Discussion Paper

What can we ask respondents in terms of data donation? The case of energy usage

Jeldrik Bakker
Maaike Kompier
Vera Toepoel

January 2026

Abstract

In official statistics, there is a growing demand for more frequent and granular household energy data, including end-uses and a distinction between household and transport-related electricity. In this mixed-mode feasibility pilot we assess one possible implementation of energy data donation, using consumer-grade smart-meter dongles combined with respondent diaries. Participants installed a HomeWizard P1 dongle and kept an eight-day diary on occupancy and appliance use. We processed 15-minute exported meter readings (electricity import/export, gas; per-phase power where available), derived quarter-hourly consumption, and linked diaries for interpretation. Installation was largely straightforward; respondents valued real-time feedback, but diaries were burdensome. Quarter-hourly data reveal daily patterns and, together with diaries, allow recognition of some high-load appliances. Disaggregation is limited when activities overlap, and hybrid heating requires contextual cues. In households with solar panels, combining net import with power readings helps to reveal masked background loads. Energy data donation via consumer-grade dongles is feasible and acceptable but, at 15-minute resolution, insufficient for automated device-level statistics. We recommend (i) a privacy-by-design approach controlled by the National Statistical Institute, (ii) a rotating panel with onboarding and a short stabilisation period, and (iii) higher-frequency sensing and human-in-the-loop validation to reduce burden and enable scalable production statistics.

Keywords: smart surveys, energy, smart meters, energy data donation, official statistics, P1 meter

Contents

1 Introduction

Both national and international stakeholders have a strong and growing interest in improving the frequency, granularity, and scope of household energy data in view of the energy transition. Energy statistics are essential to monitor progress of energy saving measures and policies. These statistics are part of the European Statistical System (ESS), with some outputs mandatory and others highly relevant to policy and research. In terms of granularity, ESS requirements demand disaggregation of household energy use into final uses such as heating, cooling, cooking, lighting, and water heating. Additionally, it is mandatory to distinguish energy used for transport (e.g. electric vehicles, battery systems) from general household consumption. In terms of broadening, very little statistical information currently exists on household energy production via solar panels or use of heat pumps, even though these are rapidly becoming more common. To assess the impact hereof on the total energy use, detailing of energy statistics is crucial. One way to obtain such detailed data is energy data donation: asking sampled households to share the detailed consumption data their own smart meters already record. This paper reports a feasibility pilot of one specific implementation of this approach.

In both academic research and industry, smart meter data is an emerging and promising research area with many different application areas (Wang et al. 2019). Especially, the disaggregation of total power consumption into its individual sub-components has received a lot of attention (for a review see Kaselimi et al. 2022). Advanced signal processing and machine learning techniques have been combined in non-intrusive load monitoring methods with which the first steps to real-life application scenarios are being made. This research area is maturing, yet still several issues regarding reliability, scalability, robustness, explainability, representation and measurement bias need to be addressed before the methods can be implemented (Kaselimi et al. 2022).

National statistical institutes from Austria, Denmark, Estonia, Sweden, Italy and Portugal have been working together in the ESSnet Big Data project on smart meter data analytics (2016-2018). In the follow-up ESSnet Big Data II project, Denmark and Sweden continued to work on electricity statistics of households from smart meter data (2018 – 2021) (Eurostat 2021). Ireland and the United Kingdom have both examined the feasibility of classifying households using smart meter data (Anderson and Newing 2015; Carroll et al. 2018), but to the authors' knowledge, no National Statistical Institute has yet implemented highly detailed energy data from smart meters in the production of their energy statistics.

The goal of this study was to explore the feasibility of one concrete way of collecting household energy data through smart technologies: a mixed-methods feasibility pilot in which households donate consumer-grade smart-meter data alongside a self-reported diary. Rather than comparing alternative collection methods, we examine in depth what the possibilities and draw-backs are from this particular setup. Specifically, in this study we investigate:

- The technical feasibility of collecting energy data through a smart meter dongle.
- The integration of passively collected energy data with self-reported information such as energy diaries and questionnaires.
- The conceptual accuracy of energy donations in relation to official statistical needs.
- The potential for disaggregated statistics, such as separating household energy use from transport-related energy (e.g. electric vehicle charging).
- The practical and ethical implications of implementing smart features at scale.

Thirteen participants were recruited and asked to participate in a pilot. The pilot combined three elements:

- Passive energy data collection via the HomeWizard P1 dongle and app.
- A short questionnaire capturing background information and perceptions.
- A self-reported diary logging presence and appliance usage over eight days.

This pilot focused on electricity and gas, with particular attention to variation between households, so that non-technical and non-energy-minded households were also represented. The pilot allowed us to explore installation barriers, privacy perceptions, data quality, and the potential of combining passive measurements with diary data to distinguish energy usage by device and context. This information is needed before upscaling to a larger population-based pilot.

This study was carried out within the Smart Survey Implementation project of Eurostat (Eurostat 2023), which explored how "smart features" such as sensor data, app-based participation, and data donations can be implemented in official statistics. The pilot reported here is the energy data donation case study within that wider project. We make this distinction explicit because the broader project frames energy data donation as one form of a smart survey, whereas the present pilot is a small, exploratory feasibility study of a potential implementation.

2 Theoretical Background

2.1 Smart Surveys and Data Donation with Smart Meter Data

Traditionally statistics were made from register data and survey data. With the advent of the digital age, an increasing amount of digital data is being generated, for example by the sensors in your phone, appliances that you wear, or sensors in your house. This opens the possibility to enhance statistical production, with respect to quality and level of detail. Schouten et al. (2025) have coined the term Smart Surveys for surveys that combine these new forms of data with traditional surveys. Data Donation is one form of smart surveys in which sampled individuals are asked to share personal data they have collected or have access to. The European Data Act states that individuals are allowed to access and share the digital data on their behaviour with third parties (European Parliament and Council 2023). This institutional and conceptual setting provides the background for the present study; the pilot itself is a first exploration of the feasibility of the data donation element.

Data from smart meters is read by the network operators to monitor the stability of the electricity and gas grid and detect issues faster. Users typically cannot retrieve detailed time-series via these network operators, yet they can install an energy consumption manager to get insights in their own data. These energy consumption managers exist in various forms: for example, a display, a smart thermostat, an app or an online tool. For some of these, users need to plug in a dongle in the P1 port of their smart meter to be able to retrieve the data. This can be done using commercial third parties' products or via a self-programmed home automation system. In this pilot, we focus on smart meter data collected via the installation of a dongle of a third party (HomeWizard) in the P1 port of the smart meter.

2.2 Energy Statistics

Currently, Statistics Netherlands is adapting its energy statistics to move from publication based on data received yearly via the network operators to more accurate and faster publication based on the monthly data from smart meters (CBS 2025). This is possible due to the wide adoption of smart meters in the Netherlands, with a penetration of over 90%. With the monthly readings from these smart meters, the first steps towards more frequent, granular and accurate energy statistics have been taken. Yet, further disaggregation is desirable and possible with the actual measurement frequency of smart meters. This opens the opportunity to collect data on energy consumption with a high frequency level, directly from respondents.

