



Proceeding Paper

# Electricity Consumption Forecasting for Out-of-Distribution Time-of-Use Tariffs <sup>†</sup>

Jyoti Narwariya <sup>\*†</sup>, Chetan Verma <sup>\*‡</sup>, Pankaj Malhotra, Lovekesh Vig, Easwara Subramanian and Sanjay Bhat

TCS Research, New Delhi 110 001, India; malhotra.pankaj@tcs.com (P.M.); lovekesh.vig@tcs.com (L.V.); easwar.subramanian@tcs.com (E.S.); sanjay.bhat@tcs.com (S.B.)

\* Correspondence: jyoti.narwariya@tcs.com (J.N.); verma.chetan@tcs.com (C.V.)

† Presented at the AAAI Workshop on Artificial Intelligence with Biased or Scarce Data (AIBSD), Online, 28 February 2022.

‡ These authors contributed equally to this work.

**Abstract:** In electricity markets, electricity retailers or brokers want to maximize profits by allocating tariff profiles to end-consumers. One of the objectives of such demand response management is to incentivize the consumers to adjust their consumption so that the overall electricity procurement in the wholesale markets is minimized, e.g., it is desirable that consumers consume less during peak hours when the cost of procurement for brokers from wholesale markets are high. We consider a greedy solution to maximize the overall profit for brokers by optimal tariff profile allocation, i.e., allocate that tariff profile to a consumer that maximizes the profit with respect to that consumer. This, in turn, requires forecasting electricity consumption for each user for all tariff profiles. This forecasting problem is challenging compared to standard forecasting problems due to following reasons: (1) the number of possible combinations of hourly tariffs is high and retailers may not have considered all combinations in the past resulting in a biased set of tariff profiles tried in the past, i.e., the retailer may want to consider new tariff profiles that may achieve better profits; (2) the profiles allocated in the past to each user is typically based on certain policy, i.e., tariff profile allocation for historical electricity consumption data is biased. These reasons violate the standard IID assumptions as there is a need to evaluate new tariff profiles on existing customers and historical data is biased by the policies used in the past for tariff allocation. In this work, we consider several scenarios for forecasting and optimization under these conditions. We leverage the underlying structure of how consumers respond to variable tariff rates by comparing tariffs across hours and shifting loads, and propose suitable inductive biases in the design of deep neural network based architectures for forecasting under such scenarios. More specifically, we leverage attention mechanisms and permutation equivariant networks that allow desirable processing of tariff profiles to learn tariff representations that are insensitive to the biases in the data and still representative of the task. Through extensive empirical evaluation using the PowerTAC simulator, we show that the proposed approach significantly improves upon standard baselines that tend to overfit to the historical tariff profiles.

**Keywords:** out-of-distribution generalization; forecasting; temporal bias; permutation equivariance; optimization



**Citation:** Narwariya, J.; Verma, C.; Malhotra, P.; Vig, L.; Subramanian, E.; Bhat, S. Electricity Consumption Forecasting for Out-of-Distribution Time-of-Use Tariffs. *Comput. Sci. Math. Forum* **2022**, *3*, 1. <https://doi.org/10.3390/cmsf2022003001>

Academic Editors: Kuan-Chuan Peng and Ziyang Wu

Published: 8 April 2022

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



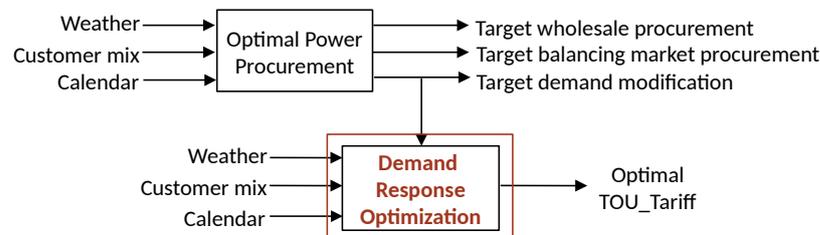
**Copyright:** © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

A smart grid consists of multiple types of entities such as those involved in generation, distribution, and consumption (smart appliances and buildings). One of the aims of a smart grid is to manage electricity demand in an economical manner via integration and exchange of information about all entities involved. For the customers or the end-consumers as well as the electricity distributing agencies or the electricity brokers, it offers the flexibility to choose/allocate among dynamically changing tariffs to meet certain objectives, e.g.,

minimize electricity bill for customers, maximize profit for retailers, etc. However, meeting such objectives is challenging due to dynamics of the market, e.g., changing wholesale electricity prices, supply–demand fluctuations, etc.

As depicted in Figure 1, a broker typically performs three functions: (1) purchase or sell power to its subscribers or customers in the retail market, (2) purchase or sell power in the wholesale market, and (3) rectify any supply–demand imbalance within its portfolio through the balancing market. In this work, we consider a simplified setting where the broker performs the following two functions: (1) sell power to those customers in the retail market who are electricity consumers, and (2) purchase power in the wholesale market. Typical examples of consumers include offices, housing complexes, hospitals, and villages. Furthermore, we focus on only those subset of consumers who have a *shiftable* load component in their total or aggregate consumption in addition to the traditional fixed or non-shiftable load, i.e., the consumption (e.g., appliance usage) at an hour that cannot be moved to another hour. This shiftable load can be shifted from the originally preferred hour to another hour in the day if the tariff for the latter is lower. The broker may want to encourage such a behavior, known as demand response management [1], to maximize profit or balance demand–supply.



**Figure 1.** Various aspects and objectives in an electricity markets. In this work, we focus on a sub-problem related to allocation of optimal time-of-use tariff (TOU Tariff) to each customer.

In this work, we consider the following out-of-distribution generalization problem: given historical aggregated consumption of consumers to tariff profiles allocated to them, forecast the aggregated consumption for new tariff profiles. These new tariff profiles are part of the electricity broker or retailer’s plan to explore new profiles to further improve the profits. This is different from standard forecasting problems as the exogenous variables (tariff profiles) at test time are different from the exogenous variables at train time. Furthermore, the allocation of tariff profiles in the past is not random, so the data is biased in the sense that, for different consumer personas, not all historical tariff profiles would have been tried. We note that the logic based on which the consumers respond to tariff profiles is consistent irrespective of the tariff profile. We propose to capture that logic in the neural network by using permutation equivariant networks and attention mechanisms.

