Virtual Assistants for Energy Efficiency: Real World Tryouts

Francisco Santos Fialho
Instituto Superior Técnico
Lisbon, Portugal
francisco.santos.fialho@tecnico.ulisboa.pt

Sabrina Scuri
Politecnico di Milano
Interactive Technologies Institute, LARSyS
Milan, Italy
sabrina.scuri@polimi.it

Lucas Pereira
Instituto Superior Técnico
Interactive Technologies Institute, LARSyS
Lisbon, Portugal
lucas.pereira@iti.larsys.pt

Nuno Jardim Nunes
Instituto Superior Técnico
Interactive Technologies Institute, LARSyS
Lisbon, Portugal
nunojnunes@tecnico.ulisboa.pt

ABSTRACT

The negative effects of climate change are calling for new approaches to promote energy efficiency and the use of renewable energy sources at multiple scale levels. As virtual assistants are becoming a common household item, recent studies have looked at integrating IoT and virtual assistants for energy management purposes. Despite the prominence of these works, a critical gap in the current body of research is the almost absence of real-world implementations covering different sectors of society. To address this gap, we developed the PowerShare Virtual Assistant (VA), a voice-based eco-feedback system. The paper presents results from the real-world deployment of the PowerShare VA in three distinct sectors - 1) residential, 2) commerce, and 3) industry. By looking at the human response to our system in different daily life scenarios, we aim to contribute to future research on using VA in the context of energy efficiency.

CCS CONCEPTS

• **Human-centered computing** → *HCI design and evaluation methods*; *Ubiquitous and mobile devices*; **Personal digital assistants**.

KEYWORDS

virtual assistant, energy efficiency, real-world, artificial intelligence

ACM Reference Format:
Francisco Santos Fialho, Sabrina Scuri, Lucas Pereira, and Nuno Jardim Nunes. 2023. Virtual Assistants for Energy Efficiency: Real World Tryouts. In *ACM International Conference on Information Technology for Social Good (GoodIT '23), September 06–08, 2023, Lisbon, Portugal.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3582515.3609548

GoodIT '23, September 06–08, 2023, Lisbon, Portugal

ACM ISBN 979-8-4007-0116-0/23/09…$15.00
https://doi.org/10.1145/3582515.3609548

1 INTRODUCTION

The climate crisis is one of the greatest threats facing humanity. As the mitigation of climate change has become a global priority, countries all over the world (including Non-Member States [7]) are implementing regulatory mechanisms to help meet the energy-saving targets proposed by the EU Energy Efficiency Directive [35]. Reaching these targets requires making significant changes in the energy market and implementing energy efficiency actions at different scale levels, from buildings (e.g., energy certifications [7]) to urban scale (e.g., smart cities [30]). In this scenario, recent advances in energy monitoring and control technologies, combined with the opportunities offered by the use of Information and Communications Technologies (ICTs), are paving the way for the development of new tools and solutions to increase energy efficiency, reduce consumption, and even promote the use of renewable energy sources (RES).

With the evolution of technology in the virtual assistant (VA) field, many home assistant solutions are being developed to bring people closer to this new technology, creating products that can be integrated into their homes. Home assistants have many functionalities that allow users to automate their homes using smart devices, ultimately helping the user in their daily life. The most famous home assistants on the market are from the top three tech companies, Amazon, Google, and Apple. Each has its own virtual assistant, Alexa[1], Google Assistant[2], and Siri[3].

The home assistant market is continuously growing, with an estimated 123.5 million adult users in the US alone by the end of 2022" [20]. This represents a great opportunity for systems to grow and extend their functionalities to home assistants. Consequently, several projects have emerged in multiple areas, particularly in energy efficiency [17, 18, 36]. Still, despite the prominence of studies on virtual assistants for energy management, a critical gap in the current body of research is the absence of real-world implementations covering different sectors of society since most of the surveyed works have been developed in residential settings under controlled or simulated environments [18, 23].

Instead, this paper reports on the development and deployment of the Power Share Virtual Assistant (VA) in the real world in three

[1]Amazon Alexa, https://alexa.amazon.com/
[2]Google Assistant, https://assistant.google.com/
[3]Apple Siri, https://www.apple.com/siri/

distinct sectors: 1) residential, 2) commercial, and 3) industry. The Power Share VA combines a smart-meter solution and a Google Assistant application to inform the user about their energy consumption in their home or facility, bringing users closer to virtual assistants while making their interaction seamless and convenient on a daily basis. Being an eco-feedback system [10], the Power Share VA provides users with energy-related data to encourage more efficient behaviors. In particular, it allows users to:

- Explore their energy consumption and production (available to owners of energy generation assets only) data;
- Access simulation data of potential solar PV production for their location;
- Receive on-demand personalized information via email.

To the best of our knowledge, this is the first work that compares the human response to virtual assistants in different daily life scenarios in the context of energy efficiency. By having these three distinguished types of environments, it is possible to understand the differences in how the users interact with the system and how it can be improved to fit best in each scenario.

