Summary

This paper concerns a CO2e estimation and visualisation system for food procurement tailored to Aalborg Municipality, and contributes to a current research gap in sHCI for supporting organisations and public institutions in implementing sustainable food practices. The system itself relies on municipal invoice data from food procurement to estimate CO2e emissions, which is used for generating CO2e reports for kitchens in the municipality. The system is designed through the use of user-centred design in collaboration with a stakeholder, who is the sustainability controller of Aalborg Municipality. The system aims to address the shortcomings of the current system used by the municipality. Their current system is slow and lacks granularity in terms of splitting up the CO2e emissions into smaller categories. The reason the current system is slow is due to the current generation of CO2e emissions, as the sustainability controller currently manually generates reports for each kitchen, which is a slow process. The administrator system described in this paper aims to address these shortcomings by automating report generation and transparency in terms of the estimation of the CO2e emission.

As mentioned, the paper follows the iterative process that is used during the user-centred design process. During this project, we had 3 iterations of user-centred design. Iteration 1 builds upon the knowledge we got from our previous work, which is described in [14]. We, therefore, began Iteration 1 by designing a wireframe prototype, which we would use during Interview 1, with the sustainability controller, as a starting point for a discussion about how an administrator system for CO2e emission estimation and visualisation would fit into their daily work. After gaining more information about how the sustainability controller in the municipality spends their time during work, it was clear our wireframe prototype would need some refinement to provide better support for the work they perform. Therefore, we began Iteration 2 by gaining knowledge about how such a system could support them. After gaining this knowledge, we made some requirements based on the information gained during Interview 1. Based on these requirements, we developed a Blazor web application, which is the administrator system described throughout this paper. The application is a fully functional web app, where the sustainability controller can: Upload invoices containing food procurement data where the system can identify products from the invoice and map them to an entity in an LCA database using an LLM, view uploaded invoices already in the system, view graphs displaying CO2e emissions estimations for kitchens in the municipality, view a graph displaying the organic share of the products procured by the kitchens and view a graph displaying the total emission for all food products procured by all kitchens in the municipality. The system is also able to generate targeted CO2e reports for all kitchens in the municipality, where the system automatically anonymises all other kitchens in the report except for the targeted kitchen. After finishing the development of the administrator system, we conducted an evaluation of the system with the sustainability controller, where we uncovered some minor shortcomings in the system. Some of these shortcomings were in regards to the CO2e reports, where the sustainability controller had some feedback that would significantly improve the use of them; they also had some other minor suggestions to the system, which were in relation to the use of some of the pages. Most of the suggestions the sustainability controller provided were implemented during Iteration 3, where we addressed the shortcomings of the CO2e reports and implemented other minor changes that would enhance the use of the system. Some of the suggestions were not implemented, as they either were out of the scope of this project or due to the time constraint of the project.

The implications of the study in the field of sHCI and in a broader context are discussed, in which we argue for the aspects of the system that promote scalability of the system, while providing generic support for kitchens in municipalities throughout Denmark in terms of effectively implementing simplistic visualisations with a necessary level of granularity. Furthermore, we discuss the benefits and disadvantages of the design choices that we have made, therein how the use of a user-centred design over another design methodology helped us design a system that met the needs of stakeholders. Lastly, we discuss privacy concerns that may arise through the use of the system, and propose ideas for data minimisation, among other potential requirements for future iterations of the project.

Overall, the system proved useful for the sustainability controller in Aalborg Municipality in terms of being very efficient in generating targeted CO2e reports for the kitchens. The accuracy of the mapping was also high, which we uncovered during testing of the system. Some of the features in the system are, however, lacking; for instance, the sustainability controller suggested that the system could have a forecasting mechanism that could be used to determine whether the municipality was on track to hit their goal of minimising their CO2e emissions on a municipal scale [1]. This feature is not implemented due to time constraints, but should be implemented in future iterations of the project.

Designing a CO2e Emissions Estimation and Visualisation System for Municipal Food Procurement: A Case Study from Aalborg Municipality

Marcus Hasselgaard Skifter Jensen
Aalborg University
mhsj20@student.aau.dk

Tobias Martin Pihl
Aalborg University
tpihl20@student.aau.dk

ABSTRACT

The food sector is a major contributor to global greenhouse gas emissions. This paper presents the design, development, and evaluation of a CO2e emissions estimation and visualisation system tailored for Aalborg Municipality's food procurement practices. Grounded in user-centred design and informed by Life Cycle Assessment (LCA) principles, the system addresses key shortcomings of the municipality's current estimation system. The solution leverages invoice data, semantic mapping via an LLM, and structured categorisation to provide accurate emissions estimation data across kitchens in the municipality. It includes features for visualising emissions by food category, normalising emissions per person served or per DKK spent, and generating anonymised, kitchen-specific CO2e reports. Through iterative prototyping and stakeholder feedback, the system was refined to support strategic decision-making and provide insight into intervention areas. The resulting tool significantly improves the efficiency and relevance of emissions reporting, offering a scalable model for public sector sustainability initiatives.

General Terms

Sustainable HCI (sHCI), User-Centred Design (UCD), Life Cycle Assessments (LCA), Greenhouse gases (GHG), CO2 equivalent (CO2e), Low fidelity (Lo-fi), High fidelity (Hi-fi).

Keywords

Sustainability, food products, food procurement, sHCI, CO2e emissions, administrator system.

1. INTRODUCTION

In recent years, efforts to reduce CO2e emissions have been a focus for governments and organisations, therein municipalities [1], promoting a more sustainable future. Key initiatives include ambitious climate policies and advancements in green technology [1]. However, meaningful reductions require ongoing innovation, collaboration across sectors, and a commitment to transformative change. By focusing on critical areas such as transportation, agriculture, buildings, and waste management, while enforcing stricter sustainability measures, the global community, including the EU, can move closer to a low-carbon economy, fostering long-term environmental and economic resilience [9].

The food sector plays a significant role in global greenhouse gas (GHG) emissions, contributing approximately 26% of the total emissions [10]. The CO2e emissions in the food sector arise from various stages, including livestock, crop production, land use, and supply chains (which include processing, transport, packaging, and retail). Notably, livestock production alone accounts for about 31% of the total GHG emissions in the food sector [10].

Transitioning to more sustainable food habits can substantially reduce these emissions. Adopting plant-based diets, for instance, can lower the climate footprint associated with food consumption [9]. Additionally, local sourcing of food reduces transportation-related emissions, as consuming local food limits the distance products must travel, thereby directly decreasing GHG emissions [15]. Furthermore, reducing food waste is a crucial step towards minimising GHG emissions. In the European Union, food waste accounts for approximately 16% of the total GHG emissions from the food system[15]. In 2020, nearly 59 million tonnes of food were wasted in the EU, equating to 131 kg per person annually. This substantial amount of waste contributes significantly to environmental degradation and underlines the need for effective food waste reduction strategies [15].

Large-scale organisations, such as municipalities, are well-positioned to lead by example in reducing the food sector's climate footprint [9]. By addressing inefficiencies and promoting more sustainable practices, such as supporting plant-based diets and local food initiatives, municipalities can minimise their GHG emissions associated with food systems [9]. Integrating measures such as these into a broader sustainability framework, such as municipalities, can enhance their role in driving climate action. These strategies highlight the substantial potential for emission reductions within the food sector through comprehensive and sustainable approaches.

To further investigate this initiative, it is necessary to gain a more detailed understanding of the climate footprint at the municipal level. Specifically, by focusing on the kitchens in Aalborg

Municipality, this project aims to address key limitations in the current calculation method for food-related CO2e emissions. One of the main challenges is that their existing CO2e calculation model is based on economic value, with regard to inferring estimated weights of food products, rather than the actual weight of food products. This makes their current method sensitive to price fluctuations rather than reflecting the true environmental impact [1]. Moreover, the process itself is slow in the sense that it requires several days to compute the CO2e emissions; therefore, their current CO2e estimation model is run yearly, leading to the kitchens within the municipality not being able to view their emissions throughout the year. As a consequence of this, the staff in the kitchens are unable to reflect and adjust their food procurement based on their CO2e emissions during the year.

Additionally, individual kitchens in Aalborg Municipality currently receive only a total CO2e emission value rather than a breakdown by specific food categories. As the existing model only works with 22 broad food categories, each with its emission factor, this lack of granularity makes it difficult for kitchens to identify which food groups contribute most to their climate footprint [1]. Without access to detailed insights, kitchens face challenges in making data-driven decisions to reduce their environmental impact.

Aalborg Municipality has already recognised the need for improved data quality and is working towards refining calculations based on food weight rather than expenditure [1]. By enhancing the calculation method and providing more precise data, this project aligns with those ongoing efforts, increasing transparency and supporting more effective CO2e reduction strategies within Aalborg Municipality. This leads us to the following research question:

How can we design a tool that enables Aalborg Municipality to accurately and efficiently estimate their food-related CO2e emissions?

2. RELATED WORK

This section reviews relevant work in three key areas: Life Cycle Assessment (LCA), municipal-level emission tracking, and Sustainable Human-Computer Interaction (HCI). LCA offer a method for evaluating environmental impacts across the food life cycle. Municipal-level tools extend this analysis to help cities track and reduce emissions. Meanwhile, the HCI community has explored digital tools to raise awareness and encourage behaviour change, often focusing on individual consumers. Together, these domains inform the design of effective, scalable solutions for sustainable food practices in public and organisational settings.

2.1 Life Cycle Assessment

LCA is a widely used methodology for assessing the environmental impacts of products, processes, and systems throughout their entire life cycle. The assessment typically considers all stages, from raw material extraction, production, transportation, and use, to end-of-life disposal or recycling [5]. LCA plays a crucial role in sustainability research by providing a quantitative approach to evaluate GHG emissions, energy consumption, and resource depletion associated with various stages of the product's life cycle from "cradle to grave" [12].

In the food sector, LCA has been instrumental in identifying the key sources of environmental impact. Studies have shown that food production contributes significantly to global CO2e emissions, and with an upcoming need to increase food production by more than 60% by 2050 due to the world population increasing. Therefore, there is a serious need for change in the food production sector [16]. A previous study by Poore and Nemecek[19] found that livestock production is particularly carbon-intensive, to such a degree that meat, eggs, and dairy use approximately 83% of the available farmland in the world [19]. Additionally, it accounts for approximately 56-58% of the total CO2e emissions in the food sector, despite only providing 37% of our protein and as low as 18% of our calories [19]. Plant-based foods generally have a lower environmental footprint. This highlights the importance of dietary shifts and sustainable food production strategies in reducing the climate footprint [19].

LCA methodologies vary based on system boundaries and functional units. Attributional LCA (ALCA) focuses on static environmental impacts related to specific products or services, whereas consequential LCA (CLCA) assesses the broader economic and systemic changes from a given decision, which is a cause-effect relationship between demand and other relevant changes in supply [6]. For instance, if there suddenly is a demand for more milk, it influences other parts of the food supply chain, and could cause a higher demand in other sectors such as animal feed for the cows [6]. The choice of LCA approach significantly influences the results and their applicability in policy and decision-making [17].

A key challenge in applying LCA to food systems is the variability of data sources and the complexity of food supply chains. Many rely on LCA databases, which contain estimated emissions factors; therefore, transparency and reviews of the factors are important to allow for reproducibility and avoid hidden manipulation in terms of tweaking of the factors [12]. Moreover, factors such as land use change, soil degradation, and biodiversity loss are difficult to quantify accurately but are crucial for a more comprehensive sustainability assessment [5]. Addressing these uncertainties requires robust data collection methods and collaboration between researchers and industry stakeholders [5].

Despite these challenges, LCA remains a fundamental tool for guiding more sustainable decision-making. Municipalities and larger organisations can leverage LCA-based tools to assess the environmental impact of food procurement, identify high-emission food categories, and implement targeted strategies to reduce CO2e emissions [16]. Given the urgency of climate change mitigation, integrating LCA into public food systems is an essential measure for achieving their long-term sustainability goals [16].

2.2 Municipality Emission Tracking

Municipalities are critical actors in climate mitigation, not only due to their role in urban planning and service delivery, but also because they are in a unique position to monitor and manage local GHG emissions. While much of the literature on environmental impact assessment has traditionally focused on product-level or sector-level analyses using LCA methodologies, recent studies have begun to address the challenges and opportunities associated with emission tracking at the municipal level[5, 6, 12].

An additional contribution to this field is the FEWprint carbon accounting platform [23]. This integrated tool offers municipalities all over the world a comprehensive method to assess carbon emissions associated with the consumption of Food, Energy, and Water (FEW) resources. Designed as a spreadsheet-based framework, FEWprint calculates consumption-based footprints by analysing data from various sectors, including food consumption, thermal and electrical energy use, car fuel demand, water management, and domestic waste processing [23].

The platform comprises three key components, each corresponding to its primary objectives:

(1) Data Collection and Processing: Gathers detailed information on resource consumption across the FEW sectors, ensuring that data inputs are both comprehensive and manageable for users [23].
(2) Emission Calculation: Utilises standardised emission factors to convert consumption data into carbon emission estimates, facilitating consistent and comparable assessments [23].
(3) Scenario Analysis and Visualisation: Allows users to model various scenarios to understand potential impacts of different policy decisions or behavioural changes, aiding in strategic planning and decision-making [23].

A comparative assessment using FEWprint across six different communities revealed significant variations in total annual emissions, underscoring the platform's capability to highlight specific areas for targeted interventions. For example, the study found that emissions from the food sector ranged between 993 kg CO2e per capita per year in Amsterdam and 1,366 kg CO2e per capita per year in Tokyo, indicating both the impact of local consumption patterns and the potential for tailored mitigation strategies [23].

By providing an accessible and integrative approach to carbon accounting, FEWprint enables municipalities to identify emission hotspots and evaluate the effectiveness of various mitigation strategies. This holistic perspective is particularly valuable for urban areas aiming to transition towards more sustainable resource management practices. The platform's design addresses common challenges in municipal emission tracking, such as the complexity of the FEW nexus and the need for user-friendly tools that do not require exhaustive data inputs [23].

