
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Sierla, Seppo; Pourakbari-Kasmaei, Mahdi; Vyatkin, Valeriy

A taxonomy of machine learning applications for virtual power plants and home/building energy management systems

Published in:
Automation in Construction

DOI:
[10.1016/j.autcon.2022.104174](https://doi.org/10.1016/j.autcon.2022.104174)

Published: 01/04/2022

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Sierla, S., Pourakbari-Kasmaei, M., & Vyatkin, V. (2022). A taxonomy of machine learning applications for virtual power plants and home/building energy management systems. *Automation in Construction*, 136, Article 104174. <https://doi.org/10.1016/j.autcon.2022.104174>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



Review

A taxonomy of machine learning applications for virtual power plants and home/building energy management systems

Seppo Sierla^{a,*}, Mahdi Pourakbari-Kasmaei^a, Valeriy Vyatkin^{a,b,c}

^a Department of Electrical Engineering and Automation, School of Electrical Engineering, Aalto University, FI-00076 Espoo, Finland

^b Department of Computer Science, Electrical and Space Engineering, Luleå University of Technology, 97187 Luleå, Sweden

^c International Research Laboratory of Computer Technologies, ITMO University, 197101 St. Petersburg, Russia

ARTICLE INFO

Keywords:

Machine learning
Reinforcement learning
Virtual power plant
Home energy management system
Building energy management system
Forecast

ABSTRACT

A Virtual power plant is defined as an information and communications technology system with the following primary functionalities: enhancing renewable power generation, aggregating Distributed Energy Resources and monetizing them considering the relevant energy contracts or markets. A virtual power plant also includes secondary functionalities such as forecasting load, market prices and renewable generation, as well as asset management related to the distributed energy resources. Home energy management systems and building energy management systems have significant overlap with virtual power plants, but these bodies of research are largely separate. Machine learning has recently been applied to realize various functionalities of these systems. This article presents a 3-tier taxonomy of such functionalities. The top tier categories are optimization, forecasting and classification. A scientometric research methodology is used, so that a custom database has been developed to capture metadata from all of the articles that have been included in the taxonomy. Custom algorithms have been developed to generate infographics from the database, to visualize the taxonomy and trends in the research. The paper concludes with a discussion of topics expected to receive a high number of publications in the future, as well as currently unresolved challenges.

1. Introduction

The power grid is transitioning from a conventional centralized grid to a decentralized smart grid characterized by distributed renewable generation, electricity storage, and smart loads. Virtual Power Plants (VPPs) are systems for aggregating such distributed energy resources (DER), monetizing them as well as coping with the variability of renewable generation. Several reviews have addressed various aspects of VPPs from different perspectives. Liu et al. [1] reviewed the works on VPPs from an urban sustainability standpoint and identified some research gaps such as stakeholder involvement and urban environment dynamics. Yu et al. [2] identified that renewable generation, market prices, and load demand are the major uncertainties complicating the optimization task of VPPs. Adu-Kankam & Camarinha-Matos [3] studied VPPs through the lens of collaborative networks, which are heterogeneous, autonomous and distributed systems collaborating toward shared goals. Mahmud et al. [4] discussed VPPs as an enabler of an Internet of Energy consisting of prosumers. Although machine learning (ML) is a potential technology for addressing these challenges, it has not been

covered by the aforementioned recently published reviews, possibly due to the fact that an active body of research on ML applications for VPPs has emerged only over the last two years. A few very recent reviews on VPPs identify ML as an emerging technology or topic of future research in VPPs [5–7], still without performing an in-depth study of ML literature.

In this article, a VPP is defined as an Information and Communications Technology (ICT) system with the following primary functionalities: enhancing renewable power generation, aggregating Distributed Energy Resources (DERs), and monetizing the DERs considering the relevant energy contracts or markets. A VPP also includes secondary functionalities such as forecasting load, market prices and renewable generation, as well as asset management related to the DERs. Home energy management systems (HEMS) and building energy management systems (BEMS) can be seen as systems with some VPP functionalities adapted to the energy resources at a home or building context. Research in the field of HEMS and BEMS is included in this taxonomy if it fits the above definition. Examples based on ML are as follows. Maximizing PV self-consumption is a typical HEMS functionality for enhancing

* Corresponding author.

E-mail addresses: seppo.sierla@aalto.fi (S. Sierla), mahdi.pourakbari@aalto.fi (M. Pourakbari-Kasmaei), valeriy.vyatkin@aalto.fi (V. Vyatkin).

<https://doi.org/10.1016/j.autcon.2022.104174>

Received 11 November 2021; Received in revised form 8 January 2022; Accepted 12 February 2022

Available online 22 February 2022

0926-5805/© 2022 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

renewable power generation [8]. Participating in price-based demand response (DR) programs is an example of monetizing the DERs aggregated by the HEMS [8]. BEMS systems forecast building load [9] and diagnose faults in DERs [10].

Some ML-related reviews have recently been published in the HEMS and BEMS context. Wang et al. [11] focused on challenges in the deployment of ML solutions in the context of a BEMS. Fathi et al. [12] reviewed ML-based urban building energy performance forecasting applications, which is different from short-term energy consumption forecasting that would be directly relevant to a VPP. Mason & Grijalva [13] reviewed the applications of ML to control problems in building energy management. These are low-level functionalities beyond the scope of a VPP. In summary, the main focus of these reviews was not on the functionalities of VPP as defined above. The goal of this article is to fill this existing gap by presenting a taxonomy of ML functionalities within the scope of the above definition of a VPP, covering works published in the VPP, HEMS and BEMS context. Thus, this article aims to integrate relevant research being carried out in these three fields. The scope of this article is limited to the operating phase of these systems.

Most ML-related reviews in the energy domain, such as those referenced in this section, are understandable to ML practitioners. These reviews revealed that ML is emerging as a disruptive technology, so its impact and outlook needs to be understood by a broader community of energy practitioners. The primary target audience of this paper is energy practitioners who are not ML experts. Unlike the majority of reviews on ML applications to energy, this paper will not assume prior knowledge of ML, and neither will it provide any condensed presentation of ML theory. The key concepts of ML are introduced in a way that aims to be understandable to the target audience without needing to study a ML textbook. The perspective taken in this paper is the *application* of ML to realize VPP functions, rather than a detailed study of *design decisions* that were taken to implement them.

This article is organized as follows. Section 2 provides an overview of VPP, BEMS, HEMS and ML technologies. Section 3 describes the research methodology, which organizes this review as a 3-level taxonomy. Section 4 describes each tier of the taxonomy. Section 5 performs a cross-cutting analysis of the taxonomy, focusing on trends that are not evident from Section 4. Section 6 concludes the paper with a discussion of topics expected to receive significant research in the future, as well as key unresolved challenges.

2. Technology overview

2.1. Overview of virtual power plants and building/home energy management systems

The terms VPP, aggregator, HEMS (Home Energy Management System) and BEMS (Building Energy Management System) lack universally established definitions in the literature, so significant overlap exists. Fig. 1 illustrates the types of HEMS, BEMS and VPP systems investigated in this study:

- **HEMS:** The HEMS coordinates resources in a single household within the framework of an electricity contract. Examples include rescheduling smart loads to utilize cheap electricity or maximizing self-consumption of rooftop photovoltaic (PV) by means of a battery. If the HEMS is managing and aggregating several homes, it could be considered a special case of the multi-site VPP.
- **BEMS:** The BEMS is similar to a HEMS, but the scope is larger, such as a commercial building. If the BEMS involves market participation such as DR participation, it could be categorized as a single site VPP.
- **Single site VPP:** The single site VPP is similar to a BEMS but has electricity market participation capability. It may be deployed to other kinds of sites instead of buildings, consisting of one or more DERs such as the following: PV farm, wind farm, factory with smart loads, battery or electric vehicle parking lot. The DERs are geographically co-located in a single site, which may have a micro-grid (e.g. [14]).
- **Multi-site VPP:** The multi-site VPP consists of several single-site VPPs which coordinate among themselves or are coordinated by a centralized VPP. It can perform overall optimization of all of the resources at the various sites rather than local optimization, as in the case of a single site VPP. This can be done by pooling all resources regardless of their geographical location or by having a site-level local optimization function under the multi-site VPP [15]. The geographical scope of a multi-site VPP is often not explicitly defined; however, Fang et al. [15] assumed the scope of a distribution network.
- **Multi-energy VPP:** The multi-energy VPP (e.g., [16]) manages other energy carriers in addition to electricity, in particular natural gas delivered by a gas grid and heating or cooling through a hot or cold water piping network. The scope of the piping network may be on-

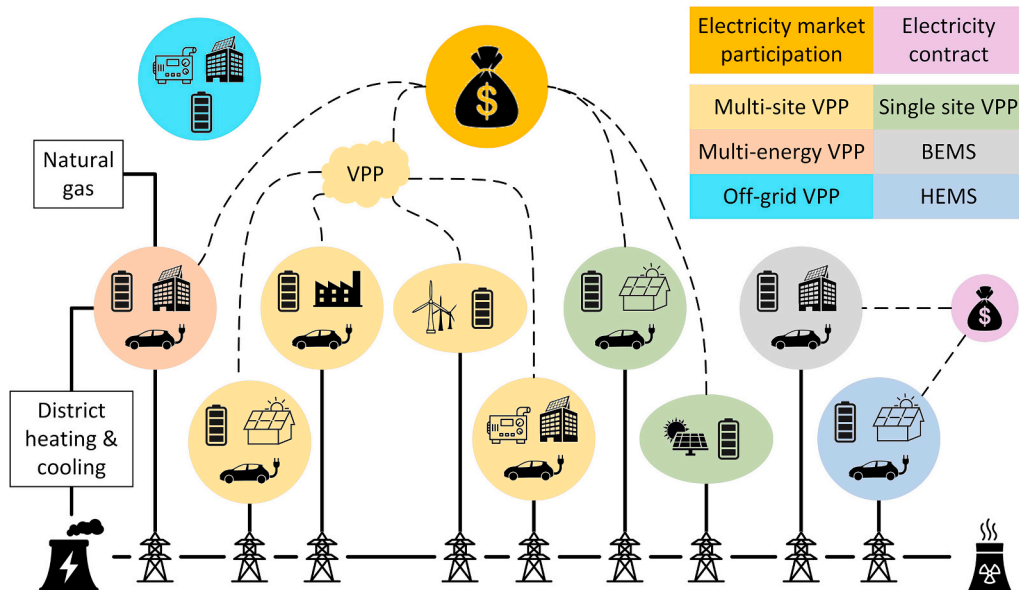


Fig. 1. Overview of different kinds of VPPs.

site in case of micro-CHP (Combined Heat & Power) generation [17], or it may be a district heating network in the case of a central CHP plant, as in Fig. 1. The potential for optimizations is highest when each energy carrier has a variable price [17].

- Off-grid VPP: The off-grid VPP does not have access to the grid, and thus there are no electricity contracts or market participation opportunities. Diesel generators are frequently available as a last resort when the other DERs cannot keep electricity production and consumption in balance [18].

Although overlapping, VPP and HEMS systems also have notable differences. This is elaborated in Fig. 2. HEMS systems are only applicable to the residential sector, whereas VPPs are also able to manage DERs in the commercial and industrial sectors. HEMS systems have IoT (Internet of Things) capabilities for interfacing to the DERs. Schedulable DERs are actively managed through these capabilities, whereas information from non-schedulable DERs can be exploited for overall optimization of the energy consumption of the home. According to Subramanya et al. [19], VPPs do not have these capabilities and rely on subaggregators. As illustrated in Fig. 2, a HEMS could either function as such a subaggregator, or it can use its smart meter to interface to a utility, independently of a VPP, to exploit variable energy prices or participate in DR. These functionalities overlap with VPPs. However, HEMS systems do not aggregate resources beyond a single home for trading on electricity markets.

Exemplary DERs are illustrated in Fig. 1. The types of DER being managed are of interest when assessing the generalizability and applicability of the proposed solutions, as well as the areas receiving high levels of research. To support the analysis in Section 5.2, a hierarchy of DER types has been defined in Fig. 3. The scope and level of detail of this hierarchy are designed for the sole purpose of capturing the types of DERs encountered in the articles that were selected for inclusion in this taxonomy. Only DERs that can be managed by VPP, BEMS or HEMS systems are included in the hierarchy. In particular, uncontrollable loads are not in the scope of the hierarchy.

2.2. Key machine learning approaches for virtual power plants

Before presenting the taxonomy, key ML concepts are presented here without assuming prior knowledge of ML from the reader.

A *reinforcement learning agent* is trained by interacting with a model of the DERs. This model is known as the *environment*. The agent is trained to optimize one or more objectives such as energy cost, PV self-consumption, battery degradation or the quality of the indoor environment. Thus, reinforcement learning is frequently used as a ML-based multi-objective optimization technique. The basic problem formulation involves a single agent. The *state* of the system captures information from the environment that is relevant for the agent to make its decision, and it can include, for example, energy prices, battery state of charge or load forecasts. The agent selects an *action*, which may be, for example, a bid on an energy market or a command to charge or discharge a battery. The action impacts the environment, which returns a *reward* to the agent. The reward quantifies how beneficial the action was. By repeated interactions with the environment, the agent learns to take actions that result in good rewards both in the short and long term.

In some cases, a single agent problem formulation may not be adequate. In the case of DERs, several buildings or several microgrids, there may be several owners or operators involved, who may want to optimize their own gain rather than the system level gain. In these cases, the *multi-agent reinforcement learning (MARL)* formulation is applicable. Each DER, building or microgrid can be managed by one self-interested agent performing local optimization, while MARL ensures that system-level optimization objectives are met. Usually, MARL involves each agent seeing the state information of all of the other agents. Notable exceptions are settings that involve bidding, in which it is not permissible for a bidder to see the state of the other bidders. The reward can be separate for each agent, or a single reward for the entire system may be used. In some cases, several types of agents are defined, with different formulations for the state, action and reward. System-wide optimization targets can be explicitly specified either by using a single system-wide reward or by defining an upper-level agent to perform the system-level optimization.

Regression approaches in ML determine a mapping from one or more

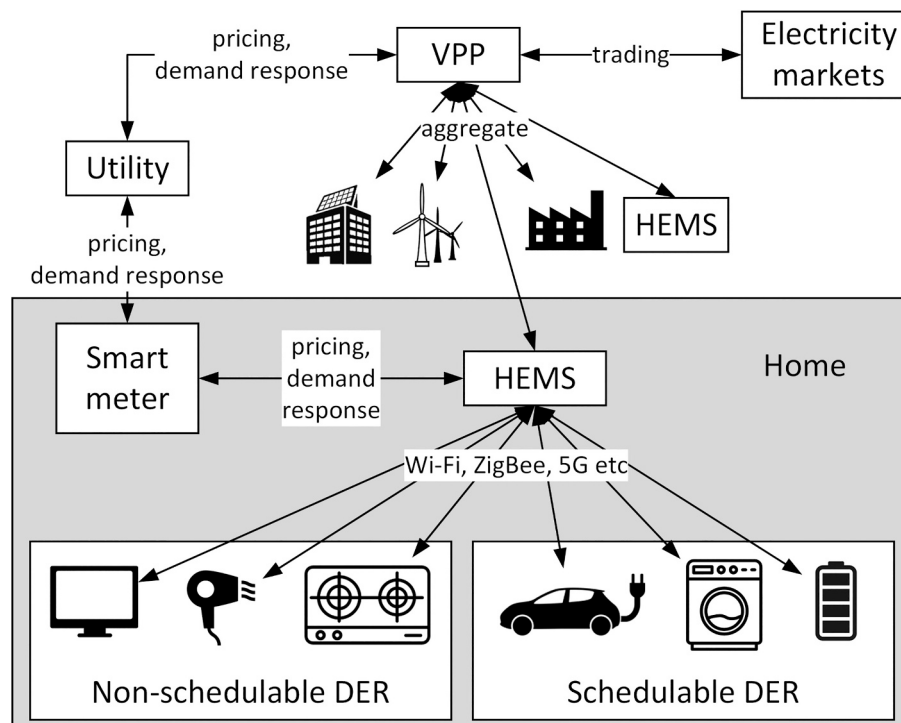


Fig. 2. Similarities and differences between VPP and HEMS systems.

