

RESEARCH ARTICLE

Open Access



# CLear: An adaptive continual learning framework for regression tasks

Yujiang He\* and Bernhard Sick

## Abstract

Catastrophic forgetting means that a trained neural network model gradually forgets the previously learned tasks when being retrained on new tasks. Overcoming the forgetting problem is a major problem in machine learning. Numerous continual learning algorithms are very successful in incremental learning of classification tasks, where new samples with their labels appear frequently. However, there is currently no research that addresses the catastrophic forgetting problem in regression tasks as far as we know. This problem has emerged as one of the primary constraints in some applications, such as renewable energy forecasts. This article clarifies problem-related definitions and proposes a new methodological framework that can forecast targets and update itself by means of continual learning. The framework consists of forecasting neural networks and buffers, which store newly collected data from a non-stationary data stream in an application. The changed probability distribution of the data stream, which the framework has identified, will be learned sequentially. The framework is called CLear (**C**ontinual **L**earning for **R**egression Tasks), where components can be flexibly customized for a specific application scenario. We design two sets of experiments to evaluate the CLear framework concerning fitting error (training), prediction error (test), and forgetting ratio. The first one is based on an artificial time series to explore how hyperparameters affect the CLear framework. The second one is designed with data collected from European wind farms to evaluate the CLear framework's performance in a real-world application. The experimental results demonstrate that the CLear framework can continually acquire knowledge in the data stream and improve the prediction accuracy. The article concludes with further research issues arising from requirements to extend the framework.

**Keywords:** Continual learning, Renewable energy forecasts, Regression tasks, Deep neural networks

## Introduction

In the late 1980s, McCloskey and Cohen [1] and Ratcliff [2] observed a phenomenon where the well-learned knowledge of connectionist models is erased by new knowledge under specific conditions when the models learn new tasks successively. It is referred to as Catastrophic Forgetting or Catastrophic Interference. The challenge is supposed to be a general problem existing in different kinds of neural networks, e.g., back-propagation neural networks and unsupervised neural networks. Each neuron of a layer is connected to all neurons of the next layer in a neural network. Weights control the strength

of the connection between two neurons. A weight vector can be expressed as a column vector  $\mathbf{w} = (w_1, w_2, \dots, w_N)^T$  with  $N$  values. Some weights have a significant influence on more than one task. For example,  $w_1, w_2, w_3$  are important for task 1 and  $w_1, w_4, w_5$  are important for task 2. In this case, the overlapped weight  $w_1$  could be adjusted during learning task 2 sequentially, which is one of the main reasons for Catastrophic Forgetting.

Overcoming the forgetting problem is a crucial step in implementing real intelligence. Models require plasticity for learning and integrating new knowledge as well as stability for consolidating what models have learned previously. Excessive plasticity can cause the acquired knowledge to be erased while learning new tasks. On the other hand, successively learning new tasks can become more

\*Correspondence: [yujiang.he@uni-kassel.de](mailto:yujiang.he@uni-kassel.de)

Intelligent Embedded Systems (IES) Group, University of Kassel, Wilhelmshöher Allee 71 - 73, Kassel, Germany

challenging due to extreme stability. It is the so-called stability-plasticity dilemma [3].

Many researchers have proposed continual learning (CL) algorithms to solve the problem, such as in [4–6]. A three-way categorization for the most common CL strategies is described in [6]: (1) regularization strategies, (2) rehearsal strategies, and (3) architectural strategies. Similarly, CL strategies are grouped as (1) prior-focused approaches, (2) likelihood-focused approaches, and (3) dynamic architectures in [7]. These categorizations standardize terminologies and outline a distinct research direction for the CL community. CL algorithms have been proven successful in supervised learning and reinforcement learning to train several tasks sequentially without forgetting the acquired ones. Application scenarios cover handwriting recognition [8–10], image classification [5], sequentially learning to play games in a reinforcement learning setting [4] and much more.

To our best knowledge, some of the standard CL benchmarks are the reconstructions of well-known datasets, such as permuted MNIST [11, 12], where the CL tasks are obtained by scrambling the pixel positions in the MNIST dataset [13]. Moreover, some datasets are explicitly generated to evaluating CL algorithms, for example, CORE50 [14] for continuous object recognition.

Most of the past research focused mainly on classification tasks rather than regression tasks, where the Catastrophic Forgetting problem usually occurs as well. For example, establishing regional smart grids requires power generation and consumption forecasts. In [15, 16], neural networks are used to forecast renewable energy generation with weather prediction data. Note that weather data is non-stationary, where the probability distribution changes over time, e.g., from summer to winter. In this situation, training neural networks has to be delayed until sufficient data is collected. Otherwise, the networks could be overfitted to the limited training dataset. Furthermore, unseen situations, e.g., the extreme weather conditions, updating/damaging/ageing of generators, and climate changes, can be called special events. The neural networks have to update themselves by continually learning these situations when they appear in the application. Besides, power consumption regarding a household or a factory is also easily affected by unpredictable things, such as purchasing new equipment or hiring more employees. These factors can change the obtained mapping between inputs and outputs. Under these conditions where historical data may be private, unrecorded, or too cumbersome to be retrained, the trained models have to learn new knowledge and consolidate the previously-stored internal representations without the help of old data.

The **main contribution** of this article is to propose a framework called CLear (Continual Learning for Regression Tasks) for continually learning the identified

changes in non-stationary data streams. Moreover, the framework is tested in two sets of experiments to assess its performance and analyze its hyperparameters' effects. The CLear framework consists of neural networks for prediction and buffers for storing new data. We calculate an error between the prediction and the corresponding true value at each point in time. The new data is labeled by comparing the error to a dynamically adjustable threshold. If the error is larger than the threshold, the data is labeled as a novelty and stored in a finite novelty buffer, or else stored as familiarity in an infinite familiarity buffer. When the novelty buffer is full, updating will be triggered. The network will be retrained on the dataset in the novelty buffer using CL. The retrained network will then be tested on the familiarity dataset to evaluate how much old knowledge is retained. After updating, the threshold needs to be re-estimated for the following learning step. The re-estimation depends on the performances of the updated network on the dataset of both buffers. Afterwards, the buffers will be emptied. Updating will be repeated until the novelty buffer is filled again.

The remainder of the article will review the literature regarding CL algorithms and applications. Then we will outline the proposed framework's fundamental structure and give an insight into the experimental details. Furthermore, we will analyze the experimental results. This article will conclude with our findings and provide an outlook for future research.

### Related work

This section will start with a brief overview of the recent academic literature regarding approaches and experimental setups designed for CL classification applications.

The changes in data or goals can be defined as new tasks in the CL community. For example, a model is expected to learn new instances of the same class while retaining its knowledge about the previous instances, or to learn new instances of different classes without losing accuracy on previous classes, or to learn new instances of the known and unknown classes. These are defined as different CL scenarios in [14]. In both [6, 7], CL algorithms are categorized into three groups in a similar way:

- *Prior-focused approaches* denote that the posterior probability of  $N$  tasks is a product of the likelihood of the  $N$ th task and the posterior probability of the first  $N - 1$  tasks, as

$$P(\Theta|D_1, \dots, D_N) = \frac{P(D_N|\Theta)P(\Theta|D_1, \dots, D_{N-1})}{P(D_N|D_1, \dots, D_{N-1})}. \quad (1)$$

As a regularization, the posterior probability of the first  $N - 1$  tasks is added in the loss function to avoid changing the weights that are important for previous

tasks. Well-known prior-focused algorithms include Elastic Weight Consolidation (EWC) [4], Synaptic Intelligence (SI) [5], Variational Continual Learning [8]. In [9], a generalization of EWC++ and SI was proposed, which is referred to as the RWalk algorithm.

- *Likelihood-focused approaches* require a subset of randomly selected samples from the original dataset of the previous  $N - 1$  tasks, or a dataset generated by a generative network that has learned the tasks, see in [7, 10, 17].
- *Dynamic architectures* enable neural networks to learn CL tasks sequentially by adjusting the networks' architecture for specific applications. Progressive Networks, Learning Without Forgetting (LWF), and Less-Forgetting Learning have been introduced in [18–20], respectively.

The above algorithms have been evaluated in multi-task scenarios with the reshaped versions of famous datasets, e.g., MNIST and CIFAR-10/CIFAR-100. In these scenarios, a model learns a new, isolated task in a sequence while remembering how to solve the learned tasks. However, there are no class overlaps among the different tasks. For example, in [5] the MNIST dataset is split into five tasks, one of which contains two labels (two digits). The model can classify data to the correct group only if the information regarding the current task is given. In this case, the model learns how to solve a series of discrete tasks rather than keep learning knowledge to address incremental problems. The experimental setups and the datasets do not allow for a fair comparison among the CL algorithms. Lomonaco et al. [14] create CORE50 specifically for single-incremental-task scenarios, which can be seen as a test benchmark for continuous object recognition. Similarly, the iCubWorld benchmark [21] is designed for robotic vision challenges, where comparison among various CL approaches is feasible.

Besides, continual learning should be considered in regression as well. In [22], He et al. propose two CL application scenarios for establishing regional smart grids: the task-domain incremental scenario and the data-domain incremental scenario. The scenarios are applicable for forecasting power, including renewable energy generation and power consumption in the middle-/low-voltage grid. Moreover, performances of four CL algorithms (EWC, Online-EWC, SI, and LWF) are evaluated concerning accuracy, forgetting ratio, and training time in the two scenarios. However, prior knowledge about new tasks is given in their experimental setup, which means that models know when new tasks will occur without novelty detection. Therefore, this setup is incompatible with the real world.

