Designing AI-driven Meal Demand Prediction Systems

Alicia Cabrejas Leonhardt, Maximilian Kalff, Emil Kobel, and Max Bauch

Hasso Plattner Institute, University of Potsdam, Potsdam, Germany
{alicia.cabrejas-leonhardt, maximilian.kalff, emil.kobel, max.bauch}@student.hpi.uni-potsdam.de

Abstract. Meeting today's demand forecasting challenges requires more than traditional methods. Economic uncertainty, evolving sustainability goals, and breakthroughs in artificial intelligence are reshaping how organisations plan for the future. Following the Design Science Research approach, we present our findings on the design of demand forecasting architectures. Based on the requirements we identified through interviews with various stakeholders, we were able to conclude nine design principles. All nine principles are consolidated in an architecture to demonstrate their feasibility. The findings presented contribute to the ongoing discourse on demand forecasting based on artificial intelligence, both in terms of its current state and future potential.

Keywords: meal demand prediction, forecasting methodology, customer choice behaviour, supervised machine learning, design science research

1 Introduction

The growing awareness of sustainability and the advancing technological capabilities have reinforced the need to act more responsibly, both ecologically and economically. However, achieving waste reduction targets and cutting emissions requires a precise understanding of future demand. Enterprises must navigate a delicate balance: while surplus inventory leads to additional costs for unused materials, insufficient supply risks lost sales and dissatisfied customers. Optimising this planning process is therefore critical, and recent advances highlight the potential of Artificial Intelligence (AI) to improve demand forecasting and decision-making (Hübner et al. 2024). In particular, factors such as the shelf life of products (Aishwarya et al. 2020) and long delivery times demand special attention and accurate forecasting to ensure sustainable and economical operations.
Despite its growing relevance, academic contributions in the specific context of in-flight catering remain limited. Recent research on food demand prediction in general (Aishwarya et al. 2020) and demand forecasting in bakeries (Hübner et al. 2024) highlights the promising role of AI in demand prediction. To asses its current relevance in industry, we reviewed the annual sustainability reports from six big airlines (RyanairGroup 2023, Portugal 2023, easyJet 2023, Iberia 2023, AmericanAirlines 2023, BritishAirways 2023) and found only one (AmericanAirlines 2023) that had done a pilot project using AI.
According to previous research, historical, internal, and external data are relevant to predict demand more accurately (Aishwarya et al. 2020). Integrating external factors such as weather conditions or demographic information is particularly promising, as it

20th International Conference on Wirtschaftsinformatik
September 2025, Münster, Germany

may allow for more precise predictions compared to those merely relying on historical data. However, identifying all relevant factors and combining them is challenging (Liman 2016). These external insights would be especially beneficial to the commercial aviation industry, which has to dispose of all the food taken on board.

In the following, we will show what we have learned from our analysis of in-flight catering and what can be concluded for other sectors where demand forecasting is also an important tool. The research question is: What are the requirements and design principles for meal demand forecasting systems driven by AI, and how does a potential system architecture look?

To get familiar with the problem space, we conducted interviews with representatives related to meal demand prediction for in-flight catering. From these interviews, we identified 11 design requirements (DRs) that represent the need for change in forecasting methods, which serve as the basis for our work. Those DRs resulted in nine design principles (DPs), building on recent work on demand forecasting. To validate our assumptions, we spoke to experts and followed an iterative process, which enabled us to present an architecture for meal demand predictions. The goal of the process is to present our findings and to contribute to the discussion on what future demand prediction tools may look like and what needs to be considered in order to make them as efficient as possible.

The paper is structured as follows. Following a brief overview of related work, the research design section explains the Design Science Research (DSR) method that we applied, as described by Peffers et al. (2007). Based on the requirements, the design principles are developed, leading to an architecture. The previous findings are then evaluated and discussed, leading to a conclusion.

2 Related Work

2.1 Demand Prediction

Demand forecasting involves estimating the future demand of customers based on available information about the goods and services that companies currently produce or plan to produce (Çetinkaya & Erdal 2019). This process often relies on quantitative forecasting techniques, which focus on predicting future trends using limited data (Miller et al. 1991). When making demand predictions, the following aspects must be taken into account: the broader economy, infrastructure, the social environment, the legal framework, the market, actions by the firm, actions by those offering competing and complementary products, and actions by others, such as unions and lobby groups (Armstrong & Green 2005). In the context of revenue management, having precise demand data is crucial for the effectiveness of various booking and pricing strategies (Haensel et al. 2011). Consequently, demand forecasting is probably the most critical element in production planning and control (Lee & Seo 2001).