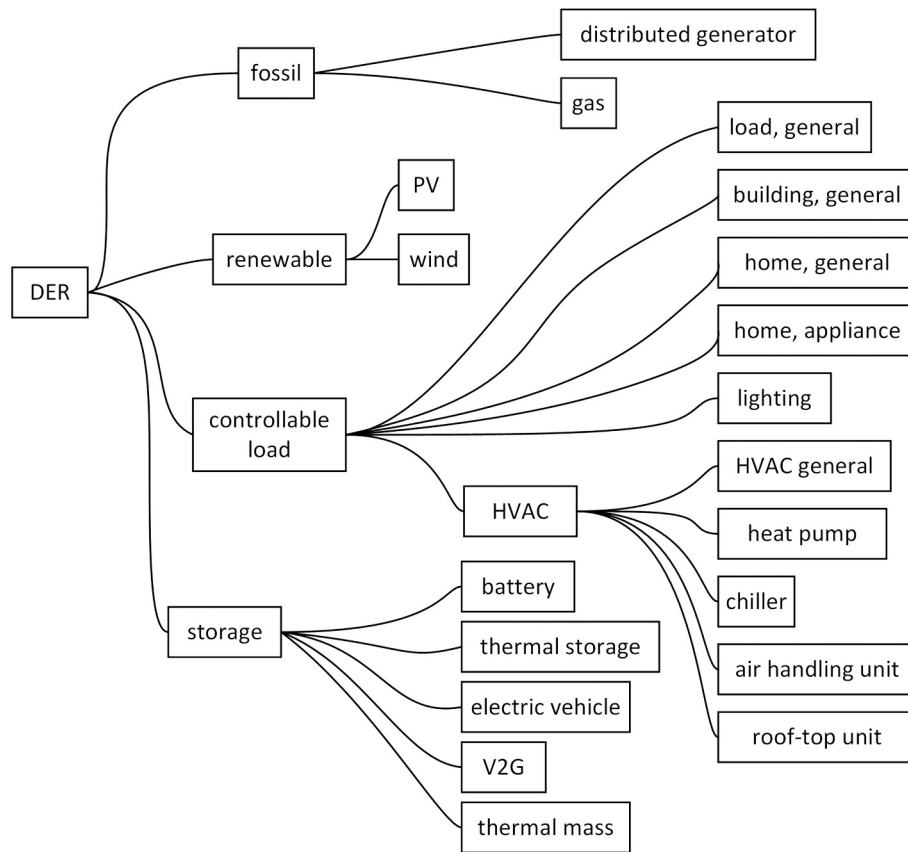


Fig. 3. Hierarchy of DERs.

independent variables to a dependent variable. In VPPs, time-series forecasting is a typical regression problem. For example, in building energy consumption, examples of independent variables are occupancy, indoor temperature and outdoor weather, and the dependent variable to be forecasted is energy consumption. The *supervised learning* technique uses a historical dataset of the independent values as well as the correct corresponding values for the dependent variable. These correct values are known as *labels*. This training set is used to fit a ML model that predicts the dependent variable when given as input the values of the independent variable. After the fitting process, the model is able to make accurate predictions with values of independent variables that were not present in the training set.

Classification methods are used to select a category from a finite set of predefined categories. In this taxonomy, classification is applied to detect or diagnose a fault state of a DER. If suitable labelled training data is available, supervised learning methods can be used to diagnose the type of fault. If labelled training data is not available, *unsupervised learning* methods can detect anomalies but cannot diagnose the specific type of fault. *Semi-supervised* learning can be used to make the best use of a training set in which only a minority of the data is labelled.

2.3. Overview of power system impacts of VPP, BEMS and HEMS systems

The ambitious goal toward carbon-neutrality has motivated researchers to focus more on energy sector coupling, and to this end, electrification plays a key role. However, electrification brings more challenges to the power system, and if not addressed adequately, it may ruin the whole decarbonization plan. VPPs, as stated earlier, can strategically provide remedies for such challenges by enhancing power system flexibility and consequently offering a better hosting capacity for renewable energy, more effective DER aggregation, richer demand response programs via HEMSs or BEMSs and deferring heavy

investments [20]. In the literature, some works have investigated the applications and effectiveness of using the VPP concept. A novel approach for the energy management system of a smart power grid was proposed by Azimi et al. [21], in which the VPP was formed by comprising fossil-fuel- and wind-based power plants, PVs, electric vehicles (EV), and DR programs. In this work, the authors showed that by considering DR in VPP, not only the operation costs will be reduced but also enhancements in peak shaving and EV's parking installations are obtained. Liu et al. [22] investigated a server client-based VPP and revealed its effectiveness in providing technical support for massive loads participating to the spot market and consequently empowering power grid regulation by considering different stages such as access, operation-planning, and settlement. Alemany et al. [23] emphasized the role of VPPs in decentralized management of DERs, especially when the interaction between the distribution system operator (DSO) and the transmission system operator (TSO) is considered. Similarly, Gorostiza et al. [24] proved that forming a VPP as clusters of EVs results in a smooth TSO-DSO interaction and consequently enhances the frequency support. Naughton et al. [25] used the concept of a VPP to facilitate the participation of DERs in the electricity market and demonstrated the capability of VPPs in supporting local networks by providing voltage and reactive power provision. Moreover, HEMSs can positively contribute to the operation-planning decisions for small-scale [26] and large-scale systems [27], aiming at increasing the flexibility of power systems as well as enhancing social welfare. In the literature, there are some works addressing the role of VPP and HEMS on power system operation and planning. Lana et al. [27] considered a low voltage DC network connected to a medium-voltage AC network and showed the importance of HEMS in providing flexibility and supporting the functionality of a network-scale battery storage network as a part of VPP. Lotfi et al. [28] studied the possible aggregation of electric vehicles and photovoltaic-equipped parking lots as a VPP. It was shown that considering HEMSs

in such VPP configuration significantly decreases the active power losses and the total energy supplied by the upstream grid without sacrificing the end-user comfort. The concept of VPP, comprising a sequential DR, was proposed by Gong et al. [29], and as a result, the ramp rate and peak power were reduced at the distribution level, while the HVAC system contributed in guaranteeing the end-users comfort level. Luo et al. [30] discussed the vision of building a VPP on distribution networks considering a central BEMS. Such a BVPP can provide adequate flexibility to support distribution system operation and enhance the effectiveness of DR, a large community of net zero energy homes, and their interconnection with the grid. Rosato et al. [31] stated the importance of ML-based approaches to develop a powerful management approach for energy clusters such as energy communities and VPP, and a reliable decision-making tool for practical power systems.

3. Research method

The search queries are presented in Table 1. One approach would have been to combine these into a single search string. In practice, it was observed that such a complex query missed some relevant papers that were found by one of the basic queries. Thus, each of the queries was performed separately in each of the databases in Table 1. The results of each query were sorted in order of relevance and the first 100 were studied manually. The number of selected papers is presented in Fig. 4. Further, conference papers were excluded from the IEEEXplore search to focus on the highest quality works. The search was limited to papers published since 2013.

The criterion for selecting an article for inclusion in the taxonomy was that it described a ML-based solution that is applicable for realizing a VPP functionality, as defined in Section 1, regardless of whether the authors use the term VPP. The selection criteria are elaborated as follows:

- The presented ML solution must be directly applicable for the purpose of implementing a VPP function. For example, day-ahead and week-ahead energy consumption forecasts are key functionalities of a VPP. However, predicting the annual energy consumption of a building [32] or predicting the Building Emission Rate (BER) resulting from retrofitting actions [33] are not generally considered to be VPP functionalities.
- Articles that discuss machine learning methods, for example in the literature review [34] or future work section [35], but which did not implement a ML solution, were not selected. Review papers were not selected; however, the most relevant review papers have been cited in Section 1.
- Articles on ML applications to low-level control (e.g. [36–38]) were not selected.
- This taxonomy covers ML applications for realizing functionalities of a VPP. ML applications for modelling electricity markets involving VPPs (e.g., [39]) are outside of the scope of this article.
- This taxonomy is only concerned with operating a VPP. Planning stage problems such as dimensioning, configuring and placing DERs [40] or selecting suitable DR customers [41] are out of scope.

Table 1
Search details.

Query	Search string	Science Direct	IEEEXplore
query1	“building energy” AND learning	4853	606
query2	“home energy” AND learning	1296	182
query3	“virtual power plant” AND learning	530	15
query4	“building” AND “reinforcement learning”	5896	227
query5	“home” AND “reinforcement learning”	2373	45
query6	building AND learning AND (diagnosis OR detection)	54,645	1292

- VPP forecasting functionalities included in this taxonomy have been restricted to forecasts that are directly applicable to the decision-making of a VPP, namely forecasting of prices, load and generation. Indirectly relevant forecasts such as building occupancy forecasting [42] are excluded. As VPPs do not act on a very short time step, works on very short forecasting intervals such as one minute [43] are excluded. Similarly, optimization approaches at such very short intervals [44] are excluded.
- This taxonomy only considers energy resources managed by a VPP. ML applications for managing other kinds of energy resources such as grid equipment [45] or batteries in IoT sensor nodes [46] are out of scope.

Several recent reviews have employed a scientometric approach to generate infographics from the metadata obtained from the publications database. Typical metadata includes the year of publication, country of affiliation, keywords, the title of the journal and cited papers. An example of such a VPP review article is [1]. This results in charts such as the number of papers per country. It would be even more interesting to obtain richer infographics from additional metadata that is specific to this taxonomy. Then it would be possible, for example, to relate the contribution of each country to each category in the taxonomy. This article employs such a scientometric research method, exploiting custom metadata collected for the purposes of this taxonomy.

The readers who are not concerned with our custom metadata management details can skip the rest of this section and proceed to Section 4. The rest of this section explains our scientometric methodology based on custom metadata that has been manually collected from the articles selected for inclusion in this taxonomy. These details are sufficient for a person with working knowledge of relational databases to repeat the methodology. However, the presentation is aimed at a general audience. Fig. 5 is a UML (Unified Modelling Language) class diagram specifying the database that was designed to capture this metadata. The steps related to the research methodology of collecting the data to this database are underlined, and the rest of the text is for the benefit of readers not familiar with the UML class diagram notation as it applies to relational databases.

- Each box is a *table* that defines the fields of information to be recorded from some entity. For example, the entity ‘Tier1’ is a top tier of our 3-tier taxonomy. The table has one *row* for each instance of that entity. For example, our taxonomy has three tier 1 categories: Optimize, Forecast and Classify, so there is one row for each of them.
- The line between ‘Tier1’ and ‘Tier2’ is a *relationship*. The notation 1 and * indicates a *one-to-many* relationship: a ‘Tier1’ category can have any number of ‘Tier2’ categories related to it, and a ‘Tier2’ category must be related to exactly one ‘Tier1’ category. Similarly, a ‘Tier2’ category can have any number of ‘Tier3’ categories related to it, and a ‘Tier3’ category must be related to exactly one ‘Tier2’ category. The result is a hierarchical 3-tier taxonomy. A tier 2 category was subdivided to tier 3 categories only if a meaningful subdivision emerged from the reviewed papers. The rows of these tables and their relationships are configured before adding the articles in Fig. 4 to the database, thus first establishing the structure of the 3-tier taxonomy.
- Each of the selected articles is inserted as a row into the ‘Article’ table. The ‘title’ and ‘year’ are fields that are filled for each article. Additionally, an article can be linked to any number of ‘Tier3’ categories and each ‘Tier3’ category can be related to any number of articles. Thus, there is a *many-to-many* relationship between ‘Tier3’ and ‘Article’, which is implemented with the *junction table* ‘Tier3_Mapping’.
- The ‘Region’ table is populated with the following entries: ‘North America’, ‘Europe’, ‘Asia-Pacific’ and ‘Rest of the World’. The ‘Country’ table is populated with the countries of affiliation of the

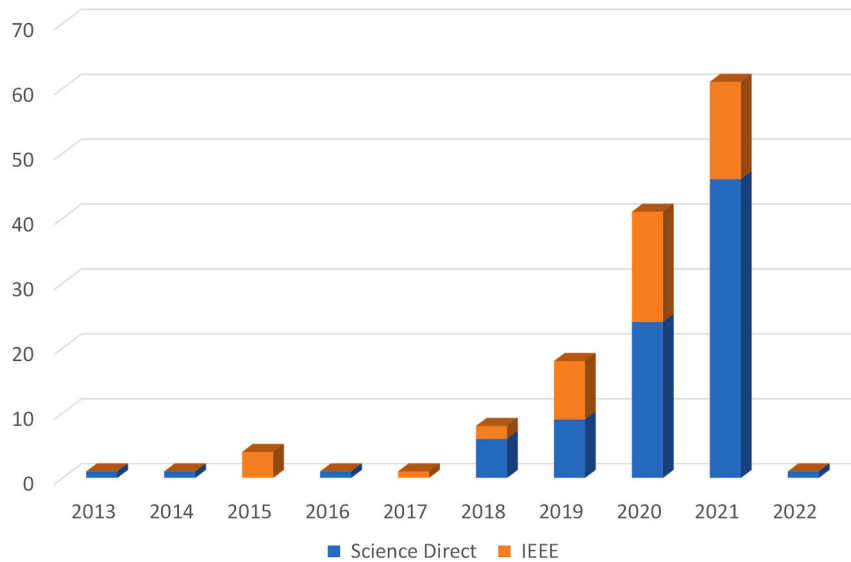


Fig. 4. Selected articles per database, by year of publication.

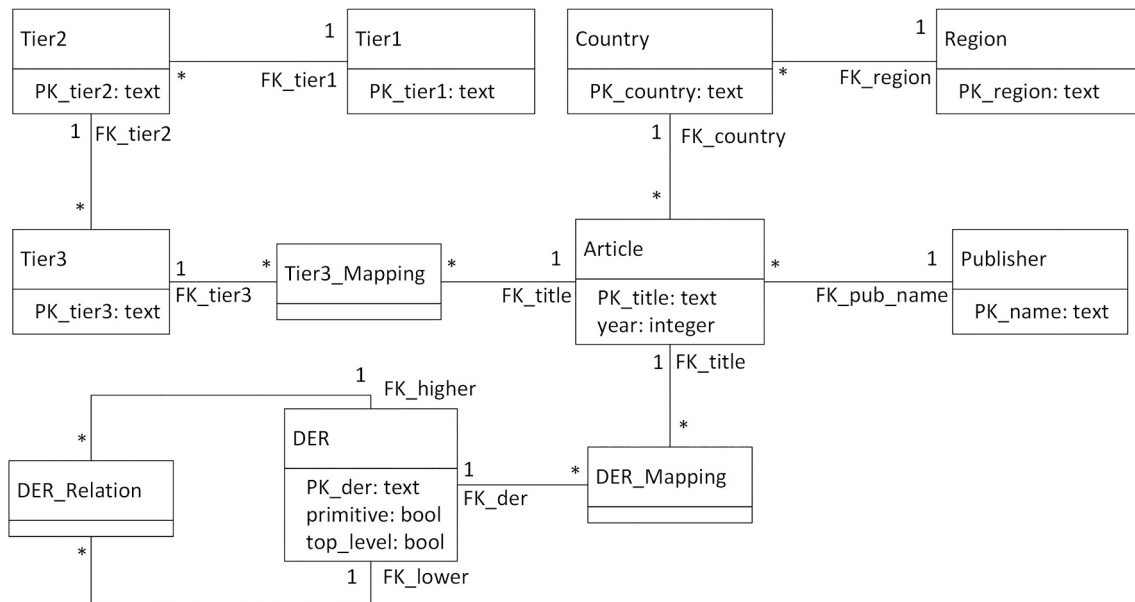


Fig. 5. Custom database for recording the metadata for this taxonomy.

first authors of the selected papers, and each country is related to one region.

- The 'Publisher' table is populated with the publishers in Table 1.
- Each article is related to one country (the country of affiliation of the first author) and to the publication database (one of the databases in Table 1).
- The DER table has one row for each DER in Fig. 3. Each DER in Fig. 3 is added as a row to the table 'DER'. Each line between DERs in Fig. 3 is added as a row to the junction table 'DER_Relation'.
- Each article in the table 'Article' is related to any number of primitive DERs by adding a row to the junction table 'DER_Mapping'.

The authors have designed algorithms and implemented them in the Python programming language for accessing the database and generating the infographics in Sections 3–6 of this article.