2.3 Sustainable HCI

In the field of sustainable HCI(sHCI), there has previously been a large focus in studies of designing systems targeted for individual consumers or consumers in social contexts with the goal of raising awareness and creating engagement in sustainable food practices, rather than a focus on designing systems for larger organisations or social contexts. Several studies focusing on sustainable food practices, such as *EcoMeal* [7] and *Habits* [3], have taken the angle of creating a connection between sustainable food practices and gamification to foster engagement and motivation for the individual consumer in understanding the environmental impact of food. EcoMeal, a mobile application, aided users in understanding the environmental impact of their daily food consumption by letting users scan barcodes on the packaging of food products [7]. By scanning the barcode of a food product, factors which had an influence on the ecological footprint of the product were revealed [7]. The study found that, after a user study with 14 participants, EcoMeal encouraged long-term usage, and the aspect of gamification increased user engagement [7]. Concerning gamification in fostering sustainable food practices, Habits combined eco-feedback technologies using a carbon calculator and gamification by implementing features such as challenges and a leaderboard for users to compare their carbon footprint against other users [3]. The study concluded that, based on a user-study with 10 participants for the gamified carbon calculator *Habits*, gamification has potential to engage users in sustainability, but mostly new users, as the application was not adaptive for users that progressed and started making larger changes towards a sustainable lifestyle through use of the application [3]. Moreover, the authors argued, based on their results, that there is a need for a social context, as users would otherwise lose interest in using the application, despite the gamification aspects [3].

Aside from studies using gamification to promote sustainable food practices, other studies, such as *Budget Your Carbon Emissions* [20] and *Visualizing Carbon Footprint from School Meals* [11], have instead focused on visualizing carbon footprints to users. The study *Budget Your Carbon Emissions* explored different designs to determine the best design for conveying a carbon budget on an individual level and thereby contextualizing a carbon budget, such as a circle, negative slope and representation of the Earth, where the visualizations would change as users continued to spend their allotted carbon budget [20]. The study found, through a user study of 10 participants, that the tool increased the awareness of users in relation to the carbon budget in relation to different aspects such as food, travel and more, which affected the carbon budget positively or negatively [20]. However, the study concluded that the proposed designs presented too much information and therefore resulted in clutter, which limited the comprehensibility of the information presented in the designs, where a simple and precise layout would have been preferred [20]. The study *Visualising Carbon Footprint from School Meals* developed a prototype for calculating and visualising GHG emissions from schools, and testing of the prototype was conducted using data for food procurement from 9 different schools in Stockholm, Sweden [11]. The study used an LCA database, which the authors of the study had developed, called *LCAFDB*, which contained CO2 data for food categories, such as "chicken" and "ox" [11]. The CO2 emission data estimated for the food products was visualised in terms of total CO2 per school, and a distribution of CO2e emissions per type of food [11]. In the study, schools were divided into two sectors: low-emission schools with an average GHG footprint of 1.6 kg CO2e/kg and high-emission schools with an average GHG footprint of 3.43 kg CO2e/kg [11]. The authors concluded that the major difference between low-emission schools and high-emission schools was that low-emission schools served almost no type of meat, namely 0.05% of the total weight of food purchased [11]. In high-emission schools, meat accounted for between 23.3% to 31.6% of the total emissions, despite that the weight of meat purchased was between 2.7% to 4.3% of the total weight of food purchased, whereas low-emission schools substituted protein-based foods such as meat with soy, legumes and mushrooms [11]. In addition to these results, the authors added that 71% of the pupils in the 9 schools were in high-emission schools [11]. The study discussed that the results presented do not initially indicate that foods such as meat should be directly replaced due to their high carbon footprint with other foods, as the GHG emissions of foods in their model are related to weight, without taking nutritional values into account [11].

In summary, previous studies highlight the importance of LCA in understanding the environmental impacts of food systems, the emerging role of municipalities in tracking and managing emissions, and the potential of sHCI tools to support sustainable behaviour change [11, 20, 7]. However, while individual-focused solutions have been widely explored, there remains a gap in designing digital tools that support larger organisations and public institutions in meeting sustainability goals, particularly in the context of food procurement [11, 20, 7]. This highlights an opportunity to develop systems that support sustainability efforts at a broader scale, particularly within institutional food procurement and municipal planning.

3. METHOD

In this paper, we conducted a single instrumental case study analysis [18] of Aalborg Municipality, in which we gathered knowledge of the current practices and solutions in Aalborg Municipality with regards to estimating CO2e emissions for food procurement. An instrumental case study is conducted to provide insight into a problem [18]. The case itself facilitates the understanding of the problem, and thereby acts as a supportive role for gaining knowledge of the problem in a broader context [18]. We employed several different methods during data gathering and analysis, where each method contributed to expanding our understanding of the problem and application domain. To conduct our study, we used the following key data collection and analysis methods: literature review, semi-structured interviews [8], user feedback analysis [8], and prototyping. Each method contributed to expanding our understanding of the problem and application domain.

First, we conducted a narrative literature review [2] to gain a broader understanding of the problem domain. A narrative literature review involves synthesising studies across multiple related, but distinct, areas and drawing connections between them to build a cohesive overview [2]. This type of review was chosen for its flexibility in exploring a broad and diverse body of literature, and when attempting to interconnect many studies on different topics [2]. Our review focused on examining existing research, case studies, and theoretical frameworks related to CO2e estimation in food procurement and sustainability management.

By reviewing prior work, we identified key challenges, potential solutions, and gaps in current methodologies. The key challenges we identified arose from prior work with the LCA methodology [6] and prior work such as FewPrint [23]. Methods such as these helped us obtain a better understanding of the problem domain when working with CO2e emissions from food products on a larger scale. Potential solutions also lead back to the prior work done with FewPrint, where the aim was to develop a larger-scale platform to track emissions for not only food, but also energy and water [23]. In our case, looking at how they handled the food emissions provided us with ideas, presented later in this paper, on how to handle food emissions on a larger scale. Additionally, the literature review provided insights into best practices and technological approaches, which helped shape the direction of our study and informed the design of our prototypes.

Second, we conducted several interviews with Aalborg Municipality, mainly with the sustainability controller. These interviews provided insights into the sustainability controller's daily workflow. The interviews also revealed the challenges the sustainability controller faces in estimating CO2e emissions from food procurement. The challenges, which the sustainability controller mentioned during our interviews, aligned with the challenges we uncovered in our literature review. In the municipality, they currently have issues with the current method that they use for estimating CO2e emissions and presenting estimated CO2e emissions, in terms of incomplete invoice data and a large maintenance overhead. Therefore, we chose to conduct qualitative interviews with the sustainability controller, compared to using a quantitative method for data collection, as the sustainability controller is our primary stakeholder.

During the interviews, we used two different types of prototypes. When selecting fitting prototypes, there are several different types of prototypes, each with its own benefits and disadvantages. Prototypes can mainly be split into two larger categories: low-fidelity (lo-fi) and high-fidelity (hi-fi) prototypes [21].

Hi-fi prototypes promote interactivity with the prototype to users, as the role of a facilitator presenting the prototype is removed [21]. Within the domain of hi-fi prototypes, there is a specialised prototype: horizontal hi-fi prototyping [21]. Horizontal hi-fi prototypes focus on the core functionality in the application without providing much depth to the lower-level details of the application [21].

Various advantages of using hi-fi prototypes rather than lo-fi prototypes are that hi-fi prototypes introduce functionality and thereby interactivity to the user, whereas the advantage of lo-fi prototypes is that they are cheap to produce [21]. By enabling interactions between users and the prototype, important usability errors and design recommendations informed by their interactions with the prototype can be discovered and corrected [21].

4. DESIGN PROCESS

In this section, we present the steps taken in designing a digital system aimed at supporting Aalborg Municipality in estimating and visualising CO2e emissions from food procurement. The structure of the section follows the iterative and exploratory nature of the design process that is involved with using a user-centred design approach. It begins with Iteration 1, which draws on insights gained from our work in the previous semester, and continues with Iteration 2 and Iteration 3. In user-centred design, designers use various investigative techniques such as surveys or interviews and generative techniques such as brainstorming with the purpose of gaining an understanding of the users' needs [13].

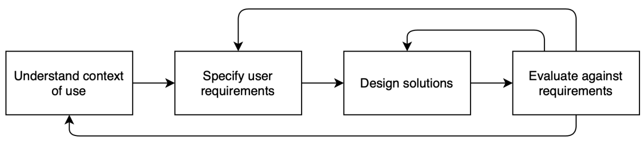

Fig. 1. User-centred design process

Following a user-centred design approach involves 4 phases that designers go through when designing a system for the stakeholders, as seen in Figure 1 [13]: understanding the context in which stakeholders would use such a system, specifying user requirements, designing solutions based on specified requirements, and evaluating the designed solutions against the specified requirements [13].

4.1 Findings from previous work

During our previous semester [14], we created a system for calculating CO2e emissions based on invoices for a single kitchen.

The system was structured in a way that a designated kitchen had to upload its invoices each time it procured food. The system would then estimate emissions based on the invoice content, and the emissions could then be used to generate simple graphs for a given selected period [14].

After we finished developing the system, we conducted a final evaluation interview where a kitchen manager and the sustainability controller of Aalborg Municipality were present. During this interview, we learned that the kitchen manager was not interested in using such a system, as they simply do not have time during their work schedule to upload their invoices. They were, however, very interested in the graphs the system was able to generate [14]. During this meeting, we learned that an administrator system, which the sustainability controller would manage and be responsible for uploading invoices, would be more suitable for Aalborg Municipality's situation. The administrator system should be similar to the old system. However, the kitchens should not be able to upload invoices, but only be able to view their emissions through the system [14].

4.2 Iteration 1: Design solutions

As we had an initial idea of how the administrator system should be structured and some unstructured requirements, through our previous work [14]. Therefore, we designed a prototype based on this knowledge. We decided to, based on the advantages of using a hi-fi prototype argued by Rudd et al. (1996) [21], and the aforementioned understanding of the problem that we had from our previous project [14], that it would be beneficial to use a horizontal hi-fi prototype for Interview 1 during this project. This prototype will serve as a reference point for our preliminary ideas for an administrative system, and it will be used for evaluating feasibility and usability during Interview 1.

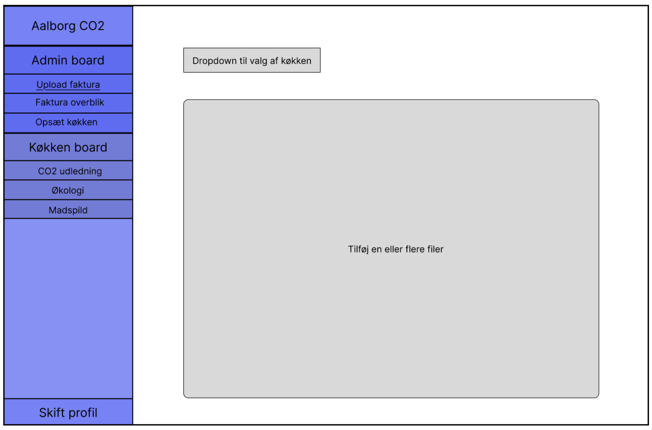

Fig. 2. Upload of invoices

The prototype contained the basic functionality that could be present in an administrator system that Aalborg Municipality would use. The administrator system should have the same functionality as our system from the previous semester, but it should be split into an administrator part and a kitchen part [14]. Therefore, the administrator system that we designed had a sidebar, as seen in Figure 2, which was split into two sections: an "Admin board" and a "Kitchen board". The reason for having two types of boards in the prototype was to make a clear distinction between the administrator, the sustainability controller, and the kitchens [14].

The "Admin board" contained 3 basic functionalities: a page for uploading invoices, an invoice overview and creating kitchens, which the administrator would be responsible for managing in terms of uploading invoices for each kitchen. The page for uploading invoices is shown in Figure 2. The page itself is simple, the user can select the designated kitchen to upload invoice data on food products for, and the option to upload one or more files containing said invoice data. The page for viewing uploaded invoices and their associated CO2e emissions is shown in Figure 3. On this page, the administrator can view the invoice data that has been uploaded for kitchens, with each uploaded invoice having an associated kitchen, total CO2e emission for the food products in the invoice data and the period for the product lines in the invoice data.

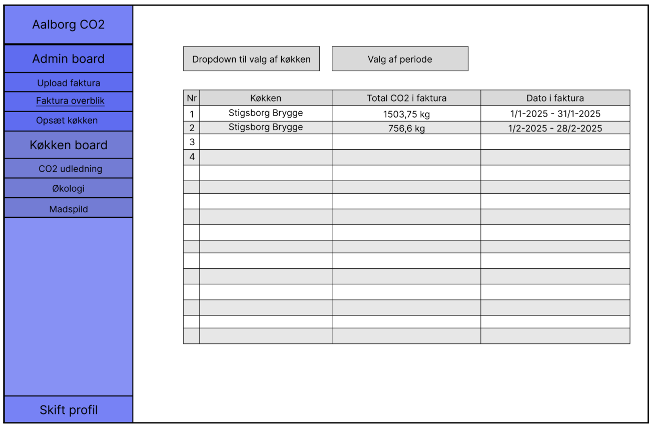

Fig. 3. Invoice overview

To avoid potential clutter when viewing the uploaded invoices, we added the possibility of selecting a kitchen and choosing a period to view uploaded invoices for. Finally, the page for creating kitchens can be seen in Figure 16, which is found in the appendix. On this page, the administrator can manage existing kitchens in the system or set up new ones. For each kitchen created in the system, there is an associated average number of people that the kitchen serves on average. This is included as the system should have appropriate means for normalising the CO2e emissions for each kitchen, given the various sizes of kitchens in the municipality. The normalisation of the CO2e emissions for each kitchen is necessary for the sustainability controller and kitchens to be able to compare their CO2e emissions to other kitchens in Aalborg Municipality. During our interview with the kitchen manager on the previous semester, we learned that comparing themselves to other kitchens is an important measurement for them, as it helps them to become better at food procurement [14].

