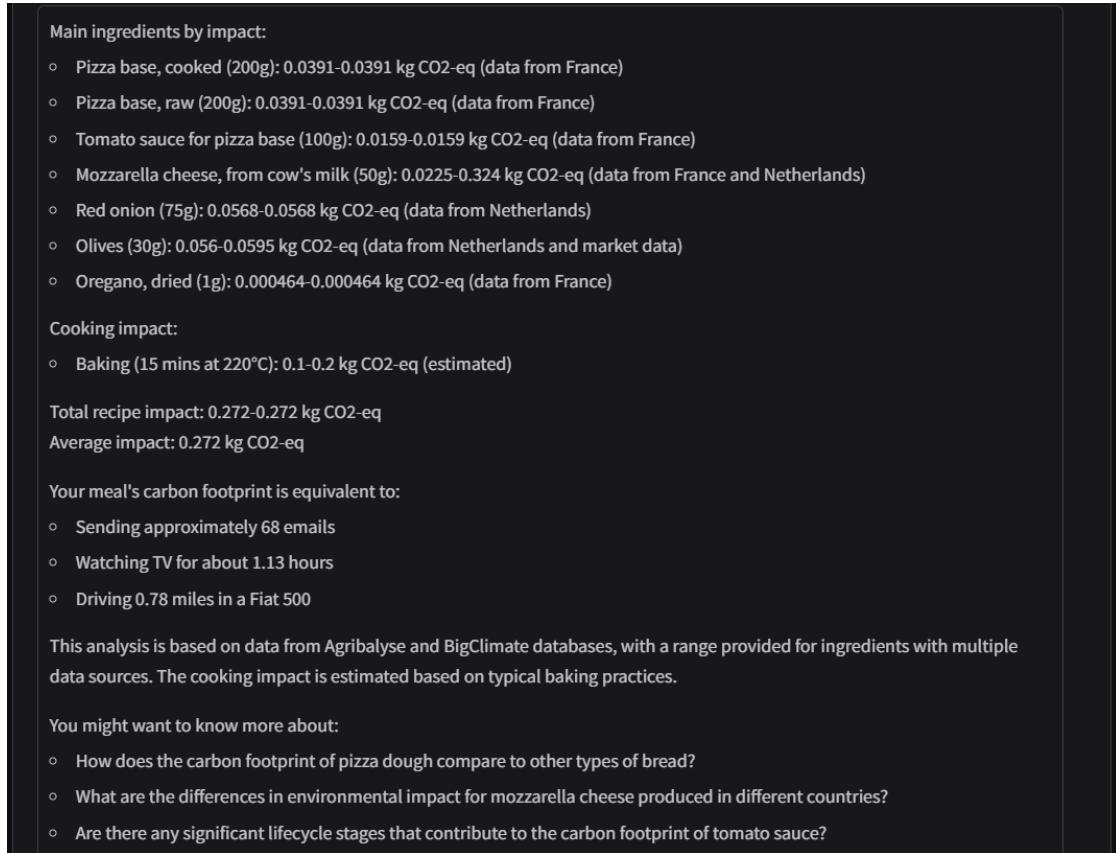


Figure 3: Initial analysis showing ranked ingredients, impact range, and real-world equivalents as a natural language summary.

sionally suggests irrelevant products (such as "tapenade" for olives or "cream sauce with herbs" for oregano), reflecting the inherent challenges of grounding natural language ingredient descriptions into standardized database terminology. While our current methodology employs a human-in-the-loop approach to mitigate this issue (by allowing users to select appropriate matches), it introduces dependency on user expertise and judgment, creating another potential source of variability in the final assessment results.