4. Hierarchical analysis of the taxonomy

Fig. 6 shows the selected papers in each of the tier 1 categories: Optimize, Forecast and Classify. A paper has been categorized under more than one category, if it presented ML solutions for each of those categories.

A sunburst chart of the 3-tier taxonomy is presented in Fig. 7. The angle corresponding to each category is in proportion to the number of papers in that category. The remainder of this section is structured according to the hierarchy in Fig. 7.

The Sankey chart in Fig. 8 maps the countries of affiliation of the first author to the tier 3 categories. For most categories, the global academic community is engaged in the work. In a few cases, a regional focus is evident: for example, although Europe contributes about 30% of the research overall, it contributes the majority of the research on solar forecasting. The contribution of the country with the largest number of publications, China, is highlighted in red. (See Fig. 8.)

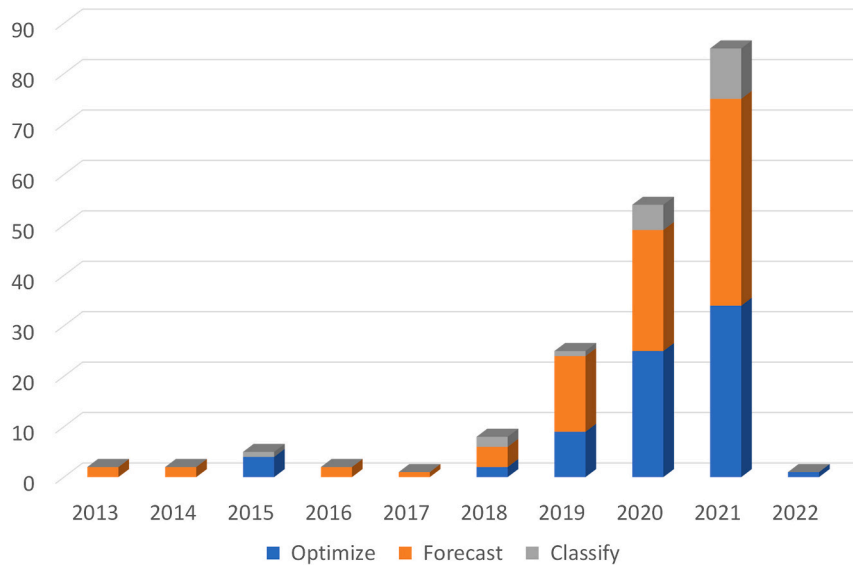


Fig. 6. Selected papers in each tier 1 category.

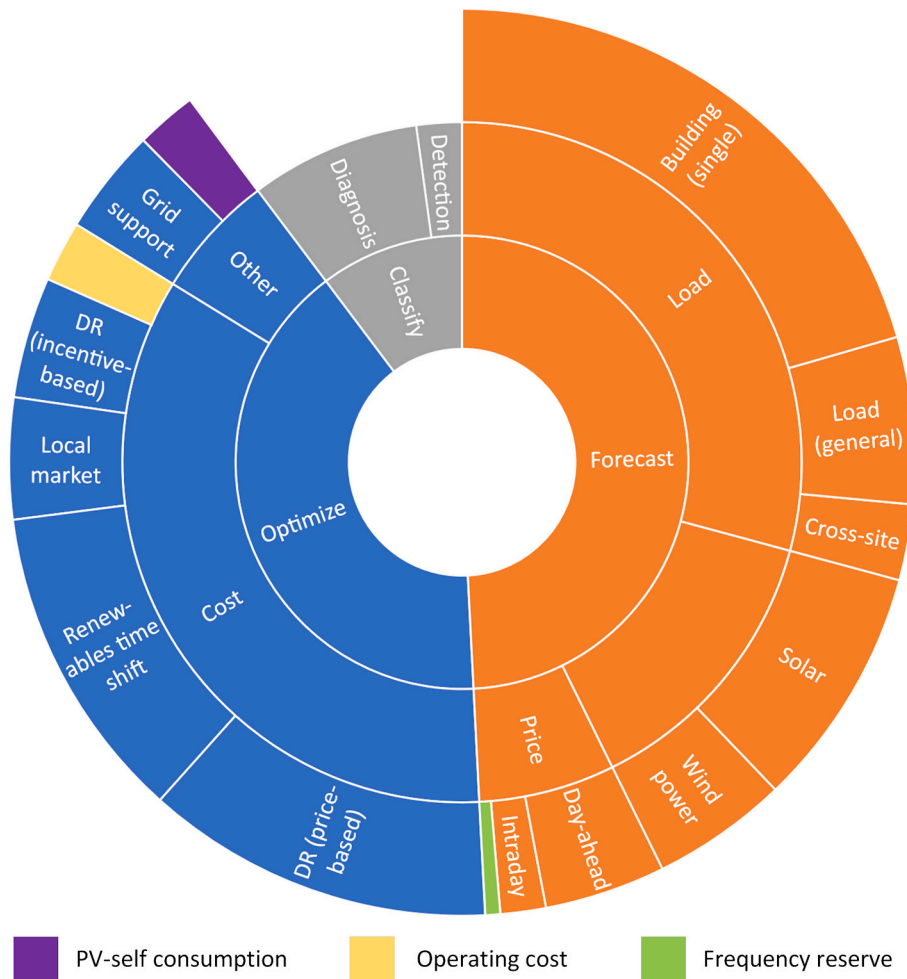


Fig. 7. Sunburst chart of the 3-tier taxonomy.

4.1. Optimize

Fig. 9 presents the breakdown of publications into tier 3 categories under 'Optimize'. The years 2019–2021 indicate a possible shift of

attention from price-based DR to incentive-based DR. A steady growth of publications in 'grid support' can be seen over the last three years. However, the 'Local market' category does not show evidence of significant growth, suggesting that the main thrust of the research remains

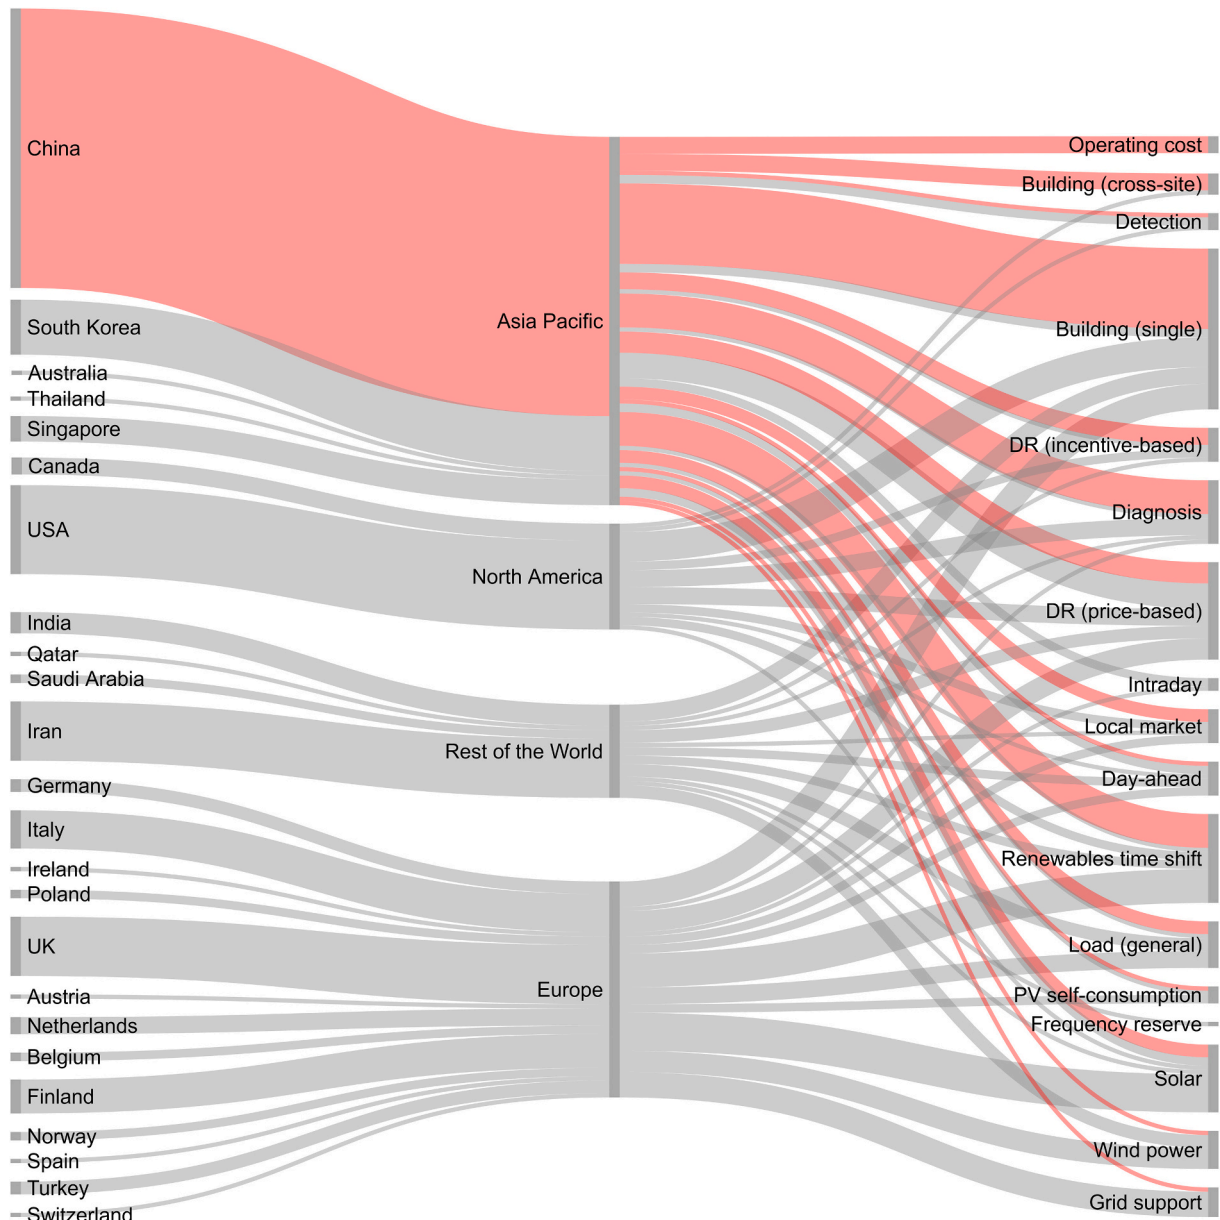


Fig. 8. Sankey chart mapping the countries of affiliation (on the left) to the tier 3 categories of the taxonomy (on the right).

on existing markets. Unless otherwise specified, all of the works in the 'optimization' category of the taxonomy use reinforcement learning.

4.1.1. Cost

4.1.1.1. Renewables time shift. If there is a possibility to sell to the grid, a storage resource can be used to shift sales of renewable energy to hours with higher prices. The basic formulation of this problem involves PV [47,48] or wind [49] and a battery, so the optimization aims to maximize profits under a variable electricity price. In a multi-energy VPP, renewable generation can be complemented with local fossil fuel generation [50,15]. The local generation can be within a building [51] or a microgrid [52]. These works assume that the utility has the same price for buying and selling electricity. Other authors use separate buying and selling prices under a Time-of-Use tariff [53,54] and a real-time pricing scheme [55–57]. Nakabi & Toivanen [14] additionally consider transmission costs. Although the majority of works use stationary batteries, Vehicle-to-Grid (V2G) [17,58] and heat storages [59] are other possible storage technologies. If the market requires that PV or wind generation

capacity be traded ahead of time, reinforcement learning can be used to optimize the operation of the battery to cope with inaccuracies in the forecast [60].

4.1.1.2. DR (price-based). Price-based DR involves curtailing or rescheduling energy consumption based on electricity prices. A utility may set the prices with the intention to reduce consumption during anticipated times of peak load [61]. The main approaches involve managing Heating, Ventilation, and Air Conditioning (HVAC) systems so that the indoor environment remains within acceptable limits [62] or managing home appliances, which impacts occupants more directly [63]. Adverse impacts to occupants are usually handled by adding a penalty term to the reward of the reinforcement learning agent.

Most price-based DR applications involve HVAC systems. Schreiber et al. [64] shift operations of a chiller to low price periods while penalizing according to deviations from an ideal temperature. Several authors minimize the electricity cost for an HVAC system, while minimizing violations of indoor temperature [65] and air quality [66] requirements. Ren et al. [8] and Li et al. [67] additionally include an EV

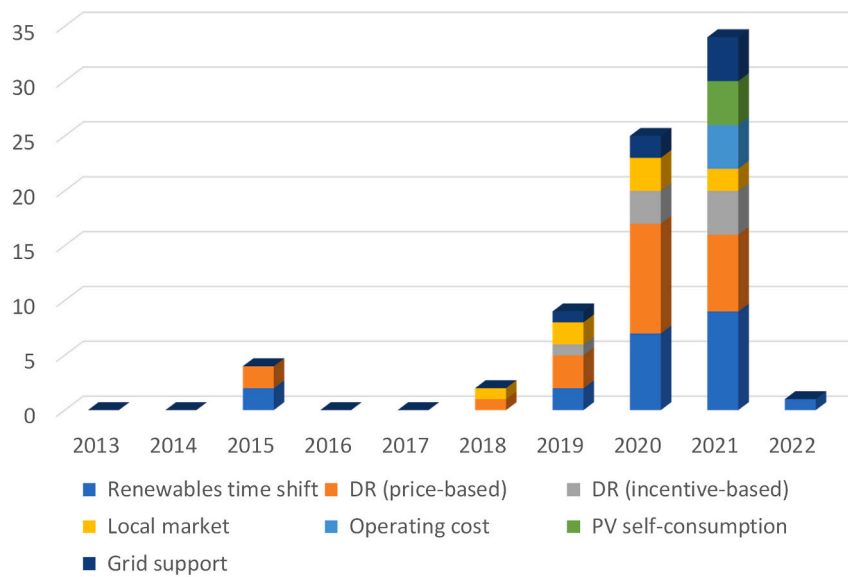


Fig. 9. Publications in the tier 1 category 'Optimize', color-coded according to the tier 3 categories under 'Optimize'.

charger and also minimize user dissatisfaction resulting from the EV not being charged as planned. Proactively pre-heating or pre-cooling before a high price period reduces HVAC electricity consumption in that period [68,69]. In all of the works involving HVAC, the thermal mass of the building is exploited as a thermal energy storage, thus avoiding the need to invest in any separate energy storage DERs. Most works assume real-time electricity pricing, but Jiang et al. [69] assume a Time-of-Use tariff and peak demand charges. Whereas the other works use reinforcement learning, Kim [70] uses supervised learning to determine an optimal operating schedule for an air handling unit. A conventional optimization framework is used to calculate the optimal schedule for a historical time period (in which there is no uncertainty), and these values are used as the labels to train the supervised learning model.

A variety of approaches are in use for capturing the user discomfort for DR involving home appliances. Chellamani & Chandramani [71] reschedule appliances to lower price hours under an Hourly Time-Of-Use tariff, and user discomfort is mitigated by learning the users' consumption patterns. Bahrami et al. [72] define discomfort as the time difference between the desirable and rescheduled operation time of the appliance. Zhang et al. [61] reschedule or curtail home appliances and model the dissatisfaction arising from these actions. Lee & Choi [73] define a separate reinforcement learning agent type for each kind of DER, so that each type of agent has its own reward function that captures user dissatisfaction in a way that is relevant for the type of DER. Xu et al. [74] take a simpler approach with one function that captures dissatisfaction, with a coefficient that can be adjusted for each appliance type. Chung et al. [75] capture the user's far-sightedness as the user's willingness to take DR actions for which the benefits are realized several hours later – however, it is debatable whether this is a relevant consideration in a system that has automated the DR decisions. Generally, customers have the option to override DR requests, and Wen et al. [76] argue that it is reasonable to assume that customers will make such decisions without considering price information. Usually, DR involves some form of peak shaving; however, Alfaverh et al. [77] additionally perform valley filling to shift consumption to times of low load.

An alternative to using reinforcement learning is to use conventional optimization techniques with historical data to obtain labels to train a supervised learning model for deciding the optimal DR actions [66].