The remainder of this paper is organized as follows: Section 2 provides an overview of existing research in IoT, VA, and voice-based assistants for energy efficiency. Section 3 describes the proposed solution and the methodology adopted for its evaluation. Results from the evaluation (quantitative and qualitative) are presented in Section 4. Finally, Section 5 presents conclusions and future work.

2 STATE OF THE ART

2.1 Internet of Things and Virtual Assistants

With the rise of smart devices and home automation solutions, home assistants like Amazon Alexa, Google Assistant, and Apple Siri - just to name a few - are becoming common household items. Home assistants offer many functionalities and can connect to smart devices, thus allowing users to automate their homes. The global home assistant market is continuously growing - according to [11], it is expected to reach approximately USD 7.8 billion by 2023 - and existing systems are quickly being upgraded to support new functionalities.

In this perspective, recent studies have looked at integrating the Internet of Things (IoT) and Virtual Assistants (VA) specifically for energy management purposes. Barman and colleagues [2] developed an IoT-based smart energy meter that monitors energy consumption and helps detect power theft. Unlike existing metering systems, their solution leverages the benefits brought by the IoT to allow for full duplex communication. It ultimately makes it easier for the end users to analyze and control their energy data. The solution is composed of five elements: (i) a Wifi module to manage the system operations, (ii) an OLED display that shows energy-related information, (iii) the energy meter, (iv) an optocoupler, and (v) a current sensor. Results from their study show that the system could identify energy consumption patterns and detect power theft and promote energy conservation.

Vishwakarma et al. [36] developed an IoT-based smart energy home automation system remotely controlling home equipment. The system was developed using Arduino NodeMcu, Adafruit, and If-This-Then-That (IFTTT) [19], and included a web-based assistant and Google Assistant. The authors highlighted the evolution and importance of the new technologies in the IoT field, which now allows older people and people with some disabilities to interact with systems with the help of virtual assistants. With a similar perspective, Isyanto et al. [21] implemented a system for people with disabilities using Google's virtual assistant, where they can control home appliances with their voice (e.g., fans, lights, and even the TV.) The authors pointed out the ability to use the VA on different devices (e.g., Google Nest and Smartphones) and Operating Systems (e.g., Android or IOS) as a main advantage since it enables it to cover a wider range of end-users.

A solution similar to the one described in the present article is the one proposed by Hadi et al. paper [16]. By leveraging the resource offered by Google, the Google Assistant and its development platform, the authors were able to create a voice-based system to control and monitor electronic appliances. As highlighted by the authors [16], Google Assistant presents several advantages, among which two are the most worth mentioning: (i) there are no memory requirements to run these applications since all the processing is made on the Google Assistant server's side, and (ii) voice-based interaction provides a more seamless user experience compared to the one offered by existing applications for energy management. The system was tested in a domestic setting. The test outcomes demonstrated the potential of voice-based energy monitoring and controlling system, as "the accomplishment pace of talk affirmation structures is 75% of success rate" [16].

2.2 Voice-based Assistants for Promoting Energy Efficiency and Sustainability

The HCI body of literature on conversational AI and VA for energy efficiency is growing rapidly.

In 2018, Gnewuch et al. [15] investigated using Conversational Agents (CAs) for energy feedback. To answer the question of how to design CAs for promoting sustainable energy use in households, the authors developed an interactive prototype of the Energy Feedback Agent - a text-based CA - which was then evaluated with industry experts in an exploratory focus group session. The responses offered by participants to the user scenarios they were presented with suggest that CA is a promising technology for energy feedback and ultimately allowed the authors to put forward a set of principles for designing such solutions.

With the development of CA technology, the research focus has quickly moved from text-based to voice-based CAs. In recent work, He and colleagues [18] investigated how proactive smart home assistants (SHAs) could be designed to nudge smart home occupants toward changing their energy-related behaviors and ultimately take energy conservation actions. Through an online experiment, the authors collected direct responses from 307 participants to energy-saving suggestions provided by an SHA in a simulated smart home scenario, demonstrating the potential of proactive SHAs in promoting energy-saving behaviors. The effectiveness of voice-based VA in stimulating energy efficiency in households is further supported by findings from a preliminary study conducted by He [17] to compare the nudging power of bi-directional communication enabled by VA with traditional text-based eco-feedback systems.

More recently, in an attempt to provide a set of guidelines for designing conversational agents to promote sustainability, Giudici et al. [14] have conducted a focus group session with experts to identify the potential areas of intervention for combining conversational technologies and digital devices to foster a more sustainable behavior in domestic spaces.

The use of VA for sustainability has been investigated in areas other than energy efficiency. For example, Esau et al. [8] investigated the use of VAs to prevent food waste and motivate conscious food handling. Likewise, in [34], the authors investigated how virtual assistants can be deployed in hotel rooms to foster sustainable tourism behavior.

2.3 Summary

Although literature seems to provide evidence of the potential of using voice-based CAs like VAs for energy management purposes, a few aspects deserving further investigation emerge from our review.