With high frequency level data, energy signals can be decomposed into individual appliance consumption patterns (Kaselimi et al. 2022; Kelly and Knottenbelt 2015; Zoha et al. 2012). This approach has been applied not only for electricity but also in water usage studies, where flow data combined with smart metering allows identification of appliance-specific water consumption (Larson et al. 2012; Marsili et al. 2024).

The required granularity to capture these dynamics typically involves sub-hourly measurements, with higher frequencies enabling more precise and accurate identification of short-duration peaks and appliance operation cycles. In scientific research, energy disaggregation is generally done using datasets with sampling frequencies ranging between 10 Hz and 100 kHz, whereas the readily installed household smart meters only measure once per second at a maximum, and many consumer apps export 15-minute aggregates (Gupta et al. 2024; Navarro et al. 2025; Young et al. 2024). The 15-minute aggregates available from consumer apps therefore set the resolution that the present pilot can work with.

3 Methods

The goal of this pilot was to explore the feasibility of one potential implementation of energy data donation: collecting detailed energy consumption data using a HomeWizard P1 meter, alongside a self-reported diary. The study assessed the technical ease of collecting passive data, the usefulness of complementary diary information for identifying the use of specific household appliances, and the feasibility of measuring energy consumption for official statistics.

The pilot is best characterised as a mixed-methods feasibility study. The smart-meter data are quantitative, and the diary and questionnaires yield structured, categorisable responses that we summarise quantitatively (for example as frequency counts). At the same time, the core of the analysis is interpretative: we use no statistical analyses but examine individual consumption curves, label peaks with the help of the diaries, and reason about what can and cannot be inferred. The interpretation of the example cases in the Results section is therefore qualitative in nature.

3.1 Participant Recruitment

The aim at this stage was to understand what the data contain rather than to generalise to a population, and thus representativeness was deliberately not a design objective. Yet, a diverse participant sample was sought to capture a range of household situations. Participants were selected to obtain diversity in gender, age, and department, increasing the likelihood of including various levels of technical skill and degrees of interest in energy. A total of 13 colleagues from Statistics Netherlands were recruited to take part in the pilot. One participant did not participate after all, and another was not able to install the dongle.

3.2 Data Collection Period

The data collection period lasted for eight consecutive days. Participants could choose when to start within a common window: everyone was asked to install the dongle and begin keeping the diary at roughly the same time, and to return their data once the eight days were complete. A delay of a few days was acceptable, so the individual measurement periods overlapped to a large extent without being perfectly aligned. Respondents were asked to keep a diary throughout this period, recording:

- Their presence at home (e.g., which parts of the day they were away or at home).
- The usage of major electrical appliances (e.g., dishwasher, washing machine, tumble dryer).
- Special activities affecting consumption (e.g., gatherings).

Upon completing the eight-day data collection, participants were instructed to export the data from the HomeWizard P1 meter and email it to the researchers. The diary was either handed in personally or emailed.

3.3 Data Collection Tools and Procedures

Respondents were provided with the HomeWizard P1 Meter and a booklet containing the pre-questionnaire, installation instructions, diary and post-questionnaire (see Supplementary Material). The pre-questionnaire collected background and household characteristics (such as gender, age, household size, housing type and energy label) together with perceptions of the home's energy efficiency, energy-saving measures already taken, and reasons for saving energy. The post-questionnaire collected feedback on the installation process and on the experience of keeping the diary. The diary itself was developed specifically for this pilot. Its main purpose was to let respondents label their own energy use over time, so that peaks in the metered data could be linked to concrete activities and appliances; the labelling was left deliberately open, so that respondents used their own descriptions rather than a fixed list.

The P1 meter is a device that connects directly to a digital smart meter via the P1 port. With the app, participants could export 15-minute intervals for electricity and gas usage, including:

- Electricity import (consumption from the grid).
- Electricity export (for households with solar panels).
- Peak usage
- Gas consumption, depending on the availability of gas readings in the smart meter.

The P1 meter was chosen for its relatively simple installation process, which, for most newer smart meters, involved plugging the dongle into the P1 port. However, one respondent with an older smart meter required an additional USB-C cable as a power source for the dongle to complete the setup.

Once connected, participants downloaded and linked the HomeWizard Energy app to their P1 meter. The app offered real-time insights into household energy usage, which many found helpful for monitoring daily consumption. At the end of the pilot, participants exported their data via the app. In order to be able to export data, respondents had to sign up for a one-month subscription. These costs were reimbursed. The exported data was then sent to the research team for analysis.

To complement the P1 meter data, participants maintained an energy diary. Each day, they noted:

1. Presence: Whether they were at home or away, allowing for an approximate link between energy usage and occupancy.
2. Appliance Usage: Which major household devices were used, along with estimated times.
3. Special Activities: Any noteworthy events that could explain lower- or higher-than- usual consumption periods.

The diary was central to the analysis rather than a secondary add-on. Because the metered data on their own only show aggregate consumption per quarter-hour, the diary provides the labels needed to interpret what those peaks represent. It is precisely the combination of passively measured consumption and respondent-supplied labels that allows us to judge the validity and accuracy of the dongle measurements and to assess how far appliance-level interpretation is possible at this resolution.

3.4 Data Handling and Preparation

Respondents submitted their data in comma-separated values (CSV) files, which is the app default. Each respondent provided a separate CSV for electricity and a separate CSV for gas, spanning the full eight-day period.

Despite using the same HomeWizard P1 meter, individual data files varied in the number of columns. This was partly due to differences in respondents' home installations and meter configurations.

- Tariff Classes: Where respondents had multiple tariff classes active (e.g. day and night tariffs), the meter recorded separate readings. The CSV included meter values per tariff class. In most cases, two tariffs were present, but there was no question included regarding potential price differences between those tariffs.
- Export to the Grid: For households with solar panels, the P1 meter recorded exported energy (i.e., electricity generated and delivered back to the grid). The data file thus contained separate columns for import (consumption) and export (production).
- Peak Wattage: For a subset of respondents the peak wattage for three "groups" was also available in the CSV. According to HomeWizard, these groups represent the three phases in an electrical installation, however, the exact household wiring details were not investigated in this study.

To facilitate analysis, several data processing steps were applied:

1. Electricity Import and Export: The original data provided meter values for each 15-minute interval. To calculate consumption, consecutive meter readings were subtracted, yielding the actual energy used (kWh) in that quarter-hour block. Where day/night tariffs existed, these readings were combined into a single "import" total for each time interval. The same approach was used for export data.
2. Peak Wattage: For respondents reporting peak usage across multiple phases (i.e. three columns of data), the values were summed into a single "peak wattage" column per time interval.

3. Gas Data: Gas readings only contained one tariff. Following the same logic as electricity import, consecutive readings were subtracted to derive quarter-hourly gas consumption.

Subsequently, the electricity and gas files were merged into a single dataset for each respondent, matched on the timestamp. To check for completeness and validity, the number of records per respondent for the 8-day period was compared with the expected minimum number of records (*4 quarters × 24 hours × 8 days = 768 records*). One respondent had less records than expected as this respondent had mistakenly exported daily totals instead of 15-minute intervals and was therefore excluded from the analysis. Some individuals surpassed the expected number of records (up to 5,280 records), because they had been using the P1 meter for a longer period and chose to export the entire data history.