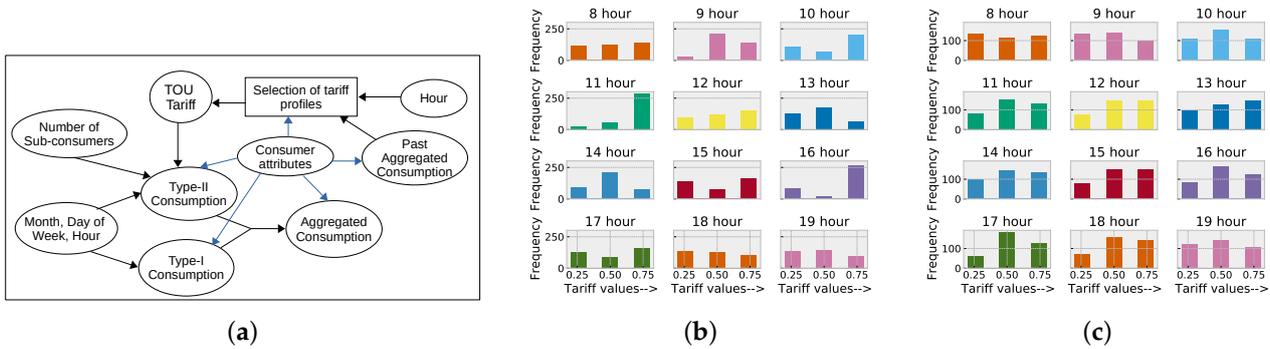
The key contributions of this work can be summarized as follows:

- We consider the problem of electricity consumption forecasting under new tariff profiles not encountered previously. This is then used for tariff profile allocation to optimize electricity broker’s profits.
- We note that the forecasting problem can be seen as an out-of-distribution (OOD) generalization problem with bias in the training data consisting of temporal and confounding bias.
- To achieve OOD generalization, we leverage the logic behind how consumers respond to tariff profiles in order to shift load, and propose a novel neural network architecture to achieve better OOD generalization.

Through empirical evaluation, we show that the proposed approach is able to improve upon vanilla methods that do not take into account suitable inductive biases guided by the knowledge of how consumers respond to tariff profiles.

## 2. Problem Formulation

The aggregated consumption  $e_{c,t} \in \mathbb{R}^+$  of a consumer  $c$  at time  $t$  has two components:  
 (1) *Type-I consumption*: this is non-shiftable consumption corresponding to the appliances that have to be used at specific hours only and cannot be shifted to alternative hour;  
 (2) *Type-II consumption*: this is shiftable component of the consumption corresponding to appliances whose usage can be planned. Refer to Figure 2a for more details.



**Figure 2.** (a) Logic for Consumption Data generation in Electricity Markets and (b,c) Hourly Tariff Rate Distributions depicting changing distribution across hours that poses generalization challenge. (a) Causal Diagram. (b) Hourly Tariff Distributions in IID Profiles depicting temporal bias ( $\mathcal{T}_{in}$ ). (c) Hourly Tariff Distributions in OOD Profiles ( $\mathcal{T}_{out}$ ).

Let  $e_{c,1:t}$  denote the time series of electricity consumption for consumer  $c$  until time  $t$ . We consider a consumer  $c \in \mathcal{C}$ , where  $\mathcal{C}$  is the set of consumers with non-zero Type-II consumption, i.e., part of their load can be shifted in response to variations in tariff across hours. Further, the  $i$ -th time-of-use (TOU) tariff profile is denoted as an ordered sequence or  $H$ -length time series of hourly tariffs  $TOU^i = TOU_1^i \dots TOU_H^i$ , where  $TOU_h^i$  ( $h = 1 \dots H$ ) denotes the tariff at hour  $h$ . In this work, we consider tariff profile with hourly rates over a day such that  $H = 24$ , without loss of generality.

Let  $\mathbf{f}_{c,1:t}$  denote all features (static or time-varying) for consumer  $c$  at time  $t$ , including e.g., past consumption time series, type of consumer (household, office, etc.), and  $\mathbf{f}_t$  denote a vector of temporal features at timestamp  $t$ , e.g., hour of the day, day of the week, week of the month, month of the year, etc. Note that  $\mathbf{f}_{c,1:t}$  refers to relevant features from entire history, but in practice, we consider a window of length  $w$  over  $t - w + 1 : t$  for deriving features at time  $t$ .

Further consider a tariff allocation policy function  $\pi$  such that

$$TOU_{c,t+\tau} = \pi(\mathbf{f}_{c,:t}, \mathbf{f}_{t+\tau}, \hat{p}_{t+1:t+H}),$$

i.e., the tariff at a future time  $t + \tau$  with  $\tau = 1 \dots H$  is decided based on consumer features at time  $t$ , the temporal features for time  $t + \tau$ , where  $\hat{p}_{t+\tau}$  denotes the estimate of electricity price  $p_{t+\tau}$  in the wholesale market at time  $t + \tau$ . Without loss of generality, we consider the scenario where  $t + 1$  corresponds to the first hour of the day, i.e., tariff profile for the next day is decided using data until the end of the current day.

Consider historical time series data  $\mathcal{D} = \{e_{c,1:t}, TOU_{c,1:t}\}_{c \in \mathcal{C}}$ , where the tariff time series are a result of sequence of tariff profile allocations over days such that any profile  $TOU^i \in \mathcal{T}_{in}$  is chosen from a fixed set of profiles  $\mathcal{T}_{in}$ .

The goal for the broker is to allocate that tariff profile  $TOU^i$  to a consumer that maximizes the gain  $G_c^i$  over the next  $H$  hours:

$$G_c^i = \sum_{t'=1}^H (TOU_{c,t+t'}^i - p_{t+t'}) \times e_{c,t+t'}. \tag{1}$$

Importantly, the electricity consumption  $e_{c,t+t'}$  at  $t + t'$  hour is a function of the entire tariff profile on that day, as the consumer could choose to shift the shiftable part of the load

from high tariff hours to low tariff hours by looking at the tariff profile allocated to the consumer at the beginning of the day.

We consider the following two scenarios depending on the tariff profiles being considered for future allocations:

**IID Scenario:** when the profiles to be allocated to the consumers in future are from the same set of profiles  $\mathcal{T}_{in}$  used historically, i.e.,  $TOU^i \in \mathcal{T}_{in}$ .

**OOD Scenario:** when the tariff profiles to be allocated to the consumers in future belong to  $\mathcal{T}_{all} = \mathcal{T}_{in} \cup \mathcal{T}_{out}$ , where  $\mathcal{T}_{out}$  is a new set of profiles not previously seen in  $\mathcal{D}$ , i.e., are out-of-distribution with respect to the training data, and not previously allocated to any consumer by the broker who wants to consider these new profiles to improve future gains, i.e.,  $TOU^i \in \mathcal{T}_{all}$ .