In particular, we couldn't find any study describing real-world implementations of voice-based energy feedback systems. Despite authors agreeing on the need for real-world deployments of these solutions [15, 18], all of the works reviewed test their prototypes through simulations [8, 14, 15, 18]. Moreover, we observe that when it comes to VAs for energy management, researchers have so far focused on the residential context only [8, 14, 15, 18, 18]. HCI researchers have long called for the need to extend research on energy feedback to the non-residential sector [4, 23]. Thus, despite the prominence of these works, we see this as a critical gap that deserves to be addressed.

3 MATERIALS AND METHODS

As stated above, the Power Share VA consists of two independent components: 1) the smart-metering infrastructure and 2) the Power Share application. These two components are described next.

3.1 Smart Metering Infrastructure

The Power Share VA was developed to be hardware agnostic when it comes to the acquisition of smart-meter data. This was achieved by separating the virtual assistant features from the energy monitoring process.

A key aspect of any energy eco-feedback system is the ability to frequently update consumption and production data, a task that, as of today, should be facilitated [1] due to the proliferation of smart meters. However, access to smart meter data is still not straightforward since many existing solutions do not offer a standard way of accessing their data through open APIs, mainly due to privacy concerns [22, 25].

A quick survey of the smart-meters market made it possible to find a few solutions with functionalities relevant to this work. For example, the Sense [24] system provides an API for data access, and it is fully integrated with Virtual Assistants such as Google Assistant and Alexa from Amazon. However, at the time of this research, the system was not compatible with the European Power Grid (Nominal frequency of 50Hz and Voltage of 230V). Another example is the Smappee meter [32], which, even though it does not offer a default integration with Virtual Assistants, is compatible with the European power grid and provides an official API for data access. Unfortunately, the access to this API is limited to one access per hour, which is not adequate for the goals of this work.

Hence, it was decided to rely on a smart-metering platform developed and tested in the scope of the Horizon 2020 SMart IsLand Energy systems (SMILE) research project [29]. This platform can monitor electric power generation from solar PV and electric energy demand using Carlo Gavazzi smart meters [13] for that effect. The data is measured locally at 1Hz and uploaded to a cloud-based server aggregated in 1-minute intervals. This is done using an Industrial Raspberry Pi (the strato-pi). An Open API is also provided to enable remote access to the stored data, with an unlimited number of API calls. Figure 1 provides an overview of the installed smart metering infrastructure. There it is possible to see that two smart meters were installed, one for household demand (1) and another to monitor the factory (2). The two smart meters are then connected to the strato-pi (3), which is responsible for reading and processing the data from the smart meters before uploading it to the online server.

3.2 Power Share Application

The Power Share application is designed to provide custom interfaces regarding energy consumption. These applications are composed of an additional entity, the fulfillment service, in other words, a web application. Power Share is then split into two applications, a Google Actions application and a web application.

3.2.1 Google Actions Application. The Google Actions app comprises two other components, a Firebase project and a Google Cloud project, linked to one another. In the Google Actions platform, we have created a project for the application and started configuring it. From here, we followed the architecture bottom-up, starting from the types moving to intents and finally to scenes.

The user interfaces are categorized into summarized and detailed dashboards, and for each type of dashboard, specific operations were created with training phrases composed of the dashboard and facility type. For example, "Show me the house daily dashboard," and "What is the restaurant complete dashboard." This composition allows the user to vary the commands as it is only required to match these two key values.

With the desired operations implemented, we created a scene for each of them. In each scene, we defined the conditions to show a specific interface. For instance, the conditions in the previous example phrases are that the facility type and the dashboard type must be present in the command given by the user. The desired interface is shown to the user if these conditions are met. Otherwise, it calls the *Match Any* scene, which is visually represented by the help page since the user asked for something the assistant does not understand. Since a big challenge of Home Assistants is that voice commands may not be well interpreted in the presence of background noises [36] this feature is particularly important since it enables the assessment of the audio recognition reliability in the Power Share VA.

3.2.2 Web Application. The web applications consist of a home screen and five dashboards. These interfaces were developed according to the Google guidelines for voice-enable applications, including the need to "design with voice-first in mind" and to "focus on one touchpoint at a time." [6]

Figure 1: Smart metering infrastructure installed in the main breaker box of one of the participants in the study. Smart matter for household demand (1), smart meter for the factory demand (2), and the strato-pi for data processing and communication (3).

- The `Home` screen introduces the application to the user with a description. It has several buttons suggesting commands to perform, and through the Assistant's voice, other commands are recommended (See Figure 2);
- The `Daily Dashboard` shows information about today's or yesterday's energy consumption and production. It also presents a results card describing the peak and low values and feedback for production and consumption. The feedback for this dashboard is made by comparing today or yesterday with the same weekday of the previous week (See Figure 3);
- The `Weekly Dashboard` shows this week's or last week's energy consumption and production. Similarly to the previous interface, here is presented a results card with the same type of information;
- The `Monthly Dashboard` provides an energy use overview throughout the current and the previous month. It also has a results card;
- The `Complete Dashboard` provides a complete summary of the daily, weekly, and monthly dashboards;
- The `Latest Dashboard` shows an overview of yesterday's and last week's dashboard.

