Human -Data Analytics Interaction through Voice Assistance in Electric Vehicle's Battery Testing

Chapter · September 2024

DOI: 10.1007/978-3-031-71633-1_20

9 authors, including:

Alexandros Bousdekis
National Technical University of Athens

Katerina Lepenioti
National Technical University of Athens

Federica Mandreoli
University of Modena and Reggio Emilia

Stefan Wellsandt
Bremer Institut für Produktion und Logistik GmbH

Human – Data Analytics Interaction through Voice Assistance in Electric Vehicle's Battery Testing

Mattheos Fikardos[1], Alexandros Bousdekis[1], Umair Haider[2], George Aristofanous[1], Katerina Lepenioti[1], Federica Mandreoli[2], Stefan Wellsandt[3], Enrico Taglini[4], Gregoris Mentzas[1]

[1] Information Management Unit (IMU), Institute of Communication and Computer Systems (ICCS), National Technical University of Athens (NTUA), Athens, Greece
[2] Department of Physical, Computer and Mathematical Sciences, University of Modena and Reggio Emilia, Modena, Italy
[3] BIBA – Bremer Institut für Produktion und Logistik GmbH at the University of Bremen, Bremen, Germany
[4] Reinova S.p.a., Soliera, Modena, Italy

Abstract. Voice assistants, alternatively mentioned as conversational agents or Digital Intelligent Assistants (DIA), represent a new form of interaction between humans and machines, providing fast, intuitive, and potentially hands-free access to systems through voice-based interaction in order to increase the efficiency of certain activities. While the literature has mainly focused on general applications of voice assistants in diverse industries, their potential in manufacturing remains mostly underexplored. This is despite the manufacturing sector is a key driver for employment and plays a critical role in economic growth. Furthermore, enabling human workers to interact with data analytics insights through voice interfaces is key to realize a human-system symbiosis in an Industry 5.0 context. However, there is limited literature regarding the data analytics potential integrated into voice assistants and related implementations due to, among others, the challenges of translating analytical results into easy-to-understand information for humans. In this paper, we face these topical issues for DIA by presenting a voice assistant equipped with data analytics functionalities to support human-machine interaction in the manufacturing sector when there are data analytics insights that are communicated to the user. We demonstrate this with a concrete instantiation in Electric Vehicle's (EV) battery testing use cases.

Keywords: Human-AI collaboration, Conversational agent, Data analytics, Machine learning, Li-Ion battery.

1 Introduction

Voice assistants, alternatively mentioned as conversational agents or Digital Intelligent Assistants (DIA), allow users to interact intuitively by using their natural language (e.g. Apple's Siri, Amazon's Alexa, IBM's Watson). These systems represent a new form of interaction between humans and machines, a paradigm shift from the current Graphical

User Interfaces (GUIs), allowing the user to interact using natural language, providing fast, intuitive, and potentially hands-free access to systems through voice-based interaction and cognitive assistance, in order to facilitate, speed up and increase the efficiency of daily activities [1–4].

In the industrial sector, the adoption of conversational agents has the potential to drive the digital transformation of organizations, to improve both customer and user-experience, and make their internal processes more efficient [5] by increasing the degree of mobility, autonomy, and independence of operators as well as by entrusting the most repetitive operations to these technologies [6]. In this sense, the development of conversational agents is focused on supporting users in interacting with machines, databases, and information systems, and completing tasks, moving towards the notion of smart operators. While the literature has mainly focused on general applications of voice assistants in diverse industries, including banking, insurance, public transportation, retail, and customer relationship management, the aforementioned potential in manufacturing is underexplored, although the manufacturing sector is a key driver for employment and plays a critical role in economic growth [2, 7].

Further, enabling human workers to interact with data analytics insights through voice interfaces is key to realize a human-system symbiosis in an Industry 5.0 context [8]. However, there is limited literature regarding the data analytics potential integrated to voice assistants and related implementations due to, among others, the challenges of translating analytical results into easy-to-understand information for humans. The synergies between voice assistance technology and data analytics to provide validation of their benefits in the real world is an important future research direction [7, 9].

In this paper, we propose an architecture for voice assistant integrated to data analytics functionalities and its application to a business case coming from battery manufacturing, one crucial manufacturing sector which can highly benefit from this kind of technology. The increasing deployment of Li-Ion batteries to satisfy the demand from the Electric Vehicle (EV) sector pushes battery manufacturers to increase their efficiency and productivity. One of the main objectives to this end is to optimize the existing battery testing procedures, which are time-consuming and costly, while they lead to high scrap rates [10]. To do this, they have started adopting emerging technologies, such as Artificial Intelligence (AI) and Machine Learning (ML) [10]. However, product testing procedures require highly manual work on the shop floor by the operators [11], thus making it inefficient to have complex visualizations on a GUI or extensive dialogues with a voice interface, especially when dealing with data analytics insights.