Previous literature on meal demand forecasting has focused primarily on the requirements of such systems before considering the impact of AI. We have listed them in Table 1. It has been noted that both overproduction and underproduction result in operational costs, affecting personnel, materials, facility management, customer satisfaction, and accountability (Miller et al. 1991). As a result, the need for accurate demand forecasting is essential, although it remains a challenging task due to daily fluctuations influenced by factors such as weather, seasonal changes, holidays, changes in customers' preferences, price changes, and local disruptions (Hübner et al. 2024, Aishwarya et al. 2020). This highlights the importance of integrating domain expertise and cultural aspects into forecasting methodologies (Armstrong & Green 2005, Çetinkaya & Erdal 2019). By including internal and external data, demand prediction could be enhanced significantly (Aishwarya et al. 2020). In order to ensure seamless data flow, the forecasting system should be integrated with the management information system (Miller et al., 1991). The accuracy level provided by the demand forecast depends heavily on the amount and quality of available data and the selected method, which should ideally be resilient to noise or variation (Lee & Seo 2001, Gorr et al. 1994). To mitigate this risk, multiple forecasting methods can be combined to reduce the error rate by 3 to 24%, compared to the average error of components, as demonstrated by a meta-analysis of 30 studies (Armstrong & Green 2005).

Table 1. Requirements for demand prediction systems mentioned in literature.

Requirement	Reference
Forecast accuracy: The system must provide precise demand forecasts to enable planning in advance, minimising costs and reducing waste.	Hübner et al. (2024)
Multifactorial forecasting method: Demand prediction must consider various factors.	Çetinkaya & Erdal (2019), Aishwarya et al. (2020), Armstrong & Green (2005)
Integrated architecture: Demand prediction must be incorporated into existing systems.	Miller et al. (1991)
Just-in-time adjustment: Demand prediction must be able to quickly react to demand fluctuations.	Haensel et al. (2011)

Traditional forecasting approaches are substantially based on personal experience and subjective market insights (Armstrong & Green 2005). Therefore, mathematical forecasting methods must be used to reduce error (Miller et al. 1991). A possible solution to this problem would be the use of AI in order to create a precise sales forecast (Rao et al. 2021, Panda & Dwivedi 2020).

2.2 The Potential of AI-driven Demand Prediction in Commercial Aviation

In contrast to regular demand prediction items, in-flight food and beverages are especially challenging products (Liman 2016). This is due to the necessity of disposing of such items after international flights. Therefore, airlines strive for a delicate balance between maintaining customer satisfaction by offering every possible meal to each passenger and potentially losing sales when giving in to financial and environmental considerations.(Hübner et al. 2024, Liman 2016).
Current techniques aim to analyse customer behaviour (van Ryzin 2005) to improve

demand prediction accuracy. With the increasing advancement of AI technologies, the application of AI-driven models to predict meal demand presents a promising opportunity (Kumar et al. 2021, Çetinkaya & Erdal 2019). Previous research has shown that several fitting algorithms can be employed in AI to generate a sales forecast (Rao et al. 2021), outperforming other prediction models (Çetinkaya & Erdal 2019). Although errors in food service operations cannot be completely eliminated, the results achieved by mathematical forecasting techniques would significantly reduce them (Miller et al. 1991), leading to less waste and enhanced competitiveness for airlines (Lee & Seo 2001). Therefore, the environmental consequences of AI are broad and considerable (Liao et al. 2021), and its economic impact is significant too.
In conclusion, accurate demand prediction models would have a significant impact on economic and environmental factors. The literature suggests that AI-driven methods will transform demand prediction due to their significant impact on precision, productivity, and efficiency (Kumar et al. 2021).

3 Research Design

To study requirements and design principles for AI-driven demand prediction, we adopted the Design Science Research approach proposed by Peffers et al. (2007). We followed the six stages of the *Problem-Centred Approach*. Given our objective of providing airlines with AI solutions that can accurately predict meal demand, this issue serves as the central context for our design efforts.
The first phase, *Problem Identification and Motivation*, defines the research problem and justifies the value of a solution (Peffers et al. (2007), p. 52). When talking to experts in the field, we observed that currently, predictions are mainly based on static models and subjective experience. Motivated to explore how airlines can use AI to enhance their demand prediction, we conducted interviews with experts to study various approaches to achieve this goal, as shown in Table 2.

Table 2. Interviewed domain experts. NoI = number of interviews.