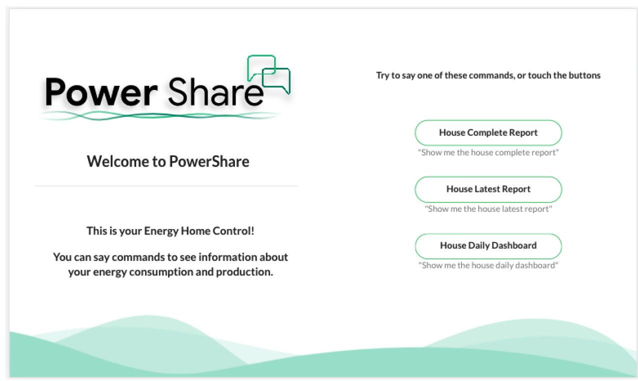

Figure 2: Overview of the Web Application: Home Screen.

Since many potential users of the Power Share VA may not own solar PV micro-production on their premises, an option was added to create a simulation of the potential for local energy production. To this end, real-time and historical solar PV estimates are obtained from Solcast [33]. If this option is not enabled, only consumption is shown in the different user interfaces.

Finally, with the user's consent, an email with information regarding the current dashboard is sent. This email contains the results of the energy usage for the specific dashboard, links to articles on optimizing energy consumption (including links to micro-production solutions from National suppliers and a link to the consumer association website), and a screenshot of the interface for future reference. Figure 4 shows an example of the email.

3.3 Evaluation Methodology

The Power Share VA system was iteratively developed and evaluated over two deployment phases. The first deployment phase, preliminary and propaedeutic to the second one, was conducted to gather feedback on the prototype design. To this aim, an initial version of the system was deployed for a period of a month at an industrial premise consisting of a factory and the owner's house. This family-owned Small Medium Enterprise (SME) produces handbags, among other fashion goods. The factory has a high energy demand, partly due to more than eighteen industrial machines and three fans.

At the beginning of the study, the family was presented with a consent form and a general explanation of the system. For the purpose of the study, a Google Home Hub was provided to participants.

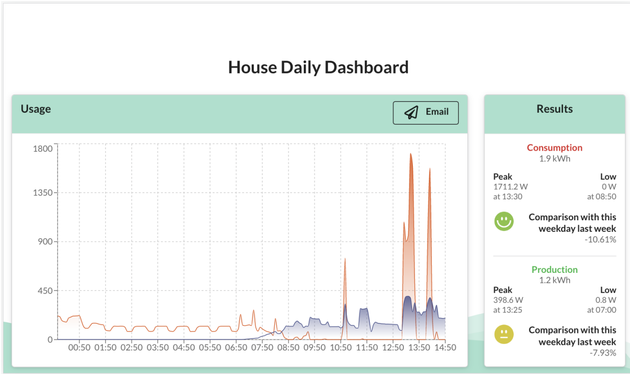

Figure 3: Overview of the Web Application: Daily Dashboard (including energy production).

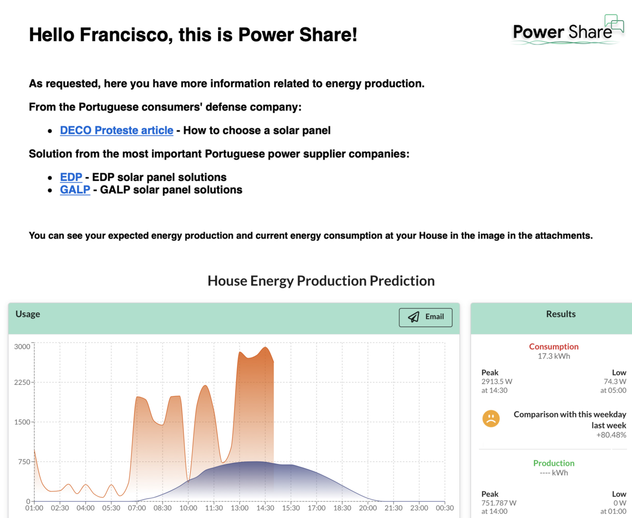

Figure 4: Email sent about the prediction of energy production.

At the end of the one-month deployment, a semi-structured interview was conducted with the system users to gather their feedback on the interface design. Revisions to some UI and UX elements were made accordingly. The refined prototype was then tested during the second deployment phase to:

- `Evaluate the usability of the system`: is intended to understand how users interact with the system;
- `Assess the information of the system`: extract from the users' assessment what is the most useful information to be presented in the system;
- `Evaluate the users' acceptance of virtual assistants`: understand the users' perspective on virtual assistants and their applications.

The second deployment was conducted in both residential (an apartment) and commercial (a restaurant) settings. Participants were recruited through convenience sampling [12] and asked to use the system for a period of two weeks. The apartment residents were a middle age couple owning a solar PV micro-production unit of 0.5kWp. As for the restaurant, it is situated on the ground floor of an urban building and does not have a solar PV installation. The restaurant offers lunch and dinner services from Tuesday to Saturday.