Finally, the system also demonstrates limitations in **handling complex analytical tasks using a single LLM**. For instance, the cooking impact estimation occasionally fails to detect implied cooking requirements, as demonstrated by the veggie pizza example, where the system incorrectly indicated "no cooking required" despite pizza preparation typically involving baking. Such commonsense violations can potentially lead to a significant underestimation of the total environmental impact of the recipe, as cooking can contribute substantially to the CF of a meal. Similarly, the generation of follow-up questions sometimes falls short of identifying the most relevant aspects for further exploration, focusing on generic queries rather than insights specific to the analyzed recipe's unique characteristics. These issues stem from the ambitious scope of responsibilities assigned to a relatively small language model (GPT-4o-mini), including ingredient extraction, impact calculation, cooking esti-

mation, visualization, data preparation, scaling real-world equivalence, and follow-up question generation. The model must simultaneously interpret heterogeneous data formats from multiple databases, perform mathematical operations, generate natural language explanations, and create structured visualization data, all within a single inference pass. This overextension of the model's capabilities could be addressed in future iterations by agentic AI, which can decompose these tasks into separate, specialized LLM calls that focus on distinct aspects of the analytical pipeline. This approach may involve employing larger and more capable models for the most complex reasoning tasks.

## Conclusions

This paper introduced an interdisciplinary methodology that bridges state-of-the-art LCA frameworks and databases with the current generation of knowledge-augmented AI. The methodology's four steps leveraged LLMs for seamless user interaction, fueled by LCA databases with contextual knowledge about food products. The methodology was designed to involve users in the ingredient selection step and to allow users to explore the data further after the summarized report was presented. The methodology is implemented in the publicly available Carbon Footprint Wizard web interface. The initial analysis of the tool's functionality showcased the power of connecting LCA insights with LLMs,

providing an intuitive entry point to complex information for lay users. At the same time, the analysis revealed challenges with missing data, inconsistencies between the databases, and violations of common sense by the LLM. These challenges point to substantial interdisciplinary obstacles that can be fruitfully explored in the future: the engineering and representation of LCA knowledge, the development of LCA data quality frameworks, and ensuring consistency in LLM behavior. Agentic AI frameworks and advanced entity resolution approaches can be employed to address these limitations. Future work should address two limitations of our approach. First, the lack of quantitative validation of our method necessitates benchmarking against expert LCAs, conducting sensitivity/error analyses, and conducting user studies to demonstrate accuracy, robustness, and usability at scale. Second, our three databases differ in geographical scope, system boundaries, and category definitions, leading to wide impact ranges and limited comparability. Thus, future work should revise the selection of LCA databases to account for their differences in coverage and focus.

## Acknowledgements

This work has been supported by the seed grant “Knowledge-based support for sustainable diet policies” awarded to the authors by the Vrije Universiteit’s Amsterdam Sustainability Institute (ASI).

## References

Abid, A.; Abdalla, A.; Abid, A.; Khan, D.; Alfozan, A.; and Zou, J. 2019. Gradio: Hassle-free sharing and testing of ml models in the wild. *arXiv preprint arXiv:1906.02569*.

Balaji, B.; Ebrahimi, F.; G Domingo, N. G.; Vunnava, V. S. G.; Faridee, A.-Z.; Ramalingam, S.; Gupta, S.; Wang, A.; Gupta, H.; Belcastro, D.; et al. 2025. Emission Factor Recommendation for Life Cycle Assessments with Generative AI. *Environmental Science & Technology*.

Berners-Lee, M. 2020. *How bad are bananas? The Carbon Footprint of Everything*. Profile Books.

Chen, Y.; Liebau, U.; Guruprasad, S. M.; Trofimenco, I.; and Minke, C. 2024. Advancing Life Cycle Assessment of Sustainable Green Hydrogen Production Using Domain-Specific Fine-Tuning by Large Language Models Augmentation. *Machine Learning and Knowledge Extraction*, 6(4): 2494–2514.

Colomb, V.; Ait-Amar, S.; Basset-Mens, C.; Gac, A.; Gaillard, G.; Koch, P.; Mousset, J.; Salou, T.; Tailleur, A.; and Van Der Werf, H. M. 2015. AGRIBALYSE®, the French LCI Database for agricultural products: high quality data for producers and environmental labelling.