In [23], Farquhar et al. conclude that an inappropriate experimental design could misrepresent the performances

of the well-known CL approaches. Therefore, they suggest five requirements for evaluating CL algorithms and demonstrate their necessities. The five requirements are: (1) cross-task resemblances; (2) shared output head; (3) no test-time assumed task labels; (4) no unconstrained retraining on old tasks; and (5) more than two tasks.

These suggestions are worth being considered in our experimental setup and inspire us to design the CLear framework. (1) Most novelties are due to changes of data  $P(X)$  or targets  $P(Y|X, \Theta)$ . The dataset of every full novelty buffer can be viewed as a new task that resembles the previous tasks. (2) The neural network outputs the power value prediction, and the new tasks will not require a change of the network's architecture. (3) The prior knowledge regarding the new task, such as when the new task appears or what the distribution of the new task is, is unknown in the application. Updating is triggered automatically only when the finite novelty buffer is filled in our experimental setup. (4) Considering that privacy laws might prohibit the long-term storage of historical datasets, we retrain the neural network only on the dataset newly collected in applications and delete it after updating. (5) More tasks will appear as the probability distribution or the mapping between inputs and outputs changes over time.

### Power forecasts using deep neural networks

At the beginning of this section, we list the chosen mathematical notations in Table 1. It can help readers to understand the following mathematical expressions. Besides, we

**Table 1** Notations

Symbol	Definition
$X$	The $N \times M$ matrix with $N$ feature samples
$Y$	The matrix with $N$ measurements associated with $X$
$\hat{Y}$	The matrix with $N$ predictions associated with $X$
$\mathbf{x}_n$	The $n$ th column feature vector, $\mathbf{x}_n^\top = X_{n,:}$
$y_n$	The $n$ th measurement, $y_n = Y_{n,:}$
$\hat{y}_n$	The $n$ th prediction, $\hat{y}_n = \hat{Y}_{n,:}$
$D$	The dataset, $D = \{(\mathbf{x}_n, y_n)   n = 1, \dots, N\}$
$\Theta_l$	The weight matrix of the $l$ th layer, $\Theta_l \in \mathbb{R}^{l-1 \times l}$
$\theta_l^j$	The $i$ th element of $\Theta_l$
$f_l$	The activation function of the $l$ th layer, $f_l: \mathbb{R} \rightarrow \mathbb{R}$
$f_\Theta(\cdot)$	The neural network with given weight matrix $\Theta$
$\mathbf{z}_l$	The output column vector of the $l$ th layer
$L$	Loss function
$p$	Probability density
$\mathcal{N}(\mu, \sigma^2)$	Gaussian distribution with mean $\mu$ and variance $\sigma^2$
$\mathbf{F}_t$	The Fisher information matrix of the $t$ th task
$F_t^i$	The $i$ th diagonal element of $\mathbf{F}_t$

use superscript  $T$  to denote the transpose of a matrix or a vector and  $T$  to denote the number of tasks.

Deep neural networks are a kind of machine learning inspired by biological neural networks to model nonlinear dependencies in high dimensional data. Compared with traditional high dimensional data reduction techniques, such as principal component analysis (PCA), the multiple deep layers of a neural network can extract representations efficiently from massive data to provide predictive performance gains.

The rest of this section will start with formulating the problem. Moreover, this section will illustrate the probability distribution changes in the experimental dataset over time and clarify how the CLear framework works in a general power forecasting workflow.

**Problem formulation**

Deep neural networks can output a prediction  $y_n$  with given a high-dimensional input  $\mathbf{x}_n$ . The goal of training is to find a mapping between  $\mathbf{x}_n$  and  $y_n$ . A general deep neural network consists of a series of hidden layers, which can be formulated as:

$$\mathbf{z}_l = f_l \left( \Theta_l^T \mathbf{z}_{l-1} \right) \tag{2}$$

with an output column vector  $\mathbf{z}_{l-1}$  of the  $l - 1$ th layer, where  $\mathbf{z}_0 = \mathbf{x}_n$ . Here  $\Theta_l^T$  denotes a transposed weight matrix, whose dimension is the dimension of the  $l - 1$ th layer by the dimension of the  $l$ th layer. A prediction  $\hat{y}_n$  of the deep neural network is then

$$\begin{aligned} \hat{y}_n &= f_\Theta (\mathbf{x}_n) \\ &= f_L \left( \Theta_L^T f_{L-1} \left( \dots \Theta_2^T f_1 \left( \Theta_1^T \mathbf{x}_n \right) \right) \right), \end{aligned} \tag{3}$$

where  $L$  is the number of layers.

Training a neural network is to minimize the defined loss function. In this section, we explain the process with an example of Mean Square Error (MSE), i.e.,

$$\mathbb{L} (Y, \hat{Y}) = \frac{1}{N} \sum_{n=1}^N \|y_n - \hat{y}_n\|_2^2. \tag{4}$$

In order to avoid overfitting, a regularization term is usually added in Eq. 4, as

$$L (Y, \hat{Y}) = \frac{1}{N} \sum_{n=1}^N \|y_n - \hat{y}_n\|_2^2 + \lambda R (\Theta), \tag{5}$$

where  $\lambda \in (0, \infty)$  is a hyperparameter that weights the contribution of penalty term  $R (\Theta)$ . Different choices for  $R (\Theta)$  can result in different solutions.

From a probabilistic perspective, according to chapter 9 in [24], we can assume that we are given an input  $\mathbf{x}_n$  and the corresponding noisy observation  $y_n = \hat{y}_n + \epsilon$ . More specifically, we assume further that this noise  $\epsilon$  follows an

independent and identical Gaussian distribution with zero mean and variance  $\sigma^2$ . Therefore, the regression problem can be considered with a likelihood function:

$$P (y_n | \mathbf{x}_n, \Theta) = \mathcal{N} (f_\Theta (\mathbf{x}_n), \sigma^2). \tag{6}$$

When we are given the datasets  $X$  and  $Y$ , Eq. 6 can be expressed as

$$\begin{aligned} P (Y|X, \Theta) &= \prod_{n=1}^N P (y_n | \mathbf{x}_n, \Theta) \\ &= \prod_{n=1}^N \mathcal{N} (f_\Theta (\mathbf{x}_n), \sigma^2), \end{aligned} \tag{7}$$

where assumed that the  $y_i$  and  $y_j$  are conditionally independent given their feature vectors  $\mathbf{x}_i$  and  $\mathbf{x}_j$ . To avoid overfitting during training, we seek parameters  $\Theta$  that maximize the posterior distribution  $P (\Theta|X, Y)$  instead of the likelihood. We can obtain the posterior distribution by applying Bayes' theorem as

$$P (\Theta|X, Y) = \frac{P (Y|X, \Theta) P (\Theta)}{P (Y|X)} \tag{8}$$

Note that the posterior distribution depends on the given  $X$  and  $Y$ . If the statistical properties of the distributions change, such as the mean or the variance, new parameter values will become optimal.

**Change in probability distribution**

A learning task can be described as changes in distributions of data  $P(X)$  or targets  $P(Y|X, \Theta)$ . In our experimental setup, meteorological features and the power measurements are viewed as the inputs and outputs of the neural network, respectively. Some meteorological features, e.g., temperature and wind direction, fluctuate yearly periodically. Also, renewable energy generation depends on meteorological conditions. Therefore, the joint probability distribution  $P(X, Y)$  changes period by period. Mathematically, the change can be expressed as follows:

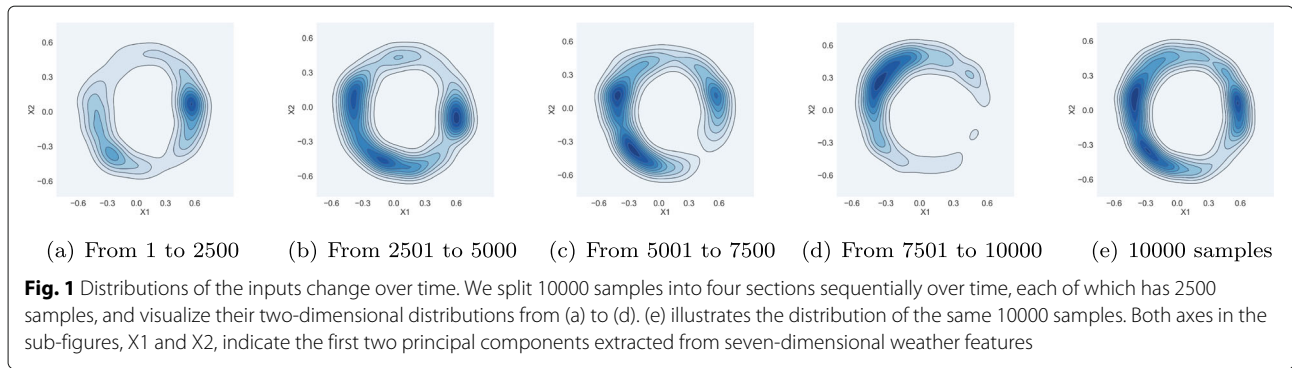
$$\exists t_0 \neq t_1 : P (X_{t_0}, Y_{t_0}) \neq P (X_{t_1}, Y_{t_1}), \tag{9}$$

where  $X$  and  $Y$  represent a batch of features and power measurements respectively, which are measured from a non-stationary data stream. If two measurement periods  $t_0$  and  $t_1$  are far apart, both distributions are different due to concept drift. The formula:

$$P (X, Y) = P (Y|X) P(X) \tag{10}$$

demonstrates that  $P(X, Y)$  is affected by the change in the probability distribution of the inputs and the obtained mapping.