Market Offering	**NoI**	**Role**	**Experience**	**Time**
Startup: Meal demand forecast platform for airlines	5	Co-Founder	1 year at a startup, 2 years at an airline innovation centre	3 h 45 min
Airline innovation centre	3	Senior Venture Development Manager	3 years at an airline innovation centre, 1 year in research	2 h
Airline	1	Data Scientists	1 year at an airline, 1 year in data analytics	30 min
University	2	Research Associate	1 year in research	50 min

The second phase, *Objectives of a Solution*, derives the solution goals from the problem definition and existing knowledge (Peffers et al. (2007), p. 55). We gathered DRs (Table 4) through 11 stakeholder interviews that serve as our design objectives and formulated user stories addressing the airlines' need for precise demand forecasts.

The third stage, *Design and Development*, consists of creating the artefact (Peffers et al. (2007), p. 55). We derived design principles (Table 5) based on Seidel et al. (2017) and chose to design an architecture for AI-driven demand prediction (Table 4).

Table 3. Domain expert interviews for evaluation.

Role	**Experience**	**Evaluated artefacts**	**Time**
Co-Founder of a startup in meal demand forecast for airlines	1 year at a startup, 2 years at an airline research centre	Design Requirements, design principles, and system architecture	80 min
Senior venture development manager at an airline innovation centre	3 years at an airline research centre, 1 year in research	Meal demand prediction prototype	30 min
Data scientist	1 year in data analytics, 2 years in research	Design Requirements, design principles	30 min
Senior researcher at large enterprise system vendor	3 years at a vendor, 9 years in research	Meal demand prediction prototype	20 min

The fourth phase, *Demonstration*, tests the artefact's ability to address the problem (Peffers et al. (2007), p. 55). Based on our proposed architecture, we developed a prototype forecasting tool as a proof of concept, which helped validate the design and generated further insights.
The fifth stage, called *Evaluation*, is to observe and measure how well the artefact supports a solution to the problem (Peffers et al. (2007), p. 56). We conducted five additional interviews, as shown in Table 3, presenting the prototype, architecture, requirements, and principles to gather expert feedback. These interviews included both participants from the first round and other experts in the field.
Finally, the *Communication Stage* aims to communicate the problem and its importance, the artefact, its utility and novelty, the rigour of its design, and its effectiveness to researchers and other relevant audiences (Peffers et al. (2007), p. 56). We aim to present and discuss our findings with experts in the information system community in both academia and practice.

4 Findings

This section presents our findings on AI-driven demand prediction in airline catering. We present key system requirements, derived from expert interviews and industry-specific challenges, seeking a balance between operational efficiency, cost savings, and passenger satisfaction. Based on these requirements, we propose design principles that guide the development of such systems and suggest a system architecture that integrates AI-based forecasting into existing airline operations.

4.1 Design Requirements

Based on our expert interviews, we have gathered 11 design requirements shown in Table 4. Requirements 1 to 3 refer to the data, requirements 4 to 6 refer to the prediction, and requirements 7 to 11 refer to the dashboard.

Table 4. Requirements for meal demand forecasting systems.

ID	Requirements	User Stories and Quotes
1	Unified and Normalised Data Infrastructure	User story: As a data scientist, I want a standardised and comprehensive data infrastructure so that I can build reliable and accurate machine learning models. Quote: "Normalising all that data is just pain." (Data Scientist, Airline)
2	Data Security, Integrity, and Compliance	User story: As an airline, I want the system to comply with security and privacy regulations to ensure the protection of sensitive data. Quote: "We are working with sensitive data and the airline takes data privacy very seriously." (Startup)
3	Integration of External Data Sources	User story: As a data scientist, I want the system to integrate with external data sources so that forecasts can account for real-world influencing factors. Quote: "External data like weather, holidays and culture-specific eating behaviour could be of great interest." (Data Scientist, Airline)
4	Long- and Short-Term Demand Forecasting	User story: As a catering manager, I want forecasts for months ahead and quick updates for sudden changes, so I can plan purchases and adapt when needed. Quote: "The caterer needs the predictions three to four months in advance, but should also be able to adjust them in case of unforeseen events." (Startup)
5	Customer Satisfaction and Meal Availability	User story: As an airline, I want to ensure that passengers can always choose from all available meal options to ensure customer satisfaction. Quote: "Happy customers are essential." (Airline Innovation Centre)
6	Budget Compliance	User story: As an airline, I want to ensure that predicted meal quantities fit within budget constraints so that costs remain controlled. Quote: "We can't afford loading only salmon. Even if it's the only meal that's wanted." (Airline Innovation Centre)
7	Intuitive Dashboard for Decision Support	User story: As a catering manager, I want a dashboard visualising meal demand trends, and recommendations to take quick and informed decisions. Quote: "They [catering managers] need a front-end displaying the predictions." (Airline Innovation Centre)
8	Transparent Data Sources	User story: As an airline, I want to see how different data sources influence predictions to prioritise those that add the most value to forecasting accuracy. Quote: "It should be clear where the predictions come from." (Startup)
9	Cost and Waste Reduction Analytics	User story: As an airline, I want detailed analytics on cost savings and waste reduction so that we can track financial and environmental impact. Quote: "The amount of unused meals and the potential cost savings would be interesting." (Airline Innovation Centre)
10	CO_2 Tracking for Sustainability	User story: As an airline, I want to track the environmental impact of meal demand forecasting to report CO_2 savings and meet regulations. Quote: "We envision a future where premium experiences do not cost us the world." (Startup)
11	Integration into Existing Systems	User story: As an airline, I want the forecasts to seamlessly integrate into my and the caterers' current systems to ensure an efficient workflow. Quote: "Catering Managers don't want to learn a new tool." (Airline Innovation Centre)