CONCITO. 2021. The Big Climate Database, version 1.0.

Cornago, S.; Ramakrishna, S.; and Low, J. S. C. 2023. How can Transformers and large language models like ChatGPT help LCA practitioners? *Resources, Conservation and Recycling*, 196: 107062.

Cortesi, A.; Colpaert, M.; Saint-Eve, A.; Maurice, B.; Bris, G. Y.-L.; Souchon, I.; and Pénicaud, C. 2023. Data on the Life Cycle Assessment of pizzas cooked and consumed at home taking into account the variability of consumer practices. *Data in Brief*, 48: 109143.

De Rosa, M.; and Weidema, B. P. 2019. BONSAI - Big Open Network for Sustainability Assessment Information. <https://bonsai.uno/>.

Deru, M. 2009. US life cycle inventory database roadmap (brochure). Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States).

Douze, M.; Guzhva, A.; Deng, C.; Johnson, J.; Szilvassy, G.; Mazaré, P.-E.; Lomeli, M.; Hosseini, L.; and Jégou, H. 2024. The faiss library. *arXiv preprint arXiv:2401.08281*.

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. *arXiv preprint arXiv:2404.16130*.

Finkbeiner, M. 2016. *Special types of life cycle assessment*. Springer.

Frischknecht, R.; Hofstetter, P.; Knoepfel, I.; Dones, R.; and Zollinger, E. 1994. Environmental Life-Cycle Inventories of Energy Systems. *Bundesamt für Energiewirtschaft July*.

Hertwich, E. G.; and Wood, R. 2018. The growing importance of scope 3 greenhouse gas emissions from industry. *Environmental Research Letters*, 13(10): 104013.

Hunter, J. D. 2007. Matplotlib: A 2D graphics environment. *Computing in science & engineering*, 9(03): 90–95.

Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.

Joshi, S.; Ilievski, F.; and Pujara, J. 2024. Knowledge-powered recommendation for an improved diet water footprint. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 23805–23807.

Krajnc, D.; and Glavić, P. 2003. Indicators of sustainable production. *Clean technologies and environmental policy*, 5: 279–288.

Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Kuttler, H.; Lewis, M.; tau Yih, W.; Rocktäschel, T.; Riedel, S.; and Kiela, D. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *Advances in Neural Information Processing Systems*, 33: 9459–9474.

Li, Z.; Tang, P.; Wang, X.; Liu, X.; and Mou, P. 2025. PCF-RWKV: Large Language Model for Product Carbon Footprint Estimation. *Sustainability*, 17(3).

Luo, B.; Liu, J.; Deng, Z.; Yuan, C.; Yang, Q.; Xiao, L.; Xie, Y.; Zhou, F.; Zhou, W.; and Liu, Z. 2024. AutoPCF: A Novel Automatic Product Carbon Footprint Estimation Framework Based on Large Language Models. *Proceedings of the AAAI Symposium Series*, 2(1): 102–106.

PET, P. T. 2003. Eco-profiles of the European plastics industry. *Sea*, 370000(3400000): 3800000.

Preuss, N.; Alshehri, A. S.; and You, F. 2024. Large language models for life cycle assessments: Opportunities, challenges, and risks. *Journal of Cleaner Production*, 466: 142824.

Purvis, B.; Mao, Y.; and Robinson, D. 2019. Three pillars of sustainability: in search of conceptual origins. *Sustainability science*, 14: 681–695.

Schebesta, H.; Bernaz, N.; and Macchi, C. 2020. The European union farm to fork strategy. *European food and feed law review*, 15(5): 420–427.

Tu, Q.; Guo, J.; Li, N.; Qi, J.; and Xu, M. 2024. Mitigating Grand Challenges in Life Cycle Inventory Modeling through the Applications of Large Language Models. *Environmental Science & Technology*, 58(44): 19595–19603. PMID: 39431580.