Similarly to the first deployment phase, the family and the restaurant owner received a Google Home Hub, instructions on how to use the system, and a consent form. Interactions with the application - user ID; timestamp of the interactions; features accessed; and number of unsuccessful interactions (i.e., calls to the Match Any interface / Help page) - have been electronically monitored throughout the study using Dashbot.io [5]. At the end of the study, semi-structured interviews were conducted with participants to explore their understanding, opinions, and usage patterns. Interviews started with a warm-up discussion about their previous experiences with energy feedback technologies and VAs. Other questions targeted how participants used the application (e.g., when and how often, most and less used features, usability issues faced, etc.) and their understanding and perceived usefulness of the information provided. In addition, questions related to overall willingness to use similar systems in the future and suggestions for improvements were included. Interviews lasted an average of 30 minutes and were fully recorded and transcribed.

4 RESULTS

4.1 Quantitative Data

As shown in Table 1, users preferred recent data - i.e., Daily Dashboard and Latest Dashboard - over historical data. Interestingly, only User Type (UT) 2 shows almost equal interest in both typologies of data.

Also, the time spent interacting with the system varied between UTs, with UT1 being the most active user and UT2 the least. This result, however, was not unexpected. Unlike UT1 and UT3, the restaurant has no residential purposes. Therefore, besides being occupied only part of the day, users do not have much time to interact with the system during work hours. The inherent characteristics of the three user scenarios investigated also influenced the "preferred time of use," as UT 1 and 2 tended to interact with the system during work hours.

The interaction data collected shows that some of the functionalities offered by the system were not used at all (see Figure 5 for additional details). In this regard, it's worth noting that most of the 'not-accessed' interfaces provide data that require time and effort to understand fully. UT2, for example, did not access the Predict Energy Production dashboard. This dashboard, compared to others, provides more complex information. It requires more effort from the user to analyze the data and infer whether installing an energy generation asset is worth it. On the contrary, the most called interfaces are those that provide a summary of energy-related data displayed in a visual form. Figure 5 shows the percentages of the features accessed by each user typology and a list of those not accessed.

Finally, based on the literature surveyed, we can reasonably argue that the overall system performance was good. Indeed, for all

Table 1: Summary of quantitative data collected for each User Type (UT)

	Factory with Domestic (UT 1)	Restaurant (UT 2)	Domestic (UT 3)
Most used feature	Daily Dashboard	Complete Dashboard	Daily Dashboard
Second most used feature	Send Email	Daily Dashboard	Latest Dashboard
Third most used feature	Predict Energy Production Dashboard	Weekly Dashboard	Send Email
Preferred time of use	12:00	14:00	21:00
Success rate	**93.19%**	**88.57%**	**76%**

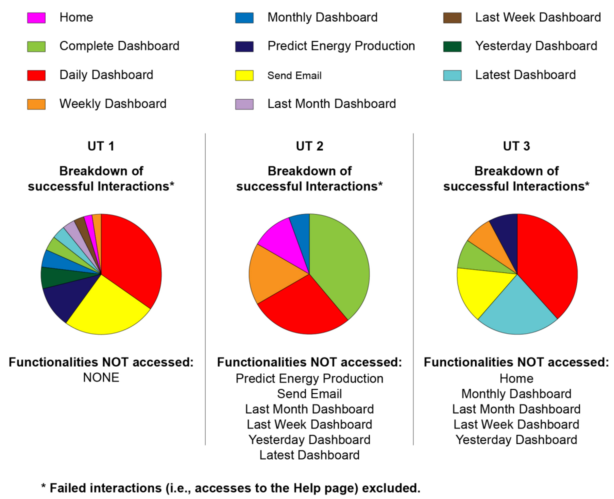

Figure 5: Breakdown of the features accessed by each user typology.

UT, the success rate - measured by the calls to *Match Any* interface - was above the 75% reported in [16].

4.2 Qualitative Data

All interviewees reported appreciating the energy feedback data provided by the system as well as the possibility of receiving the same information also via email: "*I did ask to receive data via email. It's a useful feature, and the content is easy to read*" (UT1). In particular, the opportunity to keep track of their energy-related habits and adapt them based on the feedback received was described as a useful feature: "*Thanks to the system I made a few changes. For instance, I've started to turn on the heat pump when production from solar PVs was higher*" (UT3). UT1 emphasized that factory owners could considerably benefit from using this system and ultimately save some money simply by adjusting their consumption habits: "*I think that using this system in a factory is very cost-effective!* [...] *I've started disconnecting some machines over the weekend and could see huge differences in the consumption*" (UT1).

Participants made several changes, not only to their habits - e.g., UT1 replaced all light bulbs with LED bulbs - and reported being willing to maintain them. In this regard, the *Predict Energy Production* dashboard emerged as a potential trigger for taking further actions. At the end of the study, participants that do not own any energy generation assets (UT1) said they were considering installing solar PV panels for self-consumption.