To facilitate combined plotting of import, export, gas usage, and peak wattage on a similar scale, the data were normalised per respondent. Electricity Import and Export was divided by the combined maximum value per respondent. This resulted in values ranging between 0 and 1. Similarly, gas consumption was scaled. For Peak Wattage the minimum can be below zero when more energy is created with the solar panels than is being used at that moment. Therefore, it has been divided by the maximum absolute Peak Wattage to account for the negative values. This resulted in a range between -1 and 1. This normalisation enabled straightforward visual comparisons across different types of consumption/production data, without losing relative peaks or patterns within each participant's dataset.

All diary entries were transcribed into an Excel file and subsequently imported into R. Participants recorded which appliances they used, the approximate times of use, and the duration. In total 320 time-measurements were recorded and 39 contained missing values for the time and/or duration (34 from three respondents). When the exact usage time was missing, these entries could not be plotted in time-series charts. However, they still contributed to overall counts of appliances usage.

The text descriptions of the appliances were standardised and grouped into broad categories for analysis and plotting:

- Heating/ Air conditioning (AC)
- Electric car charging
- Kitchen (general cooking appliances)
- Washing machine
- Tumble dryer
- Dishwasher
- Other (miscellaneous equipment)

Information on the number of people present in the household was extracted into a separate Excel file, split into four six-hour blocks (night, morning, afternoon, evening). This dataset was also imported into R, allowing for analysis of occupancy patterns in relation to energy consumption peaks. The questionnaire data were likewise transcribed into Excel and loaded into R. Variable formats (e.g. numeric, categorical) were verified to ensure accurate analyses.

3.5 Ethical Considerations and Privacy

Energy usage data can reveal detailed patterns about household occupancy and routines. Therefore, all participants were fully informed about the purpose of the study and the types of data collected. They provided consent to share their energy usage data and diary entries. The potential privacy implications were explained, and participants were given the option to withdraw from the study at any time. For this pilot, data transfers relied on participants exporting their usage files via the Energy app and emailing them to the research team. While this approach was adequate for a small-scale pilot, it is not suitable for large-scale population monitoring. A more secure and automated data collection infrastructure would be essential for future expansions.

4 Results

The results are organised to follow the five goals set out in the Introduction. We first describe participation and the practical experience of installation and diary keeping, which speak to the technical feasibility of collecting energy data through the dongle (goal 1). We then report perceptions of energy efficiency and the measures households take, which illustrate how passively collected data can be combined with self-reported questionnaire information (goal 2). The example patterns that follow address the conceptual accuracy of the donations and the potential for disaggregated statistics (goals 3 and 4), while observations on privacy and behavioural change feed into the practical and ethical implications of implementing such features at scale (goal 5).

4.1 Participation and Respondent Characteristics

Of the 13 participants initially recruited, 11 completed the full pilot. One participant inadvertently exported daily instead of quarter-hourly data, which limited their dataset for detailed analysis. The remaining ten provided sufficiently granular data.

Table 4.1 Demographic and household characteristics

Measure	Category	Count
Gender	Female	5
	Male	6
Age	25–35	2
	35–45	4
	45–55	3
	55–65	2
Household size	2 persons	4
	3 persons	2
	4 persons	4
	5 persons	1
Housing type	Detached	5
	Semi-detached	3
	Terraced	3
Energy label	A	1
	B/C	4
	D/E	0
	F/G	1
	Don’t know	5

As shown in Table 4.1, the respondent group was balanced in terms of gender and included variation in age, household size, housing type, and energy label, providing a broad view of energy usage contexts within a small sample.

Most participants found the HomeWizard P1 meter straightforward to install: ten of the eleven reported it was "easy" or "very easy." One respondent encountered technical issues and required an additional USB-C cable due to an older smart meter. The installation instructions were considered clear by nearly all participants.

Feedback on the energy diary was mixed. Some respondents appreciated the increased awareness it provided, particularly in identifying high-consumption appliances. However, one participant noted that logging appliance usage consistently was burdensome. No technical faults or data quality issues were reported.

4.2 Perceptions of Energy Efficiency and Measures Taken

This subsection draws on the pre-questionnaire that accompanied the diary and illustrates the second goal: integrating passively collected energy data with self-reported information. Before the start of the pilot, participants were asked about their satisfaction with their home's energy efficiency. Of the 11 who completed the study, five were (very) satisfied, four were neutral and two were dissatisfied.

Those who were satisfied generally had better energy labels or at least were aware of their label, suggesting a link between knowledge, label quality, and satisfaction. Respondents who were less satisfied often had lower energy labels or unknown labels. Since 2008, it is obligatory to have an energy label when buying a house in the Netherlands, meaning that older homes are therefore more likely to lack a label and to be less energy efficient. Respondents who reported fewer energy-saving measures frequently felt their homes were already adequately efficient, whereas those less satisfied were often unsure which measures remained viable or beneficial.

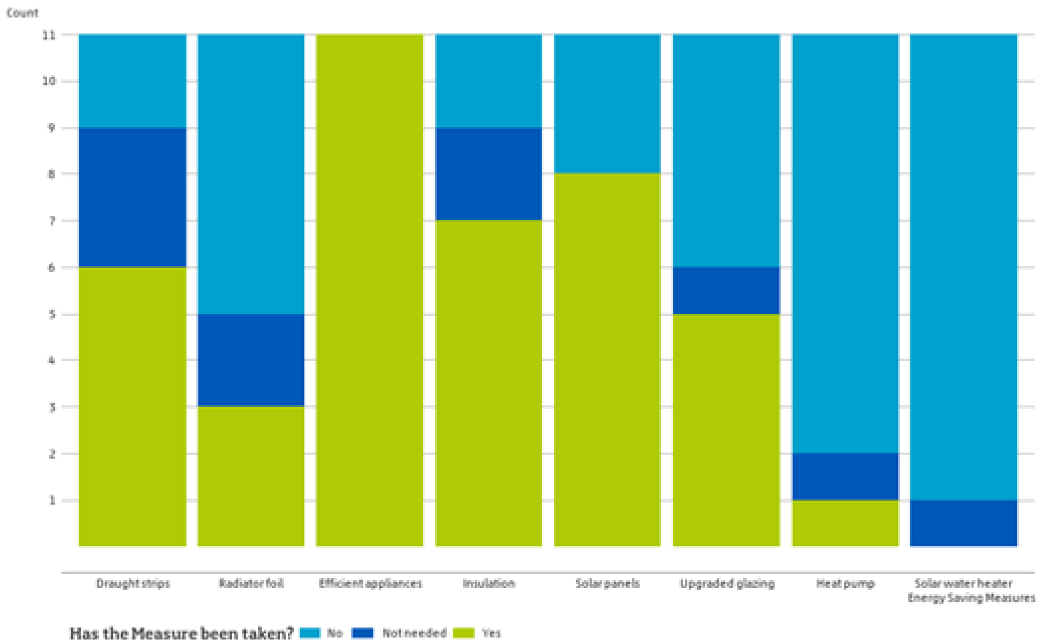

Figure 4.1 Measures taken to save energy

Energy-saving measures most frequently taken included purchasing energy-efficient appliances, installing solar panels, and improving insulation (Figure 4.1). Those who were satisfied with their energy efficiency often reported that no further measures were necessary.