### 3. Related Work

Our work relates to two bodies of literature: (1) demand response management in electricity markets and the related sub-problem of electricity consumption forecasting under exogenous variables, using reinforcement learning and deep learning methods [2–4], and (2) out-of-distribution (OOD) generalization [5–8].

There have been many studies for (1); however, to the best of our knowledge, the problem of bias in historical data in terms of the tariff profiles has been largely overlooked. We draw attention of the community working on (1) to the potential of OOD generalization by improving forecasts for previously unallocated tariffs by using the underlying structure of the problem in terms of the particular way in which consumers shift loads in response to changes in tariff. More specifically, we rely on the partial permutation equivariance property of the response to time series of tariffs.

OOD detection and generalization is an emerging area of research, and aims at improving the robustness of models to previously unseen scenarios. Many of the recent approaches for (2) rely on changes in the objective function or different training procedures. For example, the approaches based on meta-learning [9] are not applicable as there is no notion of multiple tasks. We can consider each tariff profile as a task but then the forecasting can involve different profiles in input versus output. In this work, we focus on using inductive biases in the form of the neural network architecture to improve OOD generalization. There is enough evidence to support the improvement in generalization abilities of neural networks by using the structure of the problem to introduce suitable inductive biases in the learning process. The most commonly used inductive bias is in the design of the neural network architecture motivated by the structure of the problem. Recent examples of this include using graph neural networks [10,11] and modular networks [12]. Recently, using structural biases in deep neural networks motivated by the nature of bias and the structure of the problem have been successfully evaluated for time series forecasting [13]. Data-dependent priors have been recently proposed in [14]. However, to the best of our knowledge, using consumer behavior properties for electricity time series forecasting under out-of-distribution exogenous variables to guide the design of neural network architecture has not been considered so far in the literature.

### 4. The Learning Problem

We consider a 2-step approximate solution to maximize the gain (Equation (1)):

**Step 1:** For each consumer, forecast/estimate the consumption under each potential tariff profile allocation. Given features  $\mathbf{f}_{c,1:t}$  (including  $e_{c,1:t}$ ), history of allocated tariffs  $TOU_{c,1:t}$ , and values of potential future tariff  $TOU_{c,t+1:t+H}$ , the goal is to estimate  $e_{c,t+1:t+H}$ . This can be seen as a **multi-step time series forecasting problem with exogenous variables**. We provide the details of our proposed approach for this in the next section.

**Step 2:** Compute the profit using

$$\hat{G}_c^i = \sum_{t'=1}^H (TOU_{c,t+t'}^i - \hat{p}_{t+t'}) \times \hat{e}_{c,t+t'} \quad (2)$$

for each tariff in  $TOU^i \in \mathcal{T}_{all}$  for OOD scenario ( $\mathcal{T}_{in}$  for IID scenario). Allocate the tariff profile to consumer  $c$  which results in maximum  $\hat{G}_c^i$ . Note that, in practice, the future wholesale rates  $p_{t+t'}$  ( $t' = 1 \dots T$ ) are also not known and might need to be estimated. In this work, we assume that  $p_{t+t'}$ s are known in advance or estimable accurately and focus on estimating  $\hat{e}_{c,t+t'}$ s which are the only terms controllable via  $TOU_{c,t+t'}$ s.

In summary, the tariff profile allocation policy corresponds to estimating the gain for each tariff profile for a consumer, and then allocating the profile with maximum estimated gain. We use a deep neural network based architecture as the function approximator that estimates  $\mathbb{E}[e_{c,t+t'} | TOU_{c,t+1:t+T}]$  from the data.

#### 4.1. Biased and Scarce Data

The OOD scenario is challenging as there is no historical data for the profiles in  $\mathcal{T}_{out}$ . More concretely, we consider three possible values of tariff at any time  $t$ : low (0.2), medium (0.5), and high (0.8). Therefore, there are  $3^H$  unique profiles possible. For  $H = 24$ , there can be  $\approx 3 \times 10^{11}$  profiles possible. However, in practice, the number of allocated profiles would be significantly smaller than this. In this work, we consider  $|\mathcal{T}_{in}| \in \{2, 5, 8, 10, 12, 15, 20, 30, 35\}$ , which is a range of values encountered for  $|\mathcal{T}_{in}|$  in practice. This poses serious OOD generalization challenge in estimating  $e_{c,t+1:t+T}$  for previously unseen profiles  $TOU_{t+1:t+T}^i \in \mathcal{T}_{out}$ .

We note that one peculiar type of bias that manifests in practice is the **temporal bias**: at any hour  $h$  of the day, certain values of tariff are more common than others. We explain this further using a practical scenario as depicted in Figure 2: In practice, it is common to use the following heuristic for tariff profile allocation: Keep most expensive tariff rates during peak demand periods, least expensive tariff rates during non-peak hours, and slightly cheaper (medium) rates, typically between peak and off-peak periods. Every tariff profile is curated on the basis of average aggregated consumption of each customer. High tariff is allocated when the aggregated consumption is high, and for rest of the hours, low/mid tariff are allocated. The distribution of tariff rates over hours would depend on the distribution of peak consumption across customers (refer Figure 2c). Furthermore, there is confounding bias [15] with latent consumer attributes affecting (1) past aggregated consumption which in turn affects the treatment (tariff profile allocation), and (2) the outcome (electricity consumption) in  $\mathcal{D}$  both can depend on the consumer features (refer Figure 2a). We leave the handling of confounding bias for future work, and focus on handling temporal bias in this work.

We empirically show that temporal bias poses a generalization challenge for vanilla feed-forward neural networks, and propose an attention-based architecture to deal with the same, in the next section.

#### 4.2. How Consumers Respond to Tariffs

Consider the following toy example with  $H = 6$  where there is only one tariff profile in  $\mathcal{T}_{in}$  given by  $\{HHMMLL\}$ , i.e., tariff rate is high (H) for the first two hours, medium (M) for the next two hours, and low (L) for the last two hours. Further assume that the consumer has a certain Type-II load during the 1st hour. After looking at this tariff profile, the consumer responds by shifting the load from the 1st (high tariff) hour to the 5th (low tariff) hour. Now, consider a tariff profile in  $\mathcal{T}_{out}$  as  $\{HHLLMM\}$ . Clearly, this profile is different from the profile in  $\mathcal{T}_{in}$  as the sequence of highs and lows over the hours is different. However, importantly, the underlying decision-making behavior of the consumer remains the same, i.e., shift the Type-II load from high tariff hour (1st hour in this case) to low tariff hour (3rd hour instead of 5th hour in this case). Therefore, it is still possible to forecast the behavior of the user for this OOD profile. In this work, we intend to leverage this aspect of the consumer's decision-making process that stays the same irrespective of the IID-vs-OOD profiles.