Turhan, G. D. 2023. Life Cycle Assessment for the Unconventional Construction Materials in Collaboration with a Large Language Model. In *Proceedings of the International Conference on Education and Research in Computer Aided Architectural Design in Europe*. Education and research in Computer Aided Architectural Design in Europe.

Turhan-Haskara, G. D.; and Demirbilek, A. D. 2024. Large Language Model-aided Life Cycle Analysis for Circular Bio-manufacturing. In *MATEC Web of Conferences*, volume 401, 05005. EDP Sciences.

Wang, H.; Zhang, M.; Chen, Z.; Shang, N.; Yao, S.; Wen, F.; and Zhao, J. 2024. Carbon Footprint Accounting Driven by Large Language Models and Retrieval-augmented Generation. *arXiv preprint arXiv:2408.09713*.

Wernet, G.; Bauer, C.; Steubing, B.; Reinhard, J.; Moreno-Ruiz, E.; and Weidema, B. 2016. The ecoinvent database version 3 (part I): overview and methodology. *The International Journal of Life Cycle Assessment*, 21: 1218–1230.

Yañez, P.; Sinha, A.; and Vásquez, M. 2019. Carbon footprint estimation in a university campus: Evaluation and insights. *Sustainability*, 12(1): 181.

Čuček, L.; Klemeš, J. J.; and Kravanja, Z. 2012. A Review of Footprint analysis tools for monitoring impacts on sustainability. *Journal of Cleaner Production*, 34: 9–20. Recent Cleaner Production Advances in Process Monitoring and Optimisation.

## Reproducibility Checklist

### 1. General Paper Structure

1.1. Includes a conceptual outline and/or pseudocode description of AI methods introduced (yes/partial/no/NA) **yes**

1.2. Clearly delineates statements that are opinions, hypothesis, and speculation from objective facts and results (yes/no) **yes**

1.3. Provides well-marked pedagogical references for less-familiar readers to gain background necessary to replicate the paper (yes/no) **yes**

### 2. Theoretical Contributions

2.1. Does this paper make theoretical contributions? (yes/no) **no**

If yes, please address the following points:

- 2.2. All assumptions and restrictions are stated clearly and formally (yes/partial/no) **NA**
- 2.3. All novel claims are stated formally (e.g., in theorem statements) (yes/partial/no) **NA**
- 2.4. Proofs of all novel claims are included (yes/partial/no) **NA**
- 2.5. Proof sketches or intuitions are given for complex and/or novel results (yes/partial/no) **NA**
- 2.6. Appropriate citations to theoretical tools used are given (yes/partial/no) **NA**
- 2.7. All theoretical claims are demonstrated empirically to hold (yes/partial/no/NA) **NA**
- 2.8. All experimental code used to eliminate or disprove claims is included (yes/no/NA) **NA**

### 3. Dataset Usage

3.1. Does this paper rely on one or more datasets? (yes/no) **yes**

If yes, please address the following points:

- 3.2. A motivation is given for why the experiments are conducted on the selected datasets (yes/partial/no/NA) **yes**
- 3.3. All novel datasets introduced in this paper are included in a data appendix (yes/partial/no/NA) **NA**
- 3.4. All novel datasets introduced in this paper will be made publicly available upon publication of the paper with a license that allows free usage for research purposes (yes/partial/no/NA) **NA**
- 3.5. All datasets drawn from the existing literature (potentially including authors' own previously published work) are accompanied by appropriate citations (yes/no/NA) **yes**
- 3.6. All datasets drawn from the existing literature (potentially including authors' own previously published work) are publicly available (yes/partial/no/NA) **yes**
- 3.7. All datasets that are not publicly available are described in detail, with explanation why publicly available alternatives are not scientifically satisfying (yes/partial/no/NA) **NA**

### 4. Computational Experiments

4.1. Does this paper include computational experiments? (yes/no) **no**

If yes, please address the following points:

- 4.2. This paper states the number and range of values tried per (hyper-) parameter during development of

the paper, along with the criterion used for selecting the final parameter setting (yes/partial/no/NA) **NA**