DR 1: For AI-based forecasting to be effective, the system must standardise heterogeneous data sources and normalise them into a coherent structure. Without proper data integration, inconsistencies can significantly reduce prediction accuracy and increase the complexity of model training. This underlines the need for automated data standardisation to ensure seamless integration and processing.

DR 2: As the system processes sensitive airline and passenger data, it must comply with security regulations and data protection standards. This underlines the need for access control mechanisms and compliance with airline industry regulations.
DR 3: To improve prediction accuracy, the system must incorporate external data sources such as weather, holidays, and cultural food preferences. These external factors have a significant impact on passengers' food choices and should be considered.
DR 4: An AI-driven food demand forecasting system must support both long-term planning and real-time adjustments to ensure efficient procurement and adaptability to last-minute changes. Catering managers must be able to adjust predictions to compensate for fluctuations in demand, ensuring optimal meal availability for passengers.
DR 5: Passenger satisfaction is one of the top priorities in airline catering. Therefore, it is important that the system ensures the availability of meals, seeking a balance between avoiding waste and offering the full range of meals to passengers.
DR 6: Besides reducing waste, the system must be financially viable. So, airlines can not only base meal quantities on demand; financial constraints must be considered too. This means budgeting for forecasting is needed to assure that meal allocation is cost-effective and meets the passenger's preferences.
DR 7: To effectively use forecasts, catering managers need a clear and intuitive dashboard in an accessible format. The interface should display demand trends and actionable recommendations, allowing decision-makers to take informed decisions.
DR 8: Demand forecasting relies on multiple data sources, so transparency of data usage is essential. The interface must demonstrate how data contributes to forecast accuracy.
DR 9: Beyond accuracy, the system must analyse cost savings and food waste reduction. These insights into how AI forecasting compares to traditional meal planning methods are essential to justify its implementation.
DR 10: To comply with regulations and meet waste reduction targets, airlines must actively track CO_2 emissions and sustainability metrics.
DR 11: To avoid the need for employee retraining, the system must be integrated into the existing infrastructures of airlines and catering companies. Through Application Programming Interface (API) integration, demand forecasts could be accessed through familiar interfaces while enabling advanced forecasting capabilities.

4.2 Design Principles

Derived from the requirements, we formed nine design principles, presented in Table 5. Each principle is linked to one or several requirements as shown in Figure 2.
The first DP *Data-Driven Forecasting* underlines the need to use historical and real-time data to improve forecasting. By dynamically incorporating contextual factors such as weather, major events, and booking trends, the system ensures accurate forecasts and minimises waste (DR 5, 3). Achieving both requires good data-driven predictions that take these factors into account. Requirement 3 is met by using external data sources, as real-world influencing factors can lead to a better prediction. For example, in a data-driven approach, AI models would be trained using historical demand data, real-time

adjustments based on external factors, and dynamic updates.

Table 5. Design principles for meal demand forecasting systems.