The developed voice assistant aims at tackling with human-machine interaction challenges when there are data analytics insights that are communicated to the user through an analytics service module which can effectively adapt to the user requests and a voice assistance module that is specifically designed to identify user intents and interact with the analytics service module to support the dialogue with the user. In this way, operators can get useful insights about the testing procedures and prevent overheating and potential explosions by closely monitoring the conditions of the batteries under testing. The rest of the paper is organized as follows. Section 2 discusses the related works. Section 3 presents the proposed architecture for human-data analytics interaction through voice assistance. Section 4 instantiates the proposed architecture to an EV's battery testing use case, describes its implementation, and demonstrates its results. Section 5 concludes the paper and discusses our plans for future work.

2 Related Work

The literature has mainly focused on general applications of voice assistants in diverse industries, including banking, insurance, public transportation, retail, and customer relationship management. However, their potential for manufacturing has not been investigated extensively yet [7]. In the manufacturing context, voice assistants aim at simplifying interactions between workers and complex machines, reinforcing workers' physical and cognitive capabilities, reducing the human effort involved in performing various tasks, and allowing the operators focusing on creative and value-adding activities [7], thus contributing to the realization of the Operator 5.0 concept [12]. Operators can instruct voice assistants to perform certain tasks leaving them able to focus on abstract tasks instead, and they can also obtain assistance from them in developing skills and competencies, e.g. during on-the-job training [13].

The voice assistant's tasks can be categorized into the following three macro-task types: information management (querying, retrieving, consulting, or analysing information), collaborative operations, (tasks in which voice assistants either instruct or assist human workers or assume the responsibility of completing parts of a joint task), and knowledge transfer (the involvement of voice interfaces in reducing human effort for the acquisition of skills during real-life tasks). If voice assistants access multiple information systems and provide analytics for decision-making, their use can shorten task completion time and improve task efficiency. They can also support people ubiquitously, thereby increasing flexibility [7]. However, there are several challenges that set obstacles towards this direction.

There is no agreement on the elements to be considered and developed for the creation of an industrial conversational agent, partly due to the absence of a reference standard and a general lack of mastery about their logical operation and characteristics. For example, in [2], the authors identify 5 core modules, i.e. Automatic Speech Recognition, Natural Language Understanding, Dialog Manager, Natural Language Generation, and Text-To-Speech, and two Interfaces: Conversational User Interface, and External Devices Interface. Voice assistance technology introduces complexities that can arise from intricate technical systems, for instance, accessing different user interfaces, interpreting structured, semi-structured, or unstructured data, or precise understanding of speech. In addition, there are challenges related to the integration of voice assistants into the existing systems as well as regulatory compliance [7].

At the same time, research has focused on Generative Artificial Intelligence (GAI) with the use of Large Language Models (LLM) with promising results in human augmentation, especially in industrial applications [14, 15]. Further, only a few research works describe the data analytics potential of voice assistance technology (e.g. [16–18], while the literature and the real-life applications on analytics-based voice assistance are rather limited due to the challenges related to translating analytical results into easy-to-understand information for humans [2, 7]. Existing intelligent voice assistants mainly rely on rule-based systems [2].

3 Architecture for Human - Data Analytics Interaction through Voice Assistance

In our approach, the Digital Intelligent Assistant (DIA) aims to assist technicians and engineers in a laboratory setting where machines are running tests on various devices, referred to as Device Under Test (DUT), by serving data analytics upon user request via voice. Real-time status updates and analytics are provided through the voice interface to augment the user's tasks with safety and efficiency.

The system architecture, depicted in **Fig. 1**, consists of three main modules: Voice Assistant, Analytics Service, and the Use case Infrastructure. Interaction between the user and the DIA is managed through the Voice Assistant module. This module uses Mycroft as a natural language interface, prioritizing user privacy and customizability, and identifies relevant Mycroft skills. A "Mycroft skill" is a Python program with specific functions and language-specific training data for intent classification and entity extraction.

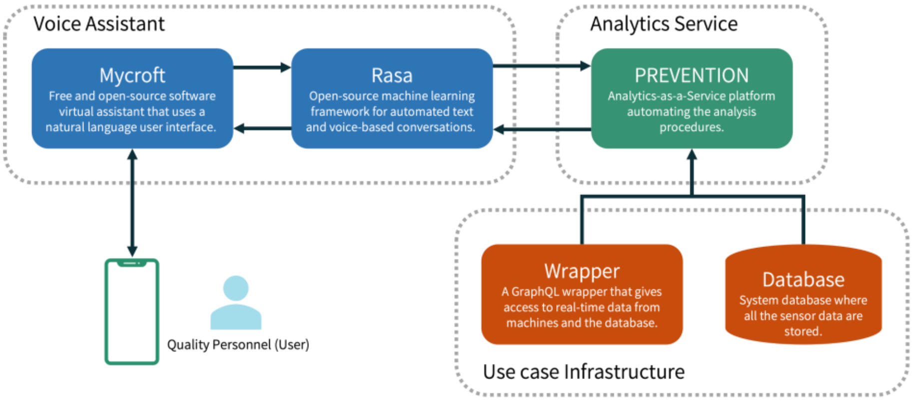

Fig. 1. DIA architecture for human - data analytics interaction through voice assistance.