- 4.3. Any code required for pre-processing data is included in the appendix (yes/partial/no) **NA**
- 4.4. All source code required for conducting and analyzing the experiments is included in a code appendix (yes/partial/no) **NA**
- 4.5. All source code required for conducting and analyzing the experiments will be made publicly available upon publication of the paper with a license that allows free usage for research purposes (yes/partial/no) **NA**
- 4.6. All source code implementing new methods have comments detailing the implementation, with references to the paper where each step comes from (yes/partial/no) **NA**
- 4.7. If an algorithm depends on randomness, then the method used for setting seeds is described in a way sufficient to allow replication of results (yes/partial/no/NA) **NA**
- 4.8. This paper specifies the computing infrastructure used for running experiments (hardware and software), including GPU/CPU models; amount of memory; operating system; names and versions of relevant software libraries and frameworks (yes/partial/no) **yes**
- 4.9. This paper formally describes evaluation metrics used and explains the motivation for choosing these metrics (yes/partial/no) **NA**
- 4.10. This paper states the number of algorithm runs used to compute each reported result (yes/no) **NA**
- 4.11. Analysis of experiments goes beyond single-dimensional summaries of performance (e.g., average; median) to include measures of variation, confidence, or other distributional information (yes/no) **NA**
- 4.12. The significance of any improvement or decrease in performance is judged using appropriate statistical tests (e.g., Wilcoxon signed-rank) (yes/partial/no) **NA**
- 4.13. This paper lists all final (hyper-)parameters used for each model/algorithm in the paper's experiments (yes/partial/no/NA) **NA**

## Prompts Used in the System

### Ingredient Extraction Prompt

Extract the ingredients and their quantities from the following user message.

If the quantities are provided in units other than grams (such as tablespoons, cups, or pieces), convert them to a reasonable weight in grams based on typical measurements.

You can directly change milliliters to grams.

Provide the ingredients and quantities in the format required by the function 'process\_ingredients'.

**User Message:** {user\_message}

3. After impacts, calculate equivalent activities:  
- Use total average (Z kg CO<sub>2</sub>-eq) with these references:
  - Sending an email = 0.004 kg CO<sub>2</sub>-eq
  - Web search on a laptop = 0.0007 kg CO<sub>2</sub>-eq
  - Watching TV 42-inch plasma = 0.24 kg CO<sub>2</sub>-eq
  - Driving 1 mile Fiat 500 = 0.35 kg CO<sub>2</sub>-eq
  - 3-minute shower = 0.09 kg CO<sub>2</sub>-eq
  - Charging phone daily = 0.003 kg CO<sub>2</sub>-eq
  - Using laptop 1 hour = 0.05 kg CO<sub>2</sub>-eq
  - Hand washing dishes = 8.0 kg CO<sub>2</sub>-eq

4. List some follow up questions:

- Generate 3-4 most relevant follow-up questions based on:
  - Market share of ingredients
  - Notable data variations between countries
  - Interesting lifecycle patterns
  - Potential impact reduction opportunities

### Result Generation Prompt

You will receive two inputs:

- user\_message: Original recipe query
- results\_text: Impact data from databases

#### DATA FORMAT OVERVIEW:

- BONSAI: Shows market/production data with country shares. Use market total impact value.
- Agribalysse: French data with lifecycle stages. Use total impact value.
- BigClimate: Country-specific data with indirect land use. Use total impact value.