ID	Design Principle	Application	DR
1	**Data-Driven Forecasting:** The system should leverage historical and real-time data to continuously improve predictions.	Train AI models using historical consumption data and external factors to improve prediction accuracy. Adjust forecasts dynamically based on real-time updates.	5,3
2	**Automated Data Integration:** The system should automatically aggregate, clean, and integrate data from multiple sources to enhance prediction quality.	Establish automated pipelines to pull and process external data in real time, ensuring forecasts are continuously updated and refined.	1,2, 3
3	**Composability and Modular Design:** The system design should be modular and adaptable to different industries.	Design a modular system with APIs, allowing different industries to integrate demand forecasting models without disrupting existing IT infrastructure.	11
4	**Explainability and Transparency:** The system should provide clear insights into how predictions are achieved and allow human oversight.	Display key influencing factors behind AI-generated forecasts and provide historical performance tracking for improvements.	7,8, 6
5	**Sustainable and Waste-Minimising Design:** The system should focus on minimising food waste while ensuring optimal meal availability.	Optimise stock levels within budget constraints to balance meal availability with minimal food waste. Track real-time stock levels and integrate redistribution options.	9,5, 6,10
6	**Real-Time Responsiveness:** The system should react dynamically to sudden changes in demand due to operational disruptions.	Automatically adjust demand forecasts based on last-minute disruptions (e.g. event cancellations). Notify supply chain teams of real-time adjustments to prevent over- or understocking.	4
7	**Long-Term and Iterative Planning:** The system should generate forecasts for extended periods of time (months) to support procurement, logistics, and strategic planning.	Generate meal demand forecasts for extended periods of time to optimise procurement and logistics. Continuously refine predictions by incorporating real-time updates and new data sources.	4,5
8	**Privacy and Data Security:** The system should ensure that all customer and operational data is handled securely and in compliance with regulations.	Apply anonymisation techniques when processing customer meal preferences. Ensure compliance and industry-specific regulations while securing API connections and data integrity.	2
9	**User-Centric Interface:** The system should provide intuitive dashboards that support decision-making for catering managers and logistics teams.	Provide intuitive dashboards that allow catering managers to view demand trends, receive AI-driven recommendations, and manually adjust predictions when necessary.	7,8

The second DP is *Automated Data Integration*, which highlights the importance of a data pipeline connecting different sources to improve forecast quality. Automatic integration, including data aggregation and cleaning, enables an optimised workflow (DR 1, 2 and 3). It is important that the data is prepared before processing (DR 1). Data

security also plays a role, as the automated pipeline works with sensitive data (DR 2).
The third DP is *Composability and Modular Design*. This principle addresses the challenge of designing systems that can seamlessly integrate into existing IT infrastructures across various industries. A modular architecture allows companies to adapt the system to their specific needs without the need for extensive customisation (DR 11). For example, this can be done using APIs.
Fourth, we propose the DP of *Explainability and Transparency*. The system should provide clear insights into how forecasts are generated and allow users to understand the factors that influence them. This enables decision-makers to build confidence in the system and make informed adjustments where necessary (DR 6, 7 and 8). The process can be supported by a transparent dashboard (DR 7). Requirement 8 is fulfilled because the system provides an overview that can be used to prioritise data sources. It is also important to guarantee budget compliance (DR 6). For example, highlighting key influencing factors ensures that stakeholders can understand the rationale behind recommendations.
The fifth DP *Sustainable and Waste-Minimising Design* emphasises the importance of balancing the availability of meals with minimising food waste and meeting environmental goals. (DR 5, 6, 9 and 10). Both DR 5 and 9 imply a balance between customer satisfaction and carefully monitoring costs and waste.
Real-Time Responsiveness is the sixth DP and highlights dynamically adjusting forecasts to sudden changes in operational conditions (DR 4). For example, catering services could update meal quantities based on event cancellations.
The seventh DP *Long-Term and Iterative Planning* focuses on generating forecasts for longer periods of time to facilitate logistic planning. As new data becomes available, forecasts are iteratively refined, ensuring they remain accurate and actionable (DR 4 and 5). Long-term forecasting (DR 4) is key to efficiency and is required to ensure meal availability (DR 5). An airline, for example, can use long-term forecasts while adjusting meal orders continuously based on passenger data.
Eighth, we identified *Privacy and Data Security* as a DP. It focuses on ensuring that all customer and operational data is handled securely and compliantly. By anonymising sensitive information and securing API connections, the system ensures data integrity and compliance. This principle addresses DR 2, which ensures compliance with existing security regulations. For example, meal preferences and booking data should be processed in a way that prevents unauthorised access or misuse.
Finally, a *User-Centric Interface* is essential. This DP underscores the need for intuitive dashboards for catering managers and logistics teams. The system should provide clear visualisations of demand trends, actionable AI-driven recommendations, and options for manual adjustments. This principle is derived from DR 7 and 8. In both cases, there is a need for a dashboard to make informed decisions and prioritise the best data sources through a transparent design. For example, catering managers could easily adjust meal orders or review forecast details in a user-friendly interface, reducing cognitive load in high-pressure situations.