Further, consider five ways to process the sequence of tariff rates (Figure 3):

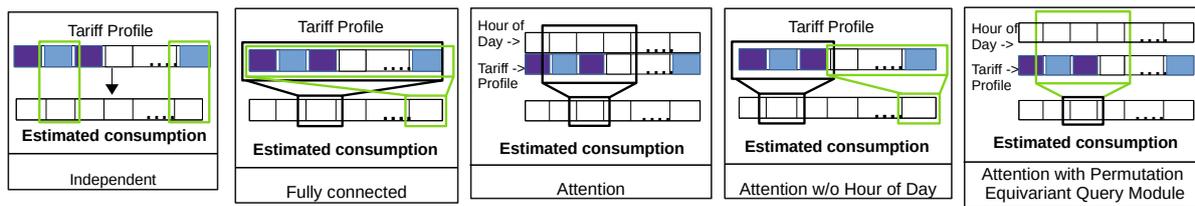


Figure 3. How different methods process the sequence of tariff rates.

- **Independent processing:** Here, the tariff at each hour is processed independently [16,17] and used to estimate the consumption at that hour. Of course, since the consumer’s decision making is based on comparison of tariff rates across hours, such a processing of tariff profiles will not be effective.
- **All considered together or fully connected:** Here, tariffs at all hours (the entire tariff profile) are processed simultaneously, e.g., through a fully connected layer in a feed-forward neural network. We argue that such processing of tariff profiles will be able to effectively learn a good function approximator for the profiles in  $\mathcal{T}_{in}$ . However, it will be highly biased to the profiles in  $\mathcal{T}_{in}$  since it does not effectively learn the way consumers are processing the tariff rates for shifting the loads. This leads to biased tariff profile processing modules due to the temporal bias in the historical profiles, as discussed above.
- **Focusing on relevant information or Attention:** Here, the tariffs rates in a day are considered as tokens and hours of a day are used as a positional information. This information is processed through a self-attention layer. We argue that such processing of tariff profiles will mimic the logic of how consumers respond to a tariff profile. However, it will be biased towards the profiles in  $\mathcal{T}_{in}$  since the tariffs and hour of the day are correlated (due to temporal bias in the historical tariff profiles).
- **Permutation Equivariance:** As discussed earlier, permutation equivariance is an important aspect of the consumer decision-making logic. To mimic the same in the processing of tariffs by the neural networks, we expect that if trained on one of the tariff sequences, say, *HHMMLL* in the earlier example), it should perform equally well on other sequence (i.e., *HHLLMM*). In other words, processing of tariffs by neural networks should be Permutation Equivariant. We propose two ways to achieve approximate permutation equivariance:
  - **Attention w/o Hour of Day (Att.-HOD):** As explained above, the standard self-attention method can mimic the logic of how consumers respond to tariffs, but due to temporal bias in the data, the attention method does not generalize well to  $\mathcal{T}_{out}$ . We propose a simple variant that does not take HOD as input in the self-attention module to obtain the permutation equivariance property.
  - **Attention with Permutation Equivariant Query Processing Module (Att.+PE):** Here, the tariff rates in a day are considered as a set and processed in such a way that ordering of the tariff rates does not matter, i.e., the processing is permutation equivariant [18,19].

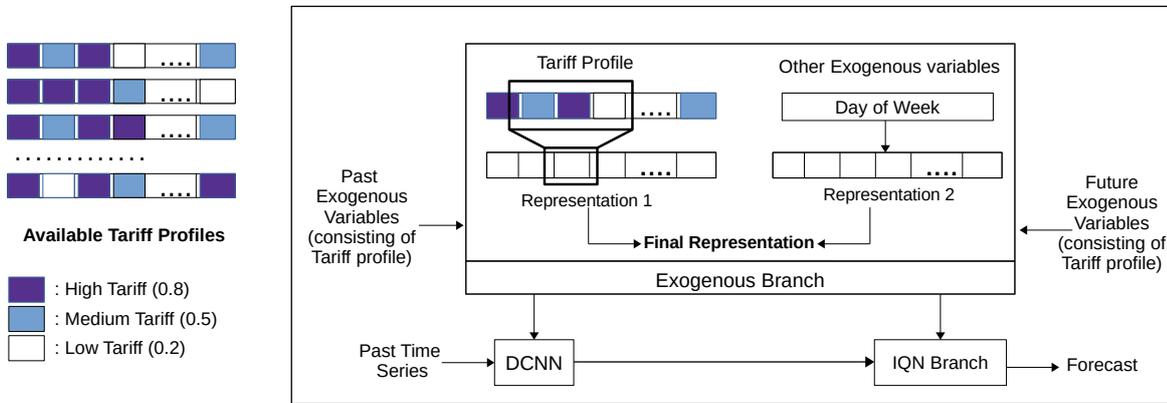
In the next section, we explain how we achieve permutation equivariance while forecasting the consumption given a consumer’s consumption history, sequence of past tariff profiles, and a future tariff profile.

### 5. Forecasting Architecture

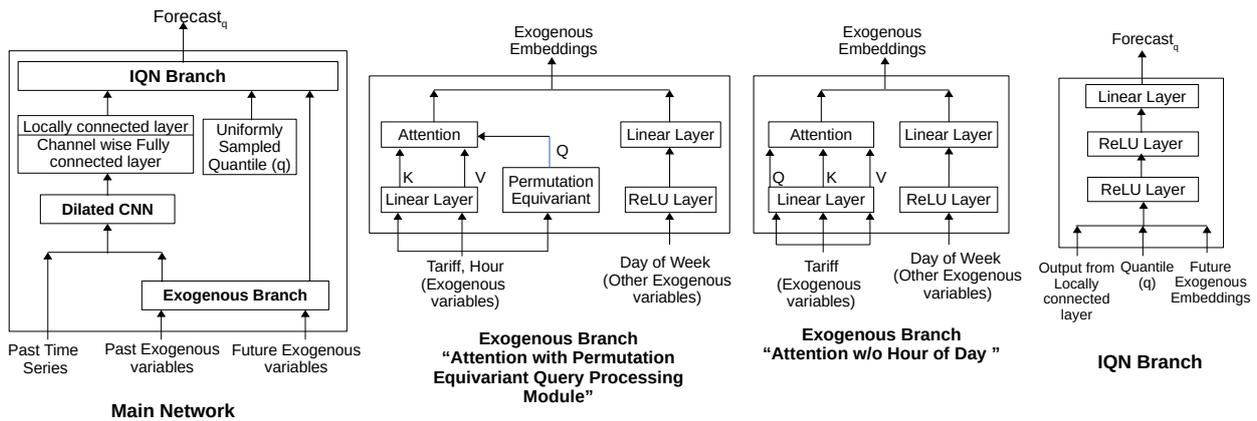
Consider the consumption history of a consumer along with past allocated tariffs to be a time series of vectors  $\mathbf{f}_{1:t}$  including dimensions for past aggregate consumption and past allocated tariff rates  $\{e_{1:t}, TOU_{1:t}\}$ , and the candidate tariff profile for the next  $H$  hours to be  $TOU_{t+1:t+H}$ . The goal is to estimate  $e_{t+1:t+H}$  while ensuring permutation equivariance in processing  $TOU_{1:t+H}$  in the sense of [19], e.g., if the output of processing  $\{TOU_1, TOU_2, TOU_3\}$  is  $\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3\}$ , then the output of processing a permutation