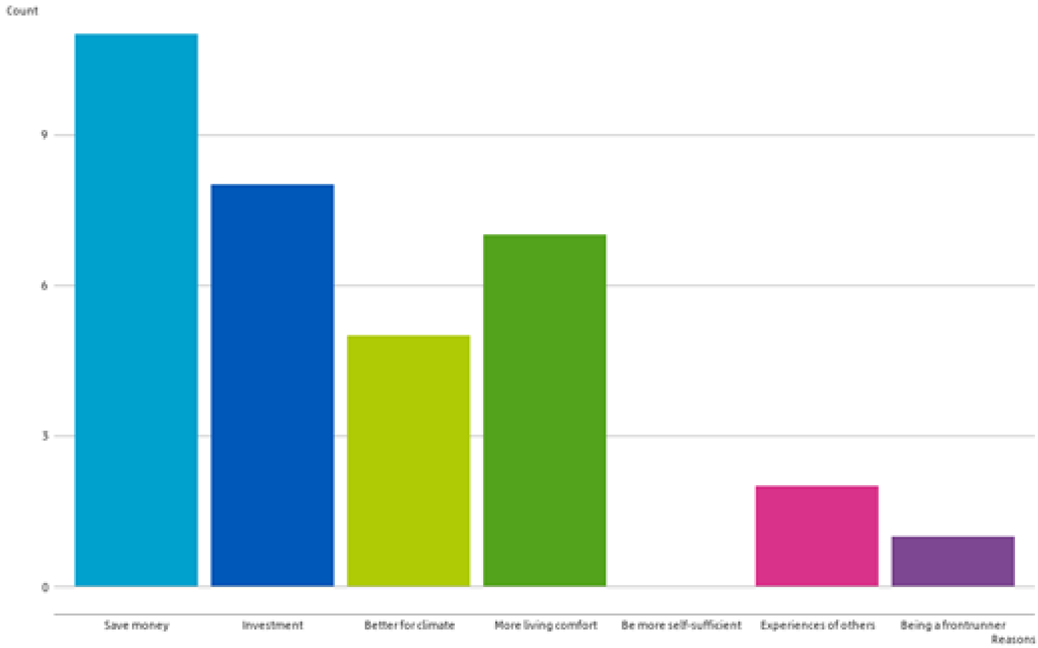

Figure 4.2 Reasons to save energy

The most frequently cited reasons for undertaking these measures were reducing costs, an investment, improving living comfort and it being better for the climate (Figure 4.2). Being self-sufficient, a frontrunner or because of the experiences of others seem to play a smaller role.

4.3 Insights and Example Patterns

During the eight days of the study, participants logged the use of 320 energy-consuming appliances. Cooking and washing machines were most frequently reported, followed by dryers and dishwashers (Figure 4.3). Note that entries marked as “cooling” actually referred to heating via air-conditioning systems during the winter period.

To illustrate the potential and limitations of quarter-hourly energy data, several examples from across the pilot participants will be discussed in the following paragraphs. These examples reflect common patterns as well as exceptional or complex cases, and help clarify what kinds of insights are feasible with the current setup.

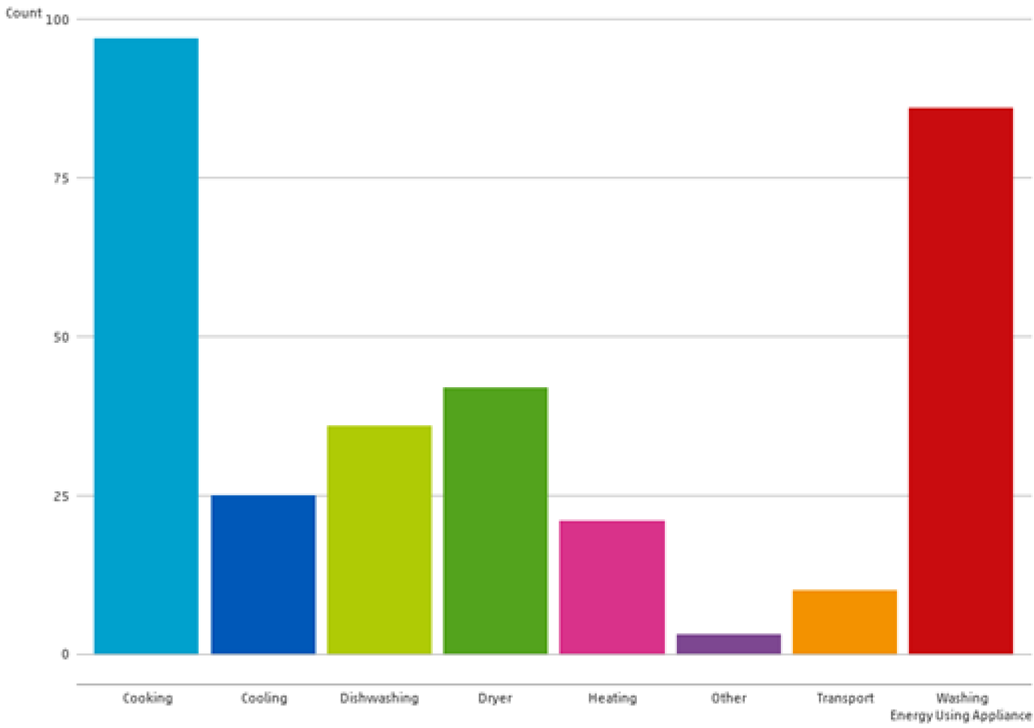

Figure 4.3 Total reported frequency of energy-using appliances according to the diaries

4.3.1 Basic energy use with clear appliance activity

The first example (Figure 4.4) shows a relatively straightforward case where energy use can be matched to specific appliances using the diary. Peaks in electricity usage align well with the start of washing machine, dryer, and dishwasher cycles. This is in line with expectations, as the initial phase of these appliances—particularly water heating—often requires the most energy. Cooking activity during dinner also corresponds to a visible peak. Gas consumption shows a strong peak in the morning, likely linked to heating, and smaller peaks at other moments, including during cooking. This example demonstrates how basic disaggregation is possible when a limited number of devices are used, their operation does not overlap and when diary information is available.