4.1.1.3. DR (incentive-based). Incentive-based DR involves a financial incentive mechanism that allows finer control of the users' DERs. Only a minority of works consider the industrially accepted international

standard OpenADR 2.0, which standardizes many aspects of automating incentive-based DR requests [79]. A straightforward approach is a contract between the consumer and the utility, permitting the utility [80] or VPP [81] to send requests to consumers to limit their energy consumption in a specified time window. In a commercial building, Zhang et al. [82] ensure that, upon receiving a DR signal, the required power reduction is achieved by separately adjusting air conditioners in different parts of the building to minimize discomfort. In a residential setting, incentive-based DR gives tools to provide unique incentives for each customer [83,84]. Kumari & Tanwar [85] have a reinforcement learning agent that sets monetary incentives in real-time for the users, while taking into account the wholesale market price. Kuang et al. [86] point out that psychological factors will be significant if customers are directly involved in the decision-making, so they define a risk attitude to capture this in the reward function of the reinforcement learning agent.

4.1.1.4. Local market. MARL is a popular approach for defining novel local energy markets. Zhang et al. [87] manage queueing at an EV charging station with an auction market operated by the station. Zhu et al. [88] operate an auction market for microgrids to aggregate the DERs across these microgrids. Each microgrid has its own agent with its own reward, aiming to maximize the total revenue for the microgrid. The system-level reward is the sum of these rewards. In contrast, Samadi et al. [89] define different types of self-interested agents for consumers, energy suppliers and batteries, and an upper-level energy management agent that trades with the lower-level agents and maximizes system-level reward. With goodwill from the utility, a local market can use the utility grid to realize physical energy trading [90,91]. Zhou et al. [92] define a local market for an energy community, in which local PV production can be sold at a price that is higher than the Feed-in-Tariff and lower than the retail price. The works in this category do not discuss obstacles related to a possible lack of cooperation from the utility, and neither do they motivate their work from the perspective of the utility.

4.1.1.5. Operating cost. A few miscellaneous works were found that minimize some aspects of operating cost but do not fit under renewables time shift, DR, or local markets. Wei et al. [93] exploit a battery at home with a real-time electricity pricing contract, charging during low prices and powering home appliances with the battery during high prices to minimize the electricity bill. Fang et al. [94] propose a concept that is otherwise similar to renewables time shift, but they assume constant

electricity prices and seek to minimize the cost of interaction with the distribution network, without elaborating how this cost is formed. Zhao et al. [95] minimize the operating cost of a multi-energy VPP managing office buildings that can be heated either by district heating or an electricity-powered boiler under a Time-of-Use electricity tariff. The district heating network operator benefits from a more stable load to the network, but the financial compensation mechanism to the VPP is unclear. Qin et al. [96] minimize the operating cost of an off-grid VPP with PV, battery storage and flexible loads, with fossil fuel-based distributed generators available if the load cannot be covered with the green DERs.

4.1.2. Other

4.1.2.1. PV self-consumption. Maximizing the self-consumption of local PV involves directing surplus generation to a storage resource, from which it can then be used later when the consumption exceeds PV generation. The storage can either be a stationary battery [8], thermal storage [97] or electric vehicle battery [98].

4.1.2.2. Grid support. A few works support the utility grid without specifying a mechanism for financial compensation for these services. Mbuwir et al. [99] use household batteries for congestion management in a microgrid, which they define as minimizing the net power exchange with the utility grid. Qiu et al. [58] apply MARL to operate a peer-to-peer market for homes with PV, batteries, V2G and smart loads. Each home is penalized for its contribution to the community's peaks. Shang et al. [100] penalize for power tracking errors at the point of common coupling between the microgrid and distribution grid. Tchnitz et al. [101] reschedule EV charging to times of low grid load but do not compensate the EV owners for this inconvenience. Totaro et al. [18] help an off-grid VPP to cope with a battery failure and to increase production from a fossil fuel generator without manual intervention. In a district heating grid, Solinas et al. [102] shave peak loads for the central CHP plant while keeping indoor temperatures within limits.

4.2. Forecast

As presented in Fig. 7, forecasting applications can be categorized under the tier 2 categories of market price forecasting, renewable generation forecasting and load forecasting. The tier 3 categories under these and the numbers of publications in each of them are presented in Fig. 10. As can be seen, steady growth is noted under most categories. However, although forecasting of the loads of single buildings in the

category 'Building (Single)' continues to attract a large amount of research, the exploitation of data across buildings for forecasting purposes in the category 'Building (cross-site)' has not yet emerged as a growing body of research. As is discussed further in Section 4.2.1.2, there are strong motivations for cross-site applications when the limitations of real-world data availability are taken into account. However, the reviewed research is not yet aiming at online deployments to real buildings, so these data availability issues are not yet a concern. The data for 2021 shows that the forecasting of load, in general, is attracting faster growth than the forecasting of load in buildings. This could be an indication of a shift of focus from BEMS and HEMS to VPPs. The latter are concerned with loads that can be abstracted and pooled so that several types of DER can be traded on energy markets as one large resource. This is discussed further in Section 5.2.

With one exception, the forecasting works referenced in this section employ supervised learning. However, the reinforcement learning technique that has been mainly used for optimization has also been applied for time series forecasting, so that the reward to be minimized is the difference between the prediction and the actual value [103].

4.2.1. Load

4.2.1.1. Building (single). Forecasting building energy consumption is a popular research problem. An hourly resolution is most common, although other resolutions such as half-hourly [104] and 15-min interval [105] are encountered. The forecasting is usually done day-ahead. The most common input data is weather data and historical energy consumption data [106], with some authors including additional data such as building automation system sensor data. Several authors present a solution based on data collected from physical buildings [107–113,9,114–118]. To select the most relevant nonredundant sets of sensor data, Eseye et al. [119] propose a feature selection method and validate its generality for building load forecasting. Tian et al. [120] generate additional artificial training data with a distribution that is similar to the sensor data from the building. The usage of such raw data has a bias, with the majority of training samples coming from working days, so Wang et al. [121] train separate models for working and rest days, whereas Zhang et al. [122] use a clustering decision tree algorithm to identify the different operating conditions to obtain a multi-model predictor that overcomes this problem. However, a bias can exist for various operating conditions, so Zhang & Wen [123] use active learning, in which the building energy management systems setpoints are actively changed to obtain the desired data from different operating conditions.

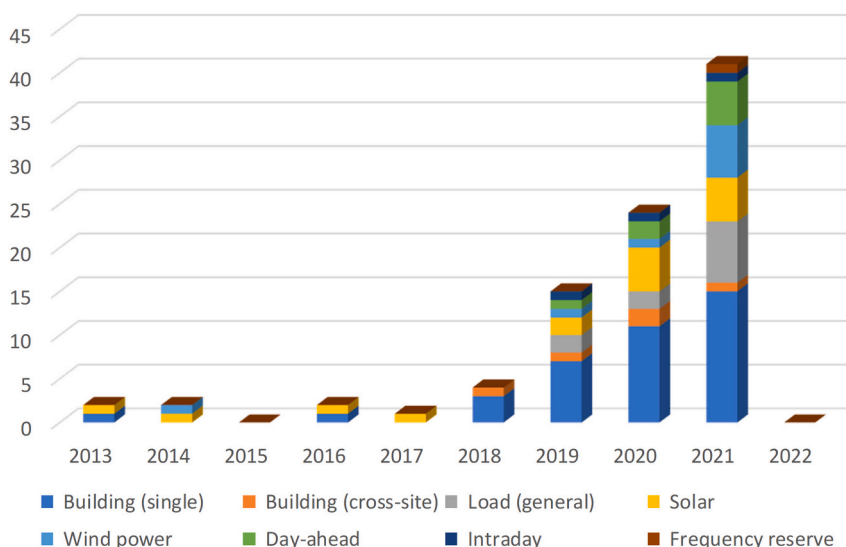


Fig. 10. Publications in the tier 1 category 'Forecast', color-coded according to the tier 3 categories under 'Forecast'.

A practical disadvantage of this method is that building operators and inhabitants may be unwilling to allow such experimentation for data collection purposes. Such problems can be avoided by using a building energy simulator, which can generate the training data for supervised learning purposes without disturbing systems and inhabitants in a physical building [124–126]. Kim et al. [127] note that this approach has the benefit of avoiding privacy issues in collecting occupancy-related data.

If the VPP is managing individual apartments, as in a HEMS context, the load is forecasted at the level of an apartment [128–130]. The forecasting of specific types of loads is possible when data is available. Datasets from volunteer homes have been used to forecast cooling load [131] and appliances load [132]. However, deployment of the solution beyond the volunteer homes may be problematic if such data is not available at the premises. Aurangzeb et al. [133] note the computational infeasibility of training forecasting models for each home, so they cluster homes with similar consumption profiles. Thus, one ML model can be trained for each cluster.

In residential contexts, the available data is usually smart meter data, which is the aggregated consumption of all of the loads. Numerous papers have proposed disaggregation methods to extract the consumption of individual loads such as lighting or a specific type of appliance. Disaggregation in itself is not a VPP functionality. However, a VPP could perform disaggregation to obtain a richer set of input features for forecasting [134] and a HEMS could do it to forecast the consumption of specific appliances [135,136].

In addition to the above-mentioned works on electricity consumption, Eseye and Lehtonen [137] propose a forecaster for district heating energy consumption, which is of relevance to the multi-energy VPP.

4.2.1.2. Building (cross-site). Cross-building forecasting can overcome the problem of inadequate historical data at the building of interest, as most buildings have deployed smart meters only recently [138]. Transfer learning is applied such that the training of the forecaster is started with data from other buildings and finalized with the available data from the building of interest [139,140]. Ma et al. [141] specifically address different scenarios for missing data, namely: random missing, continuous missing, and large proportionally missing. Ribeiro et al. [138] remove seasonal and trend components to improve the applicability of the historical data across buildings. Xu et al. [142] use social network analysis to predict energy consumption for buildings based on nearby reference buildings for which there is an accurate machine learning forecaster.

4.2.1.3. Load (general). A typical difference between a BEMS and a VPP is that the former predicts the load of a building and the latter often predicts an aggregated load, while being agnostic to the types of DERs that comprise the load [143,31,144,87]. A VPP also handles loads not generally in the scope of BEMS, such as V2G chargers [78]. The load forecasts reviewed in this taxonomy are, in general, offline applications, so practical issues may need to be solved to obtain online applications [89], as required by a VPP.

4.2.2. Generation

Renewable generation forecasts in a VPP context are usually day-ahead, as that is usually a sufficient timeframe for making decisions on energy markets. The works reviewed in this section involve day-ahead or week-ahead prediction.

4.2.2.1. Wind. Wind forecasting capabilities include wind speed [143] and wind power generation [16,90,144,145]. Sharifzadeh et al. [146] provide a thorough discussion on the pros and cons of each approach. Ultimately, a VPP is concerned with power generation. Nabavi et al. [147] forecast wind speeds and convert these to wind power generation estimates based on data sheets of the turbine and grid-tie inverter

manufacturer. However, such datasheets are based on tests in controlled conditions of a wind tunnel, whereas the power generation of a wind farm depends on wind speed and direction at several locations and altitudes.

4.2.2.2. Solar. PV generation related forecasts are either solar irradiance forecasts [145] or PV power generation forecasts [16,90,148]. Most forecasts are hourly, although there can be great variations during a one-hour period if the cloud cover changes. Rapidly changing cloud cover can cause very rapid fluctuations in PV generation, so a longer interval can even out these fluctuations and reduce the forecasting error. Thus, Hafiz et al. [92] argue that a 15-min interval is reasonable due to considerations of data availability and computation. This is also a reasonable interval for a VPP to operate. Suresh et al. [149] forecast at 10-min intervals, and unlike most works, do not use solar irradiance forecasts but rather the measured irradiance and other locally measured data from the two previous timesteps. PV forecasts with significantly shorter forecasting intervals, known as nowcasting (e.g., [93]), have not been included in this taxonomy, since VPPs generally do not operate at such a very short timescale. Forecasting PV generation at a single site is error-prone, and improved accuracy can be achieved by sharing data among the forecasters of all the generation sites under a multi-site VPP [150,24].

4.2.3. Price

4.2.3.1. Day-ahead. In many countries, day-ahead electricity trading occurs at hourly intervals. VPPs can optimize electricity bills and revenues from generation by considering the hourly changing prices and possibilities to curtail or reschedule generation or consumption or use a local energy storage. Thus, forecasting the day-ahead price is a key capability of a VPP [143,144].

4.2.3.2. Intraday. Mashlakov et al. [90] present a forecaster that can be adapted for both day-ahead and intraday market forecasting. A VPP can trade on an intraday market to manage the effects of day-ahead decisions that turned out to be suboptimal, due to the uncertain information that was available day-ahead. Intraday markets vary between different countries. However, they often involve a price that changes every half hour [95] or hour [96]. Forecasts of such prices for the next market interval are useful for a BEMS [95] or HEMS [96] that manages flexible loads to optimize the energy bill.

4.2.3.3. Frequency reserve. Frequency reserve markets compensate DERs for standing by and modifying their power generation or consumption in response to power grid frequency deviations. Several types of reserve markets with different activation requirements and trading rules usually exist in a single region. Further, the markets differ across geographic regions and are usually operated by an ISO (Independent System Operator) or a TSO (Transmission System Operator). Sadeghi et al. [144] forecast several reserve markets operated by the California ISO.

4.3. Classify

Fig. 11 presents the research in tier 2 categories under the tier 1 category ‘Classify’. Due to the limited number of articles, the tier 2 categories were not further decomposed to tier 3 categories. The reviewed articles detect or diagnose faults that do not prevent DERs from performing its function, but which reduce their capacity to generate energy [151] or adjust consumption [97], thus reducing the capacity available for a VPP, HEMS or BEMS system.

4.3.1. Detection

All of these works in this category detect that the DER is in some kind

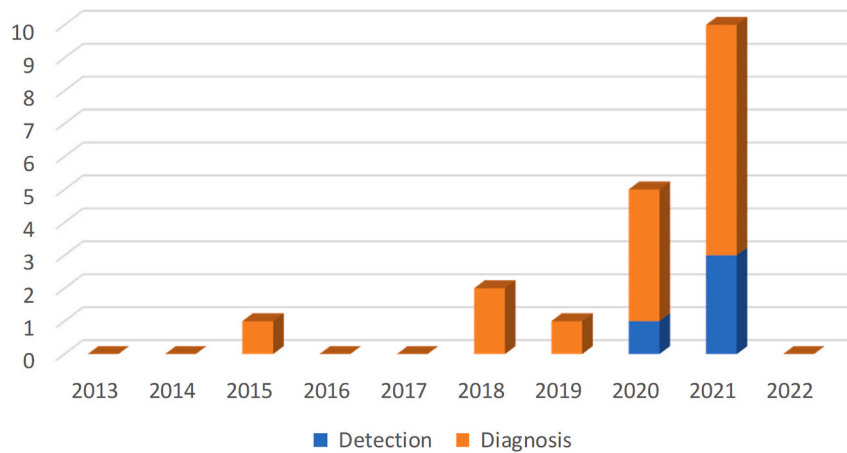


Fig. 11. Publications in the tier 1 category 'Classify', color-coded according to the tier 2 categories under 'Classify'.

of abnormal condition, but the detection methodology is not able to identify what that condition is. Choi & Yoon [25] use unsupervised learning to detect whether or not an abnormal operating condition exists in the district heating water substation system of a building. Hosseini et al. [99] perform anomaly detection for home appliances with supervised learning, due to concerns about the accuracy of unsupervised learning methods. The detection is done by observing the energy consumption. Li et al. [100] detect sensor faults in ground source heat pump systems. Yun et al. [152] perform fault diagnostics of an air handling unit with supervised learning. However, they note that in real operational systems, fault conditions may occur that do not correspond to any of the preconfigured fault types, in which case the supervised learning system is prone to perform a misclassification. To avoid this, the authors detect an undefined fault.

4.3.2. Diagnosis

Fault diagnosis methods are able to select one of a number of pre-configured categories for a DER. The categories include the normal category and several different fault categories. For example, Han et al. [10] detect faults in a water-cooled chiller and diagnose them as refrigerant leak/undercharge, condenser fouling, reduced condenser water flow, non-condensables in the refrigerant, reduced evaporator water flow, refrigerant overcharge and excess oil. Liu et al. [101] address the same problem and use the following fault categories: leakage, overcharge or reduced evaporator water flow; each fault has two severity levels. Guo et al. [102] diagnose the following faults for an air-source heat pump system: four-way reversing valve fault, the outdoor unit fouling fault, refrigerant undercharge and overcharge faults. A number of similar works exist for rooftop units [26], boilers [153] and air handling units [154]. Gharsellaoui et al. [104] diagnose faults unrelated to equipment malfunction, namely: unexpected occupancy and opening windows when the heating is on.