of the input, say  $\{TOU_2, TOU_1, TOU_3\}$ , is given by the permutation  $\{\mathbf{o}_2, \mathbf{o}_1, \mathbf{o}_3\}$  of the original output.

To achieve the above-stated goal, we consider the following modularized neural network architecture as depicted in Figures 4 and 5:



**Figure 4.** Flow diagram of “Attention w/o Hour of Day” approach. The left part of the figure indicates the variability in the tariff profiles and also some tariffs are more frequent in tariff profiles. The right part of the figure indicates flow of the inputs through the network and how the information of tariffs is consumed by the proposed approach.



**Figure 5.** Architectures contrasting “Attention w/o Hour of Day” and “Attention with Permutation Equivariant Query Processing Module” approaches.

- **Dilated Convolutional Neural Networks (DCNN)** branch for processing of past consumption time series. (Since we have large input time series ( $t = 168$  in our case), we consider 1D-Convolution Neural Networks for computational efficiency instead of Recurrent Neural Networks based architecture such as LSTMs [20].)
- **Exogenous** branch: This branch consists of **Attention with Permutation Equivariant Query Processing Module (Att.+PE)** branch for processing of tariff rates, and other modules for processing of features like hour of day, day of week, etc.
- **Implicit Quantile Network (IQN)** branch for generating the quantile estimates for future consumption.

Next, we provide details of the exogenous branch which is the key novel component of the proposed approach and helps to mitigate temporal bias.

To achieve permutation equivariance and handle temporal bias, we consider processing the tariff rates  $TOU_{t+1:t+H}$  (same processing is done for past tariffs as well) via an attention mechanism where a part of the processing is done independently for tariff at each time step  $t + t'$  ( $t' = 1 \dots H$ ) while still taking into account the global information

$TOU_{t+1:t+H}$  in order to mimic the behavior of the consumer as explained in the previous section.

More specifically, we consider key  $K$  and value  $V$  for the attention mechanism to be dependent on a single time step  $t + t'$ , while the query  $Q$  depends on the entire tariff profile  $TOU_{t+1:t+H}$  for the day. In other words,  $K_{t+t'} = f_K(TOU_{t+t'}, t + t', \theta_K)$ ,  $V_{t+t'} = f_V(TOU_{t+t'}, t + t', \theta_V)$ , and  $Q_{t+t'} = f_Q(TOU_{t+1:t+H}, \theta_Q)$ . Subsequently, the output for the part of the exogenous branch processing the tariffs at time  $t + t'$  is given by

$$\text{Att}(Q_{t+t'}, K_{t+t'}, V_{t+t'}) = \text{softmax}\left(\frac{Q_{t+t'} K_{t+t'}^T}{\sqrt{d}}\right) V_{t+t'}, \tag{3}$$

where  $d$  is the dimension of  $Q, K$ , and  $V$ . While the  $f_K$  and  $f_V$  are implemented as simple linear layers,  $f_Q$  is implemented as a permutation equivariant network as follows:

$$f(x) = \sigma(x\Lambda - \mathbf{1}\text{maxpool}(x)\Gamma) \tag{4}$$

where  $x = \text{ReLU}(TOU_{t+1:t+H}, \theta_{TOU}) \in \mathbb{R}^{H \times d}$  and  $\theta$  shared across timesteps  $t + 1 \dots t + H$ ,  $\Lambda, \Gamma \in \mathbb{R}^{d \times d'}$ , matrix of ones  $\mathbf{1} \in 1^{H \times H}$ ,  $\text{maxpool}$  is taken along columns implying that the resulting value for any timestep contains information from all timesteps and is independent of a particular timestep. In this work, we use  $d = 10, d' = 20$ .

**Objective function:** We use quantile loss for training the DCNN model given by:

$$\mathcal{L}_{\text{quantile}} = \frac{1}{b \times n} \sum_{i=1}^b \sum_{q=q_1}^{q_n} \max(q \times e^i, (q - 1) \times e^i), \tag{5}$$

where  $e^i = y^i - \hat{y}^i$  indicates the error of the forecasted consumption  $\hat{y}^i$  with respect to ground-truth consumption  $y^i$  of  $i$ -th window instance,  $b$  is the batch size and  $n$  is the number of quantiles used for training.

### 6. Experimental Evaluation

The goal is to evaluate the efficacy of the proposed approach to deal with OOD scenarios. For this, we compare the proposed approach with various baselines in the IID as well as OOD settings. We use the simulated data from a high-fidelity and popular PowerTAC (<https://powertac.org/>, accessed on 12 November 2021) [21] simulator that uses complex state-of-the-art user-behavior models and real world weather data to simulate the complex dynamics of a smart grid system.

We consider ‘Office Complex Controllable type’ consumers where consumers’ daily behavior depends on factors such as number of sub-customers, number of appliances, weather information, hour of day, month, day of week, etc. The various values these factors can take across consumers is given in Table 1.

**Table 1.** Dataset details.

S.N.	Properties of Consumers	Value(s)
1	Number of consumers	12
2	Number of sub-customers	3, 5
3	Working days	3, 4
4	Work Start hour	{8, 9, 10} (+/-) 1 h
5	Break Start hour	{13, 14} (+/-) 1 h
6	Work duration	8 (+/-) 1 h
7	Shiftable consumption( in KW)	600, 2400
8	Total data duration (in months)	6

To obtain train, validation, and test split, we divide the total data of 6 months into 4, 1, and 1 month, respectively. The time series of hourly data for each consumer is divided into

windows of length  $t = 168$  (corresponding to 7 days) with window-shift of 24 to forecast one day-head consumption, i.e., output window size is 24. We consider varying number of tariff profiles in historical data, i.e.,  $|\mathcal{T}_{in}| \in \{2, 5, 8, 10, 12, 15, 20, 25, 30, 35\}$ , and an additional set of  $|\mathcal{T}_{out}| = 40$  profiles. As the number of profiles  $|\mathcal{T}_{in}|$  in the training set increases, we expect the bias in the training data to reduce.