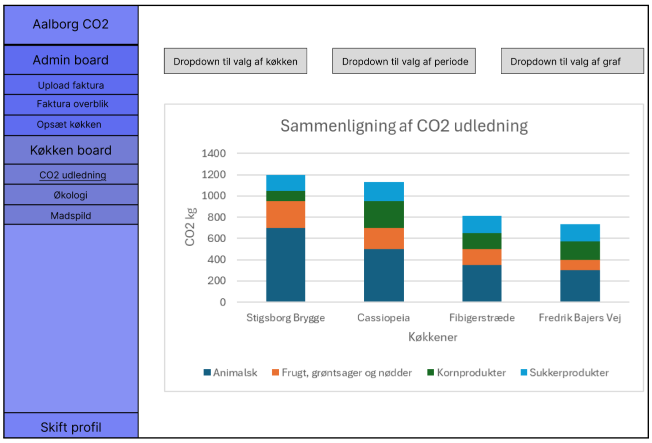

Fig. 4. Total CO2e emissions graph

For the "Kitchen board", which both the administrator and kitchens have access to, there were 3 different pages: viewing the CO2e emissions for one or more kitchens, the share of organic foods purchased by one or more kitchens and the food waste for each kitchen. The decision to include the share of organic foods purchased by kitchens was based on the interest that the municipality showed in their food strategy for 2024-2028 [1], in which they express the desire to increase the share of organic foods purchased compared to the share of non-organic foods that kitchens in the municipality purchase. The same goes for food waste, as the kitchens should have an increased focus on minimising food waste according to the food strategy [1].

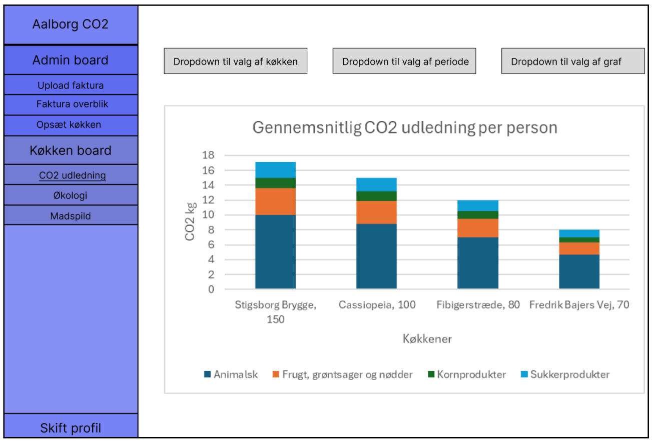

Fig. 5. CO2e emission per person served

As for the page for viewing CO2e emissions for food products purchased by kitchens, we included various graphs. These graphs included viewing a total CO2e emission for one or more kitchens, as seen in Figure 4 and an average CO2e emission per person served in the kitchen, as seen in Figure 5. The decision to include specifically these graphs was based on the last interview with the kitchen manager on our previous semester [14]. The CO2e emissions for food were sectioned into 4 different food categories, inspired by the Danish *Madpyramiden*, to convey the overall areas where kitchens had the most CO2e emissions. To keep a level of granularity for the CO2e emissions, users could, for instance, click on the category "Animalsk" on the bar chart to view the 10 most emitting food products within this category, as shown in Figure 17. This design choice was based on the interview with the kitchen manager, as the kitchen manager explained that simply viewing the general categories would not benefit them much, if they are not able to tell which products add to the specific CO2e emission of the given general category [14]. Moreover, the user has the option to choose one or more kitchens, select a period to view CO2e emissions for and choose which graph to view.

4.3 Iteration 1: Evaluate against requirements

After developing the prototype described in subsection 4.2, we conducted Interview 1 with the sustainability controller at Aalborg Municipality. The purpose of Interview 1 was to assess the degree to which the prototype aligned with the requirements derived from prior work and to validate its relevance.

The feedback obtained during Interview 1 confirmed that the system design addressed several of the core needs identified in our previous project. Most notably, the sustainability controller emphasised that two functionalities were central to their workflow:

> *"The CO2e estimation itself and targeted reports for the kitchens."*

The system was designed for estimating CO2e based on invoices, but the system was not designed for generating CO2e reports for kitchens. Therefore, these features should be a requirement for future iterations.

Regarding temporal granularity, our initial design did not have a set range, as this was something we wanted to get feedback on, whether it should be weekly, monthly or quarterly. The sustainability controller expressed that monthly intervals might lead to misinterpretation, as food purchases can fluctuate significantly between months, either due to bulk purchasing strategies or municipal events. As stated during the interview:

> *"A quarterly range will contribute to a more whole picture of each kitchen's CO2e emissions."*

Future data visualisations should support quarterly views, thereby enabling a more stable and representative analysis over time.

Another key insight related to the normalisation method. Initially, the system normalised CO2e emissions using "per person served" as the metric. Only having this metric for normalisation proved to be problematic, as noted by the sustainability controller:

> *"The kitchens do not necessarily control how many people they serve per day."*

Furthermore, he added that such a metric could be a reasonable idea; however, a normalisation using CO2e emission per DKK, which is what Aalborg Municipality currently does, would give the kitchens a better insight into their emissions. As the sustainability controller states:

> *"The kitchens can directly control which products they buy, and thereby change this metric."*

This led to a revision in the prototype, adding the metric "per DKK spent" along with the "per person served" as primary normalisation methods for CO2e metrics.

Additionally, the sustainability controller provided critical feedback regarding the level of detail in visualisations. Our initial approach of listing the 10–15 most emitting food products was considered too granular, with limited usefulness to kitchen staff. Instead, he advocated for grouping food products into predefined sub-categories—such as “Meat” or “Dairy products”—which better represent the intervention areas kitchens can act upon. This feedback directly influenced the restructuring of our food categorisation model, which in later iterations should be adjusted.

The evaluation also clarified the system’s intended user roles. It became evident that kitchens should not interact directly with the system. As the sustainability controller noted, kitchens are generally short on time. Therefore, the system should only be handled by the sustainability controller, and they should be responsible for generating CO2e reports for the kitchens, instead of providing the kitchens access to the system.

Along with the addition of CO2e reports, the feedback also confirmed the need for anonymised, comparative views across kitchens. As the controller explained:

> *“The kitchens prefer to see their CO2e emissions over a period and to be able to compare their emissions with other kitchens.”*

However, the comparison in the CO2e reports must be confidential:

> *“The names of other kitchens should be anonymised, such that the individual kitchen can only identify their emissions. However, they should be able to compare themselves to other kitchens.”*

Therefore, with the addition of CO2e reports, the system should anonymise all other kitchens than the targeted kitchen for the given report.

In summary, Interview 1 with the sustainability controller confirmed the relevance of the prototype but highlighted several areas for improvement. Key functionalities such as targeted CO2e reports were missing and should be prioritised in future iterations. The feedback emphasised the use of quarterly data intervals, normalisation by CO2e per DKK spent, and the need to group food products into actionable categories. Additionally, the system should be centrally operated, not by kitchens directly, and must support anonymised, comparative CO2e reports targeted to each kitchen.

4.4 Iteration 2: Understanding context of use

Interview 1 served both as an evaluation of our initial prototype and the starting point in the UCD process for Iteration 2. During Interview 1, it became clear that we needed a deeper understanding of the daily tasks of the sustainability controller. This realisation marked the beginning of Iteration 2, as we required the necessary information to refine the design of the prototype. To facilitate this, we prepared follow-up questions in advance, expecting the need for further inquiry beyond the prototype evaluation. These questions can be found in the Appendix in section 10.

In detail, what we wanted to learn was what basic functionalities should be present in an administrator system for estimating CO2e emissions for food in order to align the contents of the system with the tasks that the sustainability controller does daily. Additionally, we learned from the municipality’s food strategy for 2024-2028, that the municipality has had an initiative since 2013 to increase the share of organic foods purchased in the municipality [1]. Therefore, we had an interest in learning if the sustainability controller found this relevant to include in the system, as a continuous measure for the municipality to monitor their share of organic food purchased.

Regarding the task of estimating the CO2e emissions for food, a relevant aspect to consider was the food waste that each kitchen in the municipality may have. For example, in the case where food waste is repurposed as bio-fuel, this has a positive impact on the total CO2e emissions that a kitchen has [22]. For this reason, we needed to find out if it was possible to include food waste as a part of the CO2e estimations. Furthermore, for the CO2e estimations to have any value, gaining knowledge about the sustainability controller’s experiences with presenting CO2e emission data to the kitchens in the municipality was important. Specifically, we needed to learn which temporal ranges for presenting data were optimal and the preferred number of food categories, so as to avoid having an unnecessary level of detail when presenting data.

4.4.1 The basic functionality of the system.

During Interview 1, we inquired with the sustainability controller on the matter of the basic functionalities that should be present in the system.

> *”As I see it, there are two pinpoints: The CO2e estimation itself and targeted reports for the kitchens. These are the two things in my job that I spend the most time on.”*

With regard to the CO2e estimation itself, the sustainability controller does this in the following way: a single CSV file containing all product lines for the whole municipality is processed, and the CO2e emission estimation is computed. Hereafter, they produce the targeted reports for each kitchen. Moreover, we learned that the CSV files, which the sustainability controller extracts from the municipality’s database, contain information regarding which kitchen has purchased the product. Given that this is the case, we therefore learned that it was not necessary to include a drop-down button for selecting kitchens when uploading invoices, as seen in Figure 2.

The sustainability controller further elaborated on the targeted reports that they produce for the kitchens in their current solution in the municipality: the kitchens are interested in learning how much CO2e they emit based on their food purchases in a period, but not the number itself, they are more interested in graphs where they can identify food categories which have higher emission. The sustainability controller learned that the kitchens prefer to see their CO2e emissions over a period and to be able to compare their emissions with other kitchens.

Regarding the CO2e estimations for the kitchens in Aalborg Municipality, the sustainability controller currently produces bar charts for 14 kitchens in the municipality. The size of these kitchens varies heavily, as there are two major kitchens in the municipality, and the other kitchens are about 20% to 30% the size of these kitchens, which therefore needs to be accounted for in the estimation of the CO2e emissions of kitchens. The sustainability controller explains that, for the kitchens to be able to compare their CO2e emissions for a given period, the CO2e emissions for each kitchen are normalised by the amount of DKK that a kitchen has spent on its food purchases.

4.4.2 Measuring food waste and share of organic food.

With regards to the contents of the municipality's food strategy for 2024-2028 and its initiative to implement a more sustainable practice in its kitchens [1], there are two aspects related to sustainability that can be measured and reflected in the system: food waste in kitchens and the share of organic food purchased on a municipal scale.

As for implementing food waste as a feature that the municipality can track in the system, it was necessary to inquire whether or not this is something the municipality is interested in. Secondly, to implement the food waste in the system, documentation in the form of, for instance, invoices or any other type of data on food waste, was required. The sustainability controller explained the current situation in the municipality:

> *"As far as I know, there is only one kitchen that currently sorts food waste in the municipality. There is no requirement that the kitchens in the municipality sort food waste separately from other types of waste, meaning that most kitchens only have one container for waste."*

As there is no requirement for kitchens in the municipality regarding sorting food waste separately from other types of waste, this necessarily yields the fact that there are no processes in place for measuring the amount of food waste in kitchens across the municipality. Given that it is the case that there are no invoices for food waste across the municipality, this rendered the task of measuring the impact of food waste on the CO2e estimations for each kitchen impossible, although food waste is a relevant aspect to consider when estimating CO2e for food.

As stated in the food strategy [1], the municipality's goal for both increasing the share of organic food purchased across kitchens while retaining a collective organic food percentage of 64% [1]. We inquired with the sustainability controller on how organic food purchased in the municipality is currently measured:

> *"Currently, we measure the share of organic food purchased across our canteens on a weight-based approach, rather than a DKK-based approach, using the invoices from our canteens."*

The weight-based approach for measuring the organic food percentage across canteens in the municipality is simply determining whether or not a food product is organic, and annotating the weight of the product in the calculation. The DKK-based approach, as opposed to the weight-based approach, also involves determining whether or not a food product is organic and annotating the price of the product in the calculation. The sustainability controller explains the main issue with using the DKK-based approach, which is that some food products, such as different types of beef, are inherently more expensive as opposed to vegetables. As this was the case, this meant that the calculation is skewed due to the large price difference between different classes of products and therefore does not represent a true view of the organic food percentage. Due to this issue, the sustainability controller uses a weight-based approach for measuring the organic food percentage across kitchens in the municipality.

One of the main issues with the weight-based approach is the lack of detail included in the invoices with regard to the weight of the individual food product, where the weight of a product may be completely omitted in most product lines [14]. However, the municipality's current solution to this problem is to estimate the weight of products purchased by using price indices derived from the municipality's procurement contracts with each supplier. Although this is a fitting solution to the issue of omitted weight in product lines, we did, however, not have access to the municipality's suppliers or procurement contracts and can therefore not obtain these price indices. Despite this, measuring the organic food percentage in the municipality is an important aspect of the food strategy [1], and another approach must therefore be used.

4.4.3 Temporal ranges and food categories.

To ensure the value of the CO2e estimations, a critical aspect to gather knowledge about was the parameters for presenting the CO2e estimations. The estimations should be represented in fitting visualisations that provide intuitive points of improvement for transitioning to more sustainable foods in the municipality's kitchens. Specifically, the parameters we wanted to gain knowledge about were the temporal ranges for CO2e emission data and the number of food categories for displaying CO2e emission data.

About the temporal ranges, we inquired with the sustainability controller whether a weekly, monthly or quarterly range of data would be preferred and why this was the case:

> *"As for displaying CO2e emission data on a weekly or monthly range, I think there is a potential that these ranges may misrepresent the CO2e emissions for each kitchen, whereas a quarterly range will contribute to a more whole picture of each kitchen's CO2e emissions."*

The sustainability controller argued the reason why a quarterly range is better for presenting each kitchen's CO2e emissions, as some kitchens may purchase certain products in bulk in some months, if these products are on sale. When this is the case, this may mean that the kitchen purchases fewer food products the next month. Considering the case of using a monthly range for displaying CO2e emission data, this may lead to a skewed representation of the kitchen's CO2e emissions. Moreover, the sustainability controller explained that a kitchen may procure food for public events in the municipality, which could also lead to perceived increased CO2e emissions. However, the argument can be made that if a quarterly range is used may lead to some quarters having a higher perceived CO2e emission than other quarters due to different factors such as events in the municipality and discounted prices on food products. Therefore, either a monthly or quarterly range will be used to account for this.

As for the number of food categories for displaying the CO2e emission data, we had prepared an idea, before Interview 1, of how the data could be presented. The basis of this idea was to use the categories within the Danish *Madpyramiden*, which is a model used as dietary advice in Denmark. From *Madpyramiden*, we identified 4 categories to use as the main categories that covered the food products for which we would be estimating CO2e emissions. These 4 categories were the following: "Animal products", "Fruits, vegetables and nuts", "Cereals" and "Sugar products". For each of these 4 categories, the underlying 10-15 most-emitting food products in their respective categories would be presented.