Figures 1 and 2 illustrate the change regarding  $P(X)$  and  $P(Y|X)$  based on data from a European wind farm dataset.



A sample of the dataset contains seven-dimensional meteorological features and a scalar power value. The first two principal components are extracted from the seven features using PCA and labeled as  $X_1$  and  $X_2$ . 64.19% of the overall variance is explained by the two components. We split 10000 samples from the dataset into four sections sequentially over time, each of which has 2500 samples. The distribution of each section regarding  $X_1$  and  $X_2$  is plotted in Fig. 1. Similarly, the distributions of four sections regarding  $X_1$  and the power  $Y$  are shown in Fig. 2. The distributions of  $P(X_1, X_2)$  and  $P(X_1, Y)$  with the same 10000 samples are plotted in Figs. 1(e) and 2(e) in comparison to other sub-figures.

The two-dimensional probability distribution of  $P(X_1, X_2)$  with 10000 samples shows a circular shape, as shown in Fig. 1(e). From Fig. 1(a) to (d), the center of the contour line regarding each section moves clockwise along the circle. This movement illustrates the periodic change in the  $P(X)$  over time. Similarly, we can observe that the center of the distribution  $P(Y|X_1)$  moves from right to left over time, as shown in Fig. 2.

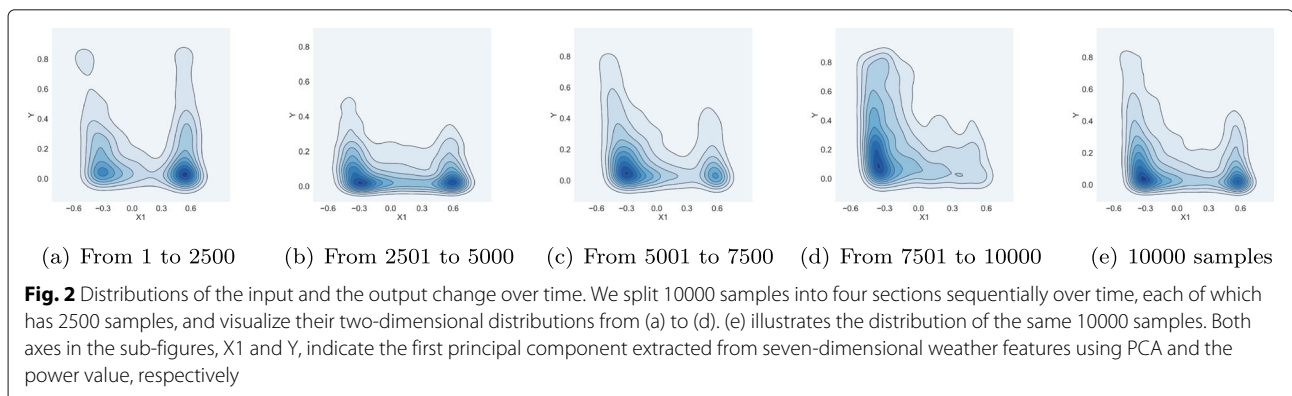
Combined with Eq. 9 and the observations, we conclude that a periodic change exists in the dataset. The model needs to be updated in applications if it was pre-trained only on a limited dataset.

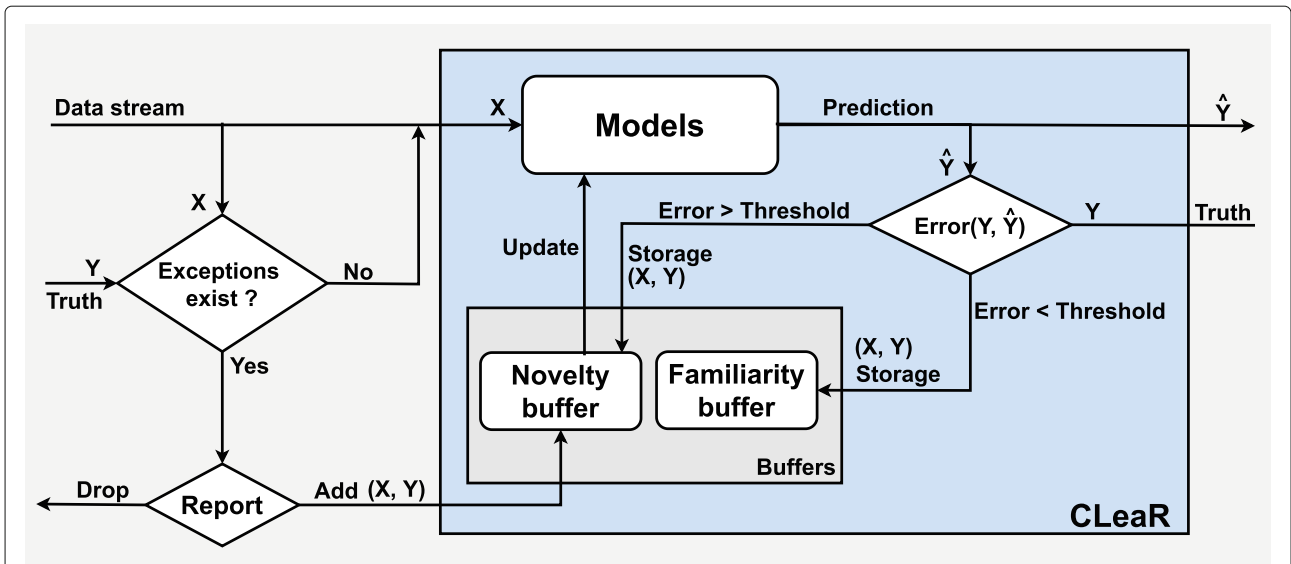
### Power forecasting workflow

We suggest that a general power forecasting workflow should comprise four phases: (1) reporting exceptions, (2) predicting targets, (3) storing data, and (4) updating models, as shown in Fig. 3.

We define that an exception deviating from the expected model is an outcome from an unknown process. For example, wind power generators are automatically shut down for protection under extreme weather conditions, such as typhoons or storms. Exceptions have to be identified first as the data arrives. If it exists in the data stream, the exceptions need to be reported and processed manually. Although the dataset used here has been cleaned in the preprocessing phase, these exceptions might appear in applications. Labeling and learning exceptions are one of the research challenges in the field of active learning, which is beyond the scope of this article. It can be further researched in the future.

The CLear framework contains the models for prediction and the buffers for storage. Once the measurements  $Y$  are available, the corresponding data is labeled and stored into the two buffers depending on the preset threshold and the error. We choose the MSE here for supervised regression tasks, but the method can be customized in other scenarios, e.g., probabilistic forecasts. The data in



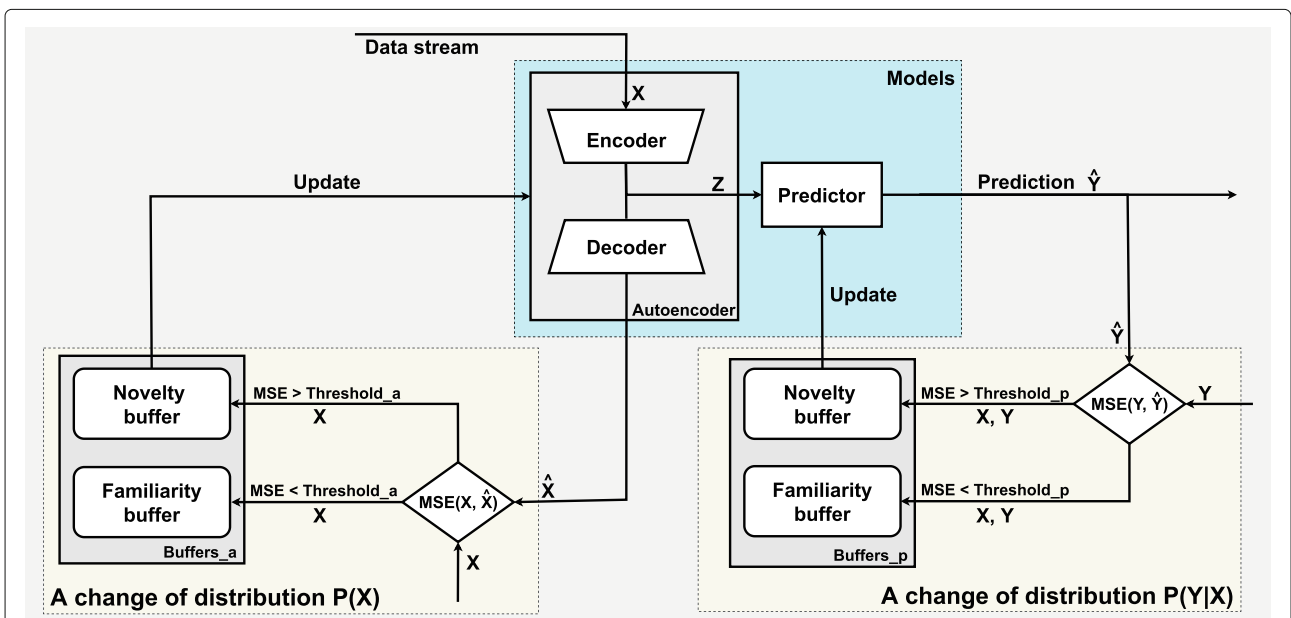


**Fig. 3** A general workflow of power forecasts comprises four phases: (1) reporting exceptions, (2) predicting targets, (3) storing data, and (4) updating models

the novelty buffer covers the change of probability distribution detected in the data stream. It is used for retraining the models when the update is triggered. The data in the familiarity buffer has information that the models are familiar with. It can be used for testing whether the models still retain the old knowledge after updating. Updating models can be considered as accumulations of knowledge for improving prediction accuracy.