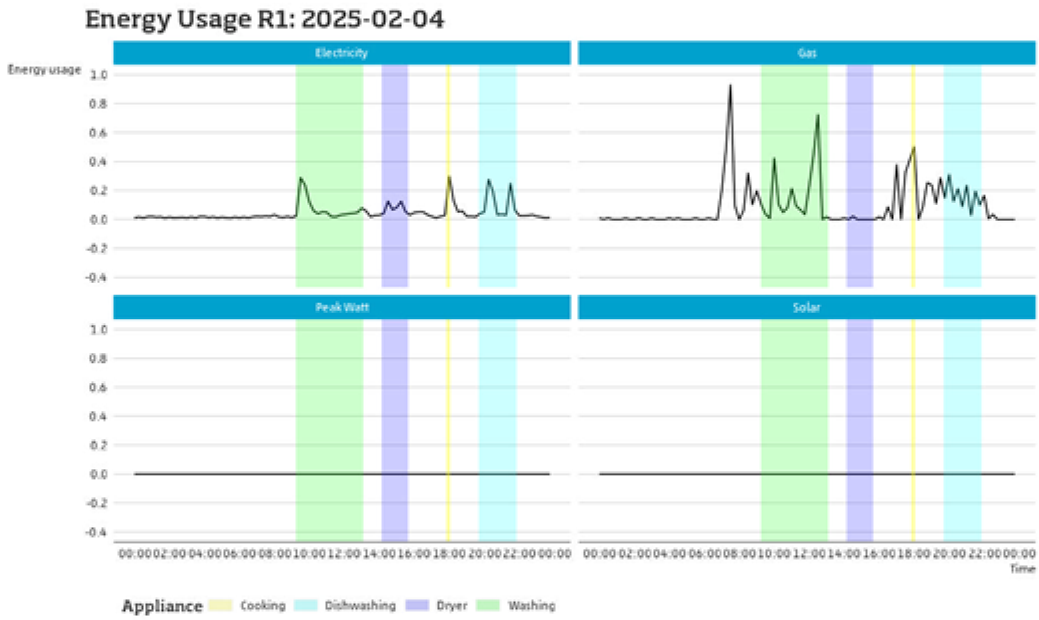

Figure 4.4 General energy usage example for gas and electricity

4.3.2 Energy usage with solar production and peak power visibility

In Figure 4.5, solar panel output and peak power data were also available. Solar energy was produced during the day, while most energy consumption occurred in the morning and evening, confirming a temporal mismatch that could overload the energy grid in decentralised systems. Notably, some periodic consumption—likely from a refrigerator or freezer—was only visible in the peak power data, not in the net energy import. During daylight, such appliances were likely powered directly by solar generation. This highlights the added value of peak power metrics for identifying hidden or background loads, especially in solar-enabled households.

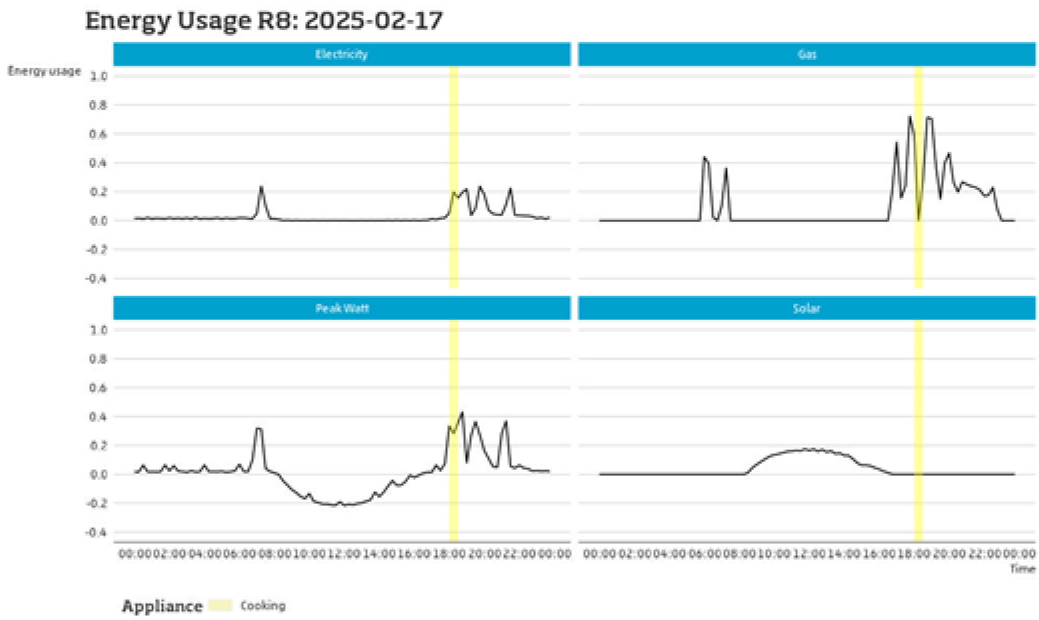

Figure 4.5 General energy usage example for gas, electricity, peak power and solar energy

4.3.3 Complex overlapping energy use patterns

Figure 4.6 presents a much more complex picture, where many appliances were used throughout the day. Even with detailed diary data, it becomes almost impossible to attribute energy use to individual devices due to overlapping activities. This example shows the practical limits of both self-reporting and the data with a 15-min frequency. It also illustrates the cognitive burden placed on participants when asked to track their energy use manually, especially if multiple appliances are used concurrently or repeatedly. These types of patterns suggest that fully manual methods are not sustainable for long-term or large-scale measurement efforts.

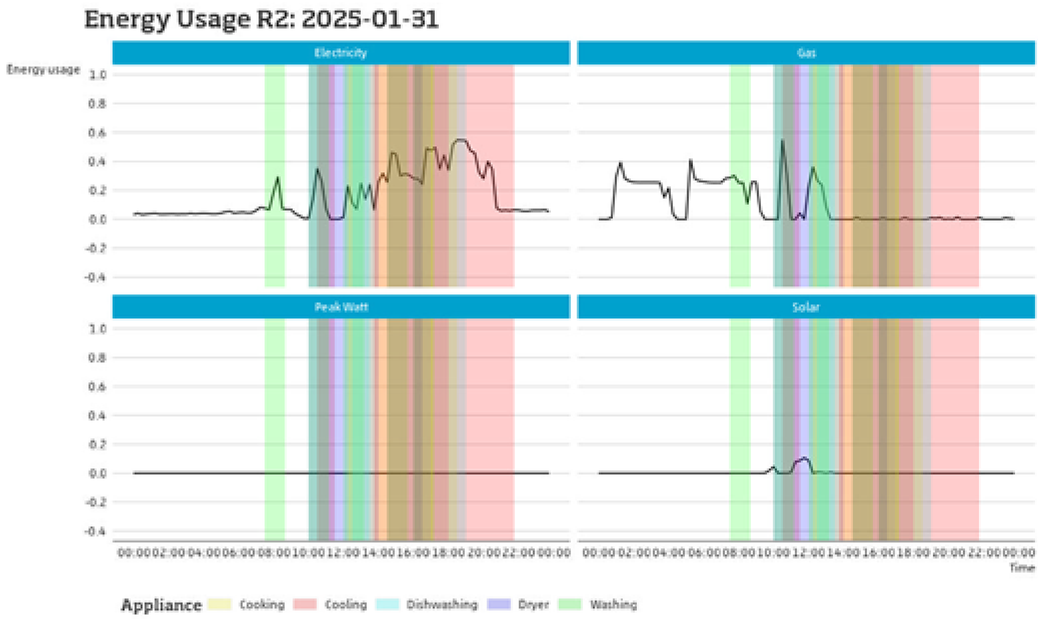

Figure 4.6 Example of complexity of energy usage

4.3.4 Heating with air-conditioning systems

A more specific use case is shown in Figure 4.7, where an air-conditioning unit was used for heating during a cold January day. This activity is clearly visible in the electricity data and aligns well with diary reports. However, without contextual information, such as temperature data, it would be difficult to determine whether the system was used for heating or cooling. This example underlines the need for integrating external contextual data, such as weather conditions, to improve the interpretation of energy usage.