Before we presented our idea for the number of food categories and the level of detail included in this idea to the sustainability controller, as seen in Figure 4, 5 and 17, which is in the Appendix, we wanted to understand the current approach in the municipality for presenting CO2e emission data to its kitchens. Currently, the sustainability controller produces targeted reports on each kitchen's CO2e emissions. In these targeted reports, they present a total CO2e emission for each kitchen in a bar chart for a given period, where the total CO2e emission is normalised by DKK spent. The total CO2e emissions for kitchens are grouped into the bar chart, anonymising all other kitchens than the kitchen the targeted report is for, such that kitchens can compare their performance to other kitchens.

Furthermore, the sustainability controller explained that, beyond the bar chart containing the total CO2 emission for each kitchen, the report also contains an individual bar chart for each kitchen. In this bar chart, the 22 food categories the municipality uses [1] are presented, with the total CO2e emission distributed across these categories depending on the products that the kitchens have bought in the given period.

Regarding our idea of basing the food categories on *Madpyramiden*, as seen in the categories used in Figure 4, the sustainability controller thought this was an interesting idea and has the potential to provide a great general overview of CO2e emissions. However, the sustainability controller argued that presenting the 10-15 most-emitting food products, as seen in Figure 17, that the kitchen has bought in a given period is too granular. Additionally, the sustainability controller, through his own experience, explained that the kitchens are not interested in knowing the specific food products, but rather some predefined food categories. For instance, meat such as pork, ground beef, lamb, and so on would have its predefined category, "Meat". Likewise, food products such as milk, sour cream, cheese and so on would have the category "Dairy products". This, according to the sustainability controller, provides immediate "levers" that a kitchen can manage to implement a more sustainable practice concerning its CO2e emissions. "Levers", in this context, are intervention areas in regards to their food procurement practices. This definition of the term "levers" will be used throughout the rest of the paper.

4.5 Iteration 2: Specifying user requirements

Based on Interview 1, which formed the basis for the phase of understanding users' context in UCD, we learned about the context in which the sustainability controller and other stakeholders in the municipality may use an administrator system for measuring CO2e emissions for kitchens. Additionally, from previous studies as elaborated in subsection 2.3, we gained knowledge on methods to improve granularity. From this, we formed the following requirements:

(1) The system must be able to process one or more invoice files at a time.(subsubsection *4.4.1*)

(2) The system must be able to produce different types of graphs, consisting of several food categories, for estimated CO2e emissions on both a kitchen and municipal scale. (subsubsection *4.4.1* & subsubsection *4.4.3* & [11])

(3) The system must be able to produce targeted reports for each kitchen in the municipality. (subsubsection *4.4.1*)

(4) The targeted reports produced by the system must contain graphs for CO2e emissions and anonymise all other kitchens but the kitchen for which the report is produced. (subsubsection *4.4.3*)

(5) The targeted reports must also contain both total emissions and normalised data to enable a kitchen to compare its performance to other kitchens. (subsubsection *4.4.3*)

(6) The CO2e emission graphs should allow for clicking each general category to further explore their given sub-categories. ([11])

(7) The system must be able to produce a graph displaying the estimated organic food percentage for multiple kitchens. (subsubsection *4.4.2*)

4.6 Iteration 2: Design solutions

The following section describes the prototype developed based on the requirements specified in subsection 4.5. The section is structured to follow the flow of the prototype, beginning with the architecture of the administrator system, basic overview of the system, then describes the uploading of invoices page, afterwards it describes the choices made about the invoice overview tab and ends with a discussion about the choices made for the graph designs.

4.6.1 System architecture.
The administrator system is split up into three main parts: The frontend which serves as the user interface, is primarily built using Blazor, but it does also contain some JavaScript function for handling DOM manipulation which i not supported in Blazor; The backend which performs data processing and data storage, this is primarily built in C#; and the LLM Integration Layer, which handles semantic mapping between invoice entries and the LCA product database, which is Den Store Klimadatabase [4]. This architectural overview is illustrated in Figure 6.

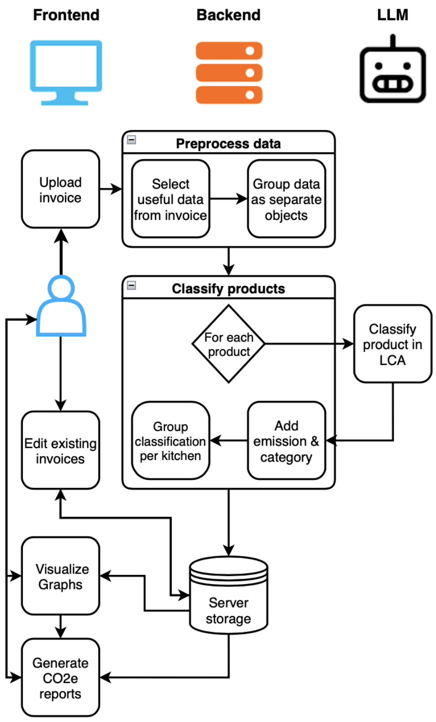

Fig. 6. System architecture of the administrator system

From the user's perspective, the system offers four primary actions when opening the page. uploading a new invoice, editing existing invoices, visualising graphs, and generating CO2e emission reports. The choices a user has will be further explained in subsubsection *4.6.2*; this section will also elaborate on why the user has these specific choices on the page. The editing of existing invoices is covered in subsubsection *4.6.4*, the visualizations in form of graphs is elaborated in subsubsection *4.6.5* and the generation of CO2e emission reports is presented in subsubsection *4.6.6*

The backend is implemented using a modular architecture to ensure scalability and maintainability. It handles two key responsibilities:

(1) Data preprocessing pipeline: when invoice upload, which is explained in detail in subsubsection *4.6.3*, the backend extracts relevant information from the invoices and groups them on a kitchen label using a record class designed specifically for invoices received by Aalborg municipality. The records are then enriched with metadata such as kitchen name, unit name and quantities, and prices.

(2) CO2e computation and storage: The backend uses emission factors from the LCA database, Den Store Klimadatabase, and computes CO2e emissions per product and aggregated CO2e emissions. The products are classified by an LLM model. These results are stored in a NoSQL database for efficient retrieval by the Frontend.

The Backend exposes a database to the Frontend, which is used for fetching stored processed invoice data, which is then used for visualising graphs, generating CO2e reports or editing existing invoices.

The LLM component of an external API integration with Google Gemini LLM. This service is used exclusively for the semantic mapping of product descriptions from invoices to corresponding entries in the LCA database. The mapping process includes prompt engineering and response validation, using a predefined response schema. If a match cannot be determined automatically with high certainty, a fallback mechanism is triggered to maintain data integrity, and the LLM returns "unknown" as a product's description. The LLM outputs are parsed and structured into predefined product categories and used downstream in the CO2e estimation pipeline, where it is later stored in a database.

By decoupling user-facing interactions from data processing and classification logic, this architecture enables a modular system structure in which individual components can be expanded or modified without affecting the overall functionality. Furthermore, the use of LLM-assisted classification ensures a flexible and scalable solution to the challenges of unstructured product descriptions, while still allowing for transparency and manual correction where necessary.

4.6.2 System overview.

Based on Interview 1, we learned that the kitchens should not necessarily access the page, but that the sustainability controller would use it to generate CO2e emission reports that they could provide to the kitchens. Therefore, the system does not include the kitchen tab, as seen in the difference between Figure 7 and Figure 5.

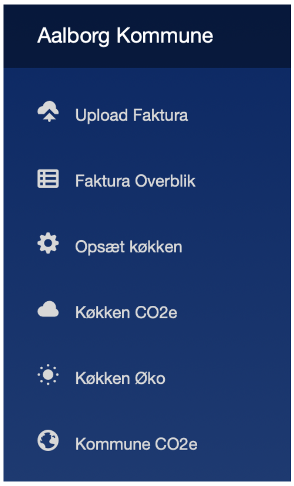

Fig. 7. The navigation menu for the administrator system

The administrator system contained the tabs "Upload Faktura", "Faktura Overblik", "Opsæt køkken", "Køkken CO2e", "Køkken Øko" and "Kommune CO2e". Each of these tabs was selected based on Interview 1. "Upload Faktura" allows the sustainability controller to upload invoices; the layout of the page can be seen in Figure 8. "Faktura Overblik" is the overview of all uploaded invoices, the page sorts the invoices per kitchen and is implemented exactly as sketched in Figure 3. The page also allows for editing of the invoices, which is to make sure that invoice lines that are wrongly classified can be changed. "Opsæt køkken" is a page for setting up kitchens in the system. This is used to register kitchens from invoices, as the invoices only contain the department, which could be specified as such in the invoice line: "Miljø og plan **". "Køkken CO2e", "Køkken Øko" and "Kommune CO2e" are all pages to view the graphs generated for the invoices uploaded to the system. "Køkken CO2e" shows 3 different graphs: a total graph showing the total CO2e emissions of each kitchen, a normalised graph showing how much CO2e emission each kitchen has per person they normally serve and a normalised graph showing how much CO2e emission each kitchen emits per DKK spent. An example of how CO2e per DKK spent looks can be seen in Figure 10.

4.6.3 Requirement 1: Upload of invoices.

Based on requirement (1), the administrator system should be able to accept CSV files for any kitchen and estimate CO2e emissions for the given CSV file provided. The prototype we developed to fulfil this requirement accepts CSV files generated from Aalborg Municipality's database, which contain invoice lines. This means that the invoice should contain specific data in order for the system to accept it.

The layout of the developed system is similar to the wireframe prototype. Code elaborating on the functionality when uploading a file can be found in Listing 1 in Appendix B. The difference between the two prototypes is the removal of the dropdown to select the given kitchen, as seen in the difference between Figure 2 and Figure 8. This is done as the sustainability controller does not generate invoices per kitchen, but rather on a municipal scale, as described in subsubsection *4.4.1*. Therefore, the system must be able to handle CSV files containing data for multiple kitchens. When a CSV file containing data for multiple kitchens has been uploaded, the system extracts relevant information from each invoice line. The relevant information in each invoice line is: invoice date, supplier name, department, food product name, total price and amount of units bought. The system can estimate CO2e emissions for each invoice line in the CSV file by extracting this information. The code explaining how the system handles the CSV data can be found in Listing 2 in Appendix B.

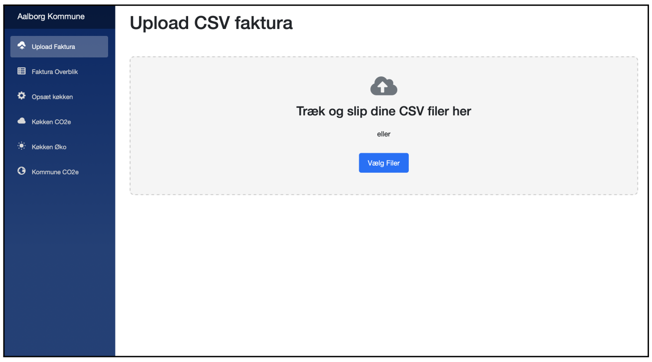

Fig. 8. The index page of the administrator system

The system can automatically detect the kitchen from the invoice line, but to estimate CO2e emissions, the food product name must first be mapped to an entity in the LCA database, Den Store Klima Database [4]. The mapping is processed through the use of artificial intelligence, which in the case of the system is provided by Google Gemini. The process of mapping is quite simple; the system provides the Google Gemini model with the given product name and the list of entities in the LCA database. Once the mapping has finished, Google Gemini passes back the name of the mapped entity in the LCA database along with a weight of the given product. If the weight is not present in the invoice line, which is often the case with vegetables in the provided data, Google Gemini estimates the weight of the product. After the mapping has finished, the system performs the following calculation to estimate the CO2e emissions of the given product:

$$CO2e \approx weight * amount * LCAemission \quad (1)$$

This process is executed for all invoice lines in the CSV file. Once this process has finished, the system groups each invoice line by kitchen names and uploads the data to a database, where it later can be seen presented in the graphs, as shown in Figure 10, or in the invoice overview, as shown in Figure 9. The code explaining how the mapping process works can be found in Listing 3 and Listing 4, which are both found in section 9 in the appendix.

4.6.4 Requirement 1: Invoice overview.

Once invoices have been uploaded to the system, they can be viewed in the invoice overview tab. The page allows for selecting a kitchen and viewing the invoices uploaded for the given kitchen. For each of the invoices, it is possible to view which products are bought on that specific invoice, which will bring up a view as seen in Figure 9.

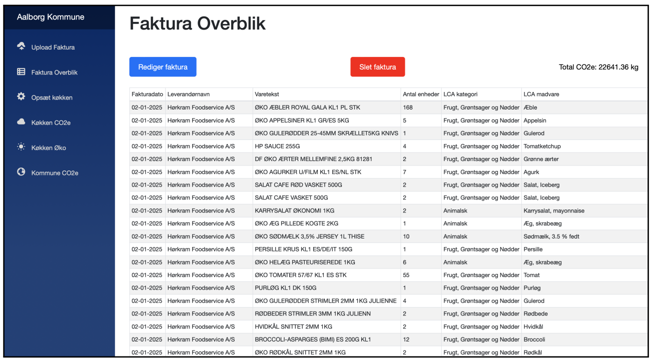

Fig. 9. The invoice overview page of the administrator system

When an invoice has been selected, the user can also edit the selected invoice. This functionality is provided such that the user can change possibly wrongly classified products. Occasionally, Google Gemini has trouble classifying products which are present in the LCA database, and it might return "Unknown". This usually occurs when there is no matching entity in the LCA database. For instance, an example is the product line "lev tom kasse", which is classified as "Unknown", given that there is no matching entity in the LCA database. This page then allows for the change of that classification, without having direct access to the database, where the invoice data is stored. The page is implemented as a convenience for the user of the system, and is not part of the requirements.

4.6.5 Requirements 2 & 6 & 7: Visualisations.

Based on requirement (2), the system must produce various number of graphs for the estimated CO2e emissions, on a kitchen and municipal scale, and based on requirement (7), a graph displaying the organic food percentage on a kitchen scale. In detail, the system has three metrics for the graphs displaying the estimated CO2e emissions: a total CO2e emission per kitchen, CO2e emission per person served in a kitchen and CO2e emission per DKK spent. Figure 10 shows how the CO2e emission per DKK spent looks in the system.