### 6.1. Baselines Considered

For comparison, we consider the following approaches all using DCNN as the core time series processing module:

- **No future exogenous variable (NoX)** is the simple univariate time series forecasting approach which uses only history of aggregated consumption without any additional future information. This can be considered as a lower bound in the sense that the network does not have access to any future tariff rates to estimate where a consumer will shift the load.
- **Independent tariff-based method (Ind.)** is an approach that treats each tariff rate independently, and uses the tariff at time  $t + t'$  to estimate the aggregated consumption at that time. Importantly, this approach has no means to capture comparison of the tariff rates in order to figure out whether the tariff at time  $t + t'$  is high or low in comparison to another timestep.
- **Fully-Connected Approach (FC)** utilizes the information of all timesteps to estimate the aggregated consumption at each timestep. As explained previously, we expect such an approach to perform well in the IID scenario but struggle in the OOD scenario where new profiles are included.
- **Permutation Equivariant (PE)** method uses only the permutation equivariance idea from our approach and ignores the attention mechanism. This method can be thought of as an ablation over our approach.
- **Attention (Att.):** This is another ablation over our approach which uses standard attention module for processing the tariffs along with hour of the day information without any permutation equivariance property.
- **Upper Bound (UB):** This is an oracle approach that assumes knowledge about the hours at which the consumer is going to shift the load. In this, a binary value indicating whether the shiftable load will be shifted to this hour or not is passed as an additional feature to the exogenous branch of the Att.+PE network.

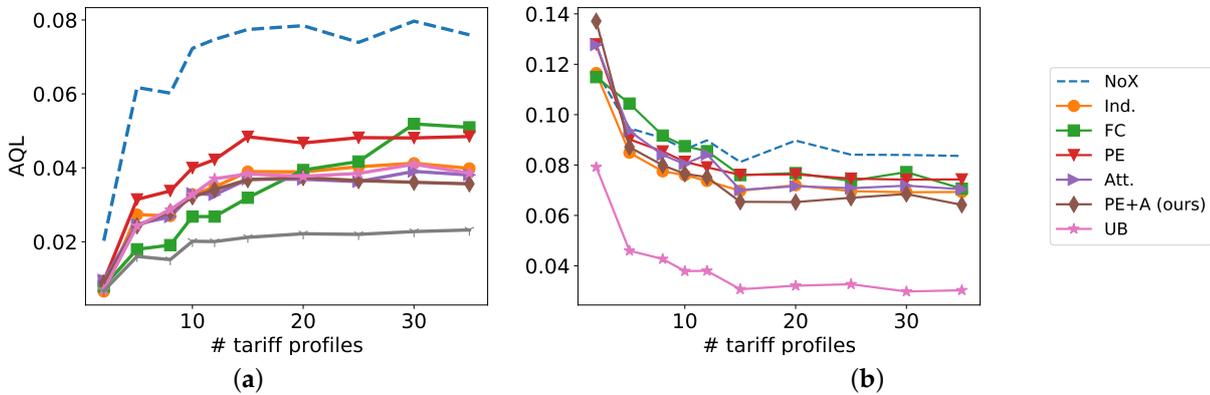
### 6.2. Hyperparameters Used

We use z-normalized consumption time series. DCNN has three layers with each layer having 16 convolutional filters of length 2, and dilation rate 1, 2, and 4, respectively. We use batch normalization and L2 filter regularizer ( $\lambda = 0.001$ ) for regularization purposes. ReLU layers are applied on each CNN layer. The output of the DCNN layer is processed by a channel-wise fully connected layer, which has 24 hidden units (equal to the output window size) i.e., 24, followed by locally connected layer with 10 filters which are applied at each time-step independently (filter size = 1).

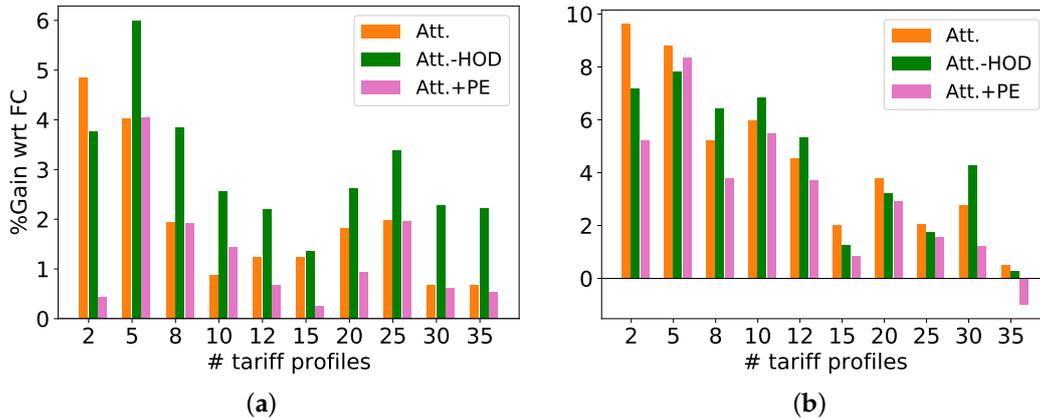
To obtain categorical feature (hour of day, day of week, month of year) embeddings and tariff rate embeddings, we use a separate feed-forward network with ReLU layer followed by linear layer, having 5 hidden units and 10 hidden units respectively. Similarly, we use 10 hidden units for each feed-forward network  $f_Q, f_K, f_V$ . Finally, the output layer is a small feed-forward network that has 2 layers followed by a linear layer having 40, 10, and 1 hidden unit, respectively. We use batch size of 16, number of epochs 200, and Adam optimizer with fixed learning rate of 0.0001 for training the neural network. During training, quantiles are sampled from uniform distribution while during validation and testing, we use three quantiles 0.1, 0.5, and 0.9. All hyperparameters were obtained via grid search based on validation quantile loss on the IID set.

### 6.3. Results and Observations

We make following key observations from the results in Figures 6 and 7:



**Figure 6.** Forecasting performance Comparison of different approaches (in terms of Average Quantile Loss). (a) IID Scenario. (b) OOD Scenario.