The Voice Assistant module translates the user's voice input into text and passes it to Rasa, a conversational AI software for building text and voice-based assistants. Rasa maps the user input to one of the defined intents and formulates a query with all the necessary variables. Rasa's responsibilities include interpreting the user input and generating queries for the Analytics Service, which responds with the required insights.

The analytics are handled by the PREscriptiVE aNalyTIcs for quality optimizatiON (PREVENTION) component, which communicates directly with the use case infrastructure to retrieve necessary data and perform analyses. The results of these analytics are encapsulated into a response and sent back to the Voice Assistant. Based on the result context, an intuitive response is formulated and given to the user in both voice and text formats.

3.1 Voice Assistant

The Voice Assistant module serves as the interface for user interaction and the link between analytics and the user. It uses Mycroft for NLP tasks to translate voice inputs to text and vice versa, and Rasa to map user inputs into intents, enabling digital assistant functionalities. This dual-framework approach ensures accurate interpretation of user queries and appropriate responses. With separate skills (Mycroft skills) for identical functions and operations, it can provide a higher level of modularity. The information flow among the voice interface components is depicted in **Fig. 2**. The Rasa server utilizes machine learning models to understand and classify user queries, ensuring responses are both contextually relevant and technically accurate. The Rasa action server is a component responsible for executing custom actions beyond the capabilities of predefined responses, allowing developers to create Python code interacting with databases, external services, or carrying out computations during conversations.

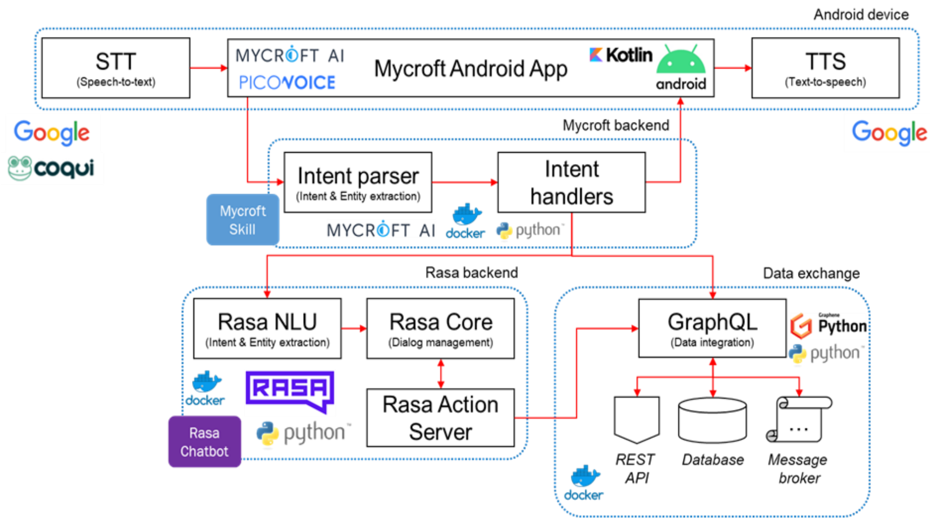

Fig. 2. The information flow among the various components of the voice interface.

3.2 Analytics Service

A crucial part of the functionalities of the described solution is the insights provided by the analytics service module of the architecture, which are communicated to the user via the voice assistant. Addressed by the PREVENTION component, the analytics module follows the Analytics-as-a-Service (AaaS) paradigm [19], serving the voice assistant with the need analytics and insights upon request. The architecture of the PREVENTION component is depicted in **Fig. 3**.

PREVENTION provides strict, yet dynamically configurable functionalities provided by the platform that augments the design and development of analytics and simplify their deployment for generating the desired results. The data analyst can make the configurations and select the building blocks for the analytics from the wide arsenal of available algorithms, ML models, and methodologies to achieve the best results based on the specific requirements and needs. Specifically, the platform can serve data-processing methods to transform the raw data, descriptive analytics by employing aggregation and group-by algorithms, as well as predictive and prescriptive analytics by implementing ML models (supervised, unsupervised, and reinforcement learning). Moreover, two additional functionalities supported by the platform are the configurable deployment of the models to continually update the produced results and a filtering mechanism that enables the users to investigate and pinpoint specific insights from the results.