4.3 Demonstration

Figure 1 presents the system architecture, demonstrating how the proposed design principles are implemented through structured data flow and component integration.
At the core of the system is a normalised database that centralises and processes internal and external data, ensuring accessibility and supporting an *Automated Data Integration* (DP 2). The Machine Learning (ML) model is a kind of AI that recognises patterns in this structured data and generates meal demand forecasts, aligning with the DP 1: *Data-Driven Forecasting*. The system also supports *Long-Term and Iterative Planning* (DP 7) by continuously exchanging data between the model and the database, ensuring *Real-Time Responsiveness* (DP 6) and improving forecast accuracy to minimise food waste, thereby promoting a *Sustainable and Waste-Minimising Design* (DP 5).

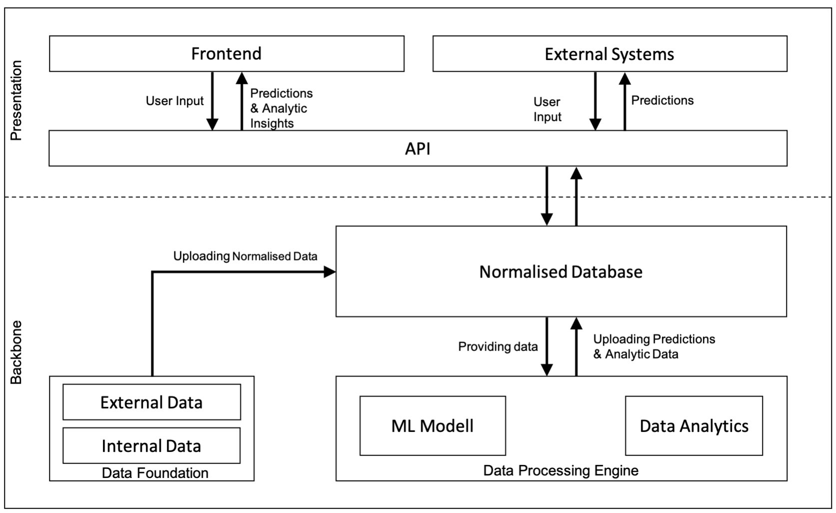

Figure 1. System architecture. The arrows indicate the data flow.

The API layer enables integration with external systems, supporting *Composability and Modular Design* (DP 3) and facilitating a *User-Centric Interface* (DP 9) through a front-end for prediction display and analytics.
To enhance *Explainability and Transparency* (DP 4), the database stores a detailed analysis, which allows users to understand and override predictions. This functionality enables a *User-Centric Interface* (DP 9) by giving users control over the forecasts, ensuring trust and human control over key decisions.
User adjustments are stored in the central database and processed by the data engine, where analytics identify patterns in manual corrections. This iterative process refines the forecasting model by addressing the discrepancies between AI predictions and real-world decisions. The updated insights are made available via the API for both the front-end and external systems, supporting continuous learning and optimisation. This contributes to a *Sustainable and Waste-Minimising Design* (DP 5) by reducing food waste.
The centralisation of data in a secure database also simplifies the implementation of

Privacy and Data Security (DP 8) guidelines, ensuring compliance and data integrity.
To validate the proposed architecture, we implemented a prototype for the entire forecasting pipeline. The system uses internal booking data and external influencing factors, such as weather conditions, public holidays, and culture-specific eating habits, which are processed through a data normalisation pipeline before being stored in the database. A machine learning model (Multi-Output Decision Tree Regressor) then generates demand predictions, which are made accessible via an API. The front-end provides real-time insights and allows users to make manual adjustments to the forecasts, ensuring control and transparency. In addition, waste reduction analysis and CO_2 tracking have been integrated, allowing airlines to monitor sustainability metrics and generate reports. This implementation successfully demonstrates the feasibility and capabilities of the proposed system and confirms compliance with the defined design principles.

4.4 Evaluation

We conducted expert interviews to evaluate our design requirements, design principles, architecture, and prototype, gaining valuable feedback. Overall, experts saw big potential in data-driven forecasting to improve prediction accuracy and emphasised the importance of incorporating external factors. As one expert noted, “It is very exciting to explore whether external factors have an impact on the prediction, and if you choose the right ones, they can have a significant impact.”
During the evaluation of the DRs, experts stressed the importance of an intuitive dashboard (DR 7) and data security (DR 2). During our conversation, one expert pointed out the aspect of budget compliance: "We can’t afford loading only salmon. Even if it’s the only meal that is wanted." Responding to this feedback, we added it as DR 6: *Budget Compliance*. During the evaluation of the design principles, experts highlighted the importance of composability and a modular design (DP 3), as this new component must be integrated into their existing system. In addition, experts pointed out the importance of explainability and transparency: “It builds trust in the system. There is a fine line between how much explanation the analysis and the model need so that it doesn’t look like a black box and giving too much unnecessary information.”. Regarding the DP 6: *Real-Time-Responsiveness*, experts noted that last-minute adjustments regarding food quantities were not feasible. Hence, we adjusted the DP accordingly, now referring only to major disruptions. Reinforcing DP 9: *User-Centric Interface*, experts praised our interface for its intuitive design and noted that this aspect was especially important to minimise onboarding time, ensuring that the system is accessible and efficient.
Lastly, we want to add that quantitative evaluation will add to future research in the field.