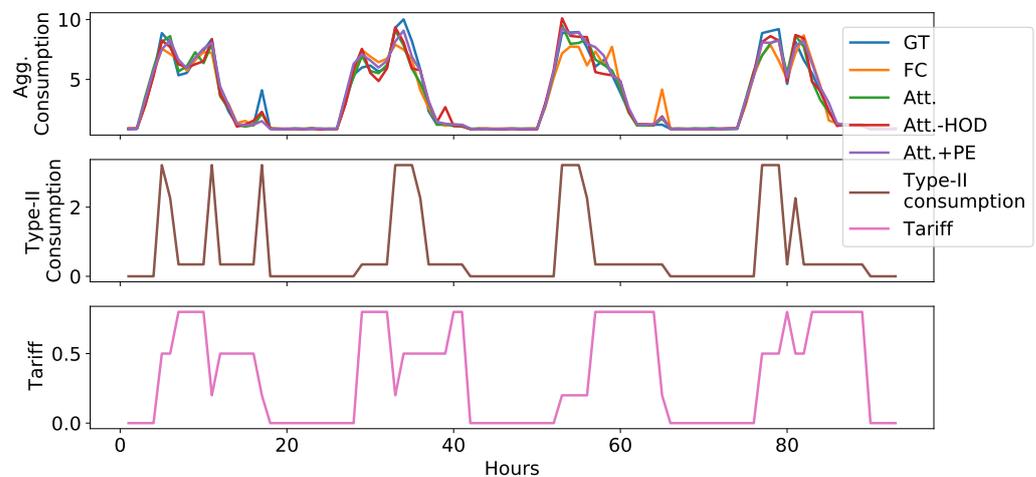


**Figure 7.** %gains of the proposed Att.+PE, Att.-HOD, and Att. approaches over the vanilla FC approach. (a) Option-1. (b) Option-2.

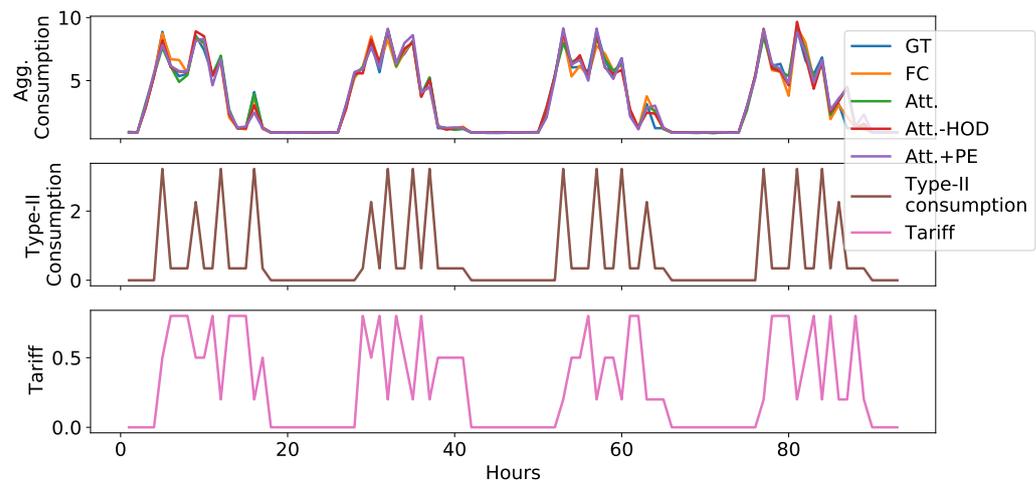
- Observations from forecasting results as shown in Figure 6:
  - In the IID scenario, the average quantile loss (AQL) for all approaches increases with increasing number of tariff profiles as the complexity of the dataset increases. The FC approach performs better than other approaches for  $|\mathcal{T}_{in}| \leq 15$ , indicating higher expressivity of the FC approach to fit to a smaller number of IID profiles, indicating potential overfitting.
  - On the other hand, for the OOD scenario, the performance of all approaches improves with increasing number of IID profiles which is expected as more IID profiles implies less bias and better generalization to OOD profiles as well. Interestingly, the FC approach which was the best approach for the IID profiles for  $|\mathcal{T}_{in}| \leq 12$ , is the worst approach (except the lower bound NoX) in the OOD setting, because it uses a fully connected layer to process the tariffs of the day, and due to temporal bias in the data, the weights of fully connected layer will try to overfit on  $|\mathcal{T}_{in}|$  and thus not generalize to OOD profiles  $|\mathcal{T}_{out}|$ . On the other hand, our proposed approaches Att.+PE and Att.-HOD are consistently better than FC for all values of  $\mathcal{T}_{in}$ , which shows that FC struggles with the temporal bias in the historical data. We also analyze that Att.-HOD as well as Att.+PE are also consistently better than Att. for all values of  $\mathcal{T}_{in}$ , which shows that permutation equivariant way of handling tariff profiles provide better generalization on OOD profiles.

- We further analyze whether the gains of Att.+PE and Att.-HOD over other methods on the OOD scenario translate into more profitable tariff profile allocation for the retailer. We compare the gain  $G$  of Att.+PE, Att.-HOD, and Att. in comparison to FC. We consider two kinds of profiles for wholesale prices  $p$ , one with two values (0.2 and 0.8, referred to as Option-1) and one with three values (0.2, 0.5, and 0.8, referred to as Option-2).
  - **Comparison with FC:** We observe that all attention-based proposed approaches Att., Att.-HOD, and Att.+PE depict significant positive gains over FC. We also observe that Att., Att.-HOD, and Att.+PE approaches have higher positive gain in fewer IID tariff profiles scenarios  $|\mathcal{T}_{in}| \leq 12$  (except  $|\mathcal{T}_{in}| = 2$ , where data is too little to claim any generalization), and the gains tend to diminish as  $|\mathcal{T}_{in}|$  increases.
  - As expected, we note that it is not important that the gains in forecasting translate directly into monetary profits, as the optimization objective involves other terms such as wholesale costs  $p$ . Therefore, the best approach on forecasting (Att.+PE) in the OOD scenario is not necessarily the best approach in terms of profit always.
  - **Comparison with Att.:** For Option-1, Att.-HOD has significantly better gains than Att. for all values of  $\mathcal{T}_{in}$  except  $|\mathcal{T}_{in}| = 2$ , which shows that the permutation equivariant way of handling tariff profiles is helpful. For Option-2, the gains of Att.-HOD are better or close to the gains of Att. approach (except  $|\mathcal{T}_{in}| = 2$ ).