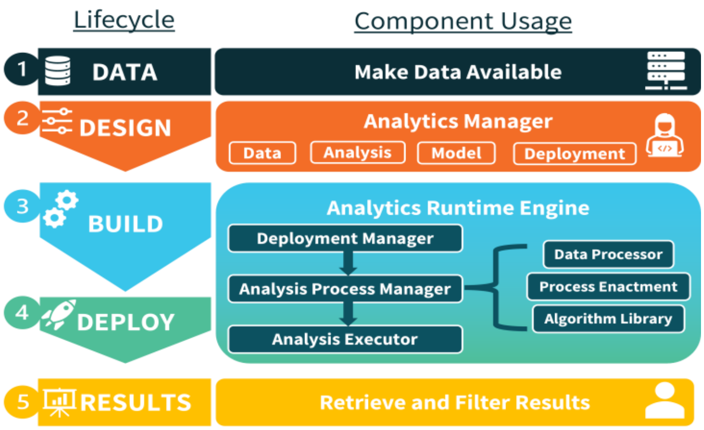

Fig. 3. The PREVENTION architecture.

The core of the platform's architecture consists of two modules: the Analytics Manager and Analytics Runtime Engine. The Analytics Manager module is responsible for the design of the analysis done manually by the data analyst. The first step for configuring the analysis is to define the entities needed to execute it, which are the Analysis, Model, Data, and Deployment. These entities are essential for the robust implementation of an analysis as they define its scope and goal (Analysis), the algorithm to be used (Model), the input raw data (Data), and how the analysis is deployed and served (Deployment).

After the configuration, the Analytics Runtime Engine kick starts that can automatically develop and deploy the respective building blocks of the analysis and produce the specified results. First, the Deployment Manager creates, schedules and manages the deployment of the analysis based on the Deployment configurations. Then the Analysis Process Manager first creates the process needed to execute the analysis, which is then instantiated by the Process Enactment sub-module. First, the analysis is compiled to prepare the needed inputs which are the data and the algorithms. The input data are produced by the Data Processor, by applying data-processing actions to the raw data to transform them into the desired format. The algorithms are implemented with the help of the Algorithms Library, giving them the configurations provided by the data analyst. As soon as the data and algorithms are ready, the analysis is executed through the Analysis Executor, which runs the pipeline created for the analysis generating the results.

The results, as well as all the algorithm instances, are then saved to the database that can be used to expose the results through an API and make their deployment independent of the building phase. Finally, the Analytics Results Query Engine is executed on top of the stored results giving the user the functionalities of the filtering mechanism. All the above functionalities and actions are communicated to the platform in the form of requests that are exposed through the API and can be divided into two groups based on their scope: (i) Analytics Management Queries: A set of CRUD queries that can be used to create, read, update, or delete the six main entities of the component's database; and, (ii) Analytics Process Queries: Requests that serve the analysis results and are enhanced with the filtering mechanism.

3.3 Use Case Infrastructure

A key ingredient that enables the utilization of the aforementioned technologies is data that are stored by the use case, usually at an on-premises infrastructure due to security factors. This data, that drives the insights generated by the analytics and the dialogue designs for the voice assistant, can be accessed either through an API or directly from the database. API access allows seamless integration with other components, while direct database access provides detailed information for data analysts to design effective analyses.

4 Application to EV's Battery Testing

The architecture presented in Section 3 has been implemented for a business case coming from an Italian SME specialized in developing, testing and validation of electric and hybrid powertrain components. In the following, we provide an overview of its working environment, describe the intents and entities defined for the specific business case, and, finally, present the flow of information for two representative case studies.

4.1 The Reinova Business Case

Reinova operates a well-equipped facility with a focus on chambers and battery testers for conducting various tests. These include Vehicle Climatic Chambers, Walk-in Big Climatic Chambers, and Walk-in Small Climatic Chambers. Each chamber transmit different sensor data including temperature, humidity, and electric power consumption. This data is used to create tests for the customers as well as several safety mechanisms that prevent dangerous situations to be occurred (e.g. fire hazards). Chamber's sensor data is continuously sent to a SQL DBMS where it is stored through a dedicated communication protocol.

Moreover, Reinova owns an important set of testing and safety instruments, the so-called "battery testers", which are devices that provide sink and source for high voltage battery packs. The output of these testers is forwarded into the test volume of the chambers, where the DUTs undergo evaluation. Similar to chambers, battery testers' information is managed via Programmable Logic Controllers (PLCs). The readings from various buttons and sensors (maintenance and emergency button, CO2/H2/O2 concentration, and others) are constantly available to the PLC, that uses them to perform certain actions and ensure safety. The architecture of Reinova's infrastructure is depicted in **Fig. 4**.

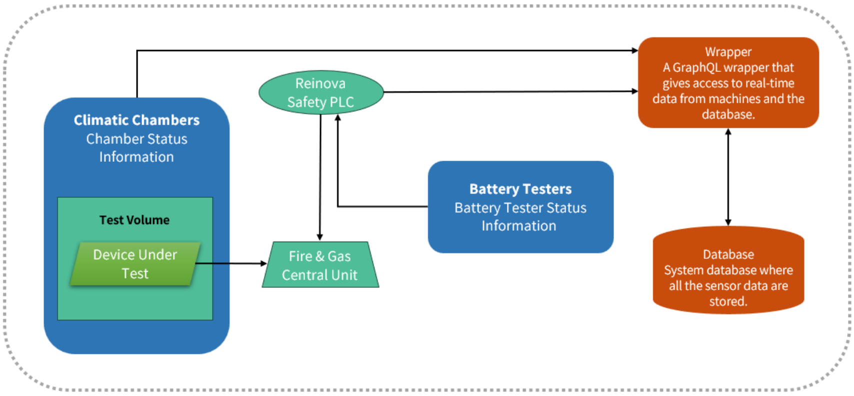

Fig. 4. The architecture of Reinova's infrastructure.