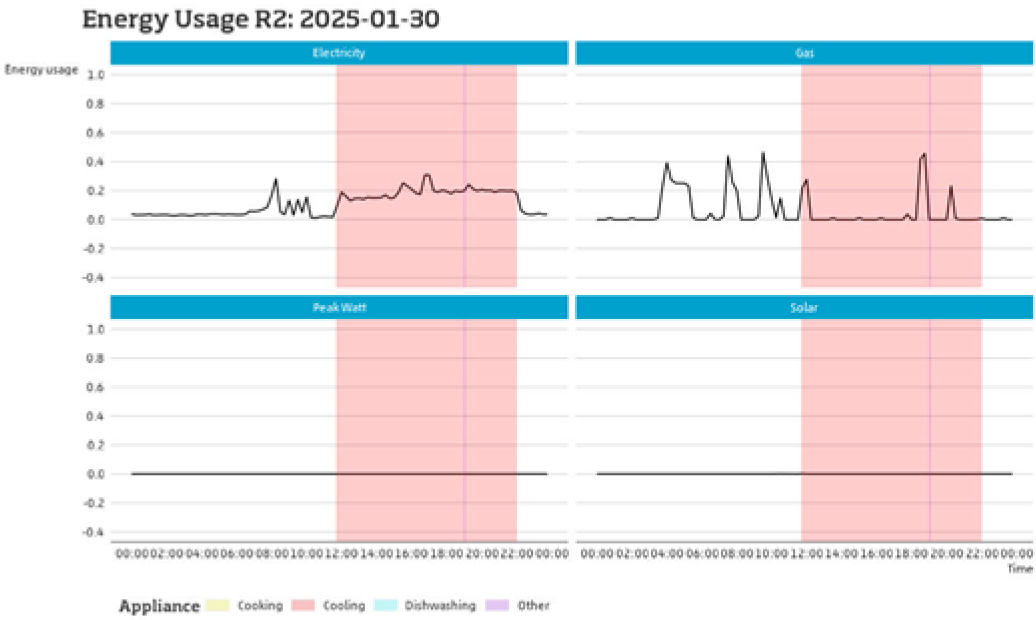

Figure 4.7 Example of energy usage of the air-conditioning system for heating with electricity

4.3.5 Combined gas and electric heating

Figure 4.8 provides an example where gas and electricity usage peaks coincide, suggesting a combined or hybrid heating system. This pattern adds another layer of complexity to disaggregation, as energy usage for a single activity (e.g. heating) may draw from multiple sources. It also shows that energy patterns do not always follow clear one-to-one relationships between devices and energy types.

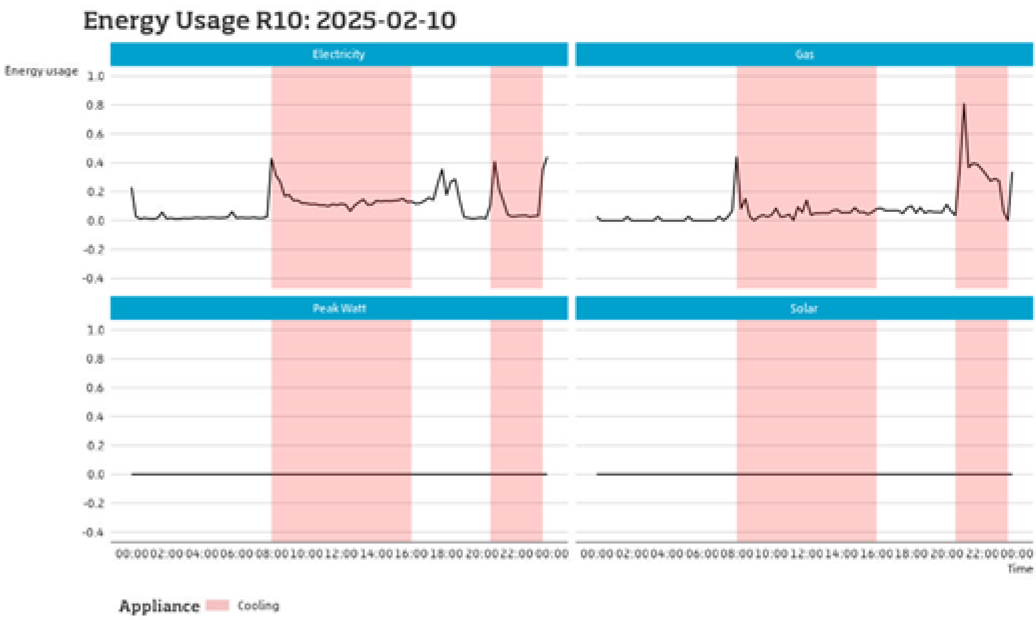

Figure 4.8 Example of energy usage of the air-conditioning system for heating with gas and electricity

4.3.6 Irregular and highly volatile usage

Finally, Figure 4.9 illustrates a case of highly irregular and unpredictable energy use. This example reinforces the idea that household energy behaviour is highly individual, and that generic disaggregation models are unlikely to perform well across all households without some form of local adaptation or calibration.

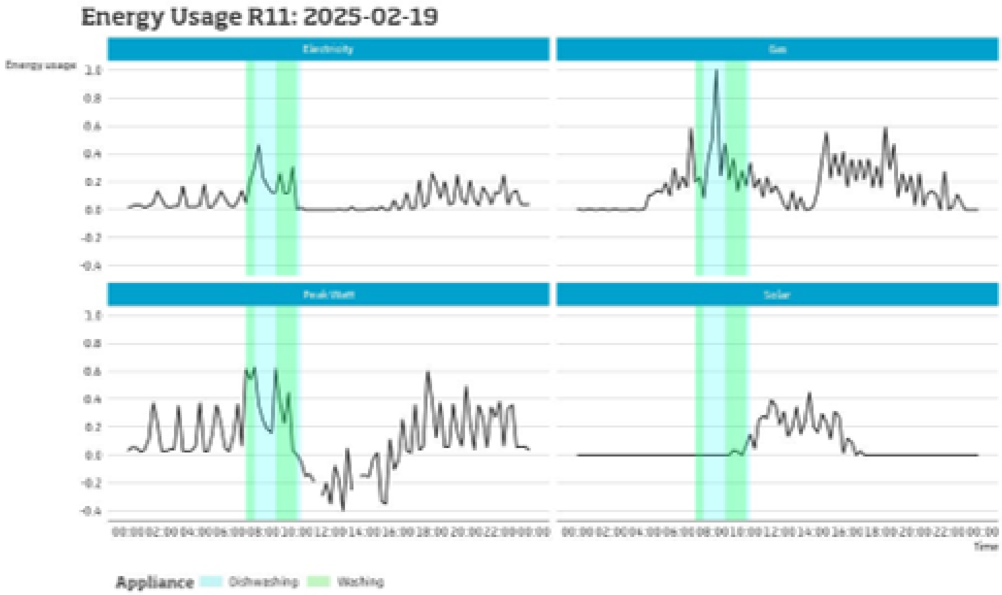

Figure 4.9 Example of the irregularity of energy data

These examples collectively show that 15-minute interval data allow for general consumption pattern analysis but lack the granularity needed to reliably identify individual appliances. The addition of solar output and peak power metrics enhances interpretability in certain contexts, especially when devices draw power directly from solar generation. However, to move towards automated disaggregation, higher-frequency data and machine learning models are essential.