**CLeaR**

The details of the CLeaR framework will be explained through the instance used in our experiments, as shown in Fig. 4. In this instance, the block **Models** contains an autoencoder and a fully-connected neural network for detecting the changes of the distributions  $P(X)$  and  $P(Y|X)$  and for forecasting power values. Data is labeled as novelty or familiarity by comparing the MSE to the



**Fig. 4** The CLeaR instance contains an autoencoder and a fully-connected neural network in the block **Models**. The  $MSE(X, \hat{X})$  is compared to **Threshold<sub>a</sub>** to detect the change of  $P(X)$ , and the  $MSE(Y, \hat{Y})$  is compared to **Threshold<sub>p</sub>** to detect the change of  $P(Y|X)$ . The novelty buffer in **Buffers<sub>a</sub>** is for updating the autoencoder and the novelty buffer in **Buffers<sub>p</sub>** is for updating the predictor. Models are evaluated with the data in the familiarity buffer

threshold. Besides, we adopt Online-EWC to update the models and adjust the threshold dynamically after each update. We suggest that the components of the CLear framework should be selected flexibly depending on the specific application scenario.

### Models

An autoencoder is a neural network that usually consists of two symmetric parts with a bottleneck between them. In an undercomplete autoencoder, the bottleneck has a smaller dimension than the input layer, which helps to extract latent representations  $\mathbf{z}$  from the input. An autoencoder can reconstruct the input at the output, rather than simply copy the input [25]. The encoder and the decoder can be formulated  $\mathbf{z}_n = f_{\Theta}(\mathbf{x}_n)$  and  $\hat{\mathbf{x}}_n = g_{\Phi}(\mathbf{z}_n)$ , where  $\Theta$  and  $\Phi$  are the parameter matrices. The optimization goal is to minimize the loss function,

$$L(\mathbf{x}_n, \Phi, \Theta) = MSE(\mathbf{x}_n, g_{\Phi}(f_{\Theta}(\mathbf{x}_n))), \quad (11)$$

by penalizing the reconstruction being different from the input. The change of distribution  $P(X)$ , as shown in Fig. 1, can be detected by the reconstruction error of the autoencoder.

At the next step, the extracted representation  $\mathbf{z}_n$  is fed into the predictor. As explained in Eqs. 4 and 5, optimizing the network in the general supervised setting is to minimize the MSE. Thus a true measurement  $y_n$  is required. The predictor, a fully-connected neural network, can be replaced by other networks, e.g., LSTM. We can also drop the predictor in applications where only the reconstruction is needed, as we will introduce in experiment 1.

### Buffers

Every neural network that needs to be updated during its application owns a limited novelty buffer and an unlimited familiarity buffer. When true target values are provided, the MSE can be used as a criterion to be compared to the preset threshold. The samples with a small MSE are stored in the familiarity buffer because the trained network has learned to cope with them before. The samples with a large MSE are stored in the novelty buffer. The network needs to be retrained based on these novelties to learn new knowledge. After updating, a validation error can be calculated using the familiarities to estimate whether the network can still retain the old knowledge acquired previously. How to deal with a poor update result remains an open question and needs to be further discussed. We empty both buffers after finishing an update.

### Threshold

Each model that needs to be updated owns a threshold. As shown in Fig. 4, **Threshold<sub>a</sub>** is for the autoencoder and **Threshold<sub>p</sub>** is for the predictor. The value of the

threshold determines how new samples are classified. The smaller the threshold, the more samples are likely to be labeled as novelties, where more well-learned knowledge will be re-learned, thus leading to unnecessary updates. A larger threshold could cause inefficient updates because too many novelties are misclassified. We suggest that the threshold value should be adjusted dynamically depending on the training results.

Figure 5 illustrates a distribution of errors for all samples after learning.  $MSE_{\min}$  refers to the minimum MSE obtained by minimizing a loss function, which can also be viewed as the mean of the distribution. The distribution has a lower mean and a lower variance indicating the model learns better on the given dataset. In this article, we adjust the threshold by

$$Threshold = \alpha \times MSE_{\min}, \text{ with } \alpha > 0, \quad (12)$$

where  $\alpha$  is a fixed threshold factor and the  $MSE_{\min}$  is re-estimated after each update.

### Update


The updating method and the trigger condition play a crucial role in the CLear framework. As mentioned in the Section about related work, retraining on old tasks should be constrained due to reasons, such as privacy protection or data storage overhead. Therefore, we adopt Online-EWC, which penalizes the loss when the overlapped significant weights are changed while learning new tasks. We adapt the notation given in [26] to explain EWC and Online-EWC in the context of the CLear framework.

### EWC

The goal of EWC is to approximate Bayesian posteriors over model parameters given tasks. In CLear, data is always split according to two kinds of tasks, the known task ( $T_k$ ) and the unknown task ( $T_u$ ).  $T_k$  refers to what the neural network has already learned. The corresponding dataset  $D_{k,T-1}$  is a combination of the datasets that were stored in the novelty and the familiarity buffers during the  $T-1$ th update. Note that novelties in  $D_{k,T-1}$  have already been learned before the  $T$ th update. Therefore, it belongs to the known task. Comparably,  $T_u$  refers to what the network will learn at the  $T$ th update. The corresponding  $D_{u,T}$  refers only to the dataset stored in the novelty buffer for triggering the  $T$ th update. Chronologically, the  $T-1$ th update occurs before the  $T$ th update. From Eq. 1, we can get the posterior after learning  $T_u$ , as:

$$P(\Theta|D_{u,T}, D_{k,T-1}) = \frac{P(D_{u,T}|\Theta)P(\Theta|D_{k,T-1})}{P(D_{u,T}|D_{k,T-1})}, \quad (13)$$

By converting Eq. 13 into a logarithm, we will get



**Fig. 5** A probability distribution of errors for the data in the novelty buffer (or training dataset) and the familiarity buffer (or test dataset).  $MSE_{min}$  indicates the minimum MSE after learning

$$\begin{aligned} \log P(\Theta|D_{u,T}, D_{k,T-1}) &= \log(D_{u,T}|\Theta) \\ &+ \log P(\Theta|D_{k,T-1}) \\ &- \log P(D_{u,T}|D_{k,T-1}). \end{aligned} \tag{14}$$

The term  $\log P(D_{u,T}|\Theta)$  is generally tractable through minimizing an MSE loss with respect to  $\Theta$  and the dataset  $D_{u,T}$ . In the case where the trained network can perform very well on the known task  $T_k$ , we have gotten the optimal parameters  $\Theta_{k,T-1}^*$ , which makes the gradient of  $-\log P(\Theta|D_{k,T-1})$  with respect to  $\Theta$  equal to 0. Therefore, the  $-\log P(\Theta|D_{k,T-1})$  can be estimated using 2nd order Taylor series around  $\Theta_{k,T-1}^*$ :

$$-\log P(\Theta|D_{k,T-1}) \approx \frac{1}{2} \Delta(\Theta)^T H(\Theta_{k,T-1}^*) \Delta(\Theta) + \text{constant}, \tag{15}$$

where  $\Delta(\Theta) = \Theta - \Theta_{k,T-1}^*$ .  $H(\Theta_{k,T-1}^*)$  is the Hessian of  $-\log P(\Theta|D_{k,T-1})$  with respect to  $\Theta$ , evaluated at the optimum  $\Theta_{k,T-1}^*$ . Furthermore, we can approximate the Hessian as:

$$H(\Theta_{k,T-1}^*) \approx N_k \cdot F(\Theta_{k,T-1}^*) + H_{prior}(\Theta_{k,T-1}^*), \tag{16}$$

where  $N_k$  is the number of samples in  $D_{k,T-1}$ , and  $F(\Theta_{k,T-1}^*)$  is the empirical Fisher information matrix on the known dataset  $D_{k,T-1}$ , and  $H_{prior}(\Theta_{k,T-1}^*)$  is the Hessian of the negative log prior with respect to  $\Theta$ .

EWC estimates the Fisher information matrix in a high-dimensional space as a diagonal matrix, i.e., the non-diagonal elements are zero [4]. The diagonal Fisher information values are denoted by  $F_k^i$  concerning the  $i$ th parameter in the network. Thus the formula of EWC is

$$\begin{aligned} \Theta_{k,T}^* &= \arg \min_{\Theta} \left\{ -\log P(D_{u,T}|\Theta) \right. \\ &\left. + \frac{1}{2} \sum_i \left( \sum_{t=1}^{T-1} \lambda_t F_{k,t}^i + \lambda_{prior} \right) (\theta^i - \theta_{k,T-1}^{*,i})^2 \right\}, \end{aligned} \tag{17}$$

where the update has been done  $T - 1$  times already.  $D_{u,T}$  is the dataset stored in the novelty buffer for the unknown task at the  $T$ th update.  $F_{k,t}^i$  is the diagonal Fisher information with respect to the parameter  $i$ , estimated by the dataset  $D_{k,t}$  stored in both buffers after the  $T$ th update.  $\theta_{k,T-1}^{*,i}$  is the optimal parameter  $i$  learned from the dataset  $D_{k,T-1}$  after the  $T - 1$ th update.  $\lambda_t$  and  $\lambda_{prior}$  are hyperparameters for EWC. Note that all of the diagonal Fisher information  $F_{k,t}$  in Eq. 17, obtained after every update, must be stored.