The DIA is set to be deployed in the laboratory environment where all the testing by technicians and engineers occurs with the aim being the streamline of two core functionalities of interest that can significantly enhance the laboratory's efficiency, responsiveness and guidance through problem resolution: *(i) the "alarm and emergency notification"* function that is designed to alert users about critical incidents, including alarms triggered during tests and emergencies related to devices, such as fires, gas leaks, sensor faults and more; and, *(ii) the "machine status request"* function that allows the users to query about the condition of a specific machine or receive an overview of the laboratory's operational status. The DIA responds with detailed information about the machine's current state (running of a test, not running, ready, emergency) alongside the status of the sub-item of the machines (sensor readings). Moreover, in the event of a malfunction of the machines, the assistant notifies the user about necessary actions to restore the machine to its operational status.

4.2 Implementation Procedure

Defining Entities. In the previous Section, we described why the intents of the voice assistant are essential to explain the entities that make it up. Entities are crucial components that the system identifies and uses to provide a context-specific response. These entities are usually used as parameters to query data from the APIs. For the intents, the entities are presented in **Table 1**.

Table 1. Definition of the entities.

Entity	Explanation
chamberId	A unique identifier for a chamber within the system. It allows the system to fetch and provide data for a specific chamber. (Example: B1, B2, B3).
batteryTesterId	A unique identifier for a battery tester. It enables the system to retrieve and communicate the battery status for the specified tester (Example: BT1, BT2, BT3).
defconValue	It is used to catch the DEFCON level of the chamber. DEFCON level refers to the status of the chamber. (Example: Ready, Alarm, Maintenance)
startDate	It is used in the show_defconlevel_date_filter intent to apply a filter in the query.
endDate	It is also used in the show_defconlevel_date_filter.
relativeDate	It is used in the show_defconlevel_relative_date_filter entity to catch relative time occurrences in the conversation related to DEFCON levels.
topValues	It is used in the show_defconlevel_sum_filter intent to catch the count.
specificDate	It is used in show_defconlevel_complex_filter intent to store a specific mention of date and time in the conversation.

Intents Understanding. An intent represents the purpose behind a user's input. In the system under examination, intents are categorized based on the type of information the user is seeking from the status monitoring. The primary intents identified in this use case are presented in **Table 2**.

Table 2. The identified intents.

Intent	Explanation
Status Query	This intent is triggered when a user inquiries about the general status of a chamber. It is designed to provide a comprehensive update on the chamber's condition based on the incoming data from the APIs.
Temperature Test	When user asks about the temperature of a specific chamber, it is invoked to fetch and relay temperature details.
Battery Status Query	It is invoked when the user wants to know the current battery status of a specific tester.
Chamber List Query	It relates to the action which is responsible for fetching the list of all the chambers.
Battery List Query	Like the chamber list intent, it relates to the action to fetch the list of battery testers available.

Criticality Status Query	It is part of the analytics component and indicates the criticality of an inquired chamber. It uses the chamberId entity to fetch the results.
Show DEFCON Levels	It is part of the analytics module. Its purpose is to handle the queries related to the DEFCON levels of the chambers.
Show Specific DEFCON Levels	It will handle the queries related to specific DEFCON levels, i.e., Alarm, Ready, and Maintenance. It uses defconValue as an entity in the conversation.
Show DEFCON Levels with Date Filter	It is used while querying DEFCON level's occurrence with a specific date and time. It can also involve the startDate and endData entities in the conversation.
Show DEFCON Levels with Relative Time Filter	It is used to resolve the relative date and time mentioned in the conversation while asking about the DEFCON levels. It uses relativeDate entity in the conversation.
Show DEFCON Levels with Count Filter	It is used to resolve the count of DEFCON level occurrences.
Show DEFCON Levels with Complex Filter	It can contain both DEFCON level and time filter in it. It can catch both defconValue and specificDate or relativeDate entities to apply the filters and to get the answers.

Available Analytics. After acquiring and analyzing the business requirements, we narrowed down to four analytics processes that had the most value for the flow of work inside the Reinova laboratories. The selected analytics processes were also defined according to the aforementioned intents and are explained in **Table 3**.

Table 3. The selected analytics processes.