As the literature on eco-feedback suggests, one of the main issues with these systems is that users normally lose interest over time [27, 28]. Results from our study suggest that voice-based eco-feedback systems might partially help overcome this issue. Interviewees reported being willing to use the application in the future (UT1), although in a more sporadic fashion (UT3). However, this is still a hypothesis that requires further research.

Finally, some recommendations for improving the system emerged from the interviews. In particular, UT3 suggested the possibility of receiving tailored notifications based on the user consumption and production data: "*It would be nice to receive a message any time production exceeds consumption. This way, I can take advantage of excess production to charge my phone or use other appliances, such as the dehumidifier*" (UT3).

5 CONCLUSIONS AND FUTURE WORK

Power Share is an integrated system that comprehends a smart meter solution, a conversational Google Assistant application, a web application, and a server. The solution was envisioned to raise people's awareness about energy, informing them about its use by using a virtual assistant and seamlessly integrating it into their daily lives. The Power Share VA provides attractive interfaces regarding facilities' energy consumption and production while offering an engaging virtual assistant experience. These interfaces include a variety of dashboards in which users can see historical and real-time data from different perspectives and even compare them to understand their energy use habits.

The system was successfully evaluated in three distinguished environments with different types of energy use, namely, houses, houses with a factory, and a restaurant, demonstrating its versatility in different scenarios.

The system was evaluated over two deployment phases, in which both quantitative and qualitative data were collected to understand the user experience. An industrial facility was involved in the first deployment phase, which lasted a month. A house in Madeira and a restaurant in Lisbon were part of the second deployment phase, which lasted two weeks.

Based on the user interaction data collected through Dashbot.io [5], it emerged that the interactions were short-lived. The system has a fast response rate, and the dashboard displays information efficiently, allowing it to be integrated into work environments such as restaurants and factories, where time is scarce. Participants clearly preferred the summarized dashboard in commercial and industrial settings as the data is more consolidated and embraces more than one scope.

Moreover, results from our test show that the real-time comparison between days, weeks, and months was one of the most praised functionalities, as it allows users to implement new measures to reduce their energy consumption and improve production. Users preferred real-time information over historical data (previous days, weeks, and months). Finally, the tests characterized the system as a system of long-term use, as people use it over time and not exhaustively over a short period. Testing the system resulted in monitoring energy use habits and creating effective ways to mitigate unnecessary energy consumption. It even boosted users' curiosity to start producing energy and finding the best solar panel solutions. Finally, the adherence to testing the system, regarding the number of interactions, aligns with the previous studies and the literature.

While the objectives of this research were achieved, at this stage, it is important to stress that the current study suffers from some limitations that should be addressed in future iterations of this work. First and foremost, it would be important to carry out a longer deployment to assess how the usage of the systems varies over time. Furthermore, while the sample spans three different sectors, it is limited to one participant per sector, which hurts generalization efforts. Furthermore, since only one Google Home device was available, running the study in parallel was impossible. This is something that should be avoided in a long-term study to remove biases due to external factors (e.g., extreme weather conditions). Moreover, another limitation to be considered is that we studied only one of the VAs currently available on the market. Future iterations of this work would benefit by including systems from other manufacturers.

Luckily, the Power Share VA was built with new technologies and procedures, namely Google Actions, which offer opportunities for improvements in its various parts. For example, in future work, it would be interesting to compare the effectiveness of a VA deployed using a stationary device such as Google Home with one that is more ubiquitous by running on a smartphone. Besides providing a much-needed comparison between the two types of VAs, the ability to use smartphones would enable the development of larger and longer studies at a much lower cost.

Likewise, the discussion with our end-users made it possible to understand that the interfaces can be improved to best fit the type of information users want to see. In other words, it should be possible to build functionalities into the VAs that are targeted at specific user segments. Given the small sample size of the evaluation, this is another example where a future study with a much larger sample and deployment duration would be very important.

In addition to the existing functionalities of the VA, the smart meter solution could be enriched by exploring the integration of Non-Intrusive Load Monitoring (NILM) [9, 26] to enable the delivery of specific information regarding individual appliances. With live data about the appliance, enhancing the VAs with warnings to adjust the demand based on external signals should be possible, allowing users to transition from passive consumers to a more active role in the smart grid [31]. Although possible from a technological standpoint, for such an approach to truly work, it will require a high degree of engagement from the end-users. As the literature suggests, one of the main challenges with the adoption of eco-feedback technologies (whether they provide aggregated or disaggregated energy data) is to sustain users' interest over time [3, 28]. As Virtual Assistants enable more natural and seamless interaction experiences, they may have great potential for promoting lasting engagement with these systems and ultimately overcoming some of the existing adoption barriers. Thus, we believe that identifying guidelines to inform the design of VAs for energy efficiency and behavior change represents another crucial direction for future research.