5 Discussion

Our DPs, derived from the DRs and evaluated by expert interviews, align with existing literature. In the following section we will discuss each design principle regarding its similarities to contributions of similar studies.
DP 1, 2, 5: There are several studies that have explored the impact of ML on food demand prediction in other contexts than aviation (Hübner et al. 2024, Aishwarya et al.

2020, Çetinkaya & Erdal 2019). In their findings, they demonstrate that *Data-Driven Forecasting* (DP1) significantly boosts forecasting accuracy. Their approach highlights the impact of external factors such as holiday periods and weather conditions, which we integrated into our prototype, among others. In addition, Armstrong & Green (2005) emphasises the need to incorporate domain knowledge. In order to facilitate this process, we emphasise *Automated Data Integration* (DP 2). In their study Hübner et al. (2024), improved forecasting accuracy led to a 30% reduction in returns, highlighting how predictive precision supports both sustainability and operational efficiency. As we found out in our expert interviews, it is important to achieve this goal within budget constraints (DP 5).

DP 3: In their work, Miller et al. (1991) describe the ideal of a forecasting system with a *Composable and Modular Design* (DP 3), which is reflected in our architecture.

DP 4: As stated by Armstrong & Green (2005), the decision makers' acceptance of the forecast defines its value. Hence, its *Explainability and Transparency* (DP 4) are crucial.

DP 6, 7: As most bookings are made within 12 weeks prior to departure (Haensel et al. 2011), ensuring the system's *Real-Time Responsiveness* (DP 6) poses a real challenge. Nevertheless, *Long-Term and Iterative Planning* (DP 7) remain important for food logistics.

DP 8, 9: We added the principles *Privacy and Data Security* (DP 8) and *User-Centric Interface* (DP 9) because experts stressed their importance during the conducted interviews.

Although our focus is on in-flight catering, the principles of meal demand forecasting apply to other sectors, including food services (Çetinkaya & Erdal 2019) and revenue management (Haensel et al. 2011), where accuracy is particularly important for managing products with short shelf lives (Aishwarya et al. 2020).

6 Conclusion

In this paper, we applied a DSR approach to derive design requirements and design principles for meal demand prediction systems. Based on an airline use case, we generalised our findings into nine design principles, following the full DSR process, including problem identification, goal definition, design and development, demonstration, evaluation, and communication. The resulting prototypical architecture aims to enhance prediction accuracy and system integration, thereby improving efficiency, sustainability, and cost savings. Nonetheless, the focus on the airline industry can introduce industry-specific bias, even though design principles were formulated for broader applicability to meal forecasting systems. Future research should validate them across different industries, including sectors beyond aviation. Moreover, the small sample size limits the diversity of requirements captured, which could be addressed by including additional regions and organisations.

References

Aishwarya, K., Nagaraja Rao, A., Kumari, N., Mishra, A. & Rashmi (2020), 'Food demand prediction using machine learning', *International Research Journal of Engineering and Technology* .

AmericanAirlines (2023), 'Sustainability report 2023', https://s202.q4cdn.com/986123435/files/images/esg/American-Airlines-Sustainability-Report-2023.pdf.

Armstrong, J. & Green, K. (2005), 'Demand forecasting: Evidence-based methods', *SSRN Electronic Journal* .

BritishAirways (2023), 'Annual report and accounts year ended 31 december 2023', https://www.iairgroup.com/media/2h0hgqdq/british-airways-annual-report-and-accounts-2023.pdf.

easyJet (2023), 'Making low-cost travel easy, annual reports and accounts 2023', https://www.easyjet.com/ejcms/cache/medialibrary/Files/Sustainability/43627_easyJet_Annual-Report_Web.pdf.

Gorr, W. L., Nagin, D. & Szczypula, J. (1994), 'Comparative study of artificial neural network and statistical models for predicting student grade point averages', *International Journal of Forecasting* **10**(1), 17–34.

Haensel, A., Koole, G. & Erdman, J. (2011), 'Estimating unconstrained customer choice set demand: A case study on airline reservation data', *Journal of Choice Modelling* **4**, 75–87.