Analysis	Explanation
Group By	It groups and counts the recorded entries from the Chamber dataset based on their status (defconlevel) per day.
Duration Of Alarm	The duration of the flagged alarms are calculated. For each chamber, the start and end timestamps are recorded for each status, and the durations are calculated in two formats: i) as a timestamp (HH:MM:SS) and ii) in seconds.
Defconlevel Duration	The duration of all the chamber statuses are calculated. For each chamber, the start and end timestamps are recorded for each status, and the durations are calculated in two formats: i) as a timestamp (HH:MM:SS) and ii) in seconds.
Defconlevel Percentages	The percentages of all the chamber statuses are calculated inside a certain time frame, in our case a day. For each chamber, the start and end timestamps are recorded for each status, and the percentages are calculated in two formats: i) as a timestamp (HH:MM:SS) and ii) in seconds.

4.3 Demonstration

This Section demonstrates the flow of information for the designed analytics targeted to the use case under examination.

Status aggregation. As part of the insights provided by the analytics, aggregation algorithms can be used to group and count different aspects of the available data points. In this case, we opted to make aggregations based on the statuses of the machines that have the most importance for the users. The grouping of the data was implemented across all the different statuses and machine IDs, and the smaller time frame of recording was per day. In the example given in **Fig. 5**, the user requests the recorded statuses of Maintenance per day, in a specified time interval, while also specifying that only the top 3 records are selected. The voice assistant maps the user's input to the appropriate intent and generates the query to be sent to the analytics component, by filling in all the necessary variables. The analysis results are filtered with the given parameters and the requested insights are sent back to the voice assistant to be displayed back to the user.

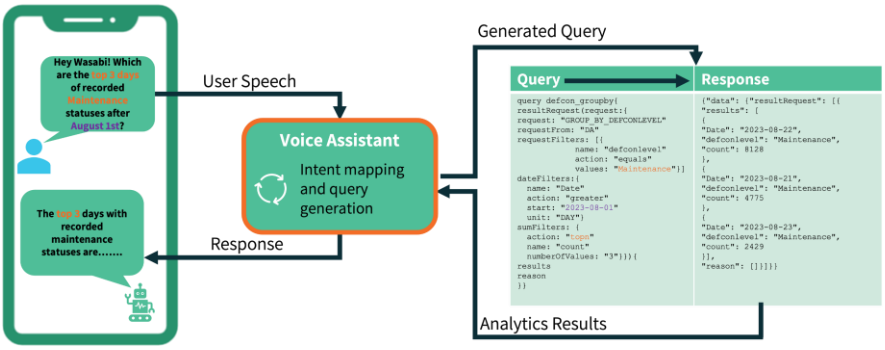

Fig. 5. Flow of information for the "status aggregation" analysis.

Status duration. The goal of this analysis is to provide the user with the aggregated durations of the different statuses recorded by the machinery sensors. The results of this analysis aggregate the status durations from all the available machinery while recording the start and end timestamps as well as the duration in seconds. Due to the immense volume of results generated, the filtering mechanism functionalities of the component can be utilized to extract specific information of interest. In the example depicted in **Fig. 6**, the user requests the status duration for a specific machinery (chamber A1) and a specific time frame (since August). Based on the input, the Voice Assistant maps the user's request to the respective intent, extracts the parameters, and automatically transforms the given time frame into filtering parameters needed for the query. After the query is generated and populated, it is sent to PREVENTION where it is processed and the specified results are returned to the Voice Assistant. With the retrieved insights, the response text is formulated and sent back to the application where it is displayed to the user as text and also narrated with voice.

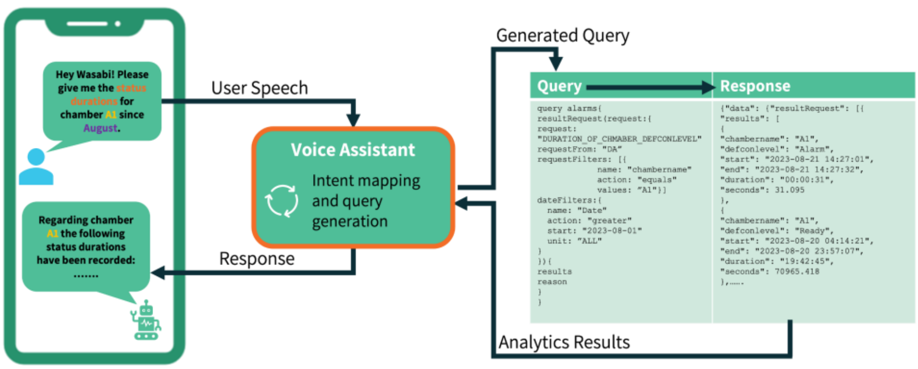

Fig. 6. Flow of information for the "status duration" analysis.

Status Percentage. A similar example to the previous one, in addition to the aggregated durations of the statuses an analysis was implemented to calculate the percentages of each status per day. Alongside the percentage for each day, the durations were also kept in order to adjust the numbers in cases where the filtering mechanism was used for a different time interval. In the dialogue of **Fig. 7**, the user requests the percentages of the status Alarm of a specific chamber A1, while also specifying to the analytics to return the results for the top 3 records for the first two months of 2024. The voice assistant maps the user's input to the defined intents and generates the query for the analytics component, which returns the requested analysis results with the specified insights.