In Figure 8, we also provide sample forecasts comparing Att., Att.-HOD, Att.+PE, and FC with the ground truth (GT) on an OOD profile, indicating better generalization ability of Att.-HOD and Att.+PE, especially around points where Type-II load gets shifted. On the other hand, all methods perform well in the IID setting as shown in Figure 9.



**Figure 8.** Sample results comparing the proposed approaches Att.-HOD and Att.+PE with FC on an OOD tariff profile. Here, GT: Ground Truth time series. FC struggles to capture the subtle changes in consumption due to shifting of load, while both Att.-HOD and Att.+PE are able to forecast better.



**Figure 9.** Sample results comparing the proposed approaches Att.-HOD and Att.+PE with FC on an IID tariff profile. Here, GT: Ground Truth time series. In IID scenario, all proposed attention-based approaches and baselines perform well.

## 7. Conclusions and Future Work

In this work, we consider the problem of demand response management from an electricity broker or retailer’s perspective. We highlight temporal bias as an issue in optimizing profits via suitable tariff profile allocations. We motivate the need for better generalization to out-of-distribution profiles, and note that this is possible by leveraging the fact that consumers respond with same logic across profiles. We propose suitable inductive biases in deep neural networks-based approach for forecasting electricity consumption in response to new tariff profiles. This takes the form of a permutation equivariance-enabled attention mechanism that can leverage the property of consumer behavior to respond in a certain way across profiles. In the future, it will be interesting to look at the generalization from the perspective of handling confounding bias as the historical profile allocation and the outcome are affected by the historical allocation policies, which in turn rely on the latent consumer attributes acting as confounders. The current optimization objective takes into account broker’s profit but ignores the cost of electricity for the end consumer—bringing this into the optimization objective is a potential next step.

**Author Contributions:** Conceptualization, methodology, resources, software, formal analysis and writing of original draft, P.M., J.N. and C.V.; Validation and data curation, J.N. and C.V.; writing—editing and review L.V., E.S. and S.B.; supervision, P.M., L.V., E.S. and S.B. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** We use simulated data from PowerTAC simulator. Further details about the data are provided in Section 6. Data is confidential, so we can not provide the simulated data.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. Siano, P. Demand response and smart grids—A survey. *Renew. Sustain. Energy Rev.* **2014**, *30*, 461–478. [[CrossRef](#)]
2. Lu, R.; Hong, S.H. Incentive-based demand response for smart grid with reinforcement learning and deep neural network. *Appl. Energy* **2019**, *236*, 937–949. [[CrossRef](#)]
3. Lu, R.; Hong, S.H.; Zhang, X. A dynamic pricing demand response algorithm for smart grid: Reinforcement learning approach. *Appl. Energy* **2018**, *220*, 220–230. [[CrossRef](#)]
4. Yang, P.; Tang, G.; Nehorai, A. A game-theoretic approach for optimal time-of-use electricity pricing. *IEEE Trans. Power Syst.* **2012**, *28*, 884–892. [[CrossRef](#)]

5. Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, BC, Canada, 11–17 October 2021; pp. 8340–8349.
6. Arjovsky, M. Out of Distribution Generalization in Machine Learning. Ph.D. Thesis, New York University, New York, NY, USA, 2020.
7. Krueger, D.; Caballero, E.; Jacobsen, J.H.; Zhang, A.; Binas, J.; Zhang, D.; Le Priol, R.; Courville, A. Out-of-distribution generalization via risk extrapolation (rex). In Proceedings of the International Conference on Machine Learning. PMLR, Virtual Event, Switzerland, 7–8 June 2021; pp. 5815–5826.
8. Sun, Y.; Wang, X.; Liu, Z.; Miller, J.; Efros, A.A.; Hardt, M. Test-time training for out-of-distribution generalization. In Proceedings of the Eighth International Conference on Learning Representations (ICLR 2020), Virtual Conference, 26 April–1 May 2020.
9. Hospedales, T.; Antoniou, A.; Micaelli, P.; Storkey, A. Meta-learning in neural networks: A survey. *arXiv* **2020**, arXiv:2004.05439.
10. Wang, T.; Liao, R.; Ba, J.; Fidler, S. Nervenet: Learning structured policy with graph neural networks. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
11. Narwariya, J.; Malhotra, P.; TV, V.; Vig, L.; Shroff, G. Graph Neural Networks for Leveraging Industrial Equipment Structure: An application to Remaining Useful Life Estimation. *arXiv* **2020**, arXiv:2006.16556.
12. Andreas, J.; Rohrbach, M.; Darrell, T.; Klein, D. Neural module networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2016; pp. 39–48.
13. Bansal, H.; Bhatt, G.; Malhotra, P.; Prathosh, A. Systematic Generalization in Neural Networks-based Multivariate Time Series Forecasting Models. In Proceedings of the 2021 International Joint Conference on Neural Networks (IJCNN), Shenzhen, China, 18–22 July 2021; pp. 1–8.
14. Liu, T.; Lu, J.; Yan, Z.; Zhang, G. Statistical generalization performance guarantee for meta-learning with data dependent prior. *Neurocomputing* **2021**, *465*, 391–405. [[CrossRef](#)]
15. Pearl, J.; Glymour, M.; Jewell, N.P. *Causal Inference in Statistics: A Primer*; John Wiley & Sons: Hoboken, NJ, USA, 2016.
16. Salinas, D.; Flunkert, V.; Gasthaus, J.; Januschowski, T. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *Int. J. Forecast.* **2020**, *36*, 1181–1191. [[CrossRef](#)]
17. Liu, X.; Yin, J.; Liu, H.; Liu, J. DeepSSM: Deep State-Space Model for 3D Human Motion Prediction. *arXiv* **2020**, arXiv:2005.12155.
18. Zaheer, M.; Kottur, S.; Ravanbakhsh, S.; Póczos, B.; Salakhutdinov, R.; Smola, A. Deep sets. *arXiv* **2017**, arXiv:1703.06114.
19. Lee, J.; Lee, Y.; Kim, J.; Kosiorek, A.; Choi, S.; Teh, Y.W. Set transformer: A framework for attention-based permutation-invariant neural networks. In Proceedings of the International Conference on Machine Learning. PMLR, Long Beach, CA, USA, 9–15 June 2019; pp. 3744–3753.
20. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [[CrossRef](#)] [[PubMed](#)]
21. Ketter, W.; Collins, J.; Reddy, P. Power TAC: A competitive economic simulation of the smart grid. *Energy Econ.* **2013**, *39*, 262–270. [[CrossRef](#)]