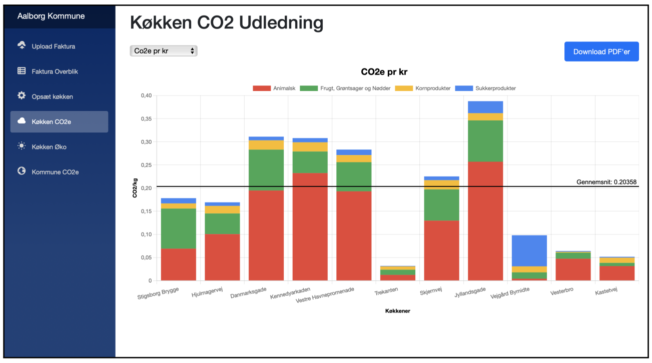

Fig. 10. A graph displaying CO2e emission per DKK spent in the administrator system.

The reason for including the metric of displaying the total CO2e emission per kitchen was to retain a fast and simple way for the sustainability controller to determine which kitchen is the largest spender in terms of CO2e, and to be able to measure this against the expectations that the municipality may have set. The metrics regarding CO2e emission per person served and per DKK spent were selected as means of normalising the estimated CO2e emissions, with the purpose of the kitchens being able to compare their performance to other kitchens. Moreover, selecting the metric of CO2e per DKK spent also provides a familiar metric that kitchens can use to measure their performance against other kitchens, as this is a metric currently used by the sustainability controller.

Additionally, beyond selecting relevant metrics for presenting CO2e emission data, it was necessary to make design decisions on what types of categories and the number of categories that should be used for food products when presenting estimated CO2e emissions. We learned, at the first meeting with the sustainability controller, that the graphs that we made last semester, for instance, the doughnut chart [14], but also the graph in Figure 17 in the appendix, consisted of too many categories. Additionally, we learned that the kitchens in the municipality are not interested in specific food products, but rather types of food. As a result, this led to an overload of information contained within the graphs, rather than providing meaningful insight into the food products that kitchens purchased. However, the sustainability controller was interested in the idea of using the Danish *Madpyramiden* as a basis for the food categories represented when displaying the estimated CO2e emissions.

Based on this and requirement (6), we proceeded to explore alternatives to retain the layer of granularity within the graphs, given that the 4 general categories, as mentioned in subsubsection *4.4.3*, do not provide much insight into the types of food kitchens purchase. Therefore, we identified several sub-categories within *Madpyramiden* for each of the 4 general categories, which were the following:

Animal products	Fruits, vegetables and nuts	Cereals	Sugar products
Beef	Fruits	Cereal	Cake
Poultry	Vegetables	Bread	Fruit concentrate
Fish	Nuts		Candy
Pork			
Lamb			
Dairy			

Table 1: General food categories and their sub-categories

Regarding the types of graphs that should be used when presenting the CO2e emissions, we decided to use bar charts to retain familiarity for the kitchens. This decision was based on Interview 1, in which the sustainability controller elaborated that they use bar charts for displaying the various food categories that kitchens procure. Figure 11 shows an example of how the graph looks, once a category has been clicked. This exact figure shows how the "Fruits, vegetables and nuts" category is split up into 3 different sub-categories: "Fruits", "Vegetables" and "Nuts".

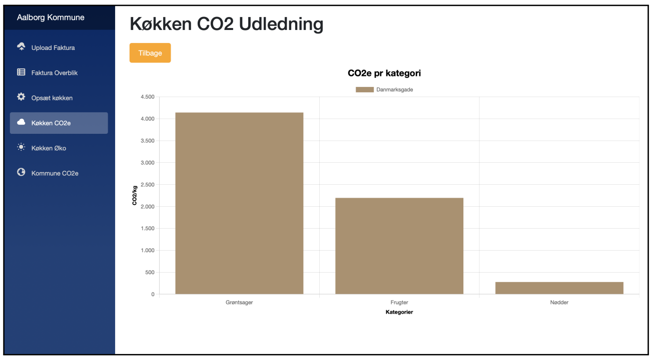

Fig. 11. A category graph in the administrator system.

Based on requirement (7), the administrator system must be able to estimate an organic percentage for a given time for each kitchen in a given time period. The organic percentage is calculated based on weight, exactly as the municipality is already calculating it. The difference is that currently, they only provide the kitchen with their exact percentage. The kitchens are therefore not able to compare their percentage with other kitchens in the municipality. Figure 12 shows how the organic percentage looks in the system, where green represents organic foods and red conventional foods. It has changed to also include the conventional percentage, as seen between the difference of Figure 12 and Figure 18 in the appendix. The graphs also contain a horizontal line, denoting the average organic percentage of the municipality.

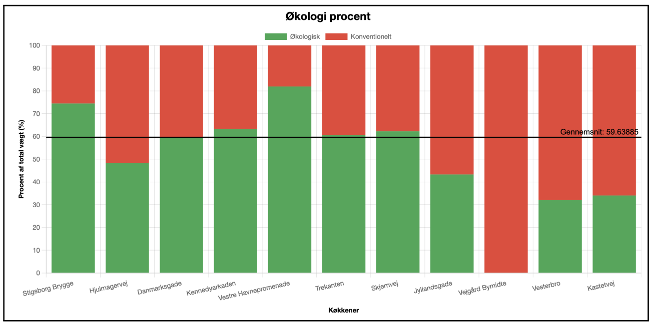

Fig. 12. A graph showing the organic percentage of the kitchens in the administrator system.

4.6.6 *Requirements 3 & 4 & 5: PDF report generation.*

Based on requirement (3), the system must be able to generate targeted reports for each kitchen in the municipality. Figure 10 shows how a graph looks, and right above the graph, there is a blue button saying "Download PDF'er"; once clicked, the system displays a loading indicator while it is generating the reports. Once the process has finished, a download begins containing a ZIP file, which has a report for every kitchen currently in the system. The invoice data that we have been given by Aalborg Municipality for this project contains 11 kitchens. Once downloaded, the ZIP file contains 11 different CO2e reports. The reports consist of all the possible graph types: Total CO2e emission per kitchen, CO2e emission per person served in a kitchen, CO2e emissions per DKK spent and all four category graphs for a given kitchen split in the categories stated in Table 1 and lastly an organic percentage graph for each kitchen. Each CO2 report thereby contains 8 different graphs.

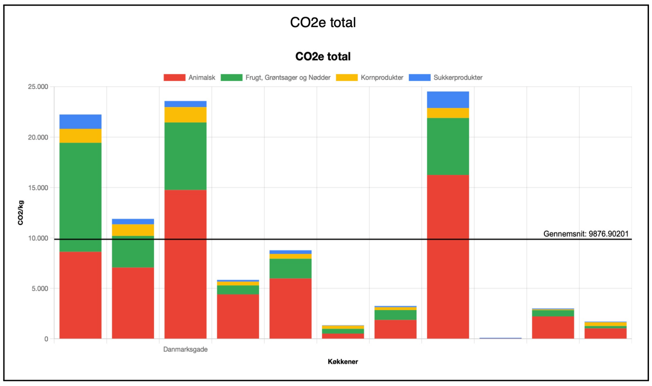

Fig. 13. An example of a graph in the CO2e report for the kitchens.

Based on requirement (4), the system should anonymise all other kitchens, and based on requirement (5), contain both total emissions and normalised emissions. Therefore, on the graphs containing multiple kitchens, the system automatically anonymize all other kitchens by only displaying the name of the kitchen targeted in the report, as shown in Figure 13, where the report in this case is generated for "Danmarksgade".

4.7 Iteration 2: Evaluate against requirements

After finishing the design and development of the administrator system prototype, we conducted a preliminary walkthrough of the system for each page with the sustainability controller. We also arranged a field deployment of the system that would last for 14 days. This field deployment was arranged, as we needed to inquire how the system would fit into the work tasks of the sustainability controller. During the field deployment, we asked the sustainability controller to collect their thoughts on each page in the system, and if they discovered any significant usability errors or shortcomings, based on questions that we provided them with, which can be found in section 11 in the appendix.

After the field deployment, we arranged Interview 2 with the sustainability controller. Prior to the meeting, we prepared some questions to structure the interview. These questions can be found in section 12 in the appendix.

The following section will follow a similar structure as subsection 4.6, where it will begin by describing an overall evaluation of the system. Afterwards, the evaluation of how the system implements each of the requirements is described in subsection 4.5.

4.7.1 Evaluation of system overview.

Overall, there were no major issues or shortcomings for any of the pages that the sustainability controller used when testing the system, but rather a few smaller changes that would improve the system. These changes will be presented later in this section. Concerning our research question, we wanted to inquire with the sustainability controller whether or not such a system would help them in their work tasks. Specifically, we inquired about the efficiency and the accuracy of the system, compared to the system that they are currently using. In terms of whether such a system would help them do their work tasks more efficiently with such a system, the sustainability controller elaborated:

> *"Absolutely. As of right now, the targeted reports that I currently make for the kitchens are made using Excel, which is a sad affair. In Excel, there is both the data processing itself, and next manually creating the targeted reports, which is done much faster in your system. So the system is a candidate for replacing the system that I currently use."*

The sustainability controller seemed to be fond of the system that we have designed, and could see a use case for the system in their work life, as the system would minimise the amount of time that they spend on estimating CO2e emissions and producing targeted reports for kitchens. In terms of the accuracy of the system, the sustainability controller compared the system to the system that they are currently using:

> *"[...]. As for the 23 food categories that I have in my current system, the vast number of food categories is irrelevant to the kitchens. So the 4 general categories that you use in your system provide a more simplistic overview and remove potential clutter that may be caused by presenting many categories at once."*

The sustainability controller, therefore, agreed with the approach of using general categories and their respective sub-categories to provide a simpler overview for the kitchens.

In general, the sustainability controller was satisfied with how the system that we have designed would help them work more efficiently with their work tasks, without compromising the accuracy of their current system.

4.7.2 Evaluation of requirement 1: Upload of invoices.

During the field deployment, the sustainability controller tried uploading some invoice data to the system for the period April 2025, they explained:

> *"I've uploaded invoices for April of 2025. While doing this, I experienced no issues."*

While they experienced no issues when using the system for uploading invoices, and the upload worked as they expected, the sustainability controller still had some suggestions on how to make the uploading of invoices even more user-friendly. They explained:

> *"The flow of uploading invoices functioned as I expected, but when I upload invoices, I have to keep the page open and wait for the mapping to finish. A better flow would be: When I upload a file, the server handles the mapping, and I could close the page while this is happening. Afterwards, the server could send me a notification that the mapping has finished."*

The flow of upload was an issue that we were aware of, but did unfortunately not have time to change. However, it would be a good addition in the future to move the processing of files from the client to a server instead. This would first of all minimise the workload of the client's computer and also make it possible for the server to handle a large queue of CSV files and inform the client once the processing has finished.

Even though the sustainability controller had a small issue with keeping the site open while the system is mapping the invoice content, the system does, however, still fulfil the requirement of being able to process one or more invoices at a time. As it is currently able to accept one or more invoices at a time and process them, even though the processing is performed client-side, rather than server-side.

4.7.3 Evaluation of requirements 2 & 6 & 7: Visualisation.
Regarding the visualizations for CO2e emissions for kitchens, as shown in Figure 10, Figure 11 and Figure 13, the sustainability controller generally found the visualisations sufficient in terms of providing the ”levers” that kitchens need to change their food procurement practices if necessary, although they proposed a few additions. They explain:

> *”I think the visualisations of CO2e emissions are really good, and are exactly what the kitchens need. If the functionality to choose dates to observe CO2e emissions for a specific period, then the visualisations would be perfect.”*

The issue of not being able to select a period to view the CO2e emissions for kitchens, as shown in Figure 13, was an obvious mistake and was a result of us focusing on implementing a sufficient amount of general categories and their sub-categories to provide the ”levers” for the kitchens. The reason for focusing on this was that we thought this to be the most important aspect of the visualisations of CO2e emissions, based on Interview 1. Considering the fact that we had only received data for the first three months of 2025, the sustainability controller informed us that being able to select a period to view CO2e emissions for was not a critical aspect. Although the sustainability controller emphasised that the level of detail contained within the visualisations was good, implementing the functionality of being able to select a period to view CO2e emissions for would be a necessary change in the future.

Additionally, we inquired with the sustainability controller whether or not they perceived the system as providing more or less information about CO2e emissions for kitchens, compared to the current system that they use. They explained:

> *”[...]. In regards to the stacked bar charts, some staff may misinterpret the visualisations as the emissions of each general category as the sum of the previous category or categories, even though this would be completely ridiculous to interpret the graph in this way.”*

As we use stacked bar charts, as shown in Figure 13, for visualizing the 4 general categories, ”Animal products”, ”Fruits, vegatables and nuts”, ”Cereals” and ”Sugar products”, and their emissions for each kitchen, rather than a bar for each general category, the sustainability controller speculated that some staff in the kitchens perhaps could perceive the different categories to be separate bars behind each other. Meaning that, for instance, on Figure 13 the blue bar, representing sugar, could be the largest bar for each of the kitchens. The sustainability controller did, however, argue that it would be ridiculous to perceive the visualisations like this. To circumvent such edge cases and provide an easier experience for some staff in navigating these visualisations, a short reading guide could be added to the visualisations to prevent such misconceptions.