Note that three respondents reported adapting their habits during the study. For example, they scheduled energy-intensive appliances for daytime hours to optimise solar panel output or to avoid overlapping multiple appliances at once. Half of the participants raised concerns about privacy, pointing out that detailed energy consumption data could reveal whether a household member was present or absent, especially in single-person dwellings.

Most households saw greater electricity consumption in the mornings and evenings, aligning with diary reports of higher occupancy during those periods (see Figure 4.10). In one notable instance, a respondent recorded an unusually high number of occupants for one evening although this effect was not clearly visible in the energy use. It seems like the number of occupants is less relevant to the energy consumption.

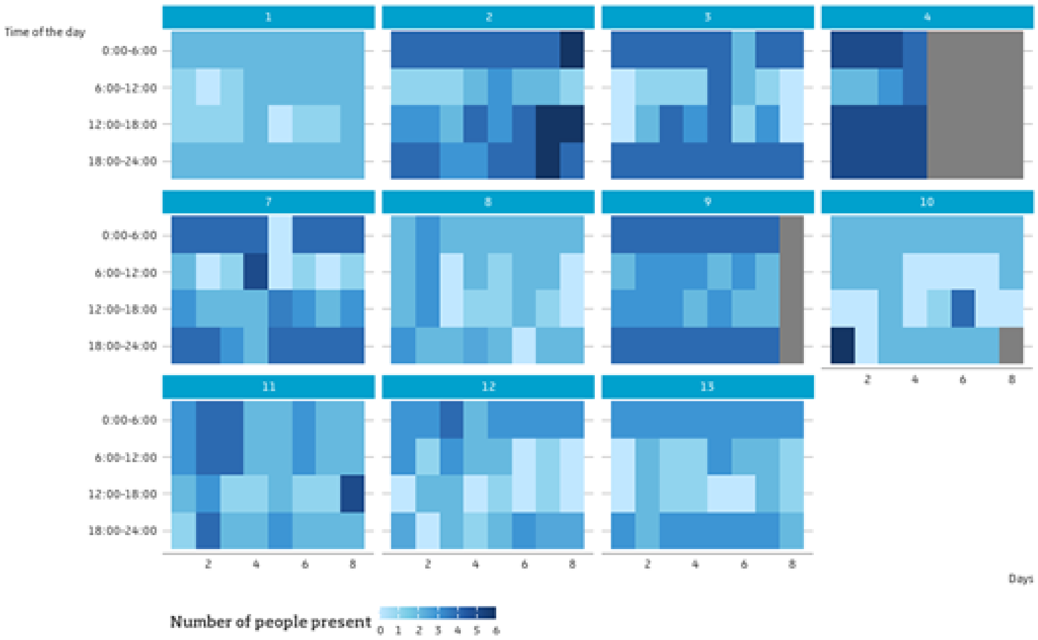

Figure 4.10 Number of people present during the pilot per respondent

5 Discussion

In this pilot we explored the feasibility of one potential implementation of energy data donation: collecting household energy data via consumer-grade smart-meter dongles, combined with a self-reported diary. This marks an important first step in establishing what such passively collected data can and cannot contribute to official statistics. We discuss the findings in the order of the five goals set out in the Introduction: the technical feasibility of the dongle, the integration of passive data with self-reported information, the conceptual accuracy of the donations, the potential for disaggregation, and the practical and ethical implications of working at scale.

The pilot demonstrated that passive data collection through a smart meter dongle is technically feasible and well-received by participants. The HomeWizard P1 dongle was relatively easy to install, and the accompanying app provided a level of real-time feedback that participants appreciated. Once installed, the HomeWizard P1 meter required little effort from participants, and the accompanying app provided real-time insights that many found both engaging and informative. However, the study also highlighted critical limitations in using this method for producing detailed energy statistics that meet current and future statistical demands.

The most significant limitation is the lack of granularity: the data collected in 15-minute intervals does not provide enough resolution to automatically detect which appliances are being used or to disaggregate energy use by type (e.g., heating, lighting, cooking). This limitation prevents the transition to fully automated, diary-free statistics, which is a key long-term goal.

Most participants managed to install the dongle, but some required assistance. This finding may have been shaped by the participant group. Recruiting participants from within Statistics Netherlands made onboarding and follow-up practical for a first pilot, but compared with the general population, these participants are likely to have above-average data literacy, greater institutional trust, and possibly different privacy perceptions. Each of these could make installation, diary keeping, and acceptance appear smoother than they would be at scale. We therefore treat the findings on burden and acceptance as indicative of what is achievable under relatively favourable conditions, and return to this point when discussing upscaling. A broader, population-based pilot will be needed to test whether the same results hold for less data-literate households.

Moreover, the current approach of manual data exports via email is not scalable. The diary element should be reduced to a minimum to limit response burden. Participants also reported increased awareness of their energy use, and some even adapted their behaviour accordingly. Privacy was a further concern. While it did not deter participation in this study, several participants expressed discomfort about how revealing energy-use patterns can be. This reinforces the importance of strong privacy-by-design measures, particularly as data frequency increases.

In this context, privacy-by-design means more than a general principle. Concretely, it implies minimising the data that leave the household in the first place, for example by aggregating or extracting features locally on the dongle so that only the information needed for statistics is transmitted rather than full high-resolution traces. It implies encrypting data both in transit and at rest, and processing them within infrastructure controlled by the statistical office rather than a commercial third party. It also implies clear data-governance arrangements: explicit purpose limitation, defined retention periods, access controls, and transparent communication to participants about what is collected and why. Making these elements explicit not only clarifies what we mean by privacy-by-design but also turns it into concrete design requirements that other national statistical institutes can adopt.

From a data integration perspective, it was possible to combine quarter-hourly gas and electricity usage with self-reported diary data on presence and device usage. This allowed for illustrative analyses of peak usage patterns and presence-activity correlations. However, despite this progress, the pilot also made clear that the data collected at 15-minute intervals is not sufficiently detailed to automatically distinguish between types of energy usage or devices. At this granularity, diaries remain essential for interpretation, limiting the scalability and automation potential of the current method.