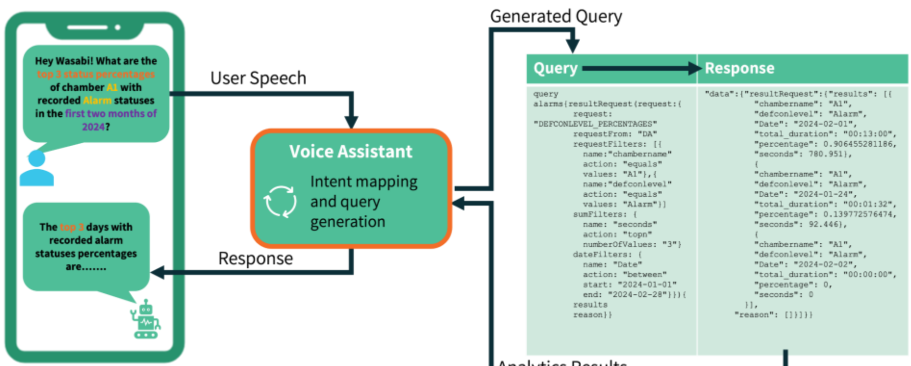

Fig. 7. Flow of information for the "status percentage" analysis.

5 Conclusions and Future Work

This paper presented a voice assistant integrated to data analytics functionalities in order to support human-machine interaction when there are data analytics insights that are communicated to the user in EV battery testing use cases. We designed and developed the architecture and we described how we implemented it in the EV battery testing use case.

In our future work, we will embed additional data analytics functionalities, in order to fulfill more business requirements, and to incorporate Large Language Models (LLM) for user interaction. Moreover, we will conduct an extensive evaluation of the proposed solution in real-life scenarios derived from the Reinova business case. This evaluation will focus on both quantitative and qualitative aspects. For the quantitative evaluation, we aim to measure the performance of the trained Rasa model in identifying intents and entities. As a part of our evaluation process, we will leverage the Rasa module's testing framework which processes test stories and produces a comprehensive set of performance metrices, i.e. precision, recall, F1 scores, confusion matrix, and confidence histogram model. For the qualitative evaluation, instead, we aim at implementing the evaluation methodology for AI-based digital assistants, proposed in [20], which consists of 4 dimensions: system usability, cognitive workload, business benefits, and AI trustworthiness. DIA testing will be performed on-premises and tailored towards tasks ypically carried out by technicians and engineers in the laboratory, which are expected to benefit from the introduction of the DIA. Standard usability evaluation surveys, such as System Usability Scale (SUS) [21] and NASA Task Load IndeX [22], will be administered to the operators after several weeks of usage to collect statistically relevant measures on system usability and cognitive workload. Business benefits will concern with Key Performance Indicators related to the performed tasks while AI trustworthiness will be measured with the Assessment List for Trustworthy Artificial Intelligence (ALTAI) questionnaire. The collected statistics will be assessed to contribute to the continuous improvement of the developed DIA.

Acknowledgments. This work is funded by the European Union's Horizon Europe project WASABI "White-label shop for digital intelligent assistance and human-AI collaboration in manufacturing" (Grant agreement No 101092176). The work presented here reflects only the authors' view and the European Commission is not responsible for any use that may be made of the information it contains.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Almansor, E.H., Hussain, F.K.: Survey on Intelligent Chatbots: State-of-the-Art and Future Research Directions. In: Barolli, L., Hussain, F.K., and Ikeda, M. (eds.) Complex, Intelligent, and Software Intensive Systems. pp. 534–543. Springer International Publishing, Cham (2020).
2. Colabianchi, S., Tedeschi, A., Costantino, F.: Human-technology integration with industrial conversational agents: A conceptual architecture and a taxonomy for manufacturing. Journal of Industrial Information Integration. 35, 100510 (2023).