Hübner, N., Caspers, J., Coroama, V. & Finkbeiner, M. (2024), 'Machine-learning-based demand forecasting against food waste: Life cycle environmental impacts and benefits of a bakery case study', *Journal of Industrial Ecology* **28**, 1117–1131.

Iberia (2023), 'Sustainability report two thousand and twenty-three', https://grupo.iberia.com/contents/archives/475/109/pdfcontent/475_109_1726662870.pdf.

Kumar, I., Rawat, J., Mohd, N. & Husain, S. (2021), 'Opportunities of artificial intelligence and machine learning in the food industry', *Journal of Food Quality* **2021**.

Lee, Y.-C. & Seo, C.-G. (2001), 'Airline in-flight meal demand forecasting with neural networks and time series models', *The Journal of Information Systems* **10**.

Liao, M., Lan, K. & Yao, Y. (2021), 'Sustainability implications of artificial intelligence in the chemical industry: A conceptual framework', *Journal of Industrial Ecology* **26**.

Liman, D. F. (2016), 'Demand forecasting for perishable commodities: A case study of inflight food demand for low cost airline'. Licensed under CC BY 4.0.

Miller, J. L., Thompson, P. A. & Orabella, M. M. (1991), 'Forecasting in foodservice: Model development, testing, and evaluation', *Journal of the American Dietetic Association* **91**(5), 569–574.

Panda, S. & Dwivedi, M. (2020), *Minimizing Food Wastage Using Machine Learning: A Novel Approach*, pp. 465–473.

Peffers, K., Tuunanen, T., Rothenberger, M. & Chatterjee, S. (2007), 'A design science research methodology for information systems research', *Journal of Management Information Systems* **24**, 45–77.

Portugal, T. A. (2023), '2023 sustainability report'.

Rao, G., Shastry, K. A., Sathyashree, S. & Sahu, S. (2021), *Machine Learning based Restaurant Revenue Prediction*, pp. 363–371.

RyanairGroup (2023), ‘Aviation with purpose, 2023 sustainability report’, https://corporate.ryanair.com/wp-content/uploads/2023/07/Ryanair-2023-Sustainability-Report.pdf.

Seidel, S., Chandra Kruse, L., Székely, N., Gau, M. & Stieger, D. (2017), ‘Design principles for sensemaking support systems in environmental sustainability transformations’, *European Journal of Information Systems* **27**.

van Ryzin, G. (2005), ‘Future of revenue management: Models of demand’, *Journal of Revenue and Pricing Management* **4**, 204–210.

Çetinkaya, Z. & Erdal, E. (2019), Daily food demand forecast with artificial neural networks: Kırıkkale university case, pp. 1–6.

7 Appendix

7.1 Design Requirements and Design Principles

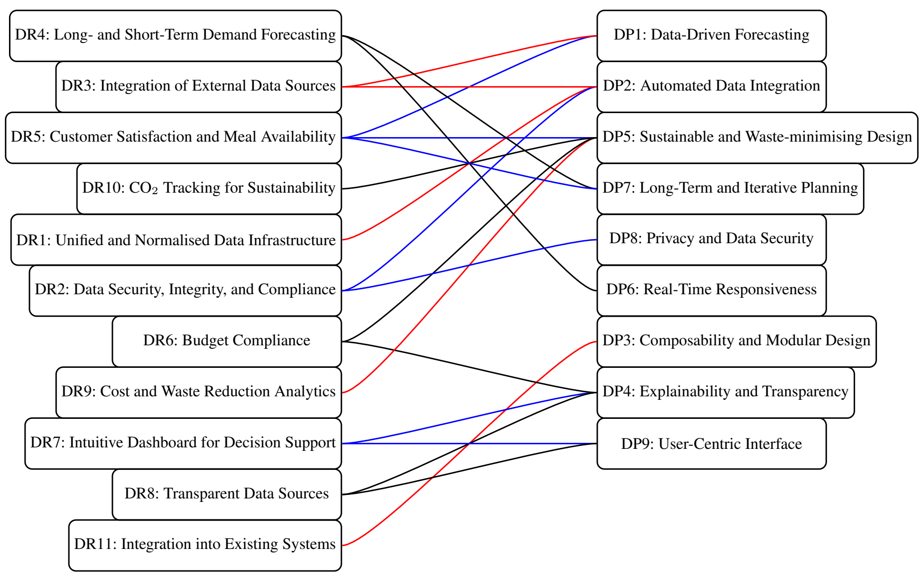

Figure 2. Mapping Design Requirements (DR) to Design Principles (DP). Coloured lines and a changed order improve readability.