Whatever the type of DER, the ML model for performing the diagnosis can be trained either with actual faults from operational DERs captured in building automation systems [105], data from fault scenarios generated with an energy simulator [106] or faults that were physically created in the experimental setup and measured by the available sensors [107]. The problem with collecting real fault data for operational DERs is that getting a large training set of data from different fault conditions is very difficult. The other approaches for artificially generating the data by inserting faults into simulators or experimental setups are labor-intensive and may not reflect real operating conditions. In order to obtain solutions that can be scaled up from individual DERs or buildings, with acceptable manual effort, transfer learning and semi-supervised learning have been used.

Transfer learning involves training a ML model with data from other DERs with rich data and then completing the training with the DER of

interest, for which limited data may be available. This has been done for chillers [101,155].

In real-world DER deployment, large quantities of unlabeled data are available from automation systems, but generally, only a relatively small quantity of this is labelled, as the labelling requires expert workforce. Several authors have argued that semi-supervised learning is well suited in this environment. Zhao et al. [151] use semi-supervised learning to generate a large, classified dataset of faults for a PV array. The PV-array's current and voltage measurements at the maximum power point are received from the inverter. In order to achieve better clustering, these are normalized using current and voltage measurements from a PV reference module. The contribution of the semi-supervised learning technique is that unlabeled datapoints within a cluster are given the same label as a labelled datapoint in that cluster. Li et al. [97] train a generative adversarial network to generate fake samples with a similar distribution as labelled as well as unlabeled real samples of HVAC data. An innovative fault diagnosis classifier training method is proposed, which combines the labelled, unlabeled and generated samples. Fan et al. [108] and Yan et al. [156] first train a ML model with the available labelled data from an air handling unit, using supervised learning techniques. The trained model is used to label the unlabeled data, which is then added to the set of labelled data, but with a lower confidence value. The training process is iterated with this expanded set of labelled data.

5. Cross cutting analysis of the taxonomy

5.1. Combination works in involving several categories of the taxonomy

The heatmap in Fig. 12 presents the number of articles that were categorized under several tier 2 categories. For example, the intersection of the row 'Forecast: Generation' and the column 'Forecast: Load' has a value of 11, meaning that 11 articles were categorized as presenting solutions for forecasting load and generation. The heatmap is symmetric across the diagonal from the top-left corner to the bottom right corner. Since the values throughout the top two rows are zero, none of the detection or diagnosis works are being exploited in conjunction with forecasting or optimization functionalities. Rows 3–5 and columns 3–5 are the load, generation and price forecasting works. The intersection of these rows and columns shows that a significant number of researchers have developed general-purpose forecasting solutions. The second row from the bottom shows that some cost optimization works are exploiting forecasts of load, generation and prices.

The heatmap in Fig. 13 is similar to Fig. 12, but at the level of tier 3. The intersection of the 'Forecast: Generation: Solar' column with the 'Forecast: Generation: Wind power' row is the brightest yellow square in Fig. 13, meaning that this was the most common combination of any two

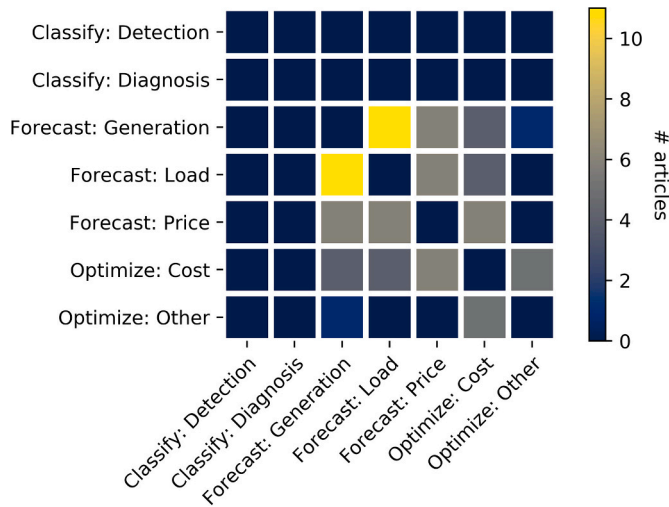


Fig. 12. Heatmap counting the number of articles that were classified under two different tier 2 categories.

tier 3 categories. This combination was encountered in the following works: [16,145], and [31,146], which additionally forecast 'Load (general)'. Another common combination is the forecasting of 'Solar' and 'Building (single)' [92,147,128,148]. Both forecasts are obviously useful for any further work on exploiting rooftop PV locally in a building. A few works forecast renewable generation, load and market prices [143,90,144], demonstrating the potential for versatility.

In Fig. 13, there are some combinations of DR or renewables time shift with forecasting approaches. Price-based DR approaches have

made use of forecasts of load [78], the day-ahead price [65] and intraday price [96]. Incentive-based DR has been optimized using forecasts of day-ahead price and load [83,84]. Wen et al. [84] additionally forecast PV generation. Harrold et al. [49] optimize renewables time shift using forecasts of day-ahead price, load, wind and PV generation.

In Fig. 13, the most common combination of two different optimization approaches is between renewables time shift and price-based DR. The basic concept involves exploiting the flexibility of DR capable loads to reschedule surplus PV sales to the grid to a more profitable market interval [55]. Liu et al. [54] make an unpractical assumption that consumption and PV generation are known exactly ahead of time. Xu et al. [74] is similar, with an hourly changing real-time price which is forecasted, so that the forecast is made available for the reinforcement learning agent performing the optimization. Systems with price-based DR and renewables time-shift capability have been additionally categorized under grid support if they perform peak shaving [56,58] and under a local market when such a market mechanism was proposed [89].

Some miscellaneous approaches combining tier 3 categories are noted as follows. When there is no possibility to sell surplus generation to the grid, one can optimize DR and PV self-consumption [8]. This approach is more effective with a PV generation forecast [73]. Price-based DR has been applied in innovative local markets for EV charging stations [87] and a VPP that aggregates several HEMS systems [75]. Shang et al. [100] propose a renewables time shift solution for microgrids, in which there are constraints for power flow at the point of common coupling (PCC) between the microgrid and the utility grid. A form of grid support has been designed that mitigates power tracking errors at the PCC.

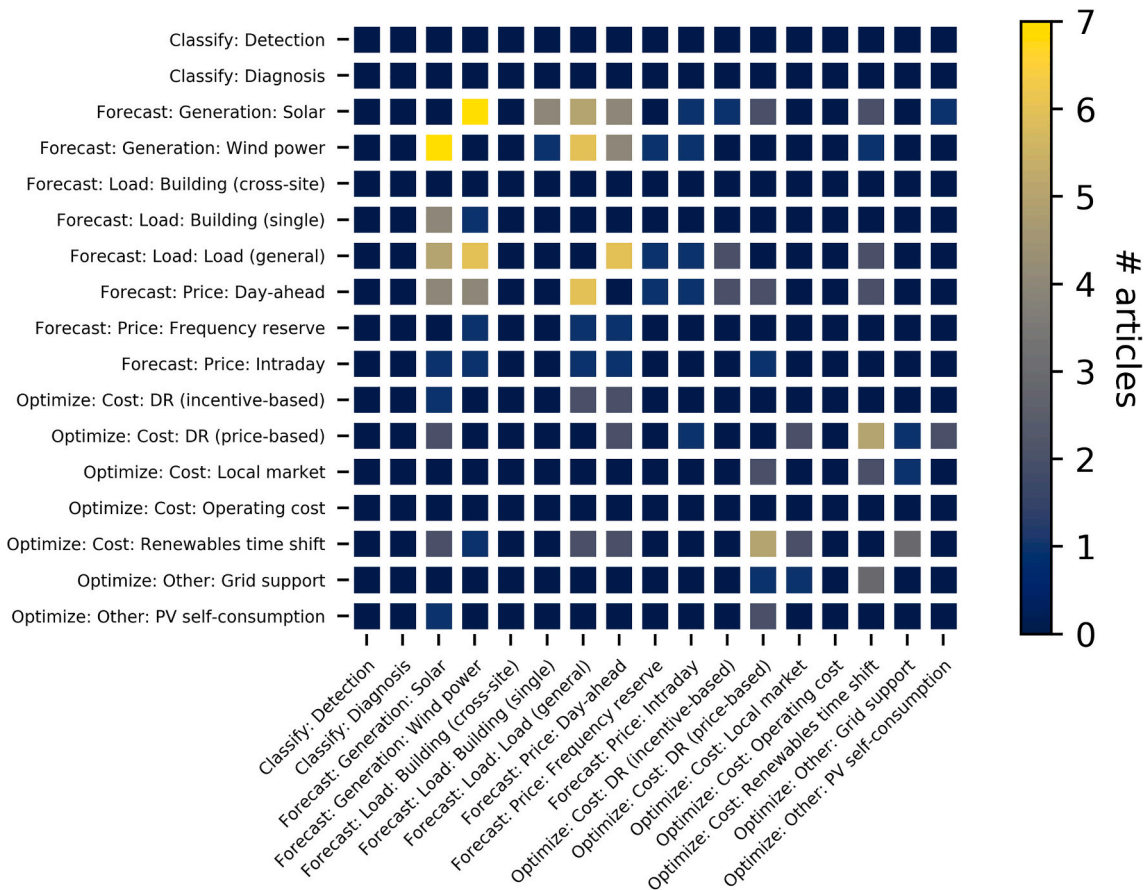


Fig. 13. Heatmap counting the number of articles that were classified under two different tier 3 categories.

5.2. Mapping of the taxonomy to types of distributed energy resources

The Sankey chart in Fig. 14 relates the tier 3 categories on the left to our hierarchy of DERs (Fig. 3). The width of each flow from a tier 3 category to a top-level DER type (fossil, renewable, load or storage) is proportional to the number of articles in that tier 3 category, which involved such a DER. Flows toward fossil fuel burning DERs are colored in gray in Fig. 14, whereas the other DERs are considered green energy solutions.

Fig. 14 reveals the level of abstraction at which loads are being modelled. The most abstract level ‘Load, general’ has been employed in a moderate number of publications. The majority of publications use a medium level of abstraction: ‘Hvac, general’, ‘Building, general’ or ‘Home, general’. A moderate number of articles were categorized under ‘Home, appliance’, which involves constraints on how the loads can be curtailed or rescheduled. A few articles addressed specific types of HVAC equipment such as chillers, air handling units, heat pumps or rooftop units. The highest level of abstraction (‘Load, general’) is required to enable a VPP to aggregate the DERs. The level of abstraction of data

modelling in a commercial VPP is discussed in detail by Subramanya et al. [19]. Sub-aggregators must handle the special requirements that constrain the use of particular types of DER. BEMS and HEMS systems are strong candidates to serve as such sub-aggregators, but further research is needed to integrate them into VPPs. Chung et al. [75] is a good example of a VPP that has been integrated into HEMS systems in a sub-aggregator role. Although a VPP needs to access DERs through an interface at a high level of abstraction, the underlying forecasting, optimization or classification solutions need to model the DERs at an appropriate level of detail. As an example of inadequate detail, a reinforcement learning environment for optimizing HVAC can define indoor temperature as a linear function of outdoor temperature and HVAC power [67,57]. This approach only maintains the average indoor temperature of the building within thermal comfort limits, whereas applications for real buildings must achieve this separately in every zone of the building. Further, it is unclear if a reinforcement learning agent trained in a linear environment could generalize to handle non-linear dynamics of real-world HVAC systems and a much larger state and action space with temperature measurements and HVAC actuators in

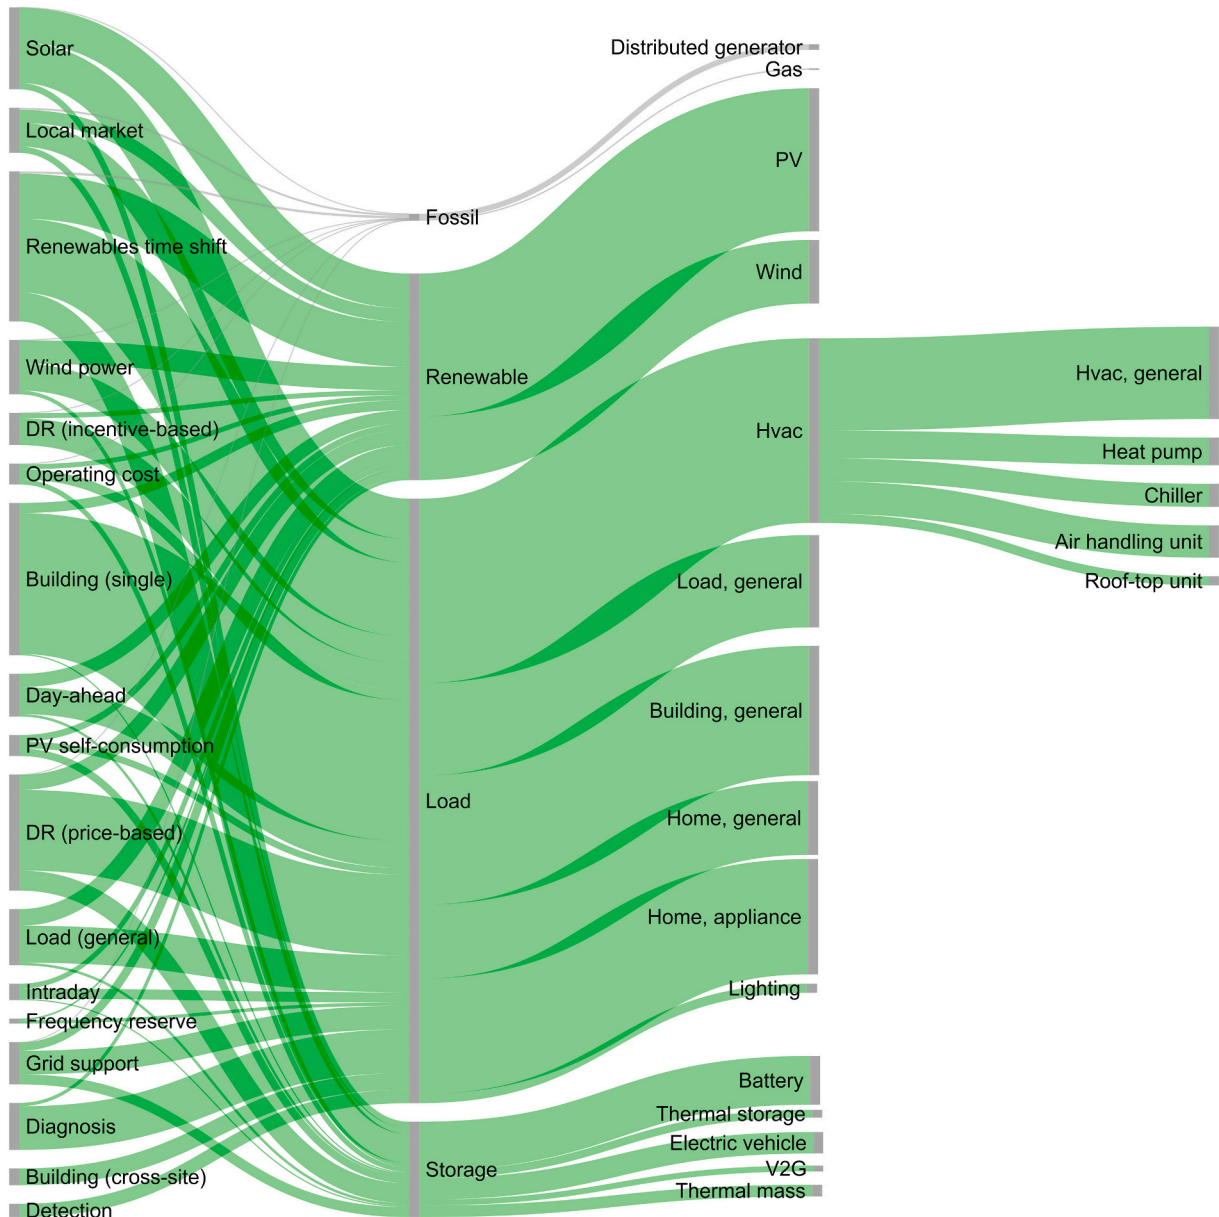


Fig. 14. Sankey chart relating the tier 3 categories (on the left) to the types of DERs encountered in this review.

several zones of the building.