#### CALCULATION RULES:

1. Per ingredient:

- Multiple databases: Use min-max range and (min+max)/2 average
- Single database: Use value for both
- Note if data isn't from target country
- Use market data from BONSAI when available

2. Recipe total:

- Range: sum(min) to sum(max)
- Average: sum(individual averages)
- Cooking analysis:
  - Determine if the recipe requires any form of cooking for preparation
  - If cooking is required:
    - Estimate cooking impact range (min-max)
    - Add cooking range to total recipe range
    - Add cooking average to total average
    - Include cooking in visualization data
  - If no cooking required: state "No cooking required (0 kg CO<sub>2</sub>-eq)"

#### OUTPUT FORMAT:

- Main ingredients by impact:
- [Ingredient] ([amount]g): [X-Y kg CO<sub>2</sub>-eq] (note if data is from a different country)
- Cooking impact:
- [Method] ([time] mins at [temp]°C): [X-Y kg CO<sub>2</sub>-eq]
- Total recipe impact: [X-Y kg CO<sub>2</sub>-eq]
- Average impact: [Z kg CO<sub>2</sub>-eq]

Your meal's carbon footprint is equivalent to:

- [Scaled comparisons]

[Brief paragraph: data sources, range explanation, cooking estimate]

You might want to know more about:

- [List the follow-up questions]

**CRITICAL:** You must return TWO components in your response as the defined "process\_impact\_results":

1. answer\_user: Follow the output format above
2. visualization\_data: Must include these exact fields:
  - ingredients: Array of ingredient names from the user recipe (including cooking if applicable)
  - impacts: Array of corresponding impact values in kg CO<sub>2</sub>-eq (use average when range exists)
  - DO NOT include total recipe impact on visualization!

After the initial response, user will continue chatting with you. Do not use function format after initial response, answer directly.

Here is user message:

{user\_message}

Here is the information from our sources:

{results\_text}

## Example Outputs from Each Step

## Product Matching

Below are the example outputs from each step of the environmental impact analysis process for a vegetarian pizza recipe.

### User Recipe Input

Could you estimate the environmental impact of my veggie pizza? Ingredients: 200g of pizza dough, 100g of tomato sauce, a handful of shredded mozzarella, half a red onion, a few olives, and a sprinkle of oregano.

### Ingredient Extraction

Ingredient	Amount (grams)
Pizza dough	200
Tomato sauce	100
Shredded mozzarella	75
Red onion	70
Olives	30
Oregano	5

Ingredient	Matching Products
Pizza dough	Bread, Mixes and doughs, Pastry, Pizza base, cooked *, Pizza base, raw *, Pizza dough, Pizza with meat, tomato and cheese, ready meals, Pizza with tomato and cheese, ready meals, Pizza, cheese and mushrooms *
Tomato sauce	Chili sauce, Cream sauce *, Juice of tomatoes, Soya sauce, Tomato sauce for pizza base *, Tomato sauce, with onions, prepacked *, Tomatoes, fresh, Tomatojuice, canned
Shredded mozzarella	Cheese from skimmed cow milk *, Cheese, semihard, mozzarella, 30 % ffdm., Mozzarella cheese, from cow's milk *, Plant-based cheese, without soybean, prepacked, shredded *, Processed cheese, Processed cheese, in slices *, Tomato, peeled, canned, Vegan cheese, grated or sliced, Whey cheese
Red onion	Garlic, Onion, cooked *, Onion, dried *, Onion, raw, Onions, chopped and deepfried, Onions, dry, Onions, shallots (green), Red onion
Olives	Oil of olives, virgin, Olive oil, Olive oil, extra virgin *, Olives, Olives, black, without stones, in brine, Olives, green, pickled, canned, Pizza, onion anchovy and black olives *, Tapenade (a puree of capers, pitted black olives, anchovy and herbs, with olive oil and lemon juice) *
Oregano	Cream sauce with herbs *, Garlic, Garlic, raw, Ginger, Olive oil, Oregano, dried *, Rosemary, dried *, Soya paste, Sunflower oil

A screenshot of the matching step in our interface is shown in Figure 4.