3. McTear, M.: Conversational AI: Dialogue Systems, Conversational Agents, and Chatbots. Springer Nature (2022).
4. Wellsandt, S., Foosherian, M., Bousdekis, A., Lutzer, B., Paraskevopoulos, F., Verginadis, Y., Mentzas, G.: Fostering Human-AI Collaboration with Digital Intelligent Assistance in Manufacturing SMEs. In: Alfnes, E., Romsdal, A., Strandhagen, J.O., von Cieminski, G., and Romero, D. (eds.) APMS. Production Management Systems for Responsible Manufacturing, Service, and Logistics Futures. pp. 649–661. Springer Nature Switzerland, Cham (2023).
5. Motger, Q., Franch, X., Marco, J.: Software-Based Dialogue Systems: Survey, Taxonomy and Challenges. ACM Comput. Surv. 55, 1–42 (2023).
6. Rooein, D., Bianchini, D., Leotta, F., Mecella, M., Paolini, P., Pernici, B.: Chatting About Processes in Digital Factories: A Model-Based Approach. In: Nurcan, S., Reinhartz-Berger, I., Soffer, P., and Zdravkovic, J. (eds.) Enterprise, Business-Process and Information Systems Modeling. pp. 70–84. Springer International Publishing, Cham (2020).
7. Zheng, T., Grosse, E.H., Morana, S., Glock, C.H.: A review of digital assistants in production and logistics: applications, benefits, and challenges. International Journal of Production Research. 0, 1–27 (2024).
8. Grosse, E.H., Sgarbossa, F., Berlin, C., Neumann, W.P.: Human-centric production and logistics system design and management: transitioning from Industry 4.0 to Industry 5.0. International Journal of Production Research. 61, 7749–7759 (2023).
9. What Is a Digital Assistant? | Oracle, https://www.oracle.com/chatbots/what-is-a-digital-assistant/, last accessed 2024/04/18.
10. Zanotto, F.M., Dominguez, D.Z., Ayerbe, E., Boyano, I., Burmeister, C., Duquesnoy, M., Eisentraeger, M., Montaño, J.F., Gallo-Bueno, A., Gold, L., Hall, F., Kaden, N., Muerkens, B., Otaegui, L., Reynier, Y., Stier, S., Thomitzek, M., Turetskyy, A., Vallin, N., Wessel, J., Xu, X., Abbasov, J., Franco, A.A.: Data Specifications for Battery Manufacturing Digitalization: Current Status, Challenges, and Opportunities. Batteries & Supercaps. 5, e202200224 (2022).
11. Lu, Y., Xu, X., Wang, L.: Smart manufacturing process and system automation – A critical review of the standards and envisioned scenarios. Journal of Manufacturing Systems. 56, 312–325 (2020).
12. Romero, D., Stahre, J.: Towards The Resilient Operator 5.0: The Future of Work in Smart Resilient Manufacturing Systems. Procedia CIRP. 104, 1089–1094 (2021).
13. Roth, E., Moencks, M., Bohné, T.: The Augmented Workforce: A Systematic Review of Operator Assistance Systems. In: Handbook of Human-Machine Systems. pp. 267–279. John Wiley & Sons, Ltd (2023).
14. Xia, L., Li, C., Zhang, C., Liu, S., Zheng, P.: Leveraging error-assisted fine-tuning large language models for manufacturing excellence. Robotics and Computer-IntegratedManufacturing.88,102728(2024).
15. Kernan Freire, S., Wang, C., Foosherian, M., Wellsandt, S., Ruiz-Arenas, S., Niforatos, E.: Knowledge sharing in manufacturing using LLM-powered tools: user study and model benchmarking. Front. Artif. Intell. 7, (2024).
16. Abner, B., Rabelo, R.J., Zambiasi, S.P., Romero, D.: Production Management as-a-Service: A Softbot Approach. In: Lalic, B., Majstorovic, V., Marjanovic, U., von

Cieminski, G., and Romero, D. (eds.) APMS Smart and Digital Manufacturing. pp. 19–30. Springer International Publishing, Cham (2020).
17. Bousdekis, A., Wellsandt, S., Bosani, E., Lepenioti, K., Apostolou, D., Hribernik, K., Mentzas, G.: Human-AI Collaboration in Quality Control with Augmented Manufacturing Analytics. In: Dolgui, A., Bernard, A., Lemoine, D., von Cieminski, G., and Romero, D. (eds.) APMS. Artificial Intelligence for Sustainable and Resilient Production Systems. pp. 303–310. Springer International Publishing, Cham (2021).
18. Wellsandt, S., Klein, K., Hribernik, K., Lewandowski, M., Bousdekis, A., Mentzas, G., Thoben, K.-D.: Hybrid-augmented intelligence in predictive maintenance with digital intelligent assistants. Annual Reviews in Control. 53, 382–390 (2022).
19. Lepenioti, K., Fikardos, M., Bousdekis, A., Apostolou, D., Mentzas, G.: Autonomous Analytics as a Service: From Vision to Action. In: Enterprise Interoperability XI: Enterprise Interoperability through Data, AI, and Robotics.
20. Bousdekis, A., Mentzas, G., Apostolou, D., Wellsandt, S.: Evaluation of AI-Based Digital Assistants in Smart Manufacturing. In: Kim, D.Y., von Cieminski, G., and Romero, D. (eds.) APMS. Smart Manufacturing and Logistics Systems: Turning Ideas into Action. pp. 503–510. Springer Nature Switzerland, Cham (2022).
21. Brooke, J.: SUS-A quick and dirty usability scale. Usability evaluation in industry, 189(194), 4-7 (1996).
22. Colligan, L., Potts, H. W., Finn, C. T., Sinkin, R. A.: Cognitive workload changes for nurses transitioning from a legacy system with paper documentation to a commercial electronic health record. International journal of medical informatics, 84(7), 469-476 (2015).