Two boundary conditions should be kept in mind when reading these results, and both point towards what a follow-up study should address rather than undermining the present findings. First, this was a small-scale, purposively selected pilot with 11 complete cases. This is entirely appropriate for prototyping and for surfacing the practical issues that a larger study must anticipate, but it means the results characterise what is possible rather than how common each pattern is in the population; establishing the latter is a task for a subsequent population-based pilot. Second, the pilot relied on a commercial third party for data collection. This was a practical choice for a first feasibility test, and it usefully demonstrates where the data governance boundaries lie: for routine statistical production, the data pipeline should be under the control of the statistical office to ensure transparency, reproducibility, and long-term stability. Both points therefore inform the design of the next phase rather than detracting from what this pilot has shown.

To advance energy data donation from a conceptual pilot to a robust, production-ready approach, several strategic investments and design choices are needed:

- A high-frequency follow-up pilot should be conducted to determine the optimal sensor resolution required for device-level disaggregation. Insights from earlier research (e.g. Froehlich et al. 2009) suggest that much finer sampling may be needed than currently used. Furthermore, this pilot should be used to develop a self-learning algorithm that reduces

respondent burden over time. Finally, the algorithm needs to be validated. Careful attention needs to be given on the detection of behavioural change due to monitoring and how to adjust for this.
– Establish a rotating energy panel with onboarding and a short stabilisation period. A dedicated onboarding phase (field support where needed) and treating the first weeks as a stabilisation period help mitigate reactivity and reduce selection bias at scale.
– Develop a hybrid system where a machine learning model makes initial predictions about appliance usage and the respondent is able to validate and adjust these using a mobile app. Occasionally, the app can prompt users to validate or correct these predictions, especially when uncertainty is high.
– A dongle and data infrastructure controlled by the National Statistical Institute must be developed to ensure end-to-end control of the data collection process, ethical compliance, and integration into existing statistical workflows. Important elements here are local data aggregation and data minimisation strategies.
– Privacy must remain central to the design. Participants need to understand why data is collected, how it will be used, and what protections are in place. Transparent communication and participatory design will be essential in maintaining trust.

The pilot has shown what is possible and what is still out of reach. Passive measurement technologies hold great promise for reducing respondent burden and enabling more timely and detailed statistics. However, to realise this potential, more research is needed. Investments are needed for purpose-built infrastructure, privacy-conscious algorithms, and robust methodological frameworks.

Taken together, our findings indicate that energy data donation can complement, rather than replace, existing survey and administrative sources today. With a pipeline that is controlled by the National Statistical Institute and designed around privacy-by-design, combined with higher-frequency sensing and light-touch human validation, energy data donation can realistically evolve from feasibility pilot to a scalable component of official household energy statistics.

Acknowledgments

The views expressed in this paper are those of the authors and do not necessarily reflect the policies of Statistics Netherlands.

Funding

This work was carried out as part of the Eurostat Smart Survey Implementation (SSI) project (Grant Agreement No. 101119594). Statistics Netherlands provided in-kind support in the form of staff time and infrastructure. No additional external funding was received for this study.

Data availability

The data underlying this study consist of household smart meter readings and detailed energy-use diaries from a small group of Statistics Netherlands employees. Due to privacy and confidentiality requirements under the Statistics Netherlands legal framework, the raw data cannot be shared publicly. Aggregated outputs, analysis code, and supplementary materials related to this study may be made available upon reasonable request to the authors, subject to approval by Statistics Netherlands.

References

Anderson, B. and A. Newing (2015). *Using Energy Metering Data to Support Official Statistics: A Feasibility Study*. Office for National Statistics.

Carroll, P., T. Murphy, M. Hanley, D. Dempsey, and J. Dunne (2018). "Household classification using smart meter data". In: *Journal of Official Statistics* 34.1, pp. 1–25.

CBS (2025). *Verbeterd inzicht in energieverbruik van woningen met slimmemetergegevens*. Statistics Netherlands. URL: `https://www.cbs.nl/nl-nl/longread/discussionpapers/2025/verbeterd-inzicht-in-energieverbruik-van-woningen-met-slimmemetergegevens`.

European Parliament and Council (2023). *Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data (Data Act)*. Official Journal of the European Union.

Eurostat (2021). *ESSnet Big Data II: Grant Agreement Number: 847375 – 2018-NL-BIGDATA*. Eurostat.

– (2023). *Smart Survey Implementation: Grant Agreement Number: 101119594*. Eurostat.

Froehlich, J., E. Larson, E. Saba, T. Campbell, L. Atlas, J. Fogarty, and S. Patel (2009). "HydroSense: Infrastructure-Mediated Single-Point Sensing of Whole-Home Water Activity". In: *Proceedings of the 11th International Conference on Ubiquitous Computing*.

Gupta, P. S., S. S. Alam, Y. Zhang, and M. A. Islam (2024). "Towards Non-Intrusive Real-Time Monitoring of Behind-the-Meter Residential Distributed Energy Resources". In: *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*, pp. 428–433.

Kaselimi, M., E. Protopapadakis, A. Voulodimos, N. Doulamis, and A. Doulamis (2022). "Towards trustworthy energy disaggregation: A review of challenges, methods, and perspectives for non-intrusive load monitoring". In: *Sensors* 22.15, p. 5872.

Kelly, J. and W. Knottenbelt (2015). "The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes". In: *Scientific Data* 2, p. 150007. DOI: `10.1038/sdata.2015.7`.

Larson, E., J. Froehlich, T. Campbell, C. Haggerty, L. Atlas, J. Fogarty, and S. N. Patel (2012). "Disaggregated water sensing from a single, pressure-based sensor: An extended analysis of HydroSense using staged experiments". In: *Pervasive and Mobile Computing* 8.1, pp. 82–102.

Marsili, V., F. Mazzoni, S. Alvisi, and M. Franchini (2024). "From pressure to water consumption: Exploiting high-resolution pressure data to investigate the end uses of water". In: *Water Resources Management* 38.13, pp. 4969–4985.

Navarro, V. M., M. Barragán, R. Nieto, J. Ureña, and ⍰. Hernández (2025). "HIFDA: High-Frequency Electrical Voltage and Current Signals from Household Appliances". In: *Scientific Data* 12.1, p. 527.

Schouten, B., P. Lugtig, and A. Luiten (2025). "Can Smart Surveys Have a Positive Business Case? An Evaluation Based on Three Case Studies". In: *Journal of Official Statistics* 41.2, pp. 547–568.

Wang, Y., Q. Chen, T. Hong, and C. Kang (2019). "Review of smart meter data analytics: Applications, methodologies, and challenges". In: *IEEE Transactions on Smart Grid* 10.3, pp. 3125–3148. DOI: `10.1109/TSG.2018.2818167`.

Young, T. L., J. Gopsill, M. Valero, S. Eikevag, and B. Hicks (2024). "Comparing four machine learning algorithms for household non-intrusive load monitoring". In: *Energy and AI* 17, p. 100384.

Zoha, A., A. Gluhak, M. A. Imran, and S. Rajasegarar (2012). “Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey”. In: *Sensors* 12.12, pp. 16838–16866.

Colophon

Publisher
Statistics Netherlands
Henri Faasdreef 312, 2492 JP The Hague
www.cbs.nl

Prepress
Statistics Netherlands, Grafimedia

Design
Edenspiekermann

Information
Telephone +31 88 570 70 70, fax +31 70 337 59 94
Via contact form: www.cbs.nl/information