## Carbon Footprint Data

The results text is a structured compilation of environmental impact data for each selected product from the databases (e.g., AgriBalyse, BigClimateDatabase, BON-SAI). It includes sections for each ingredient, with detailed data such as total impact and contributions from various lifecycle stages. Below is an excerpt for the 'pizza dough' ingredient:

### Results for selected most similar items to 'pizza dough':

#### Agribalyse database results for 'Pizza base, raw' (DATA FROM FRANCE):

- Impact for 200 grams: 0.0391 kg CO<sub>2</sub>-eq
- Data quality rating: 2.3277205962237506
- Agriculture impact for 200 grams: 0.00267 kg CO<sub>2</sub>-eq, Percentage: 6.8%
- Processing impact for 200 grams: 0.00128 kg CO<sub>2</sub>-eq, Percentage: 3.3%
- Packaging impact for 200 grams: 0.00121 kg CO<sub>2</sub>-eq, Percentage: 3.1%
- Transportation impact for 200 grams: 0.000711 kg CO<sub>2</sub>-eq, Percentage: 1.8%
- Retail impact for 200 grams: 0.000381 kg CO<sub>2</sub>-eq, Percentage: 1.0%
- Consumption impact for 200 grams: 0.0014 kg CO<sub>2</sub>-eq, Percentage: 3.6%

#### BigClimateDatabase results for 'Pizza dough' in Netherlands:

- Impact for 200 grams: 0.241 kg CO<sub>2</sub>-eq
- Agriculture impact for 200 grams: 0.174 kg CO<sub>2</sub>-eq, Percentage: 72.2%
- Indirect Land Use Change impact for 200 grams: 0.011 kg CO<sub>2</sub>-eq, Percentage: 4.6%
- Food processing impact for 200 grams: 0.00465 kg CO<sub>2</sub>-eq, Percentage: 1.9%
- Packaging impact for 200 grams: 0.0354 kg CO<sub>2</sub>-eq, Percentage: 14.7%
- Transport impact for 200 grams: 0.0155 kg CO<sub>2</sub>-eq, Percentage: 6.4%
- Retail impact for 200 grams: 0.000312 kg CO<sub>2</sub>-eq, Percentage: 0.1%

Similar data is provided for all ingredients (tomato sauce, shredded mozzarella, red onion, olives, oregano) in the recipe, detailing impacts from relevant databases.

A visual summary of the analysis results provided by our interface is shown in Figure 5.

## Natural Language Report

An example report in our interface is shown in Figure 3. An example of a follow-up user query and a system response is shown in Figure 6.

### Main ingredients by impact:

- Pizza base (200g): 0.0331–0.241 kg CO<sub>2</sub>-eq (note: data from France and Netherlands)
- Tomato sauce (100g): 0.0159 kg CO<sub>2</sub>-eq (note: data from France)
- Mozzarella cheese (75g): 0.0338–0.486 kg CO<sub>2</sub>-eq (note: data from France and Netherlands)
- Red onion (70g): 0.053 kg CO<sub>2</sub>-eq (note: data from Netherlands)
- Olives (30g): 0.056–0.0595 kg CO<sub>2</sub>-eq (note: data from France)
- Oregano (5g): 0.00232 kg CO<sub>2</sub>-eq (note: data from France)

### Cooking impact:

- No cooking required (0 kg CO<sub>2</sub>-eq)

### Total recipe impact: 0.13942–0.84032 kg CO<sub>2</sub>-eq

### Average impact: 0.48987 kg CO<sub>2</sub>-eq

### Your meal's carbon footprint is equivalent to:

- Sending approximately 122 emails
- Watching TV for about 2 hours
- Driving 1 mile in a Fiat 500 approximately 1.4 times

The data sources include AgriBalyse and BigClimate databases, with a range calculated from the minimum and maximum impacts of the ingredients. No cooking is required for this recipe, so the total impact reflects only the ingredients.

### You might want to know more about:

- What are the market shares of the ingredients used in this recipe?
- How do the impacts of these ingredients vary between different countries?
- Are there any lifecycle patterns that stand out for the ingredients?
- What are some potential opportunities to reduce the carbon footprint of this recipe?

# The Carbon Footprint Wizard 🌱

Recipe Input   **Product Selection**   Chat with Assistant

## 2) Select Your Products

Select the most similar products for each ingredient listed. If multiple options are relevant, please choose all that apply to ensure accuracy.