A mention of a specific type of DER does not guarantee that the reported performance can be realized in a real building. For example, works on heat pumps do not capture the non-linear phenomena in compressors, expansion valves, condensers and evaporators [60,99]. In particular, compressors are the key electricity-consuming component in common HVAC equipment such as heat pumps and chillers. Generally, the reviewed literature does not distinguish between kW of power consumption of the compressors and the kW of cooling or heating provided to the building. None of the papers tried to optimize the load so that the compressor could operate at a high coefficient of performance.

Fig. 14 shows that storages and renewable generation resources are all modelled as a specific type of DER. Differences in fidelity have been observed related to how battery degradation has been modelled and included in the optimization. Renewable generation optimization approaches do not discuss how the proposed optimization may interfere with maximum power point tracking, possibly causing adverse effects on energy extraction.

According to Fig. 14, the majority of the DERs investigated in the articles of this taxonomy are located in buildings. Building automation systems often employ a 15 min control time step. Thus, a RL agent can only take actions and gain experiences at this interval. Such an interval is very long compared to other major applications of RL such as robotics and self-driving cars. Several of the articles studied in this paper avoid this problem by using a building energy simulator as the RL environment, running a sufficient number of episodes in this environment. An unresolved problem with this approach is that there are challenges in keeping the simulator updated, for example when building occupancy patterns change or when retrofits to the building are made. A topic for further research would be to investigate the use of a building digital twin to overcome these issues.

6. Conclusion

6.1. Areas expected to receive significant research in the future

The focus of the reviewed research according to the six types of VPP in Fig. 1 is as follows. A significant body of research exists for HEMS and BEMS. The boundary of single-site and multi-site VPPs is not clear cut, as many articles are not explicit about whether all the DERs are in the scope of a microgrid or utility grid. A small body of research is emerging for multi-energy VPPs with connections to district heating grids or gas grids. Only two publications were found for off-grid VPPs [96,18].

Forecasting of load, generation and prices has established itself as an active area of research. In the absence of benchmarks, it is difficult to identify which works are contributing performance improvements to the state of the art. An emerging body of research is combining such forecasts with optimizations using reinforcement learning. Reinforcement learning is emerging as a potentially disruptive technology for optimizing HEMS, BEMS and VPP systems. Its strength is in obtaining heuristic solutions to complex optimization problems under uncertainty. The further development and increased penetration of renewable energy and intelligent loads and storages is likely to generate a greater need for such solutions, so both the forecasting and optimization trends identified in this taxonomy are expected to attract a large amount of publications in the future.

6.2. Key unresolved challenges

Significant unaddressed potential issues have been identified, should the approaches reviewed in this paper eventually be deployed to real buildings:

1. Simplifications are commonly made in the environments and state and action spaces of reinforcement learning agents. It is unclear if the agents that have been trained in these environments could generalize

to the operating environment in a real building, which many sensors and actuators.

2. Forecasting solutions generally are customized for a single building or load, and require significant labelled training data for supervised learning. A few authors note that this approach is not scalable, due to lack of data and the high level of manual work. To overcome these problems, they propose transfer learning or grouping buildings to clusters that can be serviced with a single ML model.
3. Forecasting and optimization solutions tacitly make the impractical assumption that DERs are always in full health. Although the research in the classification branch of the taxonomy provides solutions to detect and diagnose DERs that perform sub-optimally, these are not exploited in the optimization and forecasting branches.
4. The majority of the optimization research is under the tier 2 category 'Cost', which fits well with the purpose of a VPP to monetize DERs. The works under the tier 2 category 'Other' are vague about who would invest in the solutions and who would get the financial benefits. These are areas in need of further research, before the solutions are ready for deployment.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This research was supported by Business Finland grant 7439/31/2018.

References

- [1] C. Liu, R.J. Yang, X. Yu, C. Sun, P.S.P. Wong, H. Zhao, Virtual power plants for a sustainable urban future, *Sustain. Cities Soc.* 65 (2021) 102640, <https://doi.org/10.1016/j.scs.2020.102640>.
- [2] S. Yu, F. Fang, Y. Liu, J. Liu, Uncertainties of virtual power plant: problems and countermeasures, *Appl. Energy* 239 (2019) 454–470, <https://doi.org/10.1016/j.apenergy.2019.01.224>.
- [3] K.O. Adu-Kankam, L.M. Camarinha-Matos, Towards collaborative virtual power plants: trends and convergence, *Sustain. Energy Grid Netw.* 16 (2018) 217–230, <https://doi.org/10.1016/j.segan.2018.08.003>.
- [4] K. Mahmud, B. Khan, J. Ravishanker, A. Ahmadi, P. Siano, An internet of energy framework with distributed energy resources, prosumers and small-scale virtual power plants: An overview, *Renew. Sust. Energ. Rev.* 127 (2020) 109840, <https://doi.org/10.1016/j.rser.2020.109840>.
- [5] N. Naval, J.M. Yusta, Virtual power plant models and electricity markets - a review, *Renew. Sust. Energ. Rev.* 149 (2021) 111393, <https://doi.org/10.1016/j.rser.2021.111393>.
- [6] E.A. Bhuiyan, Md.Z. Hossain, S.M. Mueen, S.R. Fahim, S.K. Sarker, S.K. Das, Towards next generation virtual power plant: technology review and frameworks, *Renew. Sust. Energ. Rev.* 150 (2021) 111358, <https://doi.org/10.1016/j.rser.2021.111358>.
- [7] H.M. Rouzbahani, H. Karimipour, L. Lei, A review on virtual power plant for energy management, *Sustain. Energy Technol. Assess.* 47 (2021) 101370, <https://doi.org/10.1016/j.seta.2021.101370>.
- [8] M. Ren, X. Liu, Z. Yang, J. Zhang, Y. Guo, Y. Jia, A novel forecasting based scheduling method for household energy management system based on deep reinforcement learning, *Sustain. Cities Soc.* 76 (2022) 103207, <https://doi.org/10.1016/j.scs.2021.103207>.
- [9] M. Rätz, A.P. Javadi, M. Baranski, K. Finkbeiner, D. Müller, Automated data-driven modeling of building energy systems via machine learning algorithms, *Energy Build.* 202 (2019) 109384, <https://doi.org/10.1016/j.enbuild.2019.109384>.
- [10] H. Han, Z. Zhang, X. Cui, Q. Meng, Ensemble learning with member optimization for fault diagnosis of a building energy system, *Energy Build.* 226 (2020) 110351, <https://doi.org/10.1016/j.enbuild.2020.110351>.
- [11] Z. Wang, J. Liu, Y. Zhang, H. Yuan, R. Zhang, R.S. Srinivasan, Practical issues in implementing machine-learning models for building energy efficiency: moving beyond obstacles, *Renew. Sust. Energ. Rev.* 143 (2021) 110929, <https://doi.org/10.1016/j.rser.2021.110929>.
- [12] S. Fathi, R. Srinivasan, A. Fenner, S. Fathi, Machine learning applications in urban building energy performance forecasting: a systematic review, *Renew. Sust. Energ. Rev.* 133 (2020) 110287, <https://doi.org/10.1016/j.rser.2020.110287>.

- [13] K. Mason, S. Grijalva, A review of reinforcement learning for autonomous building energy management, *Comput. Electr. Eng.* 78 (2019) 300–312, <https://doi.org/10.1016/j.compeleceng.2019.07.019>.
- [14] T.A. Nakabi, P. Toivanen, Deep reinforcement learning for energy management in a microgrid with flexible demand, *Sustain. Energy Grid Netw.* 25 (2021) 100413, <https://doi.org/10.1016/j.segan.2020.100413>.
- [15] D. Fang, X. Guan, L. Lin, Y. Peng, D. Sun, M.M. Hassan, Edge intelligence based economic dispatch for virtual power plant in 5G internet of energy, *Comput. Commun.* 151 (2020) 42–50, <https://doi.org/10.1016/j.comcom.2019.12.021>.
- [16] X. Kong, J. Xiao, D. Liu, J. Wu, C. Wang, Y. Shen, Robust stochastic optimal dispatching method of multi-energy virtual power plant considering multiple uncertainties, *Appl. Energy* 279 (2020) 115707, <https://doi.org/10.1016/j.apenergy.2020.115707>.
- [17] M. Ahrarinouri, M. Rastegar, A.R. Seifi, Multiagent reinforcement learning for energy management in residential buildings, *IEEE Trans. Ind. Inform.* 17 (2021) 659–666, <https://doi.org/10.1109/TII.2020.2977104>.
- [18] S. Totaro, I. Boukas, A. Jonsson, B. Cornélusse, Lifelong control of off-grid microgrid with model-based reinforcement learning, *Energy* 232 (2021) 121035, <https://doi.org/10.1016/j.energy.2021.121035>.
- [19] R. Subramanya, M. Yli-Ojanperä, S. Sierla, T. Hölttä, J. Valtakari, V. Vytkin, A virtual power plant solution for aggregating photovoltaic systems and other distributed energy resources for northern European primary frequency reserves, *Energies* 14 (2021) 1242, <https://doi.org/10.3390/en14051242>.
- [20] J. Zapata Riveros, K. Bruninx, K. Poncelet, W. D'haeseleer, Bidding strategies for virtual power plants considering CHPs and intermittent renewables, *Energy Convers. Manag.* 103 (2015) 408–418, <https://doi.org/10.1016/j.enconman.2015.06.075>.
- [21] Z. Azimi, R.-A. Hooshmand, S. Soleymani, Optimal integration of demand response programs and electric vehicles in coordinated energy management of industrial virtual power plants, *J. Energy Storage* 41 (2021) 102951, <https://doi.org/10.1016/j.est.2021.102951>.
- [22] J. Liu, K. Wang, L. Wang, Y. Li, Data interaction method of virtual power plant participating in spot market for massive load, in: 2021 IEEE International Conference on Advances in Electrical Engineering and Computer Applications (AEECA), IEEE, 2021, pp. 120–125, <https://doi.org/10.1109/AEECA52519.2021.9574377>.
- [23] J.M. Alemany, B. Arendarski, P. Lombardi, P. Komarnicki, Accentuating the renewable energy exploitation: evaluation of flexibility options, *Int. J. Electr. Power Energy Syst.* 102 (2018) 131–151, <https://doi.org/10.1016/j.ijepes.2018.04.023>.
- [24] F. Sanchez Gorostiza, F. Gonzalez-Longatt, Optimised TSO-DSO interaction in unbalanced networks through frequency-responsive EV clusters in virtual power plants, *IET Gener. Transm. Distrib.* 14 (2020) 4908–4917, <https://doi.org/10.1049/iet-gtd.2019.1947>.
- [25] J. Naughton, H. Wang, S. Riaz, M. Cantoni, P. Mancarella, Optimization of multi-energy virtual power plants for providing multiple market and local network services, *Electr. Power Syst. Res.* 189 (2020) 106775, <https://doi.org/10.1016/j.epr.2020.106775>.
- [26] M. Pourakbari-Kasmaei, M. Asensio, M. Lehtonen, J. Contreras, Trilateral planning model for integrated community energy systems and PV-based prosumers—a Bilevel stochastic programming approach, *IEEE Trans. Power Syst.* 35 (2020) 346–361, <https://doi.org/10.1109/TPWRS.2019.2935840>.
- [27] A. Lana, T. Kaipia, A. Pinomaa, P. Nuutinen, A. Narayanan, J. Partanen, LVDC network integrated functionalities for demand response, in: 2017 IEEE Second International Conference on DC Microgrids (ICDCM), IEEE, 2017, pp. 262–269, <https://doi.org/10.1109/ICDCM.2017.8001054>.
- [28] M. Lotfi, T. Almeida, M. Javadi, G.J. Osorio, J.P.S. Catalao, Coordinated operation of electric vehicle parking lots and smart homes as a virtual power plant, in: 2020 IEEE International Conference on Environment and Electrical Engineering and 2020 IEEE Industrial and Commercial Power Systems Europe (IEEEIC/ICPS Europe), IEEE, 2020, pp. 1–6, <https://doi.org/10.1109/IEEEIC/ICPSEurope49358.2020.9160684>.
- [29] H. Gong, E.S. Jones, R.E. Alden, A.G. Frye, D. Colliver, D.M. Ionel, Demand response of HVACs in large residential communities based on experimental developments, in: 2020 IEEE Energy Conversion Congress and Exposition (ECCE), IEEE, 2020, pp. 4545–4548, <https://doi.org/10.1109/ECCE44975.2020.9235465>.
- [30] F. Luo, A. Dorri, G. Ranzi, J. Zhao, R. Jurdak, Aggregating buildings as a virtual power plant: architectural design, supporting technologies, and case studies, *IET Energy Syst. Integr.* (2021) 1–13, <https://doi.org/10.1049/esi.12047>.
- [31] A. Rosato, M. Panella, A. Andreotti, O.A. Mohammed, R. Araneo, Two-stage dynamic management in energy communities using a decision system based on elastic net regularization, *Appl. Energy* 291 (2021) 116852, <https://doi.org/10.1016/j.apenergy.2021.116852>.
- [32] R.K. Veiga, A.C. Veloso, A.P. Melo, R. Lamberts, Application of machine learning to estimate building energy use intensities, *Energy Build.* 249 (2021) 111219, <https://doi.org/10.1016/j.enbuild.2021.111219>.
- [33] S. Seyedzadeh, F. Pour Rahimian, S. Oliver, S. Rodriguez, I. Glesk, Machine learning modelling for predicting non-domestic buildings energy performance: a model to support deep energy retrofit decision-making, *Appl. Energy* 279 (2020) 115908, <https://doi.org/10.1016/j.apenergy.2020.115908>.
- [34] M. Royapoor, M. Pazhoohesh, P.J. Davison, C. Patsios, S. Walker, Building as a virtual power plant, magnitude and persistence of deferrable loads and human comfort implications, *Energy Build.* 213 (2020) 109794, <https://doi.org/10.1016/j.enbuild.2020.109794>.
- [35] A. Hany Elgamal, G. Kocher-Oberlehner, V. Robu, M. Andoni, Optimization of a multiple-scale renewable energy-based virtual power plant in the UK, *Appl. Energy* 256 (2019) 113973, <https://doi.org/10.1016/j.apenergy.2019.113973>.
- [36] A. Gupta, Y. Badr, A. Negahban, R.G. Qiu, Energy-efficient heating control for smart buildings with deep reinforcement learning, *J. Build. Eng.* 34 (2021) 101739, <https://doi.org/10.1016/j.jobbe.2020.101739>.
- [37] S. Ikeda, T. Nagai, A novel optimization method combining metaheuristics and machine learning for daily optimal operations in building energy and storage systems, *Appl. Energy* 289 (2021) 116716, <https://doi.org/10.1016/j.apenergy.2021.116716>.
- [38] A. Kathirgamanathan, E. Mangina, D.P. Finn, Development of a soft actor critic deep reinforcement learning approach for harnessing energy flexibility in a large office building, *Energy AI* 5 (2021) 100101, <https://doi.org/10.1016/j.egyai.2021.100101>.
- [39] M. Löschenbrand, Modeling competition of virtual power plants via deep learning, *Energy* 214 (2021) 118870, <https://doi.org/10.1016/j.energy.2020.118870>.
- [40] S. Jung, J. Jeoung, H. Kang, T. Hong, Optimal planning of a rooftop PV system using GIS-based reinforcement learning, *Appl. Energy* 298 (2021) 117239, <https://doi.org/10.1016/j.apenergy.2021.117239>.
- [41] H. Liang, J. Ma, Data-driven resource planning for virtual power plant integrating demand response customer selection and storage, *IEEE Trans. Ind. Inform.* (2021) 1833–1844, <https://doi.org/10.1109/TII.2021.3068402>.
- [42] Z. Deng, Q. Chen, Reinforcement learning of occupant behavior model for cross-building transfer learning to various HVAC control systems, *Energy Build.* 238 (2021) 110860, <https://doi.org/10.1016/j.enbuild.2021.110860>.
- [43] F.U.M. Ullah, A. Ullah, I.U. Haq, S. Rho, S.W. Baik, Short-term prediction of residential power energy consumption via CNN and multi-layer bi-directional LSTM networks, *IEEE Access.* 8 (2020) 123369–123380, <https://doi.org/10.1109/ACCESS.2019.2963045>.
- [44] M. Khan, J. Seo, D. Kim, Real-time scheduling of operational time for smart home appliances based on reinforcement learning, *IEEE Access.* 8 (2020) 116520–116534, <https://doi.org/10.1109/ACCESS.2020.3004151>.
- [45] H. Malik, R. Sharma, S. Mishra, Fuzzy reinforcement learning based intelligent classifier for power transformer faults, *ISA Trans.* 101 (2020) 390–398, <https://doi.org/10.1016/j.isatra.2020.01.016>.
- [46] A. Javed, H. Larijani, A. Wixted, Improving energy consumption of a commercial building with IoT and machine learning, *IT Prof.* 20 (2018) 30–38, <https://doi.org/10.1109/MITP.2018.053891335>.
- [47] G.C. Chasparis, C. Lettner, Reinforcement-learning-based optimization for day-ahead flexibility extraction in battery pools, *IFAC-PapersOnLine.* 53 (2020) 13351–13358, <https://doi.org/10.1016/j.ifacol.2020.12.170>.
- [48] W. Kolodziejczyk, I. Zoltowska, P. Cichosz, Real-time energy purchase optimization for a storage-integrated photovoltaic system by deep reinforcement learning, *Control. Eng. Pract.* 106 (2021) 104598, <https://doi.org/10.1016/j.conengprac.2020.104598>.
- [49] D.J.B. Harrold, J. Cao, Z. Fan, Data-driven battery operation for energy arbitrage using rainbow deep reinforcement learning, *Energy* 238 (2022) 121958, <https://doi.org/10.1016/j.energy.2021.121958>.
- [50] G. Ceusters, R.C. Rodríguez, A.B. García, R. Franke, G. Deconinck, L. Helsen, A. Nowé, M. Messagie, L.R. Camargo, Model-predictive control and reinforcement learning in multi-energy system case studies, *Appl. Energy* 303 (2021) 117634, <https://doi.org/10.1016/j.apenergy.2021.117634>.
- [51] S. Abedi, S.W. Yoon, S. Kwon, Battery energy storage control using a reinforcement learning approach with cyclic time-dependent Markov process, *Int. J. Electr. Power Energy Syst.* 134 (2022) 107368, <https://doi.org/10.1016/j.ijepes.2021.107368>.
- [52] C. Guo, X. Wang, Y. Zheng, F. Zhang, Real-time optimal energy management of microgrid with uncertainties based on deep reinforcement learning, *Energy* 238 (2022) 121873, <https://doi.org/10.1016/j.energy.2021.121873>.
- [53] X. Huang, D. Zhang, X. Zhang, Energy management of intelligent building based on deep reinforced learning, *Alex. Eng. J.* 60 (2021) 1509–1517, <https://doi.org/10.1016/j.aej.2020.11.005>.
- [54] Y. Liu, D. Zhang, H.B. Gooi, Optimization strategy based on deep reinforcement learning for home energy management, *CSEE J. Power Energy Syst.* 6 (2020) 572–582, <https://doi.org/10.17775/CSEEJPES.2019.02890>.
- [55] Y. Liu, C. Yuen, N. Ul Hassan, S. Huang, R. Yu, S. Xie, Electricity cost minimization for a microgrid with distributed energy resource under different information availability, *IEEE Trans. Ind. Electron.* 62 (2015) 2571–2583, <https://doi.org/10.1109/TIE.2014.2371780>.
- [56] E. Mocanu, D.C. Mocanu, P.H. Nguyen, A. Liotta, M.E. Webber, M. Gibescu, J. G. Slootweg, On-line building energy optimization using deep reinforcement learning, *IEEE Trans. Smart Grid.* 10 (2019) 3698–3708, <https://doi.org/10.1109/TSG.2018.2834219>.
- [57] L. Yu, W. Xie, D. Xie, Y. Zou, D. Zhang, Z. Sun, L. Zhang, Y. Zhang, T. Jiang, Deep reinforcement learning for smart home energy management, *IEEE Internet Things J.* 7 (2020) 2751–2762, <https://doi.org/10.1109/JIOT.2019.2957289>.
- [58] D. Qiu, Y. Ye, D. Papadaskalopoulos, G. Strbac, Scalable coordinated management of peer-to-peer energy trading: a multi-cluster deep reinforcement learning approach, *Appl. Energy* 292 (2021) 116940, <https://doi.org/10.1016/j.apenergy.2021.116940>.
- [59] J. Lago, E. Sogancioglu, G. Suryanarayana, F. de Ridder, B. de Schutter, Building day-ahead bidding functions for seasonal storage systems: a reinforcement learning approach, *IFAC-PapersOnLine.* 52 (2019) 488–493, <https://doi.org/10.1016/j.ifacol.2019.08.258>.