**Online-EWC**

The regularization term in EWC can be replaced by one Gaussian approximation to the whole posterior of all previous tasks, as proven in [26]. Based on this derivation, Online-EWC is proposed in [27], as:



$$\Theta_{k,T}^* = \arg \min_{\Theta} \left\{ -\log P(D_{u,T}|\Theta) + \frac{1}{2} \sum_i \lambda \tilde{F}_{k,T-1}^i (\theta^i - \theta_{k,T-1}^{*,i})^2 \right\}, \quad (18)$$

where  $\tilde{F}_{k,T}^i = \gamma \tilde{F}_{k,T-1}^i + F_{k,T}^i$  and  $\tilde{F}_{k,1}^i = F_{k,1}^i$ .  $\gamma$  is a hyperparameter governing the contribution of previous tasks and not larger than 1. Therefore,  $\tilde{F}_{k,T-1}^i$  contains diagonal Fisher information of all updates which precede the  $T$ th update.

By contrast with EWC, Online-EWC requires less space to store diagonal Fisher information and optima of each previous task. Another similar algorithm is EWC++ introduced in [9], where the diagonal Fisher information is defined as:  $\tilde{F}_{k,T-1}^i = \alpha \tilde{F}_{k,T-1}^i + (1-\alpha)F_{k,T}^i$  with  $0 < \alpha < 1$ . In our experiments, Online-EWC is selected for updates within the CLear framework.

#### Trigger condition

The trigger condition for updating the model is another customizable parameter. In our realization of the framework, the update will be triggered when the finite novelty buffer is filled. The smaller the size of this novelty buffer, the more easily this buffer gets full. On the one hand, fewer data in a full buffer is available for updating, which might lead to a failed update. On the other hand, more updates will result in increasing computational overhead. Moreover, note that the hyperparameter  $\gamma$  of Online-EWC is not larger than 1, which causes a gradual decay of the previous Fisher information when the number of updates increases. If the size of the buffer is too large, updating the network will be delayed until the buffer is full. The delay also results in the fact that the threshold can not be adjusted frequently.

## Experiments

In this section, two sets of experiments will be performed and analyzed. An unlabeled artificial dataset is generated for experiment 1, where we analyze the effects of two hyperparameters of CLear, the novelty buffer size and the threshold factor. We implement experiment 2 on a labeled dataset to predict wind power generation measured in 10 European wind farms [28] using supervised learning. In this experiment, we assess the CLear framework's performance in wind power prediction applications. Both experiments will adopt three evaluation metrics: fitting error, prediction accuracy, and forgetting ratio. The remainder of the section will introduce the datasets, the experimental setup, the evaluation criteria, and analyze the results.

### Dataset

**Artificial dataset** is a seven-dimensional periodic unsupervised dataset, i.e.,  $D = [\mathbf{x}_1, \dots, \mathbf{x}_7]$ . The underlying generation model is shown as follows:

$$\begin{aligned} \mathbf{x}_n(t) &\sim N(\text{mean}_n(t), \text{var}_n(t)), \\ \text{mean}_n(t) &= A_m \times \left| \sin\left(\frac{\pi \times (t + p_n)}{T}\right) \right|, \\ \text{var}_n(t) &= A_v \times \left| \sin\left(\frac{\pi \times (t + p_n)}{T}\right) \right|, \\ p_n &\sim N(0, 1), \end{aligned} \quad (19)$$

where  $x_n(t)$  is sampled from a Gaussian distribution with a time-dependent mean and variance and denotes the  $t$ th data of the sequence  $\mathbf{x}_n$ . The mean and variance are sampled from an absolute value of a sinusoidal function at the  $t$ th point in time with a period  $T$  and a phase  $p_n$ . The phase  $p_n$  is randomly generated with a standard normal distribution. Hyperparameters  $A_m$  and  $A_v$  refer to the amplitude of the sinusoidal function.


In order to generate a dataset containing daily and yearly periodical changes, we adopt two generation models,  $x_n^d(t)$  and  $x_n^y(t)$ , with two periods  $T^d$  and  $T^y$  respectively. The period  $T^d$  is 24 to simulate the fluctuation within one day (24 hours), and the period  $T^y$  of  $x_n^y(t)$  is 8760 to mimic the periodic changes over a year (8760 hours). The resulting time series data is a sum of the two models, i.e.,  $x_n(t) = x_n^d(t) + x_n^y(t)$ .

Figure 6 displays the sequence  $\mathbf{x}_1$  (marked in red) and the corresponding generating process (marked in black) in three different time slots. Note that the first 48 samples of Fig. 6(b) are displayed in Fig. 6(a). Similarly, the first 720 samples of Fig. 6(c) are also shown in Fig. 6(b). The periodical statistic fluctuation displayed in Fig. 6 can simulate features of non-stationary data, such as temperature, which usually peaks at noon and falls at night.


**Wind power dataset** [28] contains seven meteorological features and hourly averaged wind power generation data measured from European wind farms for two years in a row. The features are 24-hour-ahead meteorological forecasts using the European Centre for Medium-Range Weather Forecasts model [29], including (1) wind speed in 100m height, (2) wind speed in 10m height, (3) wind direction (zonal) in 100m height, (4) wind direction (meridional) in 100m height, (5) air pressure, (6) air temperature, and (7) humidity.

The power generation time series are normalized by the respective rated capacity of the wind farm for easy, scale-free comparison. All features are normalized to the range between 0 and 1. When no power has been generated longer than 24 hours, time points are removed in the pre-processing phase. The dropped time points can be viewed as exceptions that need to be processed manually in real-world applications.

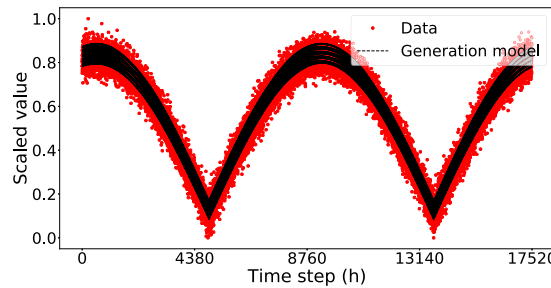
Both datasets are available on our department website [28] or by contacting the corresponding author.



(a)  $x_1$  over 48 hours (2 days)



(b)  $x_1$  over 720 hours (30 days)



(c)  $x_1$  over 17520 hours (2 years)

**Fig. 6** The three subfigures illustrate the fluctuation of the first dimension sequence  $x_1$  over 48 hours (2 days), 720 hours (30 days), and 17520 hours (2 years), respectively. The underlying generation model is marked in black, and generated data is in red. The generation model of  $x_1$  is a sum of two sub-models  $x_1^d$  and  $x_1^y$ , which have different periods  $T^d$  (24 hours) and  $T^y$  (8760 hours) respectively

Researchers can re-implement our experiments with the datasets.

**Experimental setup**

The artificial dataset and the wind power dataset are split into three subsets according to three phases: **a warm-up phase, an update phase, and an evaluation phase**. In the warm-up phase, the model is pre-trained on the first 1000 samples, containing only partial information that can describe the current task but not future tasks. The model monitors a data stream with the following 10000 samples in the update phase, simulating a theoretically infinite data stream in a real application scenario. Here, we assume that 10000 samples can provide enough information to describe the real distribution of samples because 10000

samples span 10000 hours, covering a complete period. The model will be retrained based on the novelty buffer and validated on the familiarity buffer once updating is triggered. The updated model is finally evaluated in the evaluation phase with the following 1000 samples.

We implement the following three CLear instances and a baseline model to compare their performances in experiments:

- Instance<sub>A</sub>: Its models are trained on the warm-up dataset in supervised learning mode and then directly evaluated in the evaluation phase without updating. The experimental results can reflect the disadvantages due to the lack of samples. Instance<sub>A</sub> is viewed as a **lower bound** of models' performance.

- **Instance<sub>B</sub>**: Its models are pre-trained in the warm-up phase, as Instance<sub>A</sub>. Then the models will be updated in the update phase by using **fine-tuning**. Fine-tuning allows the pre-trained models to learn a new dataset, which they were not originally trained on, by slightly adjusting all unfrozen parameters. The re-training process is monitored by early stopping with 30 epochs of patience to avoid overfitting. Early stopping will stop the re-training process if the loss is no longer decreasing in 30 epochs. The optimal models are the ones with the lowest loss before stopping. Identically, the updated Instance<sub>B</sub> will be evaluated in the evaluation phase after the update phase.
- **Instance<sub>C</sub>**: It is similar to Instance<sub>B</sub>, but uses **Online-EWC** without being monitored by early stopping for updating. It will be evaluated in the evaluation phase as other models.
- **Baseline model**: It is a common deep neural network model with the same structure as the neural network models in the above CLear instances. The baseline model is traditionally trained with 11000 samples (i.e., the samples of the warm-up phase and the update phase). Dropout layers with a dropout rate of 0.2 are used during training to avoid overfitting. It will also be evaluated in the evaluation phase.

Two sets of experiments are conducted on the artificial dataset and the wind power dataset, respectively.