**Note:** Items marked with an asterisk (\*) don't have data specific to the selected country and will use estimates from other countries.

**Pizza dough (200g)**

Bread    Mixes and doughs    Pastry    Pizza base, cooked \*    Pizza base, raw \*    Pizza dough

Pizza with meat, tomato and cheese, ready meals    Pizza with tomato and cheese, ready meals    Pizza, cheese and mushrooms \*

**Tomato sauce (100g)**

Chili sauce    Cream sauce \*    Juice of tomatoes    Soya sauce    Tomato sauce for pizza base \*

Tomato sauce, with onions, prepacked \*    Tomatoes, fresh    Tomatojuice, canned

**Shredded mozzarella (50g)**

Cheese from skimmed cow milk \*    Cheese, semihard, mozzarella, 30 % ffdm.    Mozzarella cheese, from cow's milk \*

Plant-based cheese, without soybean, prepacked, shredded \*    Processed cheese    Processed cheese, in slices \*

Tomato, peeled, canned    Vegan cheese, grated or sliced    Whey cheese

Figure 4: The product selection showing multiple database matches for each ingredient, resulting from the semantic search in the product matching phase.

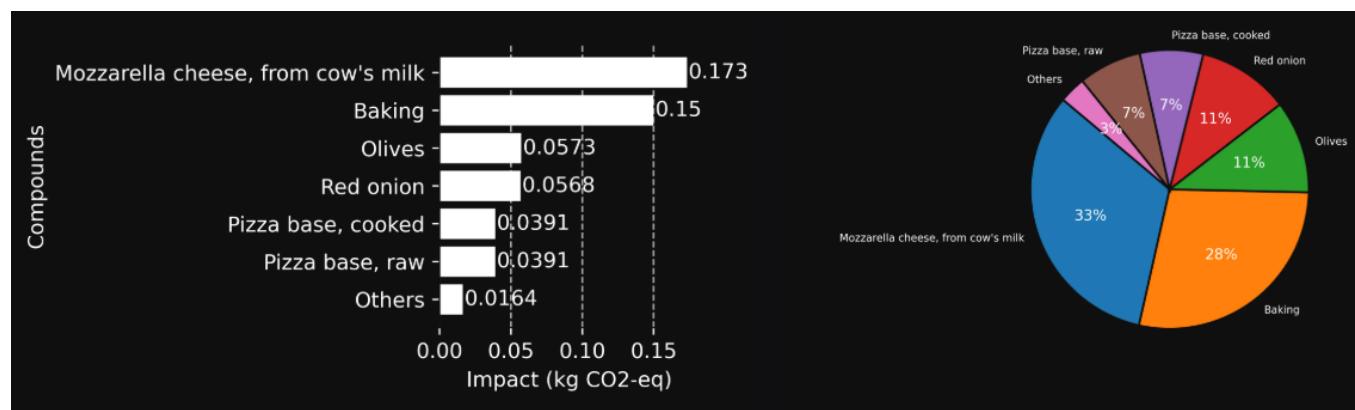


Figure 5: Carbon footprint visualization through bar and pie charts showing ingredient impacts.

Can you show me the market shares for the available products?

Here are the market shares for the available products based on the data from the BONSAI database:

1. Olives (30g):

- Market share for various countries:
- Bulgaria: 0.000202 tonnes
- Italy: 0.434 tonnes
- Germany: 0.282 tonnes
- Portugal: 0.117 tonnes
- Belgium: 0.0453 tonnes
- Greece: 0.0837 tonnes
- Spain: 0.0312 tonnes
- France: 0.00391 tonnes
- United Kingdom: 0.00337 tonnes

For the other ingredients like pizza base, tomato sauce, mozzarella cheese, red onion, and oregano, the market share data is not available in the provided databases. Therefore, I can only provide market share information for olives. If you need information on specific ingredients or alternative data sources, please let me know!

Your Message

Send

Figure 6: Follow-up conversation revealing detailed market share data for ingredients.