As for the visualisation for displaying the estimated organic food percentage for kitchens on Figure 12, regarding requirement (7), the sustainability controller proposed an unexpected change to the use of colours when representing organic and non-organic food:

> *”Regarding the choice of colours that you use for organic and non-organic food, I made the same mistake. The sticker that is used for organic food is red, which means that when the staff in the kitchens think of organic food, they associate the colour red with organic foods. Therefore, you should use the colour red for organic foods.”*

This was a surprising fact to learn, as the colour red is usually associated with something negative, and green is associated with the opposite, that is, positive. We had chosen these colours, respectively, for non-organic and organic foods, based on the well-known colour wheel. However, as this is the case for the kitchens in Aalborg Municipality, the colours used for organic and non-organic foods should be switched in the graph for the future, to avoid confusion when kitchens view the graph. For requirement (6), in terms of the sub-categories for each general category, the sustainability controller explained whether or not the graphs provided the necessary information that should be available to kitchens:

> *”The most important thing in these visualisations is that there are these levers that the kitchens can manage, meaning what can they act on, based on the knowledge provided by the visualisations? In these visualisations that you provide them with, it is, for instance, clearly specified how much they spend on animal products, such that they can reduce this if necessary.”*

The sustainability controller was generally satisfied with the sub-categories that we had included for each general category, and appreciated that we had focused on providing these ”levers” in terms of sub-categories, as this is a very important aspect of visualising CO2e emissions for the kitchens. However, in terms of the sub-categories that we had included for the general categories, the sustainability controller had a few additions:

> *”With regards to the sub-categories that are shown, if there are no products in a specific sub-category, then the sub-category should not be displayed. However, if a kitchen has not purchased any beef, then this sub-category should still be displayed, as this is a type of food that kitchens focus on a lot.”*

The reason that we had decided to display all sub-categories at all times, even though some sub-categories may not contain any products that kitchens had purchased, was to retain consistency across multiple iterations of targeted reports throughout the year. However, regarding the changes that the sustainability controller proposed to display sub-categories, the argument can be made that omitting sub-categories would remove the consistency across multiple targeted reports throughout the year, as the order of displayed sub-categories may change if no products were purchased in a specific period during a year. If this were the case, then removing said sub-category could induce confusion for the staff in the kitchens when viewing the targeted reports, as the layout of the sub-categories would change. As for the range of sub-categories, the sustainability controller had an addition to the sub-categories in the general category **Fruits, Vegetables and Nuts**:

> *”The sub-category **Legumes** should be added to the general category **Fruits, Vegetables and Nuts**, as this*

> *is a type of food that is important for the kitchens to purchase, in relation to the food strategy."*

With regard to adding the sub-category **Legumes** to the general category **Fruits, Vegetables and Nuts** and the emphasis that the sustainability controller had on the importance of this sub-category for the kitchens, this sub-category will therefore be added to the general category **Fruits, Vegetables and Nuts** in Iteration 3.

Moreover, the sustainability controller had some additional suggestions to provide the kitchens with more information, and thereby more "levers" that the kitchens can manage:

> *"For each sub-category, there could be calculated a CO2e/DKK number, such that kitchens could obtain information on which sub-category is the most efficient in terms of reducing the carbon footprint. Additionally, as for the aspect of a kitchen comparing itself to other kitchens, information could be added to each sub-category in the form of the average CO2e emissions per sub-category. This could be visualised by having an average bar, which is greyed out, behind each sub-category."*

Although these are some good suggestions for additional levers for the kitchens, adding these changes to the graphs could lead to too much information being presented in some graphs, thereby leading the staff in the kitchens to be confused when interpreting the graphs. As we focus on providing visualisations of CO2e emissions as simply as possible, while providing kitchens with enough information to have these "levers" to reduce their CO2e emissions, we will therefore not implement these changes in the graphs. In the future, these changes could be added based on a field deployment with kitchens, where it is explored if this additional information confuses the staff when interpreting the graphs.

Overall, despite some minor additions proposed by the sustainability controller, the system fulfils the requirements of producing different graphs for estimated CO2e emissions, with the additional functionality of allowing the user to click into each general food category and estimated organic food percentages for multiple kitchens.

4.7.4 Evaluation of requirements 3 & 4 & 5: PDF report generation.

Regarding whether the PDF reports from the administrator system work as expected, the sustainability controller explained:

> *"I tried generating some reports, it worked as I expected. I got a report for all 11 kitchens in a single ZIP file."*

The implementation of the CO2e reports, therefore, also fulfils the requirements. Even though the reports were generated as the sustainability controller expected, they had some criticisms of them:

> *"The graphs in the report are quite good. They are very close to what the kitchens need. What they are lacking is definitely a reading guide, which could be a simple and short explanation on what is going on on the graphs would be much better."*

This suggestion was a good addition to the reports. Currently, they only contain each of the 8 different graphs along with a headline for each of them. But as they explain:

> *"Some of our canteen personnel can get dyscalculic when looking at graphs. You sort of have to guide them in a way such as: This graph is read in this specific way, and this means this, the colours mean this and so on."*

It would therefore be a really good addition to add a short reading guide for each graph, both to make sure that the graphs are interpreted as intended, and also to highlight the important parts of the specific graph. They explain that the reading guide could just be a short paragraph, which briefly explains what is going on in the graphs. This addition was made in the system and is further elaborated in subsection 4.8. With regards to highlighting elements on the graphs, they explained:

> *"[...] To show them something on a graph, it is often easier to either make a big arrow pointing to the place they have to look or highlight in some other way."*

This aligned perfectly with the other feedback of the reports, that the reports could be better at explaining what exactly they are showing, and do it in such a way that it should not take longer than 5 seconds to figure out what the graph is showing when observing it. This addition was implemented in the system and is further elaborated in subsection 4.8. Another small suggestion the sustainability controller had was:

> *"It would also be really beneficial for the kitchens, if the reports had some sort of overview in the beginning explaining, for instance: You emit this much, you spent this amount of DKK, you emit this much per DKK and so on."*

This addition would make it really easy for the kitchen targeted in the report to look up some key numbers showing their total emission, the emission per person and their emission per DKK for the given period in the report. The reports should also contain the average emissions in each of the three categories stated, to make it easier for the kitchens to compare themselves to the rest of the municipality. Along with stating the raw emissions, adding the average emissions next to them could also be interesting. This way, it would make it even easier for the kitchens to compare themselves to the average emissions in the municipality. This addition was implemented in the system and is further elaborated in subsection 4.8. The last addition the sustainability controller suggested was:

> *"It would also be really nice for the kitchens if the reports state the 3-5 most problematic products in terms of CO2e emissions, and perhaps even have a direct replacement for the given products."*

By making such an addition, it would be possible to try and nudge the kitchens in a certain direction, by recommending products they could switch to with a lower CO2e emission. In order to explore whether such an addition affects the choices kitchens make, we would have to perform another field deployment with one or more kitchens. During this field deployment, we would focus on whether presenting the kitchens with options for supplementary products could nudge them to choose more climate-friendly products. Therefore, this is part of future implementation plans and will not be implemented during this project.

Overall, the sustainability controller was quite pleased with the implementation of the PDF report generation. Based on the feedback from the sustainability controller, the Administrator system fulfils requirements 3, 4 and 5, given that it produces targeted reports for each kitchen, anonymises every other kitchen in the report and contains both total and normalised data. The reports could, as the sustainability controller explained, be improved by adding a reading guide for each graph, by making a simple overview of the data and enhancing their use by recommending a replacement product for high-emission products. These are all good additions for the future of the administrator system, which would make it even more usable for Aalborg Municipality.

4.7.5 Key Findings.

Through our field test with the sustainability controller at Aalborg Municipality, we learned that the system was generally useful to them. They were especially fond of the PDF generation.

As explained in subsubsection *4.7.1*, the sustainability controller found the system to be efficient in terms of being able to generate reports for the kitchens much faster than their current system. Their current system for generating reports, using Excel, is manual, which means that the sustainability controller, through the use of the administrator system, would have more time to complete other work-related tasks. Regarding whether the system is more accurate than their current system, the sustainability controller found that using fewer categories is a good addition in order to avoid potential clutter that may be caused by presenting many categories at once. They, therefore, agreed with the approach the administrator system provides by using 4 general categories and their respective sub-categories.

Regarding our findings through the field deployment about the uploading of invoices, as explained in subsubsection *4.7.2*, we learned that it functioned properly. It is able to accept one or more files at a time and handle them independently. However, the sustainability controller suggested moving the handling of the file from the client to the server, which is a good addition. The current implementation was made as it is due to time constraints and due to the requirement regarding uploading of invoices, not requiring the server to handle the files. In future iterations, it should be moved to be handled by the server.

Regarding the graphs in the system, the sustainability controller found them generally good, especially because they minimise the amount of data displayed at once, by lowering the number of categories displayed. They did, however, comment that it is not possible to select dates when choosing graphs, and this should be implemented in the future. The graphs displaying organic percentage were also quite good, but they suggested swapping the colours of organic and conventional products, as argued in subsubsection *4.7.3*. This suggestion has been changed as is further described in section subsection 4.8. Overall, they informed us that the graphs generally provide the kitchens with some ”levers” they can use to change their CO2e emission.

Regarding the reports generated by the system, the sustainability controller found them good, but told us that they lacked a reading guide for each of the graphs. They also suggested making an overview page as the first page in the report, to state some of their key emissions throughout the period and the average emissions in the same period. This was implemented as described in subsubsection *4.7.4*. It was also mentioned that the graphs could be better at highlighting the targeted kitchen in the report, which could be done by using an arrow or by greying out the other kitchens in the graphs. All of these suggested changes have been implemented and are further elaborated in subsection 4.8. Additionally, they recommended that the reports could suggest supplementary products, but as argued for in subsubsection *4.7.4*, this will not be implemented, and is part of future development.

4.8 Iteration 3: Design solutions

Based on our evaluation meeting with the sustainability controller, they suggested a few changes we could make to the CO2e reports to make them even better. These included greying out other kitchens in the graphs, which contain multiple kitchens and while still having the targeted kitchen in the original colours. It also included having a reading guide for each of the graphs.

This addition was implemented in the administrator system as seen on Figure 14, where all kitchens, except for the targeted kitchen, are greyed out. The graphs were also provided with a short reading guide explaining what can be seen on the specific graphs. The difference between the previous version and the new version can be seen in the difference between Figure 13 and Figure 14.

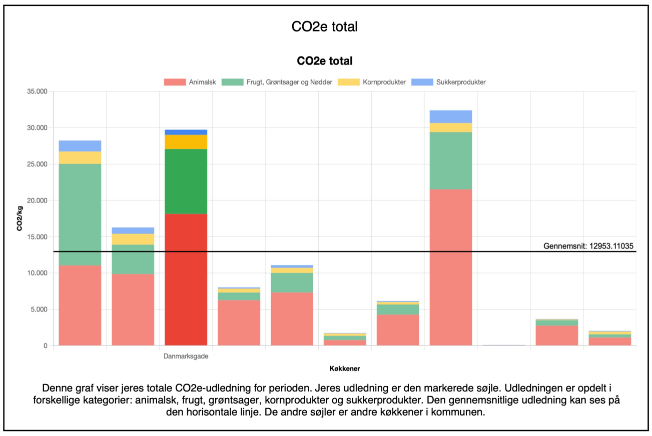

Fig. 14. Updated total CO2e emission graph in the CO2e report for the kitchen in Danmarksgade.

The sustainability controller suggested adding an overview at the beginning of the report, showing the targeted kitchen’s key emissions and the average emissions for the municipality as a whole. This has also been implemented in the administrator system, and can be seen on Figure 15.

CO2e-udlednings Rapport

Køkken: Danmarksgade

Jeres Nøgletal:

Total CO2e-udledning:	29723.64 kg
CO2e pr. person:	198.16 kg
CO2e pr. DKK:	0.28 kg

Gennemsnit for Alle Køkkener:

Total CO2e-udledning:	12953.11 kg
CO2e pr. person:	166.65 kg
CO2e pr. DKK:	0.21 kg

Fig. 15. The new overview in the CO2e report for targeted kitchen in Danmarksgade.

We have, therefore, updated the most vital parts of the system, which are the reports, to contain the extra additions suggested by the sustainability controller to make them even more useful for the kitchens in Aalborg Municipality.
Other than the report suggestions, we also added:

(1) A baseline for the total CO2e of the municipality and the ability to set it.
(2) Swapped the colours of the organic graph to fit the kitchens' perception. Now, organic is red and conventional is green.
(3) Added a period to the overview of invoices, to make it easier to find the periods covered in the system.
(4) Added legumes to the category ”Fruit, vegetables and nuts” in the CO2e graphs.

5. DISCUSSION

The following section will discuss the contribution of this project, the process we used throughout the project, privacy concerns regarding the administrator system and ends with a section discussing future steps for the project.

5.1 Our contribution

In this paper, we have presented a case study on Aalborg Municipality and the user-centred design process we followed to develop an administrator system for estimating and visualising CO2e emissions in an organisational context. The developed system was evaluated, as part of the user-centred design process, by the stakeholder of this project, the sustainability controller in Aalborg Municipality. Our aim with this project was to study how the issue of estimating CO2e emissions could be addressed in an organisational context, thereby contributing to a research gap in sHCI. While the system that we have designed is targeted at Aalborg Municipality, the design principles and architecture used in the system offer a generalisable solution which can be adopted by other municipalities in Denmark. By enabling more frequent feedback for kitchens in a municipal setting and accurate CO2e emission estimates using an LCA database[1], based on Danish products, by applying an LLM approach for invoice-based analysis, the system forms a basis for empowering municipalities in Denmark towards implementing and monitoring the efficacy of food procurement strategies. While we in this study have not been able to measure the accuracy of using an LLM approach for estimating CO2e emissions against the CO2e estimation model currently employed by the municipality, for accuracy. However, we have, during testing of the system, observed several indicators of high accuracy in terms of mapping of product names to entities in the LCA database and weight estimation of the food products. As we were not able to access any metrics for the accuracy of the model currently used by the municipality, we were unable to validate the assumed accuracy, based on our observations of our system.

Moreover, the system provides the necessary granularity for identifying high-emission food categories and various normalisation techniques for CO2e emission data, enabling kitchens in municipalities to reliably compare their initiatives to reduce carbon emissions with other kitchens. These features support greater transparency and provide actionable insights by introducing granularity to visualisations of CO2e emissions in the form of categorisations of food products. Albeit that the specific categorisations of food products used in this system are tailored to kitchens in Aalborg Municipality, making additions to these categorisations, based on the needs of future stakeholders, could assist other municipalities in moving towards sustainable food procurement practices. Through the design decisions made in this system and the emphasis on matching the requirements of stakeholders involved with UCD, the system potentially paves the way for a unified national approach for sustainable food procurement in the public sector.