- [60] B. Jiang, Y. Fei, Smart home in smart microgrid: a cost-effective energy ecosystem with intelligent hierarchical agents, *IEEE Trans. Smart Grid.* 6 (2015) 3–13, <https://doi.org/10.1109/TSG.2014.2347043>.
- [61] X. Zhang, R. Lu, J. Jiang, S.H. Hong, W.S. Song, Testbed implementation of reinforcement learning-based demand response energy management system, *Appl. Energy* 297 (2021) 117131, <https://doi.org/10.1016/j.apenergy.2021.117131>.
- [62] G. Pinto, M.S. Piscitelli, J.R. Vázquez-Canteli, Z. Nagy, A. Capozzoli, Coordinated energy management for a cluster of buildings through deep reinforcement learning, *Energy* 229 (2021) 120725, <https://doi.org/10.1016/j.energy.2021.120725>.
- [63] A. Mathew, A. Roy, J. Mathew, Intelligent residential energy management system using deep reinforcement learning, *IEEE Syst. J.* 14 (2020) 5362–5372, <https://doi.org/10.1109/JSYST.2020.2996547>.
- [64] T. Schreiber, S. Eschweiler, M. Baranski, D. Müller, Application of two promising reinforcement learning algorithms for load shifting in a cooling supply system, *Energy Build.* 229 (2020) 110490, <https://doi.org/10.1016/j.enbuild.2020.110490>.
- [65] B. Liu, M. Akcakaya, T.E. Mcdermott, Automated control of transactive HVACs in energy distribution systems, *IEEE Trans. Smart Grid.* 12 (2021) 2462–2471, <https://doi.org/10.1109/TSG.2020.3042498>.
- [66] T. Yang, L. Zhao, W. Li, J. Wu, A.Y. Zomaya, Towards healthy and cost-effective indoor environment management in smart homes: a deep reinforcement learning approach, *Appl. Energy* 300 (2021) 117335, <https://doi.org/10.1016/j.apenergy.2021.117335>.
- [67] H. Li, Z. Wan, H. He, Real-time residential demand response, *IEEE Trans. Smart Grid.* 11 (2020) 4144–4154, <https://doi.org/10.1109/TSG.2020.2978061>.
- [68] D. Azuatalam, W.-L. Lee, F. de Nijs, A. Liebman, Reinforcement learning for whole-building HVAC control and demand response, *Energy AI* 2 (2020) 100020, <https://doi.org/10.1016/j.egyai.2020.100020>.
- [69] Z. Jiang, M.J. Risbeck, V. Ramamurti, S. Murugesan, J. Amores, C. Zhang, Y. M. Lee, K.H. Drees, Building HVAC control with reinforcement learning for reduction of energy cost and demand charge, *Energy Build.* 239 (2021) 110833, <https://doi.org/10.1016/j.enbuild.2021.110833>.
- [70] Y.-J. Kim, A supervised-learning-based strategy for optimal demand response of an HVAC system in a multi-zone office building, *IEEE Trans. Smart Grid.* 11 (2020) 4212–4226, <https://doi.org/10.1109/TSG.2020.2986539>.
- [71] G.K. Chellamani, P. Venkatesh Chandramani, Demand response management system with discrete time window using supervised learning algorithm, *Cogn. Syst. Res.* 57 (2019) 131–138, <https://doi.org/10.1016/j.cogsys.2018.10.030>.
- [72] S. Bahrami, V.W.S. Wong, J. Huang, An online learning algorithm for demand response in smart grid, *IEEE Trans. Smart Grid.* 9 (2018) 4712–4725, <https://doi.org/10.1109/TSG.2017.2667599>.
- [73] S. Lee, D.-H. Choi, Federated reinforcement learning for energy management of multiple smart homes with distributed energy resources, *IEEE Trans. Ind. Inform.* 18 (2022) 488–497, <https://doi.org/10.1109/TII.2020.3035451>.
- [74] X. Xu, Y. Jia, Y. Xu, Z. Xu, S. Chai, C.S. Lai, A multi-agent reinforcement learning-based data-driven method for home energy management, *IEEE Trans. Smart Grid.* 11 (2020) 3201–3211, <https://doi.org/10.1109/TSG.2020.2971427>.
- [75] H.-M. Chung, S. Maharjan, Y. Zhang, F. Eliassen, Distributed deep reinforcement learning for intelligent load scheduling in residential smart grids, *IEEE Trans. Ind. Inform.* 17 (2021) 2752–2763, <https://doi.org/10.1109/TII.2020.3007167>.
- [76] Z. Wen, D. O'Neill, H. Maei, Optimal demand response using device-based reinforcement learning, *IEEE Trans. Smart Grid.* 6 (2015) 2312–2324, <https://doi.org/10.1109/TSG.2015.2396993>.
- [77] F. Alfaverth, M. Denai, Y. Sun, Demand response strategy based on reinforcement learning and fuzzy reasoning for home energy management, *IEEE Access.* 8 (2020) 39310–39321, <https://doi.org/10.1109/ACCESS.2020.2974286>.
- [78] H. Zhou, Y. Zhou, J. Hu, G. Yang, D. Xie, Y. Xue, L. Nordstrom, LSTM-based energy management for electric vehicle charging in commercial-building prosumers, *Journal of modern power systems and clean, Energy* 9 (2021) 1205–1216, <https://doi.org/10.35833/MPCE.2020.000501>.
- [79] A. Hassan, S. Acharya, M. Chertkov, D. Deka, Y. Dvorkin, A hierarchical approach to multienergy demand response: from electricity to multienergy applications, *Proc. IEEE* 108 (2020) 1457–1474, <https://doi.org/10.1109/JPROC.2020.2983388>.
- [80] T. Li, Y. Xiao, L. Song, Integrating future smart home operation platform with demand side management via deep reinforcement learning, *IEEE Trans. Green Commun. Netw.* 5 (2021) 921–933, <https://doi.org/10.1109/TGCN.2021.3073979>.
- [81] L. Lin, X. Guan, Y. Peng, N. Wang, S. Maharjan, T. Ohtsuki, Deep reinforcement learning for economic dispatch of virtual power plant in internet of energy, *IEEE Internet Things J.* 7 (2020) 6288–6301, <https://doi.org/10.1109/JIOT.2020.2966232>.
- [82] X. Zhang, D. Biagioni, M. Cai, P. Graf, S. Rahman, An edge-cloud integrated solution for buildings demand response using reinforcement learning, *IEEE Trans. Smart Grid.* 12 (2021) 420–431, <https://doi.org/10.1109/TSG.2020.3014055>.
- [83] R. Lu, S.H. Hong, Incentive-based demand response for smart grid with reinforcement learning and deep neural network, *Appl. Energy* 236 (2019) 937–949, <https://doi.org/10.1016/j.apenergy.2018.12.061>.
- [84] L. Wen, K. Zhou, J. Li, S. Wang, Modified deep learning and reinforcement learning for an incentive-based demand response model, *Energy* 205 (2020) 118019, <https://doi.org/10.1016/j.energy.2020.118019>.
- [85] A. Kumari, S. Tanwar, A reinforcement learning-based secure demand response scheme for smart grid system, *IEEE Internet Things J.* (2021), <https://doi.org/10.1109/JIOT.2021.3090305>.
- [86] Y. Kuang, X. Wang, H. Zhao, T. Qian, J. Wang, X. Wang, Model-free demand response scheduling strategy for virtual power plant considering risk attitude of consumer, *CSEE J. Power Energy Syst.* (2021), <https://doi.org/10.17775/CSEEJPES.2020.03120>.
- [87] Y. Zhang, Z. Zhang, Q. Yang, D. An, D. Li, C. Li, EV charging bidding by multi-DQN reinforcement learning in electricity auction market, *Neurocomputing* 397 (2020) 404–414, <https://doi.org/10.1016/j.neucom.2019.08.106>.
- [88] Z. Zhu, K. Wing Chan, S. Bu, B. Zhou, S. Xia, Real-time interaction of active distribution network and virtual microgrids: market paradigm and data-driven stakeholder behavior analysis, *Appl. Energy* 297 (2021) 117107, <https://doi.org/10.1016/j.apenergy.2021.117107>.
- [89] E. Samadi, A. Badri, R. Ebrahimpour, Decentralized multi-agent based energy management of microgrid using reinforcement learning, *Int. J. Electr. Power Energy Syst.* 122 (2020) 106211, <https://doi.org/10.1016/j.ijepes.2020.106211>.
- [90] T. Chen, W. Su, Local energy trading behavior modeling with deep reinforcement learning, *IEEE Access.* 6 (2018) 62806–62814, <https://doi.org/10.1109/ACCESS.2018.2876652>.
- [91] T. Chen, W. Su, Indirect customer-to-customer energy trading with reinforcement learning, *IEEE Trans. Smart Grid.* 10 (2019) 4338–4348, <https://doi.org/10.1109/TSG.2018.2857449>.
- [92] S. Zhou, Z. Hu, W. Gu, M. Jiang, X.-P. Zhang, Artificial intelligence based smart energy community management: a reinforcement learning approach, *CSEE J. Power Energy Syst.* 5 (2019) 1–10, <https://doi.org/10.17775/CSEEJPES.2018.00840>.
- [93] Q. Wei, Z. Liao, G. Shi, Generalized actor-critic learning optimal control in smart home energy management, *IEEE Trans. Ind. Inform.* 17 (2021) 6614–6623, <https://doi.org/10.1109/TII.2020.3042631>.
- [94] D. Fang, X. Guan, B. Hu, Y. Peng, M. Chen, K. Hwang, Deep reinforcement learning for scenario-based robust economic dispatch strategy in internet of energy, *IEEE Internet Things J.* 8 (2021) 9654–9663, <https://doi.org/10.1109/JIOT.2020.3040294>.
- [95] H. Zhao, B. Wang, H. Liu, H. Sun, Z. Pan, Q. Guo, Exploiting the flexibility inside park-level commercial buildings considering heat transfer time delay: a memory-augmented deep reinforcement learning approach, *IEEE Trans. Sustain. Energy* (2021) 207–219, <https://doi.org/10.1109/TSTE.2021.3107439>.
- [96] Z. Qin, D. Liu, H. Hua, J. Cao, Privacy preserving load control of residential microgrid via deep reinforcement learning, *IEEE Trans. Smart Grid.* 12 (2021) 4079–4089, <https://doi.org/10.1109/TSG.2021.3088290>.
- [97] P. Lissa, C. Deane, M. Schukat, F. Seri, M. Keane, E. Barrett, Deep reinforcement learning for home energy management system control, *Energy AI* 3 (2021) 100043, <https://doi.org/10.1016/j.egyai.2020.100043>.
- [98] M. Dorokhova, Y. Martinson, C. Ballif, N. Wyrsh, Deep reinforcement learning control of electric vehicle charging in the presence of photovoltaic generation, *Appl. Energy* 301 (2021) 117504, <https://doi.org/10.1016/j.apenergy.2021.117504>.
- [99] B.V. Mbuwir, F. Spiessens, G. Deconinck, Distributed optimization for scheduling energy flows in community microgrids, *Electr. Power Syst. Res.* 187 (2020) 106479, <https://doi.org/10.1016/j.epsr.2020.106479>.
- [100] Y. Shang, W. Wu, J. Guo, Z. Ma, W. Sheng, Z. Lv, C. Fu, Stochastic dispatch of energy storage in microgrids: an augmented reinforcement learning approach, *Appl. Energy* 261 (2020) 114423, <https://doi.org/10.1016/j.apenergy.2019.114423>.
- [101] F. Tuchnitz, N. Ebell, J. Schlund, M. Pruckner, Development and evaluation of a smart charging strategy for an electric vehicle fleet based on reinforcement learning, *Appl. Energy* 285 (2021) 116382, <https://doi.org/10.1016/j.apenergy.2020.116382>.
- [102] F.M. Solinas, L. Bottaccioli, E. Guelpa, V. Verda, E. Patti, Peak shaving in district heating exploiting reinforcement learning and agent-based modelling, *Eng. Appl. Artif. Intell.* 102 (2021) 104235, <https://doi.org/10.1016/j.engappai.2021.104235>.
- [103] W. Zhang, Q. Chen, J. Yan, S. Zhang, J. Xu, A novel asynchronous deep reinforcement learning model with adaptive early forecasting method and reward incentive mechanism for short-term load forecasting, *Energy* 236 (2021) 121492, <https://doi.org/10.1016/j.energy.2021.121492>.
- [104] G. Zhang, C. Tian, C. Li, J.J. Zhang, W. Zuo, Accurate forecasting of building energy consumption via a novel ensemble deep learning method considering the cyclic feature, *Energy* 201 (2020) 117531, <https://doi.org/10.1016/j.energy.2020.117531>.
- [105] Y. Chen, G. Fu, X. Liu, Air-conditioning load forecasting for prosumer based on meta ensemble learning, *IEEE Access.* 8 (2020) 123673–123682, <https://doi.org/10.1109/ACCESS.2020.2994119>.
- [106] S. Shan, B. Cao, Z. Wu, Forecasting the short-term electricity consumption of building using a novel ensemble model, *IEEE Access.* 7 (2019) 88093–88106, <https://doi.org/10.1109/ACCESS.2019.2925740>.
- [107] H. Deng, D. Fannon, M.J. Eckelman, Predictive modeling for US commercial building energy use: a comparison of existing statistical and machine learning algorithms using CBECS microdata, *Energy Build.* 163 (2018) 34–43, <https://doi.org/10.1016/j.enbuild.2017.12.031>.
- [108] Z. Ding, W. Chen, T. Hu, X. Xu, Evolutionary double attention-based long short-term memory model for building energy prediction: case study of a green building, *Appl. Energy* 288 (2021) 116660, <https://doi.org/10.1016/j.apenergy.2021.116660>.
- [109] S. Fan, Research on deep learning energy consumption prediction based on generating confrontation network, *IEEE Access.* 7 (2019) 165143–165154, <https://doi.org/10.1109/ACCESS.2019.2949030>.