ACKNOWLEDGMENTS

This work was partially funded by the Portuguese Foundation for Science and Technology (FCT) under the project EXPL/CCI-COM/1234/2021. Lucas Pereira has received funding from the FCT under grant CEECIND/01179/2017. Sabrina Scuri, Lucas Pereira, and Nuno Jardim Nunes have received funding from FCT under project UIDB/50009/2020. This work reflects only the author's views.

REFERENCES

[1] Zita Abreu and Lucas Pereira. 2022. Privacy Protection in Smart Meters Using Homomorphic Encryption: An Overview. *WIREs Data Mining and Knowledge Discovery* 12, 4 (2022), e1469. https://doi.org/10.1002/widm.1469
[2] Bibek Kanti Barman, Shiv Nath Yadav, Shivam Kumar, and Sadhan Gope. 2018. IOT Based Smart Energy Meter for Efficient Energy Utilization in Smart Grid. In *2018 2nd International Conference on Power, Energy and Environment: Towards Smart Technology (ICEPE)*. IEEE, Shillong, India, 1–5. https://doi.org/10.1109/EPETSG.2018.8658501
[3] Mary Barreto, Evangelos Karapanos, and Nuno Nunes. 2013. Why don't families get along with eco-feedback technologies? A longitudinal inquiry. In *Proceedings of the Biannual Conference of the Italian Chapter of SIGCHI*. Association for Computing Machinery, Trento, Italy, 1–4.
[4] Nico Castelli, Sebastian Taugerbeck, Martin Stein, Timo Jakobi, Gunnar Stevens, and Volker Wulf. 2020. Eco-InfoVis at Work: Role-based Eco-Visualizations for the Industrial Context. *Proceedings of the ACM on Human-Computer Interaction* 4, GROUP (2020), 1–27.
[5] Dashbot.io. 2022. . Dashbot, Inc. https://www.dashbot.io/
[6] Google Developers. 2022. Design Guidelines. https://developers.google.com/assistant/interactivecanvas/design?hl=en
[7] Mak Dukan. 2015. Climate policy info hub. https://climatepolicyinfohub.eu/energy-efficiency-policy-instruments-european-union
[8] Margarita Esau, Dennis Lawo, Thomas Neifer, Gunnar Stevens, and Alexander Boden. 2023. Trust your guts: fostering embodied knowledge and sustainable practices through voice interaction. *Personal and Ubiquitous Computing* 27, 2 (2023), 415–434.
[9] Anthony Faustine, Lucas Pereira, and Christoph Klemenjak. 2020. Adaptive Weighted Recurrence Graphs for Appliance Recognition in Non-Intrusive Load Monitoring. *IEEE Transactions on Smart Grid* 12, 20321506 (2020), 1–1. https://doi.org/10.1109/TSG.2020.3010621
[10] Jon Froehlich, Leah Findlater, and James Landay. 2010. The Design of Eco-Feedback Technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1999–2008. https://doi.org/10.1145/1753326.1753629
[11] Market Research Future. 2021. Voice assistant market size share: Growth prediction - 2030. https://www.marketresearchfuture.com/reports/voice-assistant-market-4003
[12] Alison Galloway. 2005. Non-Probability Sampling. In *Encyclopedia of Social Measurement*, Kimberly Kempf-Leonard (Ed.). Elsevier, New York, 859–864. https://doi.org/10.1016/B0-12-369398-5/00382-0
[13] Carlo Gavazzi. 2022. . Carlo Gavazzi. https://gavazziautomation.com/nsc/HQ/EN/energy_power_analyzers
[14] Mathyas Giudici, Pietro Crovari, and Franca Garzotto. 2022. CANDY: a framework to design Conversational AgeNts for Domestic sustainabilitY. In *Proceedings of the 4th Conference on Conversational User Interfaces*. Association for Computing Machinery, Association for Computing Machinery, Glasgow, UK, 1–8.
[15] Ulrich Gnewuch, Stefan Morana, Carl Heckmann, and Alexander Maedche. 2018. Designing conversational agents for energy feedback. In *Designing for a Digital and Globalized World: 13th International Conference, DESRIST 2018, Chennai, India, June 3–6, 2018, Proceedings 13*. Springer, Springer, Chennai, India, 18–33.
[16] Mokh. Sholihul Hadi, Maulana Ahmad As Shidiqi, Ilham Ari Elbaith Zaeni, Muhammad Alfian Mizar, and Mhd Irvan. 2019. Voice-Based Monitoring and Control System of Electronic Appliance Using Dialog Flow API Via Google

Assistant. In *2019 International Conference on Electrical, Electronics and Information Engineering (ICEEIE)*, Vol. 6. IEEE, Denpasar, Indonesia, 106–110. https://doi.org/10.1109/ICEEIE47180.2019.8981415

[17] Tianzhi He. 2021. *Human-Building Symbiotic Communication with Voice-based Proactive Smart Home Assistants*. Ph. D. Dissertation. Virginia Tech.

[18] Tianzhi He, Farrokh Jazizadeh, and Laura Arpan. 2022. AI-powered virtual assistants nudging occupants for energy saving: proactive smart speakers for HVAC control. *Building Research & Information* 50, 4 (2022), 394–409.