In addition to this practical contribution, our work advances the field of sHCI by shifting the design focus from individual consumers to institutional stakeholders, namely, in municipalities. Much of the prior work in sHCI, based on our research, has emphasised individual behaviour change through gamification aspects and eco-feedback mechanisms [3][7]. In contrast, our system supports sustainability efforts at an organisational level, offering a robust tool that aligns with administrative workflows and strategic goals, such as those outlined in Aalborg Municipality's food strategy [1]. We also contribute practitioner-relevant insights, demonstrating how actionable sustainability intervention areas, such as visual representations of emissions by food categories, can be effectively embedded into system design. Additionally, our work also demonstrates how norms in HCI, such as the choice of colours for positive and negative aspects, for instance, when presenting the share of organic and non-organic foods purchased, may translate differently in practice, and thereby promotes the importance of a user-centred design approach when meeting the needs of stakeholders. Furthermore, during Interview 2, we learned that there were two requirements for the system that we had not uncovered: the functionality of having a baseline and a forecasting mechanism. We believe that the reason we had not uncovered these requirements was because we had focused on the context of the problem from the sustainability controller's view, rather than also inquiring into the problem in a broader context. These insights inform future sHCI work by highlighting the importance of stakeholder alignment, data intelligibility, and strategic integration in designing for sustainability at scale.

[1] Den store klimadatabase[4]

5.2 Reflecting on the process

Throughout the project, we learned about the benefits of the methodology, which we had decided upon for the project, in practice. Namely, using techniques like prototyping and user-centred design helped us in understanding the context of stakeholders, in which stakeholders would use an administrator system for CO2e estimations, as the one that we have designed in this project. By using prototyping, we learned of potential usability errors that could be present if we had blindly followed the assumptions and ideas that we had made for our initial wireframe in Iteration 1. For instance, as described in subsection 4.2, we had initially designed the system to have a ”Kitchen board” and an ”Administrator board”. This view for the system was, however, pivoted to only be targeted at stakeholders with an administrative role in the municipality, such as the sustainability controller, as we learned that kitchens prefer concise, targeted reports. Moreover, we learned from the initial prototype we designed for Interview 1, that kitchens are more interested in general categories of food and a few sub-categories within these general categories, rather than being presented with the most 10-15 specific types of food, which a kitchen has purchased in a period. By omitting the evaluation of the design and contents of an initial prototype, usability errors and mismatch of stakeholders' requirements for the system would not have been discovered until a much later phase of the design process. For instance, substituting the user-centred design methodology used in this project with a linear and phase-gated methodology, such as the waterfall model, would have led to less user engagement and thereby led to a design which the stakeholders of the system would not have had the ability to continuously validate and refine.

Prototyping, in conjunction with using a user-centred design approach, allowed us to learn from the sustainability controller's previous experiences of working with kitchens. For instance, the aforementioned example allowed us to base our design decisions on their experience.

Although that prototyping and user-centred design helped us meet the requirements of stakeholders through evaluation of the system by the sustainability controller, we believe that having conducted a field deployment of the targeted reports with staff in kitchens in the municipality could have provided us with additional insights of what the kitchens require to make changes to their food procurement practices. By conducting such a field deployment, we would have had additional support and validation, through feedback from kitchen staff, of the design choices that we made for the visualisations of CO2e emissions and the structure of the targeted reports. However, we learned in our previous semester [14] and through our interviews with the sustainability controller during this project, that the kitchens mostly do not know what they want in terms of visualisations. Given that this is the case for the kitchens in Aalborg Municipality, conducting a field deployment, and thereby introducing more stakeholders, could have revealed any potential usability errors, but we would, most likely, based on the sustainability controller's past experiences, not have gotten any feedback for radical changes for the type of visualizations.

Overall, we are generally satisfied with the results that came of using the methodologies we selected for this project, even though making changes to the process, such as conducting an additional field deployment with the kitchens, would have provided additional validation of our design decisions for the system.

5.3 Privacy concerns

As the administrator system relies on detailed municipal food procurement data, privacy concerns must be addressed. Although the data used in this project primarily consists of food procurement records, such as product names, quantities, departments, and prices, it may still contain sensitive or semi-identifiable information. While individual kitchen staff are not directly identified in the dataset, recurring procurement patterns could, in some cases, indirectly reveal operational characteristics or purchasing behaviours of specific kitchens. This makes it essential to implement data privacy safeguards.

Furthermore, the system integrates an external LLM (Google Gemini) to perform semantic classification of food products. This means that the selected invoice information is transmitted to a third-party API. Even though the content does not include personal identifiers, transmitting any form of municipal data externally should be supported by a thoughtful review of GDPR and local data policies. For instance, it should be ensured that data processing agreements are in place and that data minimisation principles are followed, such that the system is only transmitting strictly necessary data for classification. While sending data to external LLMs enables advanced classification and potentially better user experience, this benefit must be weighed against the inherent risks of third-party data exposure, particularly in a public-sector context.

Furthermore, the system enables the generation of anonymised CO2e emission reports for individual kitchens, allowing them to benchmark their performance against other kitchens in the municipality. By anonymising the identities of other kitchens, the reports foster a constructive environment with a need-to-know structure, encouraging each kitchen to focus on its progress while being able to compare its performance to other kitchens in the municipality.

In the future of this project, it may be beneficial to explore alternative approaches that eliminate the need to transmit data to third-party services for classification purposes. One potential direction is to develop an ”in-house” classification model or to leverage models that do not retain or train on incoming data. By relying on ”in-house” or non-persistent inference models, the system could reduce its dependence on external providers, thereby improving compliance with data protection principles such as data minimisation and purpose limitation.

In summary, while the data used in this system does not contain direct personal information, its municipal nature and the use of AI classification introduce important privacy considerations. Ensuring compliance with regulatory frameworks and fostering responsible data management will be key in scaling such systems.

5.4 Future work for the system

While the current system provides improved efficiency and usability in estimating and visualising CO2e emissions from food procurement for Aalborg Municipality, several areas remain for future development. First, implementing server-side processing for invoice uploads would improve scalability and user experience, allowing asynchronous data processing and the ability to notify the user when processing has finished. Additionally, further refinement of the LLM-based product classification could enhance mapping accuracy, even though the current model performs well, as it has a high accuracy when mapping food products to their respective emissions. Incorporating a method for long-term CO2e emission forecasting would be a useful feature to implement for making the system more valuable to the municipality, as the sustainability controller explained during Interview 2. The municipality is interested in a tool that enables them to forecast future emissions to assess whether they are on track to meet their goal of reducing their CO2e emissions baseline from 2018 with 25% by 2028, and to proactively identify areas where strategic adjustments in food procurement practices may be required to stay aligned with the target baseline. Finally, restructuring the system interface into a dashboard layout with a top navigation bar could enhance usability and familiarity for Aalborg Municipality. As the sustainability controller explained during Interview 2, their existing municipal systems commonly adopt this layout, suggesting that aligning with this convention would support a more intuitive user experience.

6. CONCLUSION

This paper has presented the design and evaluation of a CO2e emission estimation and visualisation system tailored for municipal food procurement, using Aalborg Municipality as a case study. Through a user-centred design approach, we identified the key needs of the primary stakeholder, the sustainability controller, and developed a system that aligns with their daily tasks and broader municipal sustainability goals, as described in their food strategy [1]. The system enables efficient estimation of CO2e emissions based on invoice data, incorporates normalisation methods that reflect actual procurement practices, and provides targeted visualisations and anonymised reports that facilitate informed decision-making across municipal kitchens.

Our evaluation demonstrated that the system improves the efficiency of generating targeted reports for the kitchens in Aalborg Municipality, using simplified categorisations of food products for the visualisations of the CO2e emissions while providing actionable insights. The system also accurately identifies entities in an LCA database[2] based on product descriptions from invoices, and accurately estimates weights on products where they are not given in the product description. However, as we were unable to compare our model to the current model used by the municipality, we are unable to verify whether the model has made any improvements in terms of accuracy. Although the developed prototype meets the core functional requirements that we set for the system, further work could improve the system's robustness and usability in the context of usage of the prototype by Aalborg Municipality, particularly in areas such as implementing a forecasting mechanism for CO2e emissions on a municipal scale and asynchronous processing of invoice data. Overall, this project indicates the potential for future digital tools in organisational contexts, such as municipalities, to support these organisations in meeting their climate targets through more sustainable food procurement practices.

7. REFERENCES

[1] Aalborg Kommune. *Sunde og klimavenlige fødevarer - En fødevarestrategi for Aalborg Kommune 2024 - 2028*. Aalborg Kommune, 2024.

[2] Baumeister, R. F., Leary, M. R. Writing narrative literature reviews. *Review of General Psychology*, 1997.

[3] Aksel Biørn-Hansen, Cecilia Katzeff, and Elina Eriksson. Exploring the use of a carbon footprint calculator challenging everyday habits. In *Nordic Human-Computer Interaction Conference*, NordiCHI '22, New York, NY, USA, 2022. Association for Computing Machinery.

[4] Concito. Den store klimadatabse. https://denstoreklimadatabase.dk, 2025.

[5] Curran, M. A. Life Cycle Assessment: A Review of the Methodology and Its Application to Sustainability. *Current Opinion in Chemical Engineering*, 2013.

[6] Dalgaard, Randi, Jannick Schmidt, and Anna Flysjö. Generic model for calculating carbon foot- print of milk using four different life cycle assessment modelling approaches. *Cleaner Production*, 2014.

[7] Donald Degraen, Elena Werny, Marc Schubhan, Maximilian Altmeyer, and Antonio Krüger. Ecomeal: Gamified eco-feedback of food consumption using a virtual garden. CHI PLAY Companion '24, page 43–49. Association for Computing Machinery, 2024.

[8] Dennis Pagano, Bernd Bruegge. User involvement in software evolution practice: A case study. *IEEE*, 2013.

[9] Eugene A. Mohareb, Martin C. Heller, Peter M. Guthrie. Cities' Role in Mitigating United States Food System Greenhouse Gas Emissions. *American Chemical Society*, 2018.

[10] Hannah Ritchie. Food production is responsible for one quarter of the world's greenhouse gas emissions. *Our World in Data*, 2019. https://ourworldindata.org/food-ghg-emissions.

[11] Björn Hedin, Philip Claesson, and Patrik Odqvist. Visualizing carbon footprint from school meals. In *2017 Sustainable Internet and ICT for Sustainability (SustainIT)*, pages 1–3. IEEE, 2017.

[12] Hellweg, Stefanie, and Llorenç Milà I Canals. Emerging approaches, challenges and opportunities in life cycle assessment. *Science*, 2014.

[13] Interaction Design Foundation IxDF. What is User Centered Design (UCD)? https://www.interaction-design.org/literature/topics/user-centered-design, 2016.

[14] Jensen, Marcus Hasselgaard Skifter and Pihl, Tobias Martin and Kragsaa, Karl Barnholdt and Aagaard, Rune Gøttsche. Providing insight into CO2e emissions from food products for a kitchen in Aalborg Municipality, 2025.

[15] Joint Research Centre. Less food waste could bring lower EU food prices and decrease greenhouse gas emissions. https://joint-research-centre.ec.europa.eu/jrc-news-and-updates/less-food-waste-could-bring-lower-eu-food-prices-and-decrease-greenhouse-gas-emissions-2023-07-06$_en$, 2023.

[16] Notarnicola, B., Sala, S., Anton, A., McLaren, S. J., Saouter, E., Sonesson. The Role of Life Cycle Assessment in Supporting Sustainable Agri-Food Systems: A Review of the Challenges. *Journal of Cleaner Production*, 2017.

[17] Notarnicola, B., Tassielli, G., Renzulli, P. A., Castellani, V., Sala, S. Environmental Impacts of Food Consumption in Europe. *Journal of Cleaner Production*, 2019.

[18] Pamela Baxter, Susan Jack. Qualitative case study methodology: Study design and implementation for novice researchers. *TQR*, 2008.

[19] Poore, J., Nemecek, T. Reducing Food's Environmental Impacts Through Producers and Consumers. *Science*, 2018.

[2] Den store klimadatabase[4]

[20] Jayanthi Raghunathan. Budget your carbon emissions: Interactive visualisation of an individual's carbon budget, 2021.

[21] Jim Rudd, Ken Stern, and Scott Isensee. Low vs. high-fidelity prototyping debate. *interactions*, 3(1):76–85, 1996.

[22] Soleymani Angili, T.; Grzesik, K.; Salimi, E.; Loizidou, M. Life cycle analysis of food waste valorization in laboratory-scale. *Energies*, 2022.

[23] Pieter Nick ten Caat, Martin J. Tenpierik, Tithi Sanyal, Nico M.J.D. Tillie, Andy A.J.F. van den Dobbelsteen, Geoffrey Thün, Sean Cullen, Shun Nakayama, Theodora Karanisa, and Stewart Monti. Towards fossil free cities – emission assessment of food and resources consumption with the fewprint carbon accounting platform. *Cleaner Environmental Systems*, 2022.