- **Experiment 1** is based on the artificial dataset to analyze the correlation between the CLear framework’s performance and two framework-related hyperparameters, i.e., the novelty buffer size and the threshold factor. This unsupervised experiment aims to extract latent representations and reconstruct the inputs. The Instance<sub>C</sub> (only an autoencoder involved) and the baseline model (an autoencoder) are implemented.

The grid search range of the framework’s hyperparameters is shown in Table 2, including 56 available pairs of parameter values. The experiment with each pair of parameters is repeated 20 times with different initialization of the autoencoder. Identically, the baseline model is also repeated 20 times. The architecture of the autoencoder and the training setting for experiment 1 are shown in Tables 3 and 4.

- **Experiment 2** aims to evaluate all three CLear instances and the baseline model in the real-world application of wind power generation forecast based on 10 European wind farms datasets. The adopted model consists of an autoencoder and a deep neural network, as proposed in Fig. 4. Its architecture parameters are empirically selected and kept identical for a fair comparison, referring to Table 5. The parameters for the framework and Online EWC algo-

**Table 2** The grid search range of framework-related parameters for experiment 1

Hyperparameter	Value
Novelty buffer	{400, 600, 800, 1000, 1200, 1400, 1600, 1800}
Threshold factor	{0.5, 0.65, 0.8, 0.95, 1.1, 1.25, 1.4}

gorithm are selected using grid search from the range in Table 6. The training setting for experiment 2 is shown in Table 7.

**Metrics**

Models are evaluated in terms of fitting error, prediction error, and forgetting ratio.

**Fitting error** indicates how well the instance fits all seen samples after the update phase. Updating a CLear instance by mini-batch data might lead to a local minimum during the updating process. Eventually, the instance fits only a specific subset rather than all seen data. Such effects are measured by calculating the MSE on the 11000 samples, 1000 of which are from the warm-up phase, and the rest 10000 are from the update phase. Therefore, a lower fitting error reflects that more knowledge is finally accumulated.

**Prediction error** reflects the ability of CLear instances to perform predictions on previously unseen data. It is calculated with 1000 samples in the evaluation phase. The overfitting problem often exists due to neural networks’ powerful learning ability, where a test error is much larger than a training error. Regularization techniques, such as dropout and early stopping, are used to avoid overfitting during training. The prediction error can tell us whether the instance has already fallen into a

**Table 3** The architecture of the autoencoder in experiments 1

Layer	Size	Activation
Encoder-Input	7	None
Encoder-1	7×32	LeakyReLU
Encoder-2	32×16	LeakyReLU
Encoder-3	16×8	LeakyReLU
Encoder-4	8×4	LeakyReLU
Decoder-1	4×8	LeakyReLU
Decoder-2	8×16	LeakyReLU
Decoder-3	16×32	LeakyReLU
Decoder-4	32×7	LeakyReLU
Decoder-Output	7	None

Dropout layers are used with a dropout rate of 0.2. The slope parameter of the LeakyReLU is set to 0.05

**Table 4** Training setting for experiment 1

Hyperparameter	Value
Optimizer	Adam
Epochs in phase 1	512
Epochs in phase 2	512
Batch size in phase 1	32
Batch size in phase 2	16
Online EWC $\gamma$	0.9
Online EWC $\lambda$	200

Epochs denote the number of iterations for training the autoencoder. The batch size is the size of a mini-batch. Phase 1 and phase 2 indicate the warm-up phase and the update phase respectively. The experiments of training baseline model adopt the setting of the phase 1

local optimum after several updates. Also, the comparison between Instance<sub>A</sub> and Instance<sub>C</sub> in experiment 2 shows how important accumulating data and updating the models are for improving prediction accuracy, especially under the condition of a limited pre-training dataset.

**Forgetting ratio** measures how much old knowledge a model forgets after learning new tasks. In [22], He et al. compared the forgetting ratio to average test error and demonstrated that the forgetting ratio could reflect the severity of the forgetting problem. The formula is

$$\text{forgetting ratio} = \frac{\max(0, L_{\text{warm\_up}}^2 - L_{\text{warm\_up}}^1)}{L_{\text{warm\_up}}^1}, \quad (20)$$

**Table 5** Architecture parameters of the model for experiment 2

Layer	Size	Activation
Encoder-Input	7	None
Encoder-1	7×32	LeakyReLU
Encoder-2	32×16	LeakyReLU
Encoder-3	16×8	LeakyReLU
Encoder-4	8×Latent size	LeakyReLU
Decoder-1	Latent size×8	LeakyReLU
Decoder-2	8×16	LeakyReLU
Decoder-3	16×32	LeakyReLU
Decoder-4	32×7	LeakyReLU
Decoder-Output	7	None
DNN-1	Latent size×96	Tanh
DNN-2	96×64	Tanh
DNN-3	64×32	Tanh
DNN-4	32×16	Tanh
DNN-5	16×8	Tanh
DNN-Output	8×1	None

Dropout layers are used with a dropout rate selected by means of grid search. Latent size refers to the bottleneck dimension in the autoencoder. The slope parameter of the LeakyReLU is set to 0.05

**Table 6** The grid search range of hyperparameters for experiment 2

Hyperparameter	Value
Dropout	{0, 0.005}
Novelty buffer size	{512, 1024}
Latent size	{4, 5}
Threshold factor	{0.75}
Online EWC $\gamma$	{0.7, 0.8, 0.9}
Online EWC $\lambda_a$	{200}
Online EWC $\lambda_p$	{600, 800}

The parameter  $\lambda_a$  is the lambda value for the autoencoder, and  $\lambda_p$  is for the predictor. Latent size is the dimension of the bottleneck in the autoencoder

where  $L_{\text{warm\_up}}^1$  indicates the MSE on the warm-up dataset at the end of the warm-up phase and  $L_{\text{warm\_up}}^2$  indicates the error on the same dataset at the end of the update phase, and  $\max(x_1, x_2)$  returns the larger one of either  $x_1$  or  $x_2$ . The increment of the error for the same task describes the model's forgetfulness after learning new tasks. The comparison is performed only between the Instance<sub>B</sub> and the Instance<sub>C</sub> in experiment 2 because the Instance<sub>A</sub> and the baseline model are not updated.

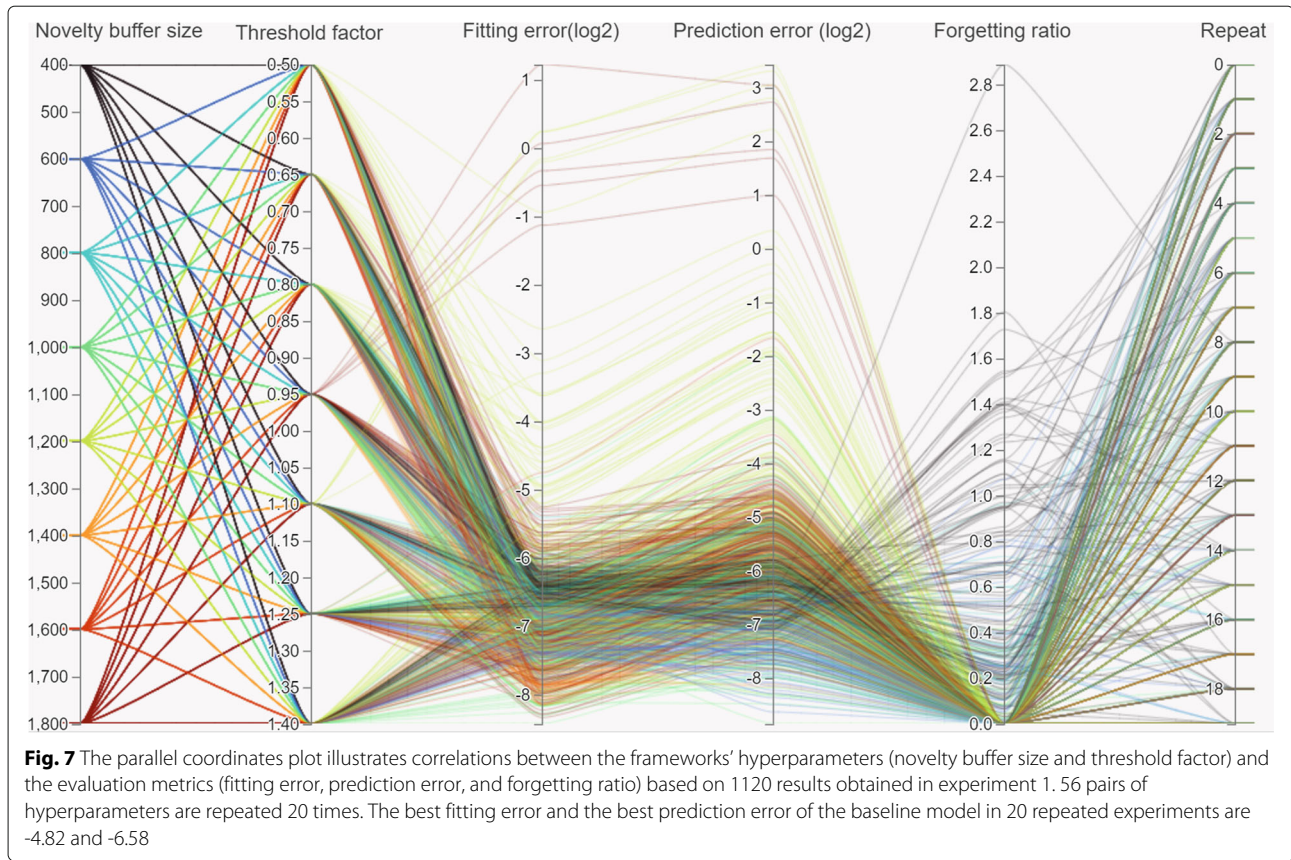
### Results of the experiment 1

Figure 7 illustrates the correlation between the frameworks' hyperparameters (novelty buffer size and threshold factor) and the evaluation metrics (fitting error, prediction error, and forgetting ratio) based on 1120 results obtained in experiment 1. **The lower the evaluation metric is, the better the model performs.** Note that the binary logarithm of the fitting error and prediction error are plotted in Fig. 7. Moreover, the binary logarithm of the best fitting error and the best prediction error of the baseline model are -4.82 and -6.58, respectively.