- [110] L. Lei, W. Chen, B. Wu, C. Chen, W. Liu, A building energy consumption prediction model based on rough set theory and deep learning algorithms, *Energy Build.* 240 (2021) 110886, <https://doi.org/10.1016/j.enbuild.2021.110886>.
- [111] T. Liu, Z. Tan, C. Xu, H. Chen, Z. Li, Study on deep reinforcement learning techniques for building energy consumption forecasting, *Energy Build.* 208 (2020) 109675, <https://doi.org/10.1016/j.enbuild.2019.109675>.
- [112] H. Lu, F. Cheng, X. Ma, G. Hu, Short-term prediction of building energy consumption employing an improved extreme gradient boosting model: a case study of an intake tower, *Energy* 203 (2020) 117756, <https://doi.org/10.1016/j.energy.2020.117756>.
- [113] X.J. Luo, L.O. Oyedele, A.O. Ajayi, O.O. Akinade, Comparative study of machine learning-based multi-objective prediction framework for multiple building energy loads, *Sustain. Cities Soc.* 61 (2020) 102283, <https://doi.org/10.1016/j.scs.2020.102283>.
- [114] N. Somu, M.R. G. Raman, K. Ramamritham, A deep learning framework for building energy consumption forecast, *Renew. Sust. Energ. Rev.* 137 (2021) 110591, <https://doi.org/10.1016/j.rser.2020.110591>.
- [115] H. Sun, W. Zhai, Y. Wang, L. Yin, F. Zhou, Privileged information-driven random network based non-iterative integration model for building energy consumption prediction, *Appl. Soft Comput.* 108 (2021) 107438, <https://doi.org/10.1016/j.asoc.2021.107438>.
- [116] Z. Wang, Y. Wang, R.S. Srinivasan, A novel ensemble learning approach to support building energy use prediction, *Energy Build.* 159 (2018) 109–122, <https://doi.org/10.1016/j.enbuild.2017.10.085>.
- [117] J. Wang, X. Chen, F. Zhang, F. Chen, Y. Xin, Building load forecasting using deep neural network with efficient feature fusion, *Journal of modern power systems and clean, Energy* 9 (2021) 160–169, <https://doi.org/10.35833/MPEE.2020.000321>.
- [118] C. Zhang, J. Li, Y. Zhao, T. Li, Q. Chen, X. Zhang, A hybrid deep learning-based method for short-term building energy load prediction combined with an interpretation process, *Energy Build.* 225 (2020) 110301, <https://doi.org/10.1016/j.enbuild.2020.110301>.
- [119] A.T. Eseye, M. Lehtonen, T. Tukia, S. Uimonen, R., John Millar, machine learning based integrated feature selection approach for improved electricity demand forecasting in decentralized energy systems, *IEEE Access* 7 (2019) 91463–91475, <https://doi.org/10.1109/ACCESS.2019.2924685>.
- [120] C. Tian, C. Li, G. Zhang, Y. Lv, Data driven parallel prediction of building energy consumption using generative adversarial nets, *Energy Build.* 186 (2019) 230–243, <https://doi.org/10.1016/j.enbuild.2019.01.034>.
- [121] Y. Wang, Z. Quan, H. Jing, L. Wang, Y. Zhao, Performance and operation strategy optimization of a new dual-source building energy supply system with heat pumps and energy storage, *Energy Convers. Manag.* 239 (2021) 114204, <https://doi.org/10.1016/j.enconman.2021.114204>.
- [122] C. Zhang, J. Li, Y. Zhao, T. Li, Q. Chen, X. Zhang, W. Qiu, Problem of data imbalance in building energy load prediction: concept, influence, and solution, *Appl. Energy* 297 (2021) 117139, <https://doi.org/10.1016/j.apenergy.2021.117139>.
- [123] L. Zhang, J. Wen, Active learning strategy for high fidelity short-term data-driven building energy forecasting, *Energy Build.* 244 (2021) 111026, <https://doi.org/10.1016/j.enbuild.2021.111026>.
- [124] M. Al-Rakhami, A. Gumaei, A. Alsanad, A. Alamri, M.M. Hassan, An ensemble learning approach for accurate energy load prediction in residential buildings, *IEEE Access* 7 (2019) 48328–48338, <https://doi.org/10.1109/ACCESS.2019.2909470>.
- [125] S. Seyedzadeh, F. Pour Rahimian, P. Rastogi, I. Glesk, Tuning machine learning models for prediction of building energy loads, *Sustain. Cities Soc.* 47 (2019) 101484, <https://doi.org/10.1016/j.scs.2019.101484>.
- [126] L. Wang, R. Kubicek, X. Zhou, Adaptive learning based data-driven models for predicting hourly building energy use, *Energy Build.* 159 (2018) 454–461, <https://doi.org/10.1016/j.enbuild.2017.10.054>.
- [127] C.H. Kim, M. Kim, Y.J. Song, Sequence-to-sequence deep learning model for building energy consumption prediction with dynamic simulation modeling, *J. Build. Eng.* 43 (2021) 102577, <https://doi.org/10.1016/j.jobe.2021.102577>.
- [128] L. Ciabattoni, G. Ippoliti, A. Benini, S. Longhi, M. Pirro, Design of a home energy management system by online neural networks, *IFAC Proc. Vol.* 46 (2013) 677–682, <https://doi.org/10.3182/20130703-3-FR-4038.00111>.
- [129] P. Nie, M. Roccotelli, M.P. Fanti, Z. Ming, Z. Li, Prediction of home energy consumption based on gradient boosting regression tree, *Energy Rep.* 7 (2021) 1246–1255, <https://doi.org/10.1016/j.egyrr.2021.02.006>.
- [130] X. Zhao, W. Gao, F. Qian, J. Ge, Electricity cost comparison of dynamic pricing model based on load forecasting in home energy management system, *Energy* 229 (2021) 120538, <https://doi.org/10.1016/j.energy.2021.120538>.
- [131] Y. Feng, Q. Duan, X. Chen, S.S. Yakkali, J. Wang, Space cooling energy usage prediction based on utility data for residential buildings using machine learning methods, *Appl. Energy* 291 (2021) 116814, <https://doi.org/10.1016/j.apenergy.2021.116814>.
- [132] D. Syed, H. Abu-Rub, A. Ghayeb, S.S. Refaat, Household-level energy forecasting in smart buildings using a novel hybrid deep learning model, *IEEE Access* 9 (2021) 33498–33511, <https://doi.org/10.1109/ACCESS.2021.3061370>.
- [133] K. Aurangzeb, M. Alhussein, K. Javaid, S.I. Haider, A pyramid-CNN based deep learning model for power load forecasting of similar-profile energy customers based on clustering, *IEEE Access* 9 (2021) 14992–15003, <https://doi.org/10.1109/ACCESS.2021.3053069>.
- [134] Y. Hong, Y. Zhou, Q. Li, W. Xu, X. Zheng, A deep learning method for short-term residential load forecasting in smart grid, *IEEE Access* 8 (2020) 55785–55797, <https://doi.org/10.1109/ACCESS.2020.2981817>.
- [135] E.U. Haq, X. Lyu, Y. Jia, M. Hua, F. Ahmad, Forecasting household electric appliances consumption and peak demand based on hybrid machine learning approach, *Energy Rep.* 6 (2020) 1099–1105, <https://doi.org/10.1016/j.egyrr.2020.11.071>.
- [136] B. Buddhahai, W. Wongseeree, P. Rakkwamsuk, An energy prediction approach for a nonintrusive load monitoring in home appliances, *IEEE Trans. Consum. Electron.* 66 (2020) 96–105, <https://doi.org/10.1109/TCE.2019.2956638>.
- [137] A.T. Eseye, M. Lehtonen, Short-term forecasting of heat demand of buildings for efficient and optimal energy management based on integrated machine learning models, *IEEE Trans. Ind. Inform.* 16 (2020) 7743–7755, <https://doi.org/10.1109/TII.2020.2970165>.
- [138] M. Ribeiro, K. Grolinger, H.F. ElYamany, W.A. Higashino, M.A.M. Capretz, Transfer learning with seasonal and trend adjustment for cross-building energy forecasting, *Energy Build.* 165 (2018) 352–363, <https://doi.org/10.1016/j.enbuild.2018.01.034>.
- [139] C. Fan, Y. Sun, F. Xiao, J. Ma, D. Lee, J. Wang, Y.C. Tseng, Statistical investigations of transfer learning-based methodology for short-term building energy predictions, *Appl. Energy* 262 (2020) 114499, <https://doi.org/10.1016/j.apenergy.2020.114499>.
- [140] X. Fang, G. Gong, G. Li, L. Chun, W. Li, P. Peng, A hybrid deep transfer learning strategy for short term cross-building energy prediction, *Energy* 215 (2021) 119208, <https://doi.org/10.1016/j.energy.2020.119208>.
- [141] J. Ma, J.C.P. Cheng, F. Jiang, W. Chen, M. Wang, C. Zhai, A bi-directional missing data imputation scheme based on LSTM and transfer learning for building energy data, *Energy Build.* 216 (2020) 109941, <https://doi.org/10.1016/j.enbuild.2020.109941>.
- [142] X. Xu, W. Wang, T. Hong, J. Chen, Incorporating machine learning with building network analysis to predict multi-building energy use, *Energy Build.* 186 (2019) 80–97, <https://doi.org/10.1016/j.enbuild.2019.01.002>.
- [143] S.S. Gougheri, H. Jahangir, M.A. Golkar, A. Ahmadian, M. Aliakbar Golkar, Optimal participation of a virtual power plant in electricity market considering renewable energy: a deep learning-based approach, *Sustain. Energy Grid Netw.* 26 (2021) 100448, <https://doi.org/10.1016/j.segan.2021.100448>.
- [144] S. Sadeghi, H. Jahangir, B. Vatandoust, M.A. Golkar, A. Ahmadian, A. Elkamel, Optimal bidding strategy of a virtual power plant in day-ahead energy and frequency regulation markets: a deep learning-based approach, *Int. J. Electr. Power Energy Syst.* 127 (2021) 106646, <https://doi.org/10.1016/j.ijepes.2020.106646>.
- [145] A. Tascikaraoglu, O. Erdinc, M. Uzunoglu, A. Karakas, An adaptive load dispatching and forecasting strategy for a virtual power plant including renewable energy conversion units, *Appl. Energy* 119 (2014) 445–453, <https://doi.org/10.1016/j.apenergy.2014.01.020>.
- [146] M. Sharifzadeh, A. Sikiroti-Lock, N. Shah, Machine-learning methods for integrated renewable power generation: a comparative study of artificial neural networks, support vector regression, and Gaussian process regression, *Renew. Sust. Energ. Rev.* 108 (2019) 513–538, <https://doi.org/10.1016/j.rser.2019.03.040>.
- [147] S.A. Nabavi, N.H. Motlagh, M.A. Zaidan, A. Aslani, B. Zakeri, Deep learning in energy modeling: application in smart buildings with distributed energy generation, *IEEE Access* 9 (2021) 125439–125461, <https://doi.org/10.1109/ACCESS.2021.3110960>.
- [148] X. Wu, X. Hu, S. Moura, X. Yin, V. Pickert, Stochastic control of smart home energy management with plug-in electric vehicle battery energy storage and photovoltaic array, *J. Power Sources* 333 (2016) 203–212, <https://doi.org/10.1016/j.jpowsour.2016.09.157>.
- [149] V. Suresh, P. Janik, J.M. Guerrero, Z. Leonowicz, T. Sikorski, Microgrid energy management system with embedded deep learning forecaster and combined optimizer, *IEEE Access* 8 (2020) 202225–202239, <https://doi.org/10.1109/ACCESS.2020.3036131>.
- [150] A. Rosato, M. Panella, R. Araneo, A. Andreotti, A neural network based prediction system of distributed generation for the management of microgrids, *IEEE Trans. Ind. Appl.* 55 (2019) 7092–7102, <https://doi.org/10.1109/TIA.2019.2916758>.
- [151] Y. Zhao, R. Ball, J. Mosesian, J.-F. de Palma, B. Lehman, Graph-based semi-supervised learning for fault detection and classification in solar photovoltaic arrays, *IEEE Trans. Power Electron.* 30 (2015) 2848–2858, <https://doi.org/10.1109/TPEL.2014.2364203>.
- [152] W.-S. Yun, W.-H. Hong, H. Seo, A data-driven fault detection and diagnosis scheme for air handling units in building HVAC systems considering undefined states, *J. Build. Eng.* 35 (2021) 102111, <https://doi.org/10.1016/j.jobe.2020.102111>.
- [153] R. Shohet, M.S. Kandil, Y. Wang, J.J. McArthur, Fault detection for non-condensing boilers using simulated building automation system sensor data, *Adv. Eng. Inform.* 46 (2020) 101176, <https://doi.org/10.1016/j.aei.2020.101176>.
- [154] S. Taheri, A. Ahmadi, B. Mohammadi-Ivatloo, S. Asadi, Fault detection diagnostic for HVAC systems via deep learning algorithms, *Energy Build.* 250 (2021) 111275, <https://doi.org/10.1016/j.enbuild.2021.111275>.
- [155] X. Zhu, K. Chen, B. Anduv, X. Jin, Z. Du, Transfer learning based methodology for migration and application of fault detection and diagnosis between building chillers for improving energy efficiency, *Build. Environ.* 200 (2021) 107957, <https://doi.org/10.1016/j.buildenv.2021.107957>.
- [156] K. Yan, C. Zhong, Z. Ji, J. Huang, Semi-supervised learning for early detection and diagnosis of various air handling unit faults, *Energy Build.* 181 (2018) 75–83, <https://doi.org/10.1016/j.enbuild.2018.10.016>.