8. APPENDIX A: PROTOTYPE

Fig. 16. A wirefram showing how creating a kitchen in the system looks

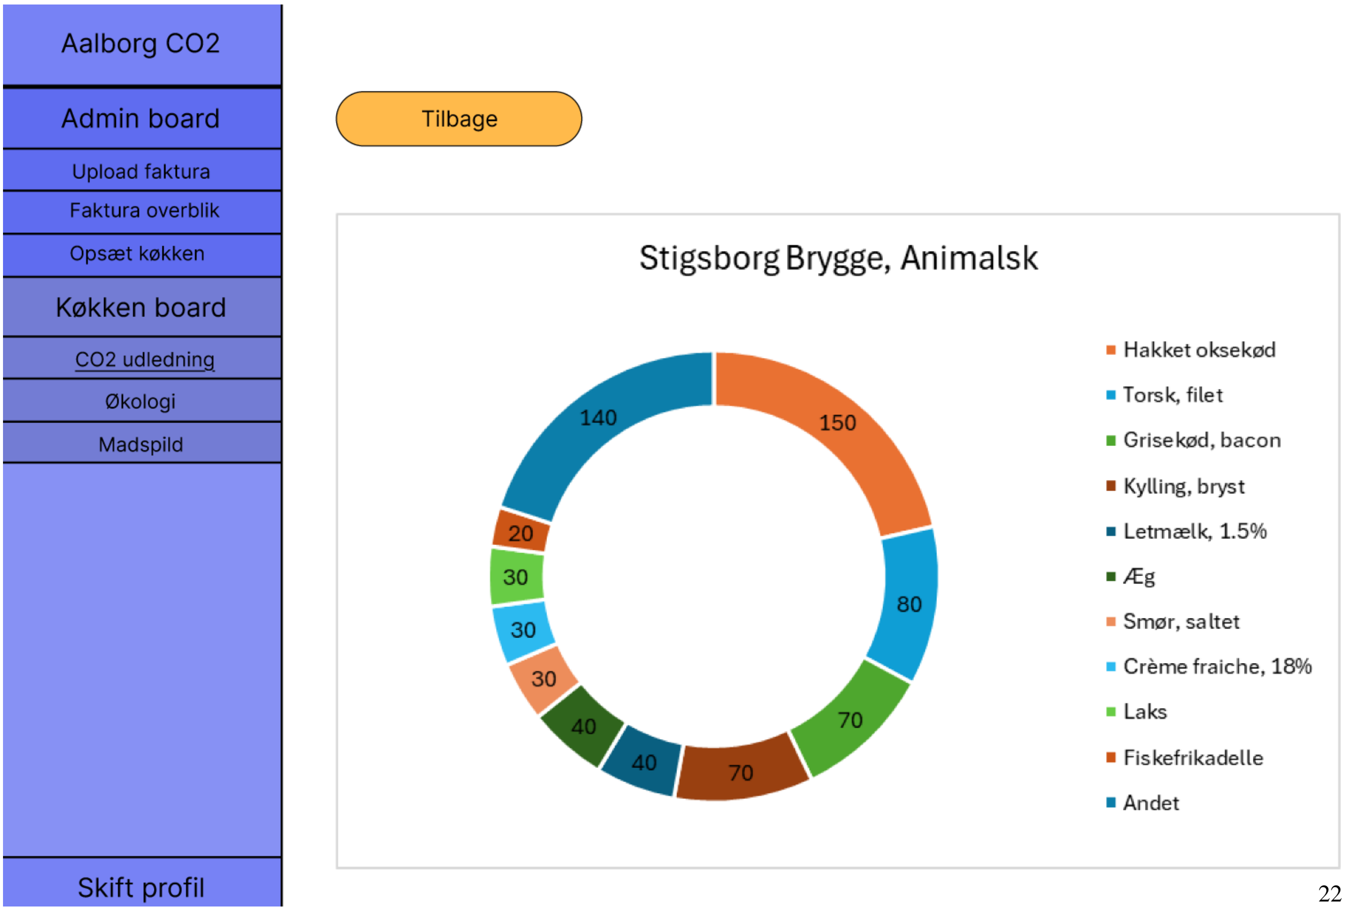

Fig. 17. A wireframe showing the 10 most CO2e emitting foods in the category Animalsk

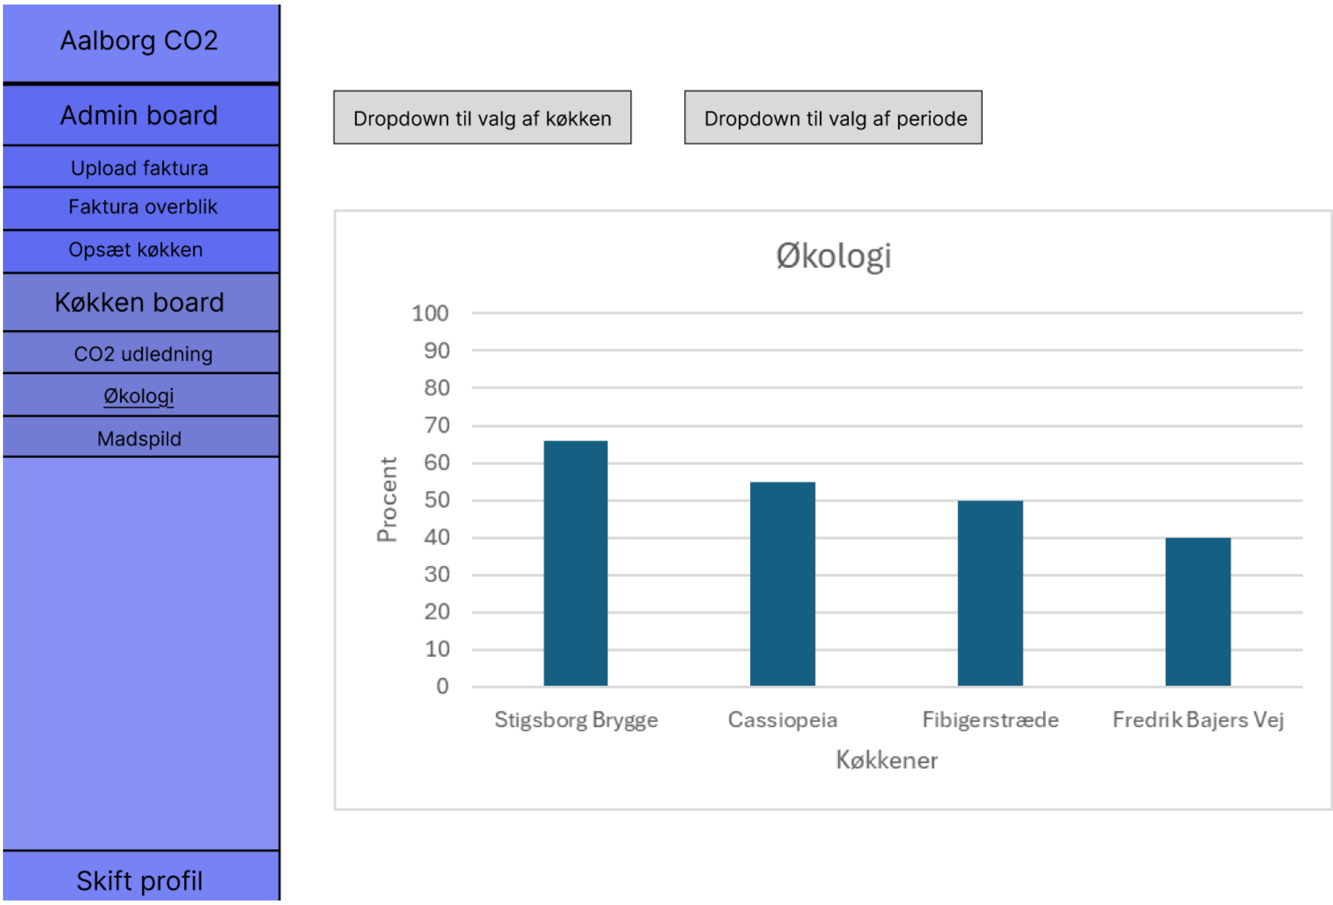

Fig. 18. A wiregram showing share of organic foods purchased in the municipality

9. APPENDIX B: CODE

```
LoadFile(Inputfiles)
{
    for each file in inputfiles
    {
        # Check whether file is a CSV file
        if (file is not .csv)
        {
            give error to user "File is not a CSV file"
            skip to next file in input
        }

        # Next line for making sure file is encoded in UTF8
        csvContent = ReadFileWithCorrectEncoding(file);

        # Parse conent from file
        InvoiceContent = ParseCSV(csvContent);

        # Parse contnet to Gemini and upload content to DB once finished
        Gemini.PromptGemini(invoiceContent);
    }
}

Listing 1: Code explaining the flow when uploading a file to the administrator system.
```

```
# Revieve file content as string
ParseCSV(string csvContent)
{
    string[] lines = csvContent.Split("\r\n");
    List<InvoiceModel> invoiceContent = [];
    string[] splitHeader = lines[0].Split(';');

    # Allocate "Fakturadato", "Leverand rnavn", "Afdeling", Varetekst", "Fakturapris", "Antal enheder" columns in file
    int Fakturadato = Array.IndexOf(splitHeader, "Fakturadato");
    int Leverand rnavn = Array.IndexOf(splitHeader, "Leverand rnavn");
    int Afdeling = Array.IndexOf(splitHeader, "Afdeling");
    int Varetekst = Array.IndexOf(splitHeader, "Varetekst");
    int Fakturapris = Array.IndexOf(splitHeader, "Fakturapris");
    int AntalEnheder = Array.IndexOf(splitHeader, "Antal enheder");

    if(Indexes not found)
    {
        #Return error to user, that file does not contain columns
        return [];
    }

    # remove header from list and create invoiceModel for each line
    foreach (string line in lines[1..])
    {
        string[] parts = line.Split(';');
        invoiceContent.Add(new InvoiceModel
        {
            Fakturadato = parts[Fakturadato],
            Leverand rnavn = parts[Leverand rnavn],
            Afdeling = parts[Afdeling],
            TotalPris = parts[Fakturapris],
            Varetekst = parts[Varetekst],
            AntalEnheder = parts[AntalEnheder]
        });
    }

    Dictionary<string, List<InvoiceModel>> groupedInvoices = [];

    # Pre-populate with all known kitchens
    foreach (var kitchen in Kitchens)
    {
        groupedInvoices[kitchen.DepartmentName] = [];
    }

    # Group invoices by kitchen
    foreach (var invoice in invoiceContent)
    {
        string department = invoice.Afdeling;
        groupedInvoices[department].Add(invoice);
    }

    # Remove empty lists
    groupedInvoices = groupedInvoices.Where(kvp => kvp.Value.Count > 0).ToDictionary(kvp => kvp.Key, kvp => kvp.Value);
    return groupedInvoices;
}

Listing 2: Code explaining the flow of extrating data from the CSV file.
```

```
# Receive products in invoices grouped by kitchens
PromptGemini(Dictionary<string, List<InvoiceModel>> GroupedKitchens)
{
    foreach (var Kitchen in GroupedKitchens)
    {
        double totalCO2e = 0;
        List<InvoiceModel> Results = [];
        foreach (InvoiceModel product in Kitchen.Value)
        {
            # Promt gemini per product in invoice
            InvoiceModel response = await ClassifyProduct(product);
            Results.Add(response);
            # Add emission for the product to the total emission
            totalCO2e += response.CO2e * response.V gtPrEnhed * response.AntalEnheder;
        }
        # Upload the invoice one prompting has finished
        await DBHandler.CreateItem("Invoices", Results);
    }
}
```

Listing 3: Code explaining the flow when prompting gemini.

```
# Recieves product as an InvoiceModel
ClassifyProduct(InvoiceModel item)
{
    # Setup model
    var model = GeminiClient.CreateGenerativeModel("models/gemini-2.0-flash");
    # System intructions parsed to the gemini model, to provide some context
    string systemInstructions = "Du skal finde p   en dansk f devare baseret p   en tekst streng. Korrektheden er vigtig, du skal give mig det der st r i tabellen og give mig en total v gt pr enhed i kg. ";
    # Setup model
    model.Config = new GenerativeAI.Types.GenerationConfig
    {
        Temperature = 0.0,
        ResponseMimeType = "application/json",
        ResponseSchema = new GeminiResponse()
    };
    model.SystemInstruction = systemInstructions;

    # Promt gemini with product and emissions data from LCA db
    var response = await model.GenerateContentAsync($"V lg den kategori fra f lgende liste der bedst matcher: '{item.Varetekst}'\n\nKategorier:\n{emissionsData}");

    # Extract result from JSON block
    var jsonBlocks = response.ExtractJsonBlocks();
    var Objects = jsonBlocks.Select(block => block.ToObject<GeminiResponse>());

    # Add the new entries to the model and return the same model
    item.LCAFood = myObjects.First().category;
    item.V gtPrEnhed = myObjects.First().weight.ToString();
    CategoryModel category = _groupService.MapToCategory(item.LCAFood);
    item.CO2e = category.emission.ToString();
    item.LCAkategory = category.category;
    item.LCAGroup = category.group;
    return item;
}
```

Listing 4: Code explaining the flow when classifying a product using gemini.

10. APPENDIX C: QUESTIONS FOR INTERVIEW 1

(1) The general design of the system (admin board, each tab)
(2) What we need to learn
   (a) How do you handle food waste (are there any invoices for the food waste that kitchens have)?
   (b) How do you measure organic products, organizational or per kitchen?
   (c) What periods for displaying the data are desired (weeks, months, quarterly)?
   (d) How many categories is best to display when presenting the data?
   (e) Basis on the Foodpyramid, with an added 4th category?
   (f) Is the Foodpyramid categories a good method for displaying the data?
   (g) More categories but limited to 10-15 of the highest emitting categories in the given invoice data?

11. APPENDIX D: QUESTIONS FOR PRELIMINARY WALKTHROUGH

Questions we will provide them that they can consider while using the system:

(1) When considering the graphs that we have implemented in the system for presenting CO2e emissions for kitchens, do they provide the necessary information that should be available to the kitchens? How and why?
(2) Is the system intuitive, both in the sense of
   (a) navigating the system
   (b) the different functionality that should be available when performing the tasks necessary to estimate CO2e emissions for kitchens?
(3) How does the system fit the work tasks you need to perform and the practices that you follow when doing these tasks?
(4) How well does the system fit into the municipality's strategy for sustainability and the work tasks you have associated with the kitchens?
(5) Are there any limitations in the system, anything that you think is crucial to have available to you in this context?

12. APPENDIX E: QUESTIONS FOR INTERVIEW 2

(1) Functionality for each of the pages. Do they provide any interesting information why/why not?
(2) Along with this, each of the pages also ask how this functionality fits with the tasks that they performs during the work. Why/why not?
(3) After the walk-through of the system, we ask the following questions and inquire about the hows and whys of his response to each of these questions.
   (a) When considering the graphs that we have implemented in the system for presenting CO2e emissions for kitchens, do they provide the necessary information that should be available to the kitchens? How and why?
      i. Do the reports containing the graphs contain better/worse information about CO2e in relation to the targeted reports that you currently produce for the kitchens in the municipality?
      ii. Do the general categories and sub-categories fit the information that the kitchens expect to be present? Are there any categories missing from the graphs?
   (b) Is the system easy and intuitive to use?
   (c) How well does the system fit into the municipality's strategy for sustainability and the work tasks you have associated with the kitchens?
   (d) Are there any limitations in the system, anything that you think is crucial to have available to you in this context?
      i. For instance, is it lacking that food waste is not an aspect present in the system? Is this an aspect that the municipality finds important to include in the future in relation to the food strategy?
(4) Related to our research question, we will ask if and why the system provides a more 1) efficient and 2) accurate approach in estimating food-related CO2e emissions in the municipality.
   (a) How is the quality of the visualizations of the CO2e emissions for each kitchen in the municipality?
      i. Are they effective in the sense of providing the “levers” that kitchens need to make changes to food they purchase?
      ii. In relation to effectiveness, do you see yourself having more time for other tasks that you perform in relation to CO2e emission estimations by using the system?
      iii. Are they accurate in relation to the current estimations that you make? Are they better or worse than estimations provided by your current system?
   (b) Does the system provide more or less information about CO2e emissions for kitchens compared to your current system? How and why?