**Table 7** Training setting for the experiment 2

Hyperparameter	Reference figure
Optimizer	Adam
Epochs <sub>a</sub> in phase 1	256
Epochs <sub>p</sub> in phase 1	256
Batch size in phase 1	32
Epochs <sub>a</sub> in phase 2	128
Epochs <sub>p</sub> in phase 2	256
Batch size in phase 2	16

The parameter epoch<sub>a</sub> denotes the number of iterations for training the autoencoder and epoch<sub>p</sub> for training the predictor. The batch size is the size of a mini-batch. Phase 1 and phase 2 indicate the warm-up phase and the update phase respectively



Regarding the metric of fitting error, 1089 of the 1120 CLear models (97.2%) perform better than the best baseline model. This result indicates that the CLear framework can accumulate knowledge continually and effectively. Besides, it shows that the CLear framework's continual fitting ability is relatively robust to both framework-related hyperparameters.

Regarding the prediction error, 360 of the 1120 CLear models (32.1%) outperform the best baseline model. 253 of the 360 models have a threshold factor greater than or equal to 0.95. 159 of the 253 models have a novelty buffer size lower than or equal to 1000. Furthermore, 455 CLear models (40.8%) perform worse than the worst baseline model, whose prediction error is -5.84. 267 of the 455 models have a threshold factor lower than or equal to 0.95. 203 of the 267 models have a novelty buffer size greater than or equal to 1200. On the one hand, the experimental results are in accordance with our expectation, i.e., continual accumulation of meaningful knowledge can improve neural networks' prediction abilities. To a certain extent, the CLear framework can even predict non-stationary data more accurately than a neural network trained with sufficient historical samples. On the other hand, our findings indicate that the CLear framework's

prediction ability is susceptible to its hyperparameter values. In the update phase, the average update frequency of the 360 models that outperform the best baseline model is 10.2 times, while the average update frequency of the 455 models that perform worse than the worst baseline model is 6.87 times. A smaller novelty buffer gets filled more easily so that updating is triggered more frequently. Besides, a higher threshold can result in that only high-entropy data is stored in the novelty buffer. We conclude that timely learning of high-entropy data can effectively improve the prediction accuracy of neural networks for non-stationary data.

We only analyze the forgetting ratio of these CLear models because the forgetting problem does not happen with the baseline model. 174 of the 1120 CLear models (15.5%) obtain a forgetting ratio greater than 0.1. 172 of the 174 models have a novelty buffer size lower than or equal to 1000. Combined with the conclusion concerning the prediction error, we find that a smaller novelty buffer can trigger updating more often, which can decrease the prediction error but also increase the forgetting ratio. The main reason can be that the hyperparameter  $\gamma$  of Online EWC is set to 0.9 (see Table 4), which leads to a decay of the Fisher information regarding the previous tasks after

each update. Therefore, how to adjust the hyperparameters and how to supervise the updating process will be one of the key points in our further research. Otherwise, it will always be faced with a trade-off between prediction error and forgetting ratio.

Figure 7 is plotted using HiPlot [30]. The raw data of the results is available by contacting the corresponding author.

### Results of the experiment 2

First, we analyze the fitting error results of the three CLear instances and the baseline model for the 11000 samples in the first two phases, see Tables 8 and 9. Table 8 shows the fitting errors regarding the weather-to-weather data, i.e., the outputs of the autoencoders, and Table 9 presents the corresponding results of the weather-to-power data, i.e., the outputs of the predictors. WF is the abbreviation of the wind farm. The values on the last row of the tables are the average results of the 10 wind farms. The best results among the three CLear instances are marked in bold. Compared to the baseline model results, Instance<sub>C</sub> can decrease the fitting error in the term of either input reconstruction or power prediction. It is similar to our observation in experiment 1. We can also observe that the average fitting error of the Instance<sub>B</sub> is slightly lower than that of the Instance<sub>C</sub> in Table 8, but rises obviously in Table 9. We infer that Catastrophic Forgetting happens on the Instance<sub>B</sub> because it applies fine-tuning that can only adapt the model to the new situations. We can obtain the conclusion regarding the existence of the forgetting problem according to the high forgetting ratio in Table 11 as well.

Table 10 shows the results of power prediction, where the Instance<sub>C</sub> clearly outperforms the other two instances.

**Table 8** The fitting errors of the **autoencoders** in the three CLear instances and the baseline model. *ae* refers to the autoencoder

Dataset	Fitting error (e-2)			
	Instance <sub>A<sub>ae</sub></sub>	Instance <sub>B<sub>ae</sub></sub>	Instance <sub>C<sub>ae</sub></sub>	Baseline <sub>ae</sub>
WF 1	2.098	<b>0.369</b>	0.502	1.505
WF 2	2.068	<b>0.348</b>	0.440	1.579
WF 3	2.292	<b>0.423</b>	0.513	1.634
WF 4	2.669	0.988	<b>0.863</b>	1.449
WF 5	1.421	0.516	<b>0.471</b>	1.549
WF 6	2.456	1.054	<b>0.761</b>	1.065
WF 7	2.129	<b>0.647</b>	0.931	1.441
WF 8	2.149	<b>0.434</b>	0.662	1.448
WF 9	2.370	<b>0.773</b>	0.793	1.340
WF 10	1.996	0.678	<b>0.411</b>	1.406
Mean	2.165	<b>0.623</b>	0.635	1.442

**Table 9** The fitting errors of the **predictor** in the three CLear instances and the baseline model


Dataset	Fitting error (e-2)			
	Instance <sub>A</sub>	Instance <sub>B</sub>	Instance <sub>C</sub>	Baseline
WF 1	4.864	3.468	<b>2.304</b>	2.574
WF 2	2.695	2.120	<b>1.167</b>	1.473
WF 3	3.517	9.172	<b>1.769</b>	2.840
WF 4	4.867	5.457	<b>3.682</b>	3.327
WF 5	10.347	12.989	<b>6.931</b>	6.835
WF 6	9.593	6.089	<b>4.225</b>	5.546
WF 7	0.867	0.838	<b>0.677</b>	0.612
WF 8	8.042	5.656	<b>3.451</b>	3.865
WF 9	4.713	5.511	<b>3.060</b>	3.405
WF 10	1.875	3.117	<b>1.025</b>	1.424
Mean	5.138	5.442	<b>2.829</b>	3.190

The baseline model results indicate that sufficient training data can enable neural network models to obtain as much meaningful knowledge as possible, which helps to improve the models' predictive ability.

Figure 8 illustrates the errors of the three instances and the baseline model over the 12000 weather-to-power samples of one wind farm. The errors are calculated after finishing the update phase. We split the samples into 12 sections, and each point in Fig. 8 refers to an MSE of 1000 samples. The average value of the points from 1000 to 11000 can be equal to the fitting error, and the value at a coordinate of 12000 is the prediction error. Compared to the other three lines, the red line is less volatile. Moreover, we can observe that the lines start to rise at a coordinate of 5000 and drop back to the original level at the end. On the one hand, this reflects that the periodic changes can

**Table 10** The prediction error of the **predictor** in the three CLear instances and the baseline model

Dataset	Prediction error (e-2)			
	Instance <sub>A</sub>	Instance <sub>B</sub>	Instance <sub>C</sub>	Baseline
WF 1	2.415	2.332	<b>1.822</b>	1.611
WF 2	0.740	0.945	<b>0.522</b>	0.502
WF 3	2.370	5.369	<b>1.208</b>	1.641
WF 4	2.124	<b>1.643</b>	2.223	1.448
WF 5	13.864	5.672	<b>5.560</b>	3.828
WF 6	8.612	8.017	<b>4.368</b>	3.123
WF 7	0.965	0.729	<b>0.587</b>	0.527
WF 8	6.129	4.969	<b>2.927</b>	2.499
WF 9	2.544	4.562	<b>1.808</b>	1.129
WF 10	1.383	3.009	<b>0.741</b>	0.866
Mean	4.115	3.725	<b>2.177</b>	1.717



**Fig. 8** Errors of the three instances about one wind farm. Each point refers to the mean error over 1000 samples

lead to these fluctuated curve shapes and influence the mapping between weather and power generation. On the other hand, in Tables 9, 10, and Fig. 8, where the Instance<sub>C</sub> outperforms the other instances, the results prove that updating is significant for predicting a non-stationary data stream, especially when the model is trained only on a limited dataset.

Table 11 presents the forgetting ratio values between Instance<sub>B</sub> and Instance<sub>C</sub> calculated after the update phase. According to the average results, we can conclude that Instance<sub>C</sub> outperforms Instance<sub>B</sub> here, for both the weather-to-weather task and the weather-to-power task. Moreover, as the conclusion of Table 9, the existence of the forgetting problem leads to the increment of the fitting error.