[19] IFTTT. 2022. . IFTTT. https://ifttt.com/explore/new_to_ifttt

[20] Insider Intelligence. 2022. *Voice Assistants in 2022: Usage, growth, and future of the AI voice assistant market.* Insider Intelligence. https://www.insiderintelligence.com/insights/voice-assistants/

[21] Haris Isyanto, Ajib Setyo Arifin, and Muhammad Suryanegara. 2020. Design and Implementation of IoT-Based Smart Home Voice Commands for disabled people using Google Assistant. In *2020 International Conference on Smart Technology and Applications (ICoSTA)*. IEEE, Surabaya, Indonesia, 1–6. https://doi.org/10.1109/ICoSTA48221.2020.1570613925

[22] Jonathan Spencer Jones. 2022. Consumer Privacy Concerns Limit Smart Meter Data Access in GB – Report. https://www.smart-energy.com/industry-sectors/smart-meters/consumer-privacy-concerns-limit-smart-meter-data-access-in-gb-report/

[23] Rebecca Afua Klege, Martine Visser, Saugato Datta, and Matthew Darling. 2022. The power of nudging: Using feedback, competition, and responsibility assignment to save electricity in a non-residential setting. *Environmental and Resource Economics* 81 (2022), 1–17.

[24] Sense Labs. 2022. *Home*. Sense. https://sense.com

[25] Eoghan McKenna, Ian Richardson, and Murray Thomson. 2012. Smart Meter Data: Balancing Consumer Privacy Concerns with Legitimate Applications. *Energy Policy* 41 (Feb. 2012), 807–814. https://doi.org/10.1016/j.enpol.2011.11.049

[26] Christoforos Nalmpantis and Dimitris Vrakas. 2018. Machine Learning Approaches for Non-Intrusive Load Monitoring: From Qualitative to Quantitative Comparation. *Artificial Intelligence Review* 52 (Jan. 2018), 1–27. https://doi.org/10.1007/s10462-018-9613-7

[27] Lucas Pereira and Nuno Nunes. 2019. Understanding the Practical Issues of Deploying Energy Monitoring and Eco-Feedback Technology in the Wild: Lesson Learned from Three Long-Term Deployments. *Energy Reports* 6 (Dec. 2019), 94–106. https://doi.org/10.1016/j.egyr.2019.11.025

[28] Lucas Pereira, Filipe Quintal, Mary Barreto, and Nuno J. Nunes. 2013. Understanding the Limitations of Eco-feedback: A One-Year Long-Term Study. In *Human-Computer Interaction and Knowledge Discovery in Complex, Unstructured, Big Data (Lecture Notes in Computer Science)*, Andreas Holzinger and Gabriella Pasi (Eds.). Springer Berlin Heidelberg, Maribor, Slovenia, 237–255.

[29] Filipe Quintal, Daniel Garigali, Dino Vasconcelos, Jonathan Cavaleiro, Wilson Santos, and Lucas Pereira. 2021. Energy Monitoring in the Wild: Platform Development and Lessons Learned from a Real-World Demonstrator. *Energies* 14, 18 (Jan. 2021), 5786. https://doi.org/10.3390/en14185786

[30] Michele Roccotelli and Agostino Marcello Mangini. 2022. Advances on Smart Cities and Smart Buildings. *Applied Sciences* 12, 2 (2022). https://doi.org/10.3390/app12020631

[31] Konrad Schmitt, Rabindra Bhatta, Manohar Chamana, Mahtab Murshed, Ilham Osman, Stephen Bayne, and Luciane Canha. 2023. A Review on Active Customers Participation in Smart Grids. *Journal of Modern Power Systems and Clean Energy* 11, 1 (Jan. 2023), 3–16. https://doi.org/10.35833/MPCE.2022.000371

[32] Smappee. 2022. *Smappee - Fueling energy efficiency for people and businesses*. Smappee. https://www.smappee.com

[33] Solcast. 2019. . Solcast. https://solcast.com/

[34] Iis Tussyadiah and Graham Miller. 2019. Nudged by a Robot: Responses to Agency and Feedback. *Annals of Tourism Research* 78 (Sept. 2019), 102752. https://doi.org/10.1016/j.annals.2019.102752

[35] European Union. 2012. DIRECTIVE 2012/27/EU OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 25 October 2012 on energy efficiency, amending Directives 2009/125/EC and 2010/30/EU and repealing Directives 2004/8/EC and 2006/32/EC. https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2012:315:0001:0056:en:PDF

[36] Satyendra K. Vishwakarma, Prashant Upadhyaya, Babita Kumari, and Arun Kumar Mishra. 2019. Smart Energy Efficient Home Automation System Using IoT. In *2019 4th International Conference on Internet of Things: Smart Innovation and Usages (IoT-SIU)*. IEEE, Ghaziabad, India, 1–4. https://doi.org/10.1109/IoT-SIU.2019.8777607