



HAL
open science

Hierarchical transfer learning with applications for electricity load forecasting.

Anestis Antoniadis, Solenne Gaucher, Yannig Goude

► **To cite this version:**

Anestis Antoniadis, Solenne Gaucher, Yannig Goude. Hierarchical transfer learning with applications for electricity load forecasting.. 2021. hal-03429702v2

HAL Id: hal-03429702

<https://hal.science/hal-03429702v2>

Preprint submitted on 19 Nov 2021 (v2), last revised 10 Nov 2022 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Hierarchical transfer learning with applications for electricity load forecasting.

Anestis Antoniadis¹, Solenne Gaucher², and Yannig Goude ^{*3}

¹Laboratoire de Mathématiques d’Orsay, Université Paris-Saclay

²Laboratoire Jean Kuntzmann, Université Grenoble Alpes, France

³EDF R&D, Saclay

Abstract

The recent abundance of data on electricity consumption at different scales opens new challenges and highlights the need for new techniques to leverage information present at finer scales in order to improve forecasts at wider scales. In this work, we take advantage of the similarity between this hierarchical prediction problem and multi-scale transfer learning. We develop two methods for hierarchical transfer learning, based respectively on the stacking of generalized additive models and random forests, and on the use of aggregation of experts. We apply these methods to two problems of electricity load forecasting at national scale, using smart meter data in the first case, and regional data in the second case. For these two usecases, we compare the performances of our methods to that of benchmark algorithms, and we investigate their behaviour using variable importance analysis. Our results demonstrate the interest of both methods, which lead to a significant improvement of the predictions.

Keywords— Demand forecasting, Semi-parametric additive model, Random forest, Transfer learning, Time series, Combining forecasts

1 Introduction

The recent abundance of electricity consumption data at more or less aggregated scales, due in part to the development of smart meters, opens up many prospects for electricity consumption forecasting (e.g. see Wang et al. (2019)). However, with these new perspectives come new challenges, among which is the question of how to include these data obtained at a finer scale (corresponding to a household or to a smaller geographical area), which can be used to create forecasts at a fine scale, into a prediction at a wider scale (at the national scale for example). In this work, we present two methods for leveraging our ability to predict a variable of interest at a finer scale, with the goal of exploiting these predictions to improve prediction at a larger scale. This problem, which consists in taking advantage of the similarities existing between forecasting problems at different scales, can be naturally formulated in the framework of *transfer learning*.

Transfer learning methods aim at transferring the knowledge acquired from solving given problems (referred to as *source* problems \mathcal{S}) to address an other problem of interest (referred to as the *target* problem \mathcal{T}). In a supervised predictive machine learning setting, the objective is to predict a variable of interest $Y^{\mathcal{T}}$ using covariates $\mathbf{X}^{\mathcal{T}}$. To do so, the learner relies on a set of observations $(\mathbf{X}_i^{\mathcal{T}}, Y_i^{\mathcal{T}})_{i \in \{1, \dots, n^{\mathcal{T}}\}}$, drawn from a joint distribution $\mathbb{P}^{\mathcal{T}}$. Popular methods rely on minimizing the empirical risk corresponding to a given loss function over a set of possible learners (e.g. tree based, neural nets, generalized additive models). When the learners and the loss are chosen appropriately, the estimated model will have good forecasting accuracy on a new dataset as long as the size of the training set $n^{\mathcal{T}}$ is sufficiently large, and the marginal and joint distributions remain unchanged in the test set. In a wide range of real-world applications, these conditions are not satisfied. Classical examples of such situations include tasks requiring massive training sets, such as computer vision or natural language processing. On the other hand, when dealing with temporal data, one may be confronted with changes in distribution resulting in large prediction errors for the target period. In these settings, it is then important to gain improvement of learning in a new task through the transfer of knowledge from a related task that has already been learned, while preserving adaptivity to the new task.

*yannig.goude@edf.fr

Transfer learning aims to tackle such a problem. It has attracted increasing attention in machine learning and has been used in many applications (see Olivas et al. (2009)). In many practical situations, a relatively small quantity of data from the target distribution \mathbb{P}^T is available. In some cases, one also has access to a larger dataset, with different distribution denoted \mathbb{P}^S , which can be used to solve a task related to the target problem. A key assumption is that \mathbb{P}^T and \mathbb{P}^S are related in a way that can be leveraged by the transfer learning method. These distributions may be defined on the same domain (in this case, the transfer learning problem is said to be homogeneous), or on different domains, in which case more complex transformations need to be developed (the transfer learning problem is then said to be heterogeneous). In this paper, we will focus on heterogeneous transfer problems, with difference between source and target in terms of features spaces, feature marginal distribution, and joint distribution. Surprisingly, although transfer learning is very popular in computer vision and text mining (see Pan and Yang (2010) and Zhuang et al. (2020) for a survey), very few developments can be found in the time series forecasting community. In Obst et al. (2021a,b) the authors propose a fine tuning approach as well as online updates to transfer information from Italian Data to French Data in order to improve electricity load forecasts during the COVID lock-down. Here the transfer is both in time (from past data (source) to future data (target)) and space (from one country to another).

In the following, we suppose to have access to two data sets

$$\mathcal{D}_T = \left(\mathbf{X}_t^T, Y_t^T \right)_{t=1, \dots, n_T} \quad \text{and} \quad \mathcal{D}_S = \left(\mathbf{X}_t^S, Y_t^S \right)_{t=1, \dots, n_S},$$

where \mathcal{D}_T is the target data set in the sense that the final objective is to forecast Y_t^T , and has underlying distribution \mathbb{P}^T . \mathcal{D}_S , with underlying distribution \mathbb{P}^S , is an auxiliary source data set sharing some common properties with \mathcal{D}_T . We then want to exploit \mathcal{D}_S in order to improve the forecast of Y_t^T .

In the present paper, we develop transfer learning methods for hierarchical prediction, that leverage data available at a fine scale to improve prediction at a wider scale. The first approach is based on the design of new features, which are then used in generalized additive models stacked with random forests. While features $\mathbf{X}_t^S = (\mathbf{X}_{t,k}^S)_{1 \leq k \leq d^S}$ and $\mathbf{X}_t^T = (\mathbf{X}_{t,k}^T)_{1 \leq k \leq d^T}$ may belong to spaces of different dimensions (respectively denoted d^S and d^T), we may assume without loss of generality that there exists a subset $C \subset \{1, \dots, d^S \wedge d^T\}$ of covariates that are common to both data sets. In the setting of electricity consumption forecasting, these common variables can be, e.g., calendar variables, or meteorological variables (at finer scale in \mathcal{D}_S , and at wider scale in \mathcal{D}_T). It is then natural to assume that these features will have similar effects on the variable of interest Y_t in the target and source dataset. To exploit this idea, we propose to learn the influence f_k of a common feature $X_{t,k}$ such that $k \in C$, using the source dataset \mathcal{D}_S . The function f_k is obtained using empirical risk minimization over a set of learners on the dataset \mathcal{D}_S . We then use f_k to generate a new feature $Z_t = f_k(X_{t,k}^T)$, which we include in the target dataset \mathcal{D}_T . When the functions f_k are learned from an auxiliary dataset \mathcal{D}_S corresponding to observations at a finer scale, adding the corresponding features to the dataset \mathcal{D}_T of observations at the wider scale allows to transfer knowledge in a hierarchical fashion. Note, however, that we can also use this technique to learn new features directly from \mathcal{D}_T . If we use different types of learner to create the features, and to make predictions from this enriched dataset, we can hope to take advantage of both types of learner. This idea motivates the stacking of generalized additive models (GAM) and random forests (RF) presented in Section 2.2.

Our empirical results suggest that the feature design method presented above, combining stacked GAM and RF, allows to improve prediction at an wider scale by leveraging knowledge acquired from data a finer scale. Unfortunately, this approach relies on knowledge acquired on a training set, which may not be relevant if a change in distribution occurs during the test period. For this reason, these methods are not adaptative to a brutal change in distribution both at the fine and the wider scale. To ensure adaptivity in our model, we propose a second transfer learning approach based on online aggregation of experts. Online robust aggregation of experts Cesa-Bianchi and Lugosi (2006) is a powerful model agnostic approach for time series forecasting. It consists in combining in a streaming fashion different forecasts (called experts) according to their past performances. When experts forecasting a variable of interest at a finer scale are aggregated so as to forecast this variable at the wider scale, this allows to transfer knowledge between these different scales. Aggregation of experts was recently applied in a forecasting competition (see Farrokhbadi et al. (2021)), where 2 of the 3 first teams (see de Vilmarrest and Goude (2021), Ziel (2021)) applied these approaches to forecast electricity load consumption during the COVID lock-down in a big city (unknown localisation). In this changing context, online aggregation allows to adapt to changes in distributions and to track the performance of the best expert. In the hierarchical context, Goehry et al. (2019) and Brégère and Huard (2021) show that aggregating experts designed on different nodes of a hierarchical partition of the data (statistical clustering based on temporal or exogenous information, spatial partition) improves forecasting performances compared with classical bottom-up approaches.

1.1 Contributions and outline of the paper

In this paper, we develop two methods for leveraging information available at a fine scale to improve prediction at a more wider scale, based respectively on *feature design* combined using *stacked generalized additive models and random forests*, and on *online aggregation of experts*. These methods are presented in Section 2, and illustrated on two real-world problems. In Section 3, we apply the first method to the problem of electricity load forecasting at national level, relying on smart meter data. In Section 4, we combine these methods to obtain adaptative methods for forecasting electricity consumption at national level during the Covid-19 pandemic period, using data available at regional level. We demonstrate the interest of our proposed approach in both cases. Our results indicate that both the stacking of GAM and RF and the use of features designed on data at finer scale lead to improvements of the forecasts at wider scale. Moreover, the use of multi-scale information transfer through aggregation of experts also increases the quality of wide-scale forecasts. Strikingly, our results indicate that in the two usecases, the presented methods can improve wide-scale predictions by using fine-scale predictions, even when no hierarchical constraints are implemented.

For reproducibility of the results, please find the code and data here: <https://github.com/SolenneGaucher/Hierarchical-Transfer-Learning-Covid19.git>.

2 Methodology and algorithms

In this section, we briefly present the different statistic tools composing our transfer learning approach: generalized additive models, random forests, and online aggregation of experts.

2.1 Generalized additive models

We chose to use generalized additive models for the learners f_k . Partially linear additive models, which are a special case of generalized additive nonparametric models (GAM), retain the parsimony and interpretability of linear models and the flexibility of nonparametric additive regression, by allowing a linear component for some predictors which are presumed to have a strictly linear effect, and an additive structure for other predictors. This choice of both linear and non-parametric components allows to reduce the degrees of freedom and to mitigate the problem known as “curse of dimensionality”.

Given observations $\{(Y_t, \mathbf{X}_t^{(1)}, \mathbf{X}_t^{(2)})\}_{t=1}^n$, where Y_t is the response at time t , $\mathbf{X}_t^{(1)} = (X_{t,1}^{(1)}, \dots, X_{t,d_1}^{(1)})^T$ and $\mathbf{X}_t^{(2)} = (X_{t,1}^{(2)}, \dots, X_{t,d_2}^{(2)})^T$ are vectors of covariates, the partially linear additive model assumes that

$$Y_t = b + \left(\mathbf{X}_t^{(1)}\right)^T \boldsymbol{\beta} + \sum_{j=1}^{d_2} f_j(X_{t,j}^{(2)}) + \epsilon_t, \quad t = 1, \dots, n, \quad (1)$$

where b is the intercept, $\boldsymbol{\beta}$ is the $d_1 \times 1$ vector of unknown coefficients for linear terms, f_j are unknown nonlinear real valued components and the ϵ_t 's are i.i.d random variables with mean 0 and variance σ^2 independent of the covariates. In order to ensure that the model is identifiable, one requires that the linear covariates are centred and that identifiability conditions $\int f_j(t)dt = 0$, $j = 1, \dots, d_2$ hold.

Such models, together with procedures that achieve estimation and simultaneous consistent variable selection, have proven their ability to cope with high level aggregate electricity data in previous work: Goude et al. (2013) applied it to national french consumption and Fan and Hyndman (2012) show their interest for regional Australia’s load forecasting. Moreover they can be applied efficiently to forecast electricity data at different level of aggregation Amato et al. (2021). In the following, GAMs are trained in R using the library `mgcv` Wood (2017).

2.2 Stacked GAM and RF

Our target models are obtained by stacking GAMs, and the correction provided by random forest regression trained on the target dataset.

Random forests (RF) are a powerful black box approach for modelling complex regression relationships (see Breiman (2001)). The very general underlying model behind a random forest regression is to assume that $y_t = g(X_t) + \epsilon_t$, where g is a generic, non-parametric function, and ϵ_t is an independent gaussian noise. Because of the generality of the model, RF necessitates very few prior knowledge of the problem. RF are obtained by aggregating an ensemble of base learners generated by applying classification and regression trees (CART, see Breiman et al. (1984)) on different subsets of the data obtained with bagging and random sampling of covariates. Thus, RF are by definition restricted to the convex envelop of the training data. This behavior prevents them from producing aberrant predictions caused by extrapolation, even when trained on very small data sets, as can typically be the case

in a transfer learning framework with small target dataset and high number of covariates Balestrierio et al. (2021). Moreover, their black box design can capture well complex non-linear interactions.

By contrast, GAM provide interpretable models and a natural way to incorporate expert knowledge into a statistical model. In addition, because of the smoothness assumptions imposed to GAM functionals, they can extrapolate out of training data. However, they only model the influence of pre-specified covariates or pairs of covariates, and can therefore fail to account for some interactions between inputs. To have the best of both worlds, we propose to stack these two approaches.

The stacked GAM and RF algorithm consists in three steps:

1. we first fit a GAM model as Equation (1) on the training set to extract the estimated GAM features $\hat{f}_j(x)$. We denote $\hat{f}(X)$ the concatenation of these effects.
2. we compute estimates of forecasting residuals (either by cross-validation, block cross-validation or forecasting errors in an online forecasting setting) denoted $\hat{\varepsilon}_t$.
3. we then fit a RF model to predict the GAM residuals: $\hat{\varepsilon}_t = g(X_t, \hat{f}(X_t)) + u_t$. The final forecasts are obtained summing GAM forecasts and the corrections provided by the RF.

For the random forest fits, we use the procedure `ranger()` from the R toolbox `ranger`. The default parameters are used (500 trees, $mtry = \sqrt{p}$, unlimited tree depth). In future work, these values could be optimized in a more refined way by combining `ranger` with procedures from the R library `caret`, at the cost of increasing the CPU time.

2.3 Online aggregation of experts

We propose here a short description of sequential expert aggregation for forecasting. A complete presentation of these methods can be found in Cesa-Bianchi and Lugosi (2006).

In this setting, data are observed sequentially. The target variable (here the electricity consumption) is supposed to be a bounded sequence $Y_1, \dots, Y_T \in [0, B], B > 0$, that we want to forecast step by step for every time t . At each time t , N experts provide forecasts of Y_t , denoted $(\hat{Y}_t^1, \dots, \hat{Y}_t^N) \in [0, B]^N$. These experts can come from a statistical model, a physical model or expert advice projection. The aggregation algorithm choses weights $\hat{p}_{j,t} \in \mathbb{R}^N$, and returns as forecast for Y_t a weighted average $\hat{Y}_t = \sum_{j=1}^N \hat{p}_{j,t} \hat{Y}_t^j$ of the N forecasts. Then, Y_t is observed and instance $t + 1$ begins. In the following, we will consider only convex aggregation, with weights $\hat{p}_{j,t}$ in the simplex of \mathbb{R}^N .

Performance of experts and aggregation forecasts are evaluated according to a convex loss function. We will consider here the square loss $\ell_t(x) = (Y_t - x)^2$. At time t , expert k suffers loss $\ell_t(\hat{Y}_t^k) = (Y_t - \hat{Y}_t^k)^2$ and the aggregation $\ell_t(\hat{Y}_t) = (Y_t - \hat{Y}_t)^2$. The purpose of expert aggregation is to minimise the total loss $\sum_{t=1}^T (Y_t - \hat{Y}_t)^2$ that can be expressed:

$$\frac{1}{T} \sum_{t=1}^T (Y_t - \hat{Y}_t)^2 \triangleq \frac{1}{T} \sum_{t=1}^T (Y_t - \hat{Y}_t^*)^2 + R_T,$$

\hat{Y}_t^* is called an oracle and can be viewed as an optimal (unknown before the forecasting run) forecast. R_T is the regret term corresponding to the error suffered by our algorithm relatively to the error of the oracle. Some algorithms are proposed to achieve low regrets. In our study we use the ML-Poly algorithm proposed in Gaillard et al. (2014) and implemented in the R package `opera` Gaillard and Goude (2016). This algorithm tracks the best expert or the best convex aggregation of experts by giving more weight to an expert that will generate a low regret. This makes this algorithm particularly interesting as no parameter tuning is needed.

3 Transfer learning for forecasting aggregated smart meter data

In this section, we will illustrate the methodology using a dataset that is commonly used for the calibration of electricity consumption forecasting models. The dataset is made up of aggregate semi-hourly consumption data of the national load for the United Kingdom, and of observations of some meteorological and calendar variables. Our goal is to forecast electricity consumption at the national level from December 2009 to August 2010 (this period will be referred to as the test set). For this purpose, we assume that we have access to national covering the period from April 2005 to November 2009 (called hereafter the learning set) and smart meter for a smaller period (from april 2009 to august 2010).

3.1 Data

3.1.1 National Data

This dataset for UK national semi-hourly electricity consumption is provided by the European Grid Standards Office (see <https://www.nationalgrideso.com/balancing-data/data-finder-and-explorer>) and covers the period between April 2005 and December 2010. We add as features the temperature data obtained from the NOAA (National Oceanic and Atmospheric Administration) ¹ for the 10 largest cities in the UK: London, Birmingham, Glasgow, Sheffield, Bradford, Liverpool, Edinburgh, Manchester and Bristol. We then compute a weighted average T_t of the temperatures recorded in these 10 stations with weights proportional to the official population of each city, and we finally perform an exponential smoothing of this weighted average with the parameters 0.2, 0.05 and 0.01.

3.1.2 Smart meters Data

This data set corresponds to smart meters data at an individual scale in the UK. This dataset has been obtained from the Energy Demand Research Project (EDRP) launched by Ofgem on behalf of the UK Government in 2007 (see AECOM (2018), Schellong (2011) and ²) where power consumption of approximately 60,000 households was collected at half hourly intervals for about two years. We consider a subset of 1925 customers from april 2009 to august 2010 located in two regions of the UK: south east (around Brighton) and north west (around Glasgow). We considered temperatures in each region, obtained from the NOAA. We add to this data set supplementary calendar covariates such as the time of year, day type and ephemeride hour along the year.

3.2 Models and forecasting

The fitting procedure used to forecast electricity consumption at national level can be described as follows.

We note a trend in the time series of consumption over the period from April 2005 to August 2010 (see figure 1). We estimate this trend in a very simple way by fitting to the series of observations a nonparametric Gaussian model $Y_t = \mu + s(t) + \varepsilon_t$, the trend $s(t)$ being represented in a cubic spline function base with a number of knots limited to three. In the following, we subtract this trend and aim at forecasting the national *de-trended* consumption, which is then given by $Y_t^c = Y_t - \hat{s}(t) - \hat{\mu}$.

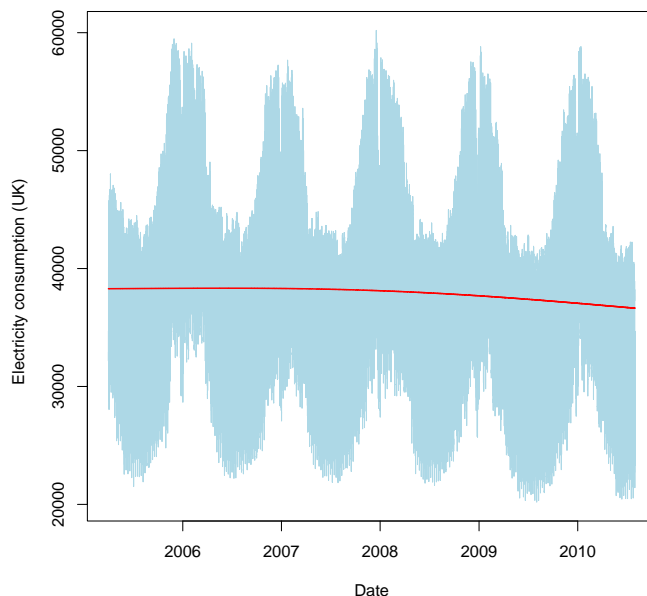


Figure 1: The national electricity consumption over the observed period and its smoothed trend (in red)

¹<https://www.noaa.gov/>

²<https://www.ofgem.gov.uk/gas/retail-market/metering/transition-smart-meters/energy-demand-research-project>

We apply the stacked GAM and RF methodology to predict national load consumption, using only data available at national level. Note that this is a special case of the general transfer learning framework, with $\mathcal{D}_{\mathcal{T}} = \mathcal{D}_{\mathcal{S}}$, and where the final forecast is obtained using RF on the data enriched with the transfer of information performed using the design of new features. We begin by fitting a semi-parametric GAM on the 82274 moments indexing the time series of national electricity consumption of the learning set. The GAM is given by

$$\begin{aligned} Y_t^c &= \sum_{j=1}^7 m_j \mathbf{I}_{\text{DayType}_t=j} + m_8 \mathbf{I}_{\text{Holiday}_t=1} + m_9 \mathbf{I}_{\text{LongWeekEnd}_t=1} \\ &+ g_1(\text{Instant}_t, \text{Temp}_t) + \sum_{j=1}^7 f_j(\text{Instant}_t) \mathbf{I}_{\text{DayType}_t=j} + s(\text{ToY}_t) \\ &+ s(\text{Temp99}_t) + \varepsilon_t \end{aligned} \quad (2)$$

where Y_t^c is the de-trended electricity load at time t , Temp_t is the weighted temperature of observation t , Temp99_t is the weighted and exponentially smoothed temperature, Instant_t is the instant of the day, an integer going from 0 to 47, DayType_t is the day of the week of observation t , Holiday_t is a binary variable indicating public holiday days, LongWeekEnd_t is a binary variable indicating the presence of a long weekend at time t , ToY_t indicates the time of year of observation t , and ε_t is a centered Gaussian noise. Each univariate smooth component of the above GAM model is fitted using regression spline functions with 40 knots (50 knots for ToY) and a tensor basis of spline functions for the interaction between time and temperature with 20 and 10 knots, respectively.

Then, once the fit (2) has been performed, we extract the estimated effects m_i , g_1 , and f_j of the features, and add them to the set of initial covariates. This enriched dataset will then be used to train the RF. Note that the initial number of covariates in the database taken into account in the Equation (2) is 7, and the number of additive components extracted by the GAM methodology is 15. We thus find ourselves, after transfer, with a sample having a number of observed covariates equal to 22. We apply the random forest stacking methodology to fit a nonparametric regression called **RF.nat.feats** on the training sample. Finally, we evaluate the prediction of this stacked GAM and RF on the test sample, and compare it to the predictions obtained with the GAM (2) model. We also compare this model to a standard regression model by random forests, denoted hereafter **RF.nat**.

In a second time, we apply the stacked GAM and RF methods to transfer information from the smart meters data. We begin by computing on the smart meter data set the total consumption of the customers, and we fit a GAM model to forecast this total. Using the same methodology than for the national data, we obtain the model presented in Equation (3). We chose to use a simpler model than the national one, because the dataset used to train it is smaller.

$$\begin{aligned} y_t &= \sum_{j=1}^7 m_j \mathbf{I}_{\text{DayType}_t=j} \\ &+ g_1(\text{Instant}_t, T_t) + \sum_{j=1}^7 f_j(\text{Instant}_t) \mathbf{I}_{\text{DayType}_t=j} + s(\text{ToY}_t) \\ &+ \varepsilon_t \end{aligned} \quad (3)$$

We then extract the features of this model as supplementary covariates, and add them to the dataset consisting of all original covariates and of the effect extracted from the GAM at national level. The dataset obtained contains 28 covariates. Finally, we use these covariates to train the stacked RF, and obtain a model called **RF.local.feats**.

Table 1: Errors of prediction for the learners **GAM.nat**, **RF.nat**, **RF.nat.feats**, and **RF.local.feats**

	GAM.nat	RF.nat	RF.nat.feats	RF.local.feats
RMSE	1409 MW	1339 MW	1214 MW	1193 MW
MAPE	2.670	2.560	2.360	2.310

An importance by permutation analysis of the variables used in the stacked RF after learning by transfer retains as most important for **RF.nat.feats** the instant of the day, followed by three terms from the national GAM modeling. For the model **RF.nat.feats**, the most important variables are the instant of the day, followed by two terms from the national GAM, and by one term from the local GAM.

By analyzing the mean absolute percentage errors (MAPE) and the root mean squared errors (RMSE) of the different methods, presented in Table 1, we see that for the British national data set, the RF are more efficient than the adopted reference model GAM. Interestingly, the stacked GAM and RF trained using only national data **RF.nat.feats** outperforms these two models. This indicates that the stacking of GAM and RF allows to obtain the best of both worlds: the RF is able to correct effects or interactions of variables (such as the instant of the day) that are not well captured by the GAM, while being robust to the large number of covariates taken as input (up to 28). Finally, the best model both in terms of MAPE and RMSE is obtained by stacking GAM and RF using effects learned from both national and smart meter data. These results underscore the value of leveraging available data at a finer scale, even when no hierarchical constraints are implemented in the algorithm.

4 Electricity load forecasting during the first Covid-19 lock-down

In this Section, we apply our methodology to short-term electricity load forecasting during the Covid-19 lock-down and post-lock-down period in France, at a resolution of half an hour and at the national level. To do so, we leverage information available at the regional level. Electricity consumption has been significantly affected by the measures taken by the government to cope with the epidemic: closures of non-essential businesses, as well as stay-at-home directives, have lead to a decrease in electricity consumption of about 10%, as well as to changes in its daily and weekly patterns (see Obst et al. (2021a) for a description of the impact of these measures on french electricity consumption). Common models trained on historical data, which rely on calendar and weather data, fail to account for these significant changes. Similarly, transfer learning methods relying on data present at a finer (e.g., regional) scale, if trained on data with different distribution than that of the target, will make poor previsions, especially if the relationship between local and global variables changes over time. Thus, these models, trained on data from the pre-pandemic period, make relatively large prediction errors on the period following the start of the lock-down. For these reasons, it is necessary to use more adaptive forecasting methods than the basic transfer of features presented above.

Transfer learning proves to be essential to address the problem of electricity load forecasting during the Covid-19 pandemic. Indeed, as the data for this period is scarce, especially since we want to make predictions from its very beginning, it is crucial to use information from the pre-pandemic period to predict power consumption during the pandemic period. To do so, we use the methods presented above to transfer information from the large data set corresponding to historical electricity consumption during the pre-pandemic period (thereafter called the source period) to improve predictions during the pandemic period (called the target period). This can be done using again stacked GAM and RF: this transfer learning algorithm allows us to rely on a GAM trained on a large set of observations of historical electricity consumption, coming from the source distribution, while correcting its error on the target using RF based on scarce observations.

On the other hand, because of the important changes in electricity consumption consecutive to the lock-down, we expect that the relationship between effects learned on regional data and national load will also change. Indeed, our studies reveal that containment measures induce changes in electricity consumption at the regional level, that however differ according to the region considered. To make use of electricity consumption data available at the regional level, we must remain adaptative to changes in distribution of both national and regional data. This is achieved by using online aggregation of experts, which allows us to combine forecasts at regional and national level in an adaptative fashion. We choose to forecast electricity consumption separately region by region using stacked GAM and RF, and then to combine the forecasts of these regional models in order to predict national electricity consumption in a hierarchical fashion. In doing so, the hierarchical model captures regional phenomena that are not apparent at a more aggregated scale, and leverages this information to improve predictions at the national level. Thus, our methods allow for transferring knowledge both at a temporal level (data from the pre-pandemic period are used to improve forecasts during the pandemic period), and at a hierarchical levels (regional predictions are used to produce forecasts at the national level).

The rest of the section is organized as follows. In Section 4.1, we present the data used to design and evaluate our models. In Section 4.2 we present the models used for forecasting electricity consumption at national and regional level. The results of our study are presented in Section 4.3: first, we compare the performances of different approaches, then we present a more detailed analysis of the stacked GAM and RF, and of the online aggregation of experts.

4.1 Data

The data are from the french TSO (Transmission System Operator) RTE. It consists of electricity consumption (in MW) at a half-hourly temporal resolution at the French national level (“Load”) and for the 12 metropolitan administrative regions (it does not include Corsica): Nouvelle Aquitaine, Auvergne Rhônes-Alpes, Bourgogne-Franche-Comté, Occitanie, Hauts-de-France, Normandie, Bretagne, Centre-Val de Loire, Île-de-France, Pays de la Loire, Provence-Alpes-Côte d’Azur, Grand Est. Our goal is to forecast the French national consumption, exploiting the regional loads information. For all the load consumption data, we compute the lags for one day and one week and denote it with the subscripts “.48” and “.336”.

Our models use the temperature and weighted temperature as explanatory variables. These variables were collected on the website of the French weather forecaster Météo France. For each region, we compute the weighted mean of meteo stations where the weights are proportional to $\exp(-dist)$ where $dist$ is the distance of the station to the barycenter of each region. Note that we use the observed temperatures instead of their predicted values in our forecast. In doing so, we cancel out the errors caused by the uncertainty of a particular weather forecast, which allows for a more precise comparison of the different models. Moreover this choice allows us to only use open data, so as to ensure reproducibility of our results.

Our models also rely on variables characterising the impact of the restrictions implemented to fight the epidemic. The first of these variables is the Oxford Covid-19 Government Response Tracker. This index, freely available at <https://www.bsg.ox.ac.uk/research/research-projects/covid-19-government-response-tracker>, aggregates indicators characterising the measures taken by governments to mitigate the epidemic in terms of containment, health and economic support. It is available at the national level. The methodology used to calculate the index and the measures on which it is based are known a few days in advance, so we assume that it is known for the day we wish to forecast. The remaining variables used to characterize the impact of lock-down measures are Google Mobility Indices. These indices are provided by Google, and obtained by aggregating geolocalisation data. They characterize the changes in frequentation of categorized places (residential, workplaces, transports, parks, grocery and pharmacy, retail and recreation). The data are freely available at <https://www.google.com/covid19/mobility/>, albeit with a little less than a week delay. Therefore, we considered lagged versions of these indicators in our prediction. The government response and mobility indices are available respectively from January and February 2020 onwards. Therefore, we do not use them as covariates in the source model, but only in the target model.

Hereafter, we call source models the models trained on historical data, collected between the beginning of 2012 to the end of August 2019. We evaluate their performances on data with the same distribution, during the pre-lock-down period ranging from September 2019 to March 15th 2020, and compare it to their performance on data from the target distribution, ranging from March 16th to September 17th. By contrast, the models specific to the lock-down and post-lock-down period, henceforth called target models, are retrained every day during this target period, so as to leverage all observations available.

Note that while the first french lock-down officially started on March 17th, we consider March 16th as the first day of the target distribution, as the electricity consumption pattern had already changed by that day.

4.2 Models

Recall that our target models are obtained by stacking generalized additive models (GAM) and an time varying random forest (RF). As presented in Section 2, the GAM is fitted on source data (pre-pandemic period). By contrast, the RF are trained in a streaming fashion on the target data (pandemic and post pandemic period).

We proceed in two steps:

1. First, we fit a GAM (described bellow) on the source data. We use it to produce predictions of the load on the target data, and compute residuals during the pandemic period, denoted $\hat{\varepsilon}_t$. Moreover, we extract the estimated GAM effects of the covariates for this period.
2. Then, each day from the first day of containment, we train a RF (described bellow) on the available data for the period of the pandemic. This RF is trained to predict the residuals $\hat{\varepsilon}_t$ of the historical GAM on the target data, and takes as input the estimated GAM effects, the original covariates and the mobility indices.

Using RF to correct the errors of the GAM during the pandemic period allows us to obtain an adaptative model able to produce predictions from the very beginning of the target period. Note that the corrections of the RF remain small compared to the predictions of the GAM: the first order of the prediction is given by the source model, trained on the large set of historical data, while the corrections learned on the scarce observations from the target dataset only provide a second order correction.

4.2.1 Generalized additive models for the pre-pandemic period

We use GAM to predict the electricity load under normal circumstances. We fit one model for each region of mainland France, as well as one at the national level, and obtain thus 13 models. To take into account the daily patterns of electricity consumption, each model is composed of 48 GAM fitted independently, and forecasting the electricity load at a given instant of the day. Thus, the 624 time series corresponding to the 48 half-hours for the 12 regions and the national level are treated independently. In order to compare the predictions, terms and errors of the models, regional and national electricity loads are normalized, that is, they are divided by their average value for the region and the half-hour considered. GAM are then fitted so as to predict this normalised load. In the following, we denote respectively by y and y^{norm} the load and the *normalized* load.

The model used to predict the electrical load for a zone z at a time t corresponding to the h -th half-hour of the day is the following:

$$\begin{aligned}
y_{z,t}^{norm} &= \sum_{i=1}^7 \sum_{j=0}^1 \alpha_{i,j}^{(z,h)} \mathbb{1}_{\text{DayType}_t=i} \mathbb{1}_{\text{DLS}_t=j} + \sum_{i=1}^7 \beta_i^{(z,h)} \text{Load.48}_{z,t} \mathbb{1}_{\text{DayType}=i} \\
&+ \gamma^{(z,h)} \text{Load.336}_{z,t} + f_1^{(z,h)}(t) + f_2^{(z,h)}(\text{ToY}_t) + f_3^{(z,h)}(t, \text{Temp}_{z,t}) \\
&+ f_4^{(z,h)}(\text{Temp95}_{z,t}) + f_5^{(z,h)}(\text{Temp99}_{z,t}) + f_6^{(z,h)}(\text{TempMin99}_{z,t}, \text{TempMin99}_{z,t}) + \varepsilon_{z,t}
\end{aligned}$$

where $y_{z,t}^n$ is the normalized electricity load at time t for the zone z ; Daytype_t is a categorical variable indicating the day of the week; DLS_t is a binary variable indicating whether t is in summer hour or winter hour; ToY_t is the time of year whose value grows linearly from 0 the 1st of January 00h00 to 1 on the 31st of December 23h30; $\text{Temp}_{z,t}$ is the temperature at time t in the zone z ; $\text{Temp95}_{z,t}$ and $\text{Temp99}_{z,t}$ are exponentially smoothed temperatures of smoothing factor 0.95 and 0.99; $\text{TempMin99}_{z,t}$ and $\text{TempMax99}_{z,t}$ are the minimal and maximal value over the day of exponentially smoothed temperature with smoothing factor 0.99; $\text{Load.48}_{z,t}$ and $\text{Load.336}_{z,t}$ are the normalized loads of the day before and the load of the week before in the zone z ; $\varepsilon_{z,t}$ is gaussian white noise. Each univariate smooth component of the above GAM model is fitted using regression spline functions with respectively 20 knots for ToY, 10 knots for Temp95 and Temp99, 5 knots for Date, and a tensor basis of spline functions for the interaction between time and temperature with 3 and 5 knots, respectively.

4.2.2 Stacked RF and stacked quantile RF

Our experts are obtained by stacking GAM and RF. Building on the results of Section 3, we choose to take as input for these RF the usual covariates, but also the GAM effects learned on the source dataset. The RF are designed to predict quantile estimates of the GAM residuals $\hat{\varepsilon}_t$. For each half-hour, and for each region and at the national level we consider 5 experts associated to quantiles 0.05, 0.1, 0.5, 0.9 and 0.95. Moreover, for each quantile prediction, we use two RFs, trained either on the residuals of the GAM for the corresponding half-hour and region (called thereafter *individual* RF), or on the residuals for the corresponding half-hour across all regions (called thereafter *common* RF). The common RF allows to efficiently learn effects of the pandemic that are common across all regions by increasing the size of the training set (namely, multiplying it by 13). The stacked experts are then obtained by adding the prediction of the GAM and the correction provided by the RF. We thus obtain 11 experts for each zone (namely, the GAM expert, as well as individual RF and common RF forecasting 5 different quantiles). Methods used to design these experts are summarized in Figure 2.

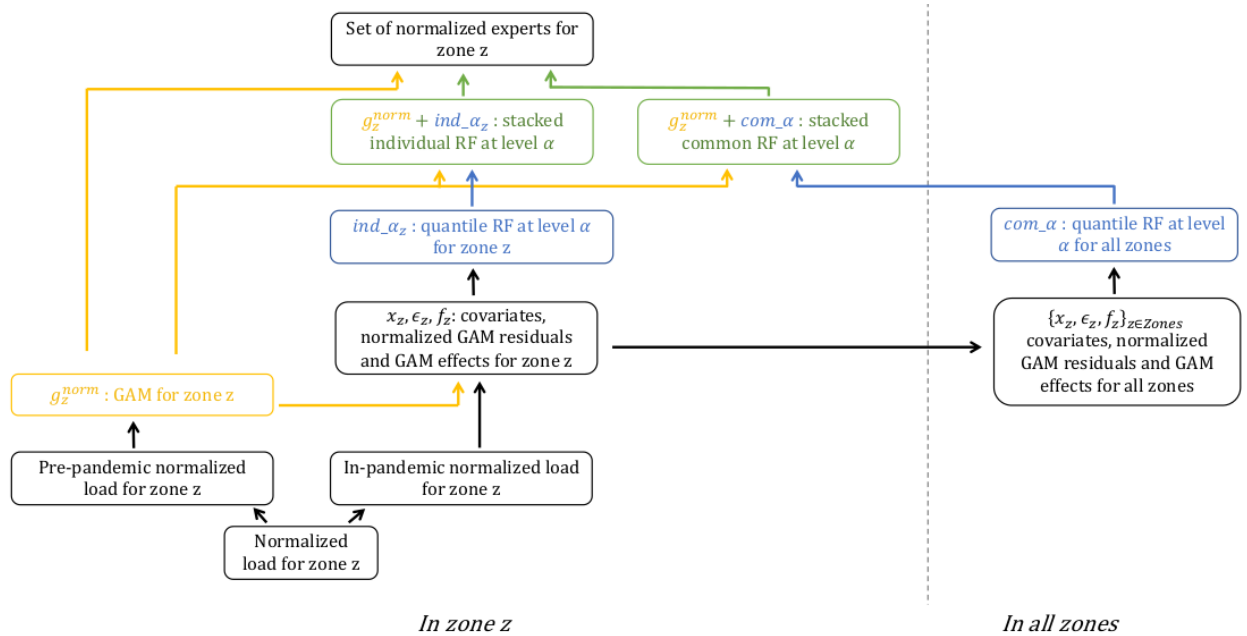


Figure 2: Experts used for predicting the normalized loads for the different regions and at national level.

Designing experts for low and high quantiles present several advantages. On the one hand, these experts, when aggregated online to track the changes in the distribution of the load using convex aggregation, are particularly relevant since there is a high probability that the real consumption falls in the convex hull of the quantile experts. On the other hand, by doing so, we obtain experts with similar behaviour across regions, that can share weights between the different regions and at the national level. Indeed, it is reasonable to assume that when an expert receives a low weight in one region, it must receive a low weight in all regions. Considering this vectorial aggregation model allows us to take advantage of the similar behaviour of the quantile experts across regions.

Variable selection for the RF We allow the RF to use many variables for their predictions, including usual calendar and weather variables, as well as mobility data, containment index, and estimated GAM effects, without knowing a priori which ones will be relevant to predict electricity consumption during the pandemic period. It is reasonable to assume that including all covariates might be detrimental to the prediction, given the high correlations between some variables and the small number of observations available to train the model, especially in the early days of lock-down. We want to take advantage of the fact that the number of observations increases rapidly, by repeating the variable selection operation several times during the pandemic period, in order to be able to enrich the model if necessary. Moreover, we expect that the relevant variables might differ from one region to the other, and so we want to perform the variable selection region by region.

To do so, we select every week the variables to be used to train the RF for the following week’s forecasts for a given region. Feature selection in RF is an ongoing field of research. State-of-the-art methods rely on ranking of variable importance measure, such as VSURF (Genuer et al., 2015), or on permutation of variables, such as Boruto (Kursa and Rudnicki, 2010). These methods suffer from their important computational cost: VSURF, in particular, is too slow to be used in our context of numerous variable selection operations. We suggest an alternative approach to determine the relevant covariates, using the technics developed for variable selection in linear regression. More precisely, we fit a linear model to predict the residuals of the GAM during the pandemic period using a LASSO penalty. Without prior knowledge on the number of covariates necessary to accurately forecast electricity load, we design three models corresponding to different number of covariates. More precisely, we use the LASSO penalty to obtain two formulas, including respectively the 5 and 10 most relevant covariates, and consider the formula taking into account all 16 available covariates. This variable selection step is repeated every week. Finally, the predictions of the 5, 10, and 16-variables RF are combined using an expert aggregation method.

Before implementing this method for all regions and all quantiles, we evaluate its interest in predicting the 0.5 quantile of national load using the available national data. We compare the MAPE of the predictor obtained using Boruto, Lasso variable selection with an aggregation step, and of the full model with 16 variables. The MAPE for the lock-down and post-lock-down periods are presented in Table 2.

Selection method	2020/03/16-2020/05/11	2020/05/11-2020/09/17
Lasso with aggregation	2.41 %	1.06 %
Boruta	2.41 %	1.03 %
Full model	2.41 %	1.03 %

Table 2: Mean Absolute Percentage Error of the stacked GAM and RF models.

The preliminary results indicate that variable selection only marginally affects the performance of the RF. This underlines the robustness of the RF against a high dimension of the inputs, event when trained on relatively small data sets. In the rest of the section, we will therefore consider the RF obtained using the full models, taking as input all covariates and estimated GAM effects.

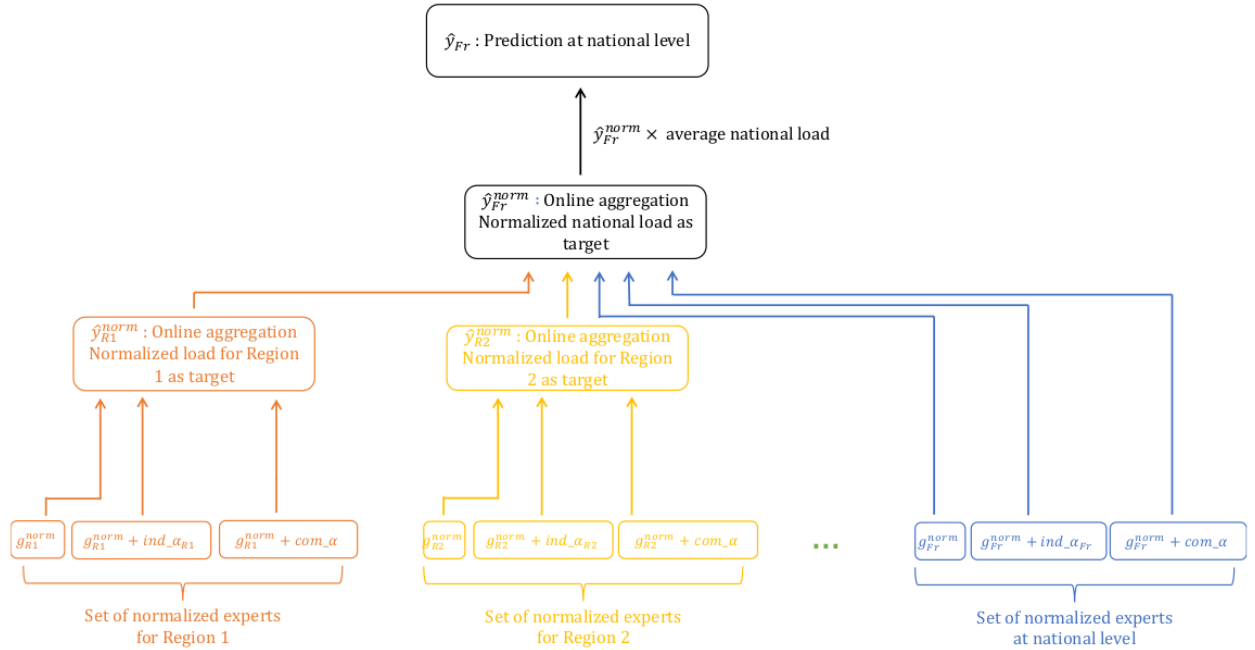
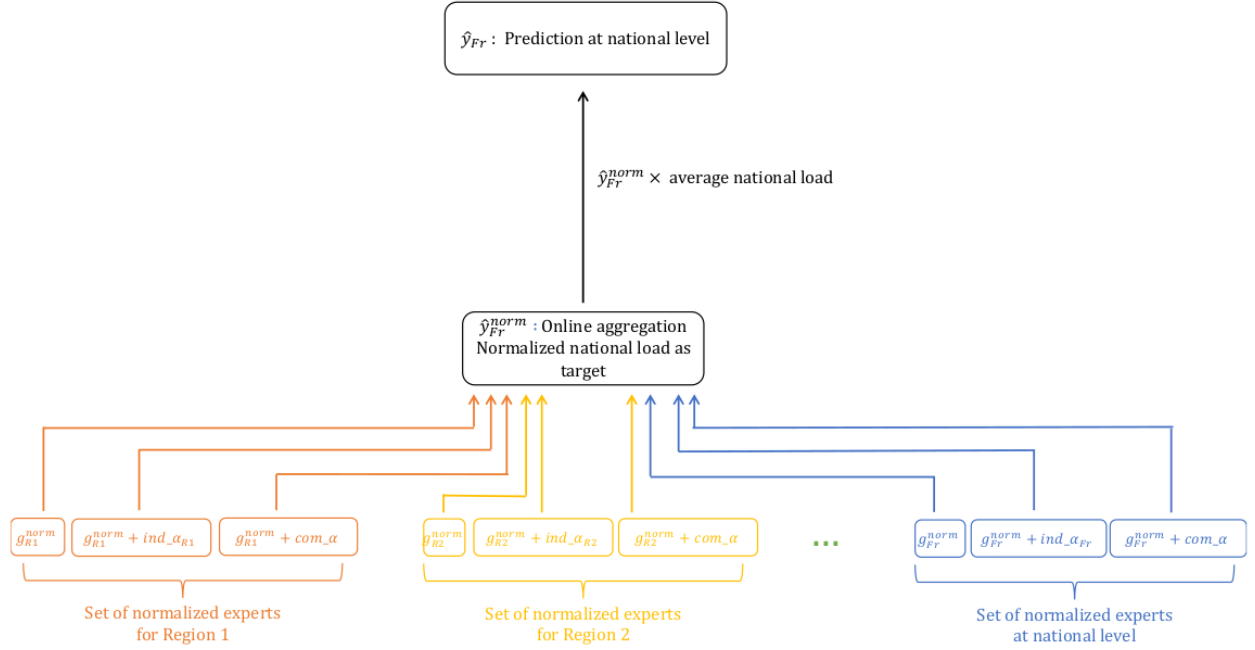
4.2.3 Aggregation

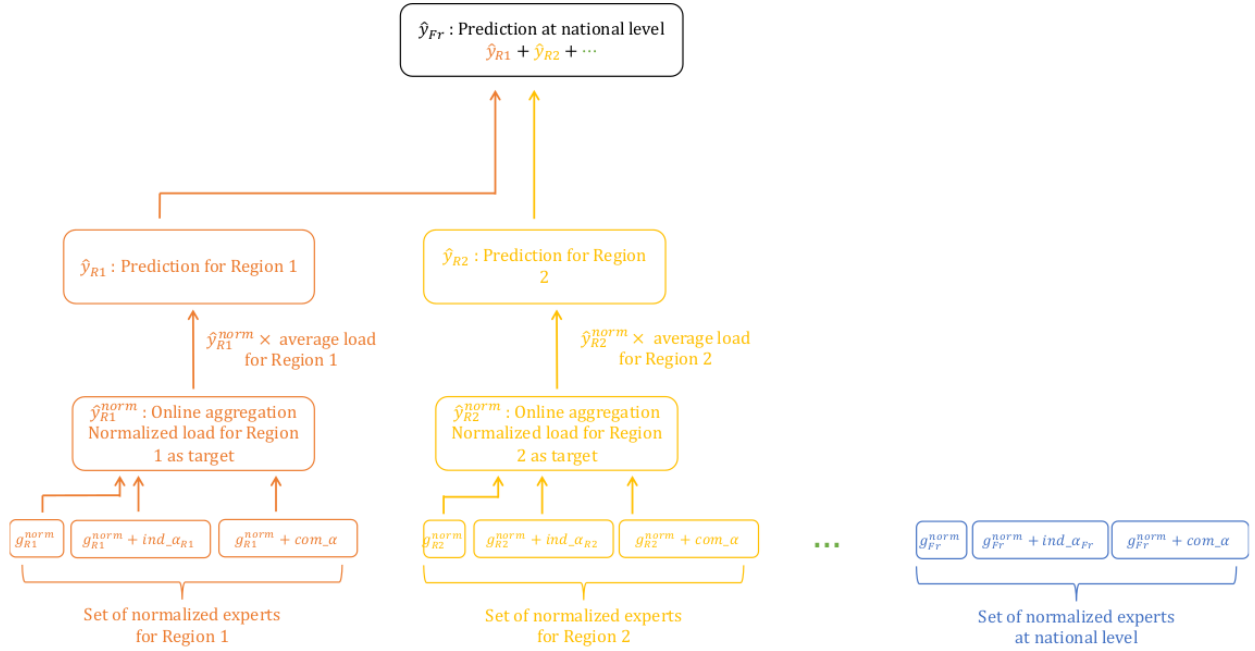
Recall that all models are fitted to forecast the electricity load for a given area, i.e. the electricity load divided by its average value for the given area and time considered. For each region, 11 experts are available: the historical GAM expert, as well as 10 experts predicting 5 different quantiles using 2 type of RF, trained on the residuals for one region or on the residuals for all regions. To provide relevant experts at the national level from regional ones, we consider different approaches for aggregating these experts:

- **Full disaggregated model:** we use the full set of scaled forecasts as experts (143 experts, 11 by regions + 11 at the national level), and the scaled national demand y_{Fr}^{norm} as our target variable. The prediction is then multiplied by the average national load to forecast the load y_{Fr} .
- **Hierarchical aggregation, scaled predictions:** as a first step, we aggregate the 11 experts in each region using the scaled regional demand for a region z , y_z^{norm} as a target, and we obtain 12 experts. Then, we aggregate these 12 experts and the quantile experts at the national level, with the scaled national demand as our target variable y_{Fr}^{norm} . The prediction is then multiplied by the average national load to forecast the load y_{Fr} .
- **Hierarchical aggregation, unscaled predictions:** we again aggregate the 11 experts in each region, using the scaled regional demand as a target, and we obtain 12 experts predicting the normalized load at regional level. Then, we multiply their predictions by the average load of the corresponding region, and sum these predictions in order to obtain a forecast of national demand y_{Fr} .

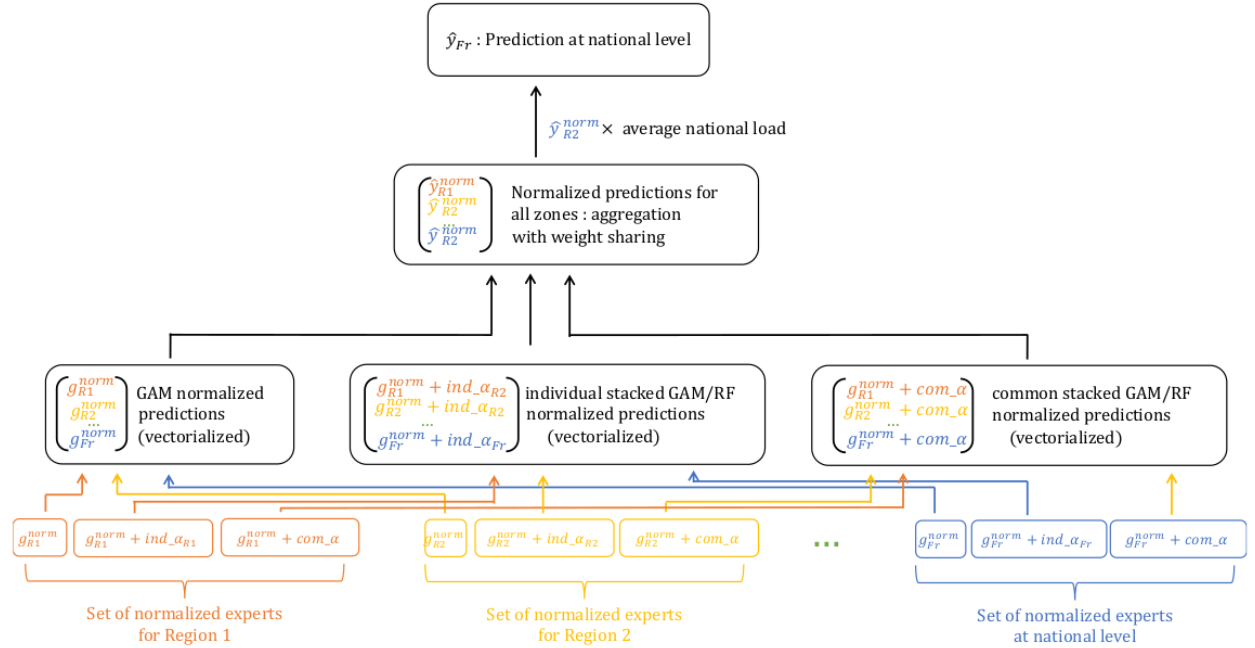
- **Vectorial aggregation:** we illustrate the possibility to share weights between the regions and at the national level. We aim at predicting the time series of the 13-dimensional vector corresponding to the scaled loads in each regions and at the national level. To do so, we aggregate 11 vectorial, 13-dimensional experts corresponding to the predictions of the GAM and of the 10 stacked RF. The prediction corresponding to the national level is then multiplied by the average national load to forecast the load y_{Fr} .

Our methods for aggregating the experts at regional and national levels are summarized in Figure 2.





(c) Unscaled hierarchical aggregation.



(d) Vectorial aggregation.

Figure 2: Aggregations models for predicting national load.

4.3 Results

In this Section, we compare the methods presented above. In a first time, we compare their performances in terms of MAPE and RMSE in Section 4.3.1. In a second time, we lead an importance by permutation analysis of the RF, which is presented in Section 4.2.2. Finally, we analyse and compare the different aggregation methods in Section 4.2.3.

4.3.1 Performances

In Table 3, we compare the MAPE and the RMSE of the 4 methods, of the GAM at the national level, and of the stacked individual RF predicting the median of the residuals at national level.

We split the test period into 3 sub-periods. In the pre-pandemic period between September 1st 2019 and March 15th 2020, only the GAM predictions are available. During this period, vectorial aggregation makes little sense since there is only one type of experts. The lock-down period ranges between March 16th and May 11th: during this period, training data are very scarce, and models must quickly adapt to a dramatic change in electricity consumption pattern. The post-lock-down period, from May 12th to September 17th, corresponds to a new change in load pattern, due to a relative rebound in activity, to which the models must adapt.

Model	2019/09/01-2020/03/15	2020/03/16-2020/05/11	2020/05/12-2020/09/17
GAM	1.36 %, 1030 MW	4.82 %, 2838 MW	1.84 %, 1045 MW
Individual stacked RF	Non applicable	2.41 %, 1813 MW	1.03 %, 592 MW
Full disaggregated	1.20 %, 910 MW	2.26 %, 1716 MW	1.09 %, 609 MW
Hierarchical aggregation scaled	1.14 %, 861 MW	2.21 %, 1648 MW	1.07 %, 609 MW
Hierarchical aggregation unscaled	1.20 %, 907	2.08 %, 1553 MW	1.02 %, 593 MW
Vectorial aggregation	Non applicable	2.56 %, 1885 MW	0.91 %, 521 MW

Table 3: Mean Absolute Percentage Error and Root of the Mean Squared Error of the stacked GAM and RF models.

4.3.2 Analysis of the stacked GAM and RF

Stacked RFs allow to quickly adapt to the change in distribution of the load during the pandemic period and to efficiently correct the errors of the historical GAM. Using a stacked RF predicting the median of the GAM residuals at national level is enough to decrease its MAPE during and after the lock-down by respectively 50 and 45%.

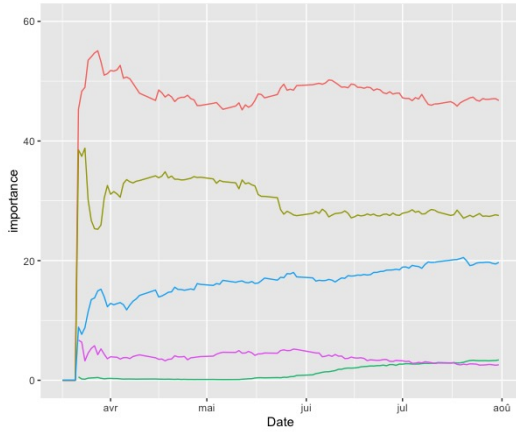
Figures 3a and 3b present the evolution of the average importance (for the different half-hours) of the variables for the stacked RF trained respectively on residuals at the national level, and on all residuals. The importance of the different variables for a given model are normalized so that their sum remains constant during the pandemic period, and equal to 100. More precisely, denoting by $I_{v,t}$ and $I_{v,t}^{normalized}$ respectively the importance and normalized importance of variable v at time t , we have at any time t

$$I_{v,t}^{normalized} = 100 * \frac{I_{v,t}}{\sum_{variables\ v'} I_{v',t}}$$

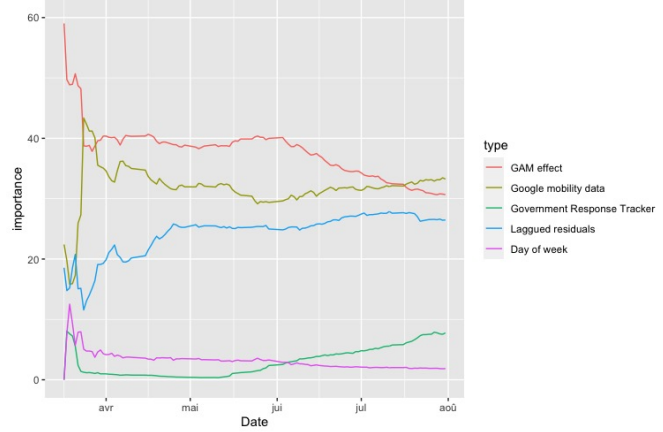
We group the variables into 5 categories: the GAM effects, the measures of mobility, the government response tracker, the lagged residuals, and the day of week. The importance of a group is simply the sum of the importances of the variables in this group. The importance of the mobility measures is detailed in Figures 3c and 3d.

We note that the effects of the GAM are among the most important covariates for predicting the GAM residuals. Using these effects as covariates allows to transfer information on the impact of weather and calendar variables learned on the large dataset of pre-pandemic observations. We observe a change in the importance of the different types of variables after the end of the lock-down, indicating that the RF is able to account for to a relative change in electricity consumption patterns. As the time passes and the size of the training set for the RF increases, relevant variables such as the Government Response Tracker, or relative occupation of residence, and grocery and pharmacies, become more important for the prediction. Conversely, spurious variables (for example, the relative frequentation of parks, highly correlated with weather) are discarded as unimportant. Interestingly, the common RF trained on residuals accross all regions detects these relevant variables more quickly than the individual RF trained solely on residuals at the national level. This highlights the interest of multi-task learning in this sparse data context.

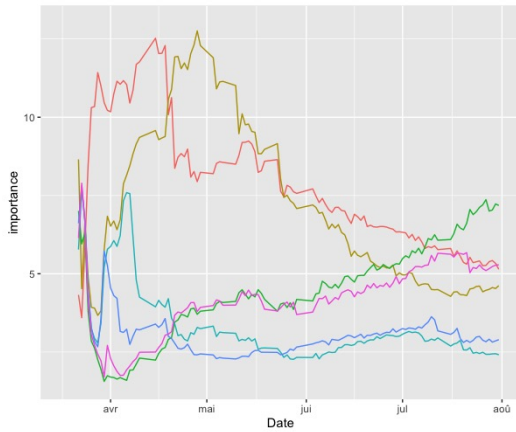
We also investigate the relative importance of the variables in the stacked RF predicting the different quantiles. More precisely, we consider the stacked RF trained on the residuals at national level for the pandemic period. We compute an importance measure of a given variable as the average increase in error in term of the pinball loss corresponding to a given quantile when the values of this variable are permuted at random (the error is computed over the training set). The importance of the different variables are normalized so that their sum is equal to 100. We compare the importance of the variables for predicting the 0.05, 0.5 and 0.95 quantiles in Figure 4. Variables important for predicting one quantile tend to be important for predicting the other quantiles. However, this is not the case for all variables: for example, the normalised load for the relative frequentation of residential place and



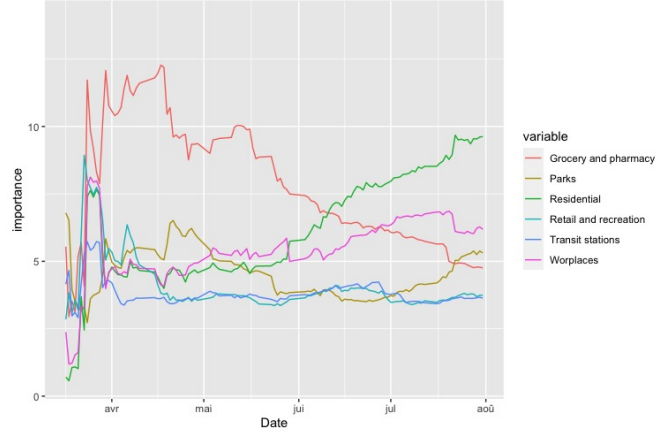
(a) Average importance of types of variables in the staked RF at national level.



(b) Average importance of types of variables in the staked RF common to all regions and the national level.



(c) Average importance of mobility measures in the staked RF at national level.



(d) Average importance of types of mobility measures in the staked RF common to all regions and the national level.

Figure 3: Evolution of the importance of the types of variables (top) and the mobility measures (bottom) in the RF trained on GAM residuals at national level (left), and on GAM residuals for all regions and at the national level.

workplaces have an outstanding importance for predicting the 0.05 and 0.95 quantiles. These variables have a very high (negative) correlation: the frequentation of workplaces is very low during the lock-down period, and remains relatively low in the post lock-down period during weekdays; on the opposite the frequentation of residential places is high during the lock-down period, and remains relatively high in the post lock-down period during weekdays. Looking at the Accumulated Local Effects of these variables, plotted in Figure 5, we see that they have a much larger impact on the prediction of the two extreme quantiles than on that of the median. However, we expect their effects to partially cancel each other out because of the correlation between these variables.

4.3.3 Analysis of online aggregation

Our results indicate that online aggregation is an efficient way to take into account information available at a finer scale. Note that the regional GAM have in average errors much larger than that of the GAM at the national level, as illustrated in Figure 7, due to larger fluctuations present at the finer scale. Interestingly, aggregating these low-accuracy models allows to obtain better performances than that of the GAM at national level, even in the pre-pandemic period, as indicated by our results in Table 3.

The fact that scaled and unscaled hierarchical aggregation obtain similar performances is somewhat counter

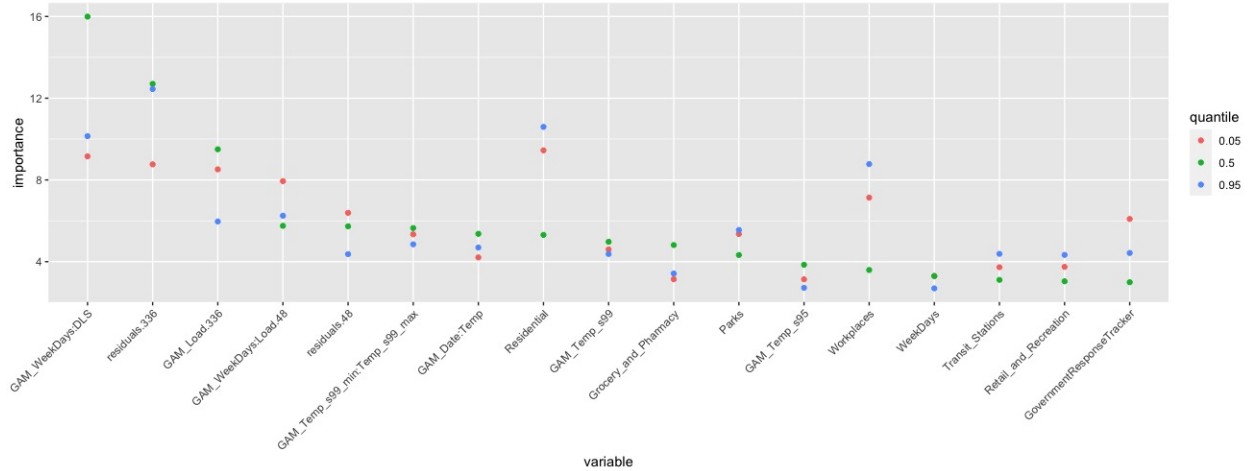
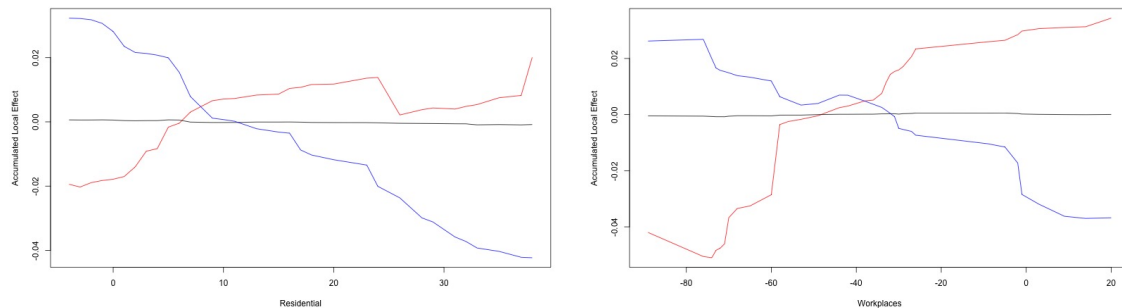


Figure 4: Importance of the variables in the stacked RF predicting the quantiles 0.05, 0.5 and 0.95 of the GAM residuals at national level. The variables “GAM_X” denotes the GAM effect corresponding to variable “X”.



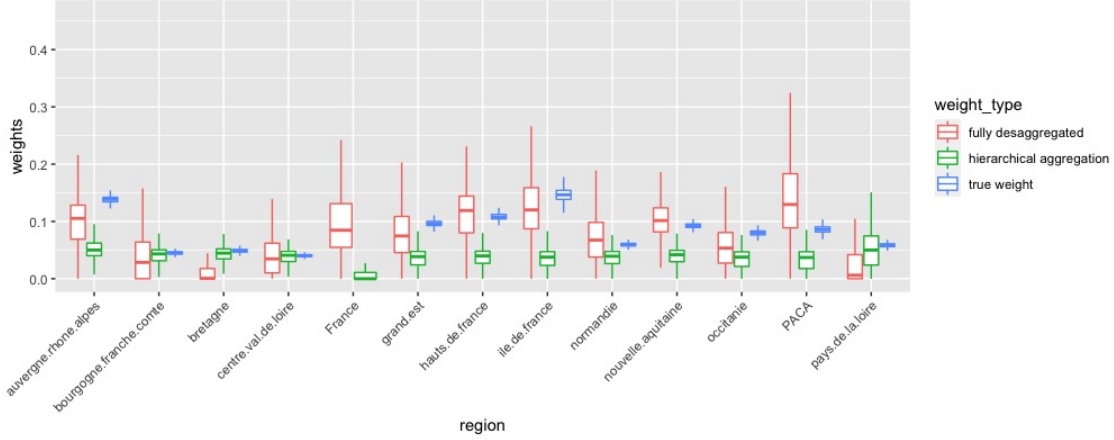
(a) Accumulated Local Effects of the relative frequentation of residential places.

(b) Accumulated Local Effects of the measure of relative frequentation of workplaces.