**Conclusions**

We believe that continual learning will be the key to the future machine’s intelligence. The non-stationary

world requires future artificial intelligence to be updated smoothly by taking account into the different data distributions but still to retain previous useful knowledge. Therefore, in this article, the proposed CLear describes the prototype structure of the continual learning based framework. It can be applied to lots of real-world projects, such as power predictions for smart grids, where prediction models have to mimic humans’ ability to acquire and transfer knowledge incrementally from new data throughout their lifespan.

The framework still needs to be improved in our future research. For example, although the cleaned exception is not considered, it is necessary to define and process such exceptions in a real application. Besides, we only use MSE to calculate the difference between predictions and measurements. However, this method might not work when the values are unavailable in unsupervised settings. Therefore, we suggest estimating the uncertainty of new predictions to detect changes in distributions, for example, using Monte Carlo dropout [31]. If the uncertainty is over a threshold, the data can be labeled as novelties. Furthermore, hyperparameters are found by grid search in the context of the current design. It is worth researching whether hyperparameters can be found by transferring relevant knowledge from a similar task dataset. In addition, we suppose that Online-EWC can be replaced by (or combined with) other CL algorithms to improve the framework. Novel analysis methods and evaluation metrics for the updated models will also be one of our main research focuses in future.

In a nutshell, we expect that the framework can be designed as a modular tool like LEGO toys. Each component of the framework is flexible and can be added, removed, replaced, or expanded. It should be possible for researchers and users to adapt the framework to a specific application scenario for achieving their own goals.

**Table 11** The forgetting ratio of Instance<sub>B</sub> and Instance<sub>C</sub>

Dataset	Forgetting ratio			
	Instance <sub>B<sub>ae</sub></sub>	Instance <sub>C<sub>ae</sub></sub>	Instance <sub>B</sub>	Instance <sub>C</sub>
WF 1	1.474	<b>0.727</b>	1.955	<b>0.994</b>
WF 2	<b>0.018</b>	0.140	1.707	<b>0.745</b>
WF 3	<b>1.205</b>	1.489	6.170	<b>0.421</b>
WF 4	2.817	<b>2.407</b>	3.371	<b>0.785</b>
WF 5	<b>0.000</b>	<b>0.000</b>	7.617	<b>2.296</b>
WF 6	0.945	<b>0.237</b>	0.969	<b>0.046</b>
WF 7	<b>0.574</b>	1.552	1.768	<b>1.253</b>
WF 8	<b>0.588</b>	1.585	5.716	<b>2.999</b>
WF 9	1.463	<b>0.848</b>	3.470	<b>1.463</b>
WF 10	4.937	<b>2.726</b>	2.754	<b>0.609</b>
Mean	1.402	<b>1.171</b>	3.550	<b>1.161</b>

### Abbreviations

CL: Continual Learning; MSE: Mean Square Error; EWC: Elastic Weight Consolidation; SI: Synaptic Intelligence; LWF: Learning Without Forgetting; Online-EWC: Online Elastic Weight Consolidation; PCA: Principal Component Analysis; AE: Autoencoder; WF: Wind farm

### Acknowledgements

Thanks to our colleagues from the Intelligent Embedded Systems group, particularly Mohammad Wazed Ali, Florian Heidecker, and Chandana Priya Nivarthi, for their helpful comments and suggestions.

### Authors' contributions

YH wrote the majority of the manuscript and was responsible for implementing the CLear and experiments. BS suggested the artificial data and the first experiment and was responsible for double-checking the manuscript. All authors read and approved the final manuscript.

### Authors' information

Not applicable.

### Funding

This work was supported within the C/sells RegioFlexMarkt Nordhessen (03SIN119) project and the Digital-Twin-Solar (03EI6024E) project, funded by BMWi: Deutsches Bundesministerium für Wirtschaft und Energie/German Federal Ministry for Economic Affairs and Energy. Open Access funding enabled and organized by Projekt DEAL.

### Availability of data and materials

The datasets generated and analyzed during the current study are available from <https://www.uni-kassel.de/eecs/ies/downloads> or the corresponding author on reasonable request.

### Declarations

#### Ethics approval and consent to participate

Not applicable

#### Consent for publication

Not applicable

#### Competing interests

The authors declare that they have no competing interests.

Received: 18 February 2021 Accepted: 28 June 2021

Published online: 16 July 2021

### References

- McCloskey M, Cohen NJ. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychol Learn Motiv.* 1989;24: 109–65.
- Ratcliff R. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychol Rev.* 1990;97(2): 285.
- Mermillod M, Bugaiska A, Bonin P. The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects. *Front Psychol.* 2013;4:504.
- Kirkpatrick J, Pascanu R, Rabinowitz N, Veness J, Desjardins G, Rusu AA, Milan K, Quan J, Ramalho T, Grabska-Barwinska A, et al. Overcoming catastrophic forgetting in neural networks. *Proc Natl Acad Sci.* 2017;114(13):3521–6.
- Zenke F, Poole B, Ganguli S. Continual learning through synaptic intelligence. *Proc Mach Learn Res.* 2017;70:3987.
- Maltoni D, Lomonaco V. Continuous learning in single-incremental-task scenarios. *Neural Netw.* 2019;116:56–73.
- Farquhar S, Gal Y. A unifying bayesian view of continual learning. *arXiv e-prints.* 2019:arXiv–1902.
- Nguyen CV, Li Y, Bui TD, Turner RE. Variational continual learning. In: *International Conference on Learning Representations.* 2018. <https://openreview.net/forum?id=BkQqOgRb>.
- Chaudhry A, Dokania PK, Ajanthan T, Torr PH. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: *Proceedings of the European Conference on Computer Vision (ECCV);* 2018. p. 532–47.

- van de Ven GM, Tolias AS. Generative replay with feedback connections as a general strategy for continual learning. *arXiv e-prints.* 2018:arXiv–1809.
- Goodfellow IJ, Mirza M, Xiao D, Courville A, Bengio Y. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv e-prints.* 2013:arXiv–1312.
- Srivastava RK, Masci J, Kazerounian S, Gomez F, Schmidhuber J. Compete to compute. In: *Advances in Neural Information Processing Systems;* 2013. p. 2310–8.
- LeCun Y, Cortes C, Burges C. Mnist handwritten digit database. ATT Labs [Online]. Available: <http://yann.lecun.com/exdb/mnist>. 2010;2:.
- Lomonaco V, Maltoni D. Core50: a new dataset and benchmark for continuous object recognition. In: *Conference on Robot Learning.* PMLR; 2017. p. 17–26.
- Gensler A, Henze J, Sick B, Raabe N. Deep learning for solar power forecasting – an approach using autoencoder and lstm neural networks. In: *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC);* 2016. p. 002858–65. <https://doi.org/10.1109/SMC.2016.7844673>.
- He Y, Henze J, Sick B. Forecasting power grid states for regional energy markets with deep neural networks. In: *2020 International Joint Conference on Neural Networks (IJCNN).* IEEE; 2020. p. 1–8.
- Shin H, Lee JK, Kim J, Kim J. Continual learning with deep generative replay. In: *Advances in Neural Information Processing Systems;* 2017. p. 2990–9.
- Rusu AA, Rabinowitz NC, Desjardins G, Soyer H, Kirkpatrick J, Kavukcuoglu K, Pascanu R, Hadsell R. Progressive neural networks. *arXiv e-prints.* 2016:arXiv–1606.
- Li Z, Hoiem D. Learning without forgetting. *IEEE Trans Pattern Anal Mach Intell.* 2017;40(12):2935–47.
- Jung H, Ju J, Jung M, Kim J. Less-forgetting learning in deep neural networks. *arXiv e-prints.* 2016:arXiv–1607.
- Pasquale G, Ciliberto C, Rosasco L, Natale L. Object identification from few examples by improving the invariance of a deep convolutional neural network. In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).* IEEE; 2016. p. 4904–11.
- He Y, Henze J, Sick B. Continuous learning of deep neural networks to improve forecasts for regional energy markets. *IFAC-PapersOnLine.* 2020;53(2):12175–82.
- Farquhar S, Gal Y. Towards robust evaluations of continual learning. *arXiv e-prints.* 2018:arXiv–1805.
- Deisenroth MP, Faisal AA, Ong CS. *Mathematics for Machine Learning:* Cambridge University Press. <https://mml-book.com>.
- Goodfellow I, Bengio Y, Courville A. *Deep Learning:* MIT Press. <http://www.deeplearningbook.org>.
- Huszár F. On quadratic penalties in elastic weight consolidation. *arXiv e-prints.* 2017:arXiv–1712.
- Schwarz J, Luketina J, Czarnecki WM, Grabska-Barwinska A, Teh YW, Pascanu R, Hadsell R. Progress & compress: A scalable framework for continual learning. In: *International Conference on Machine Learning.* PMLR; 2018. p. 4528–4537.
- Gensler A. EuropeWindFarm Data Set. <https://www.uni-kassel.de/eecs/ies/downloads>. Accessed 07 July 2021.
- ECMWF homepage. <https://www.ecmwf.int/>. Accessed 07 July 2021.
- Haziza D, Rapin J, Synnaeve G. Hiplot, interactive high-dimensionality plots. *GitHub.* 2020.
- Gal Y, Ghahramani Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: *International Conference on Machine Learning;* 2016. p. 1050–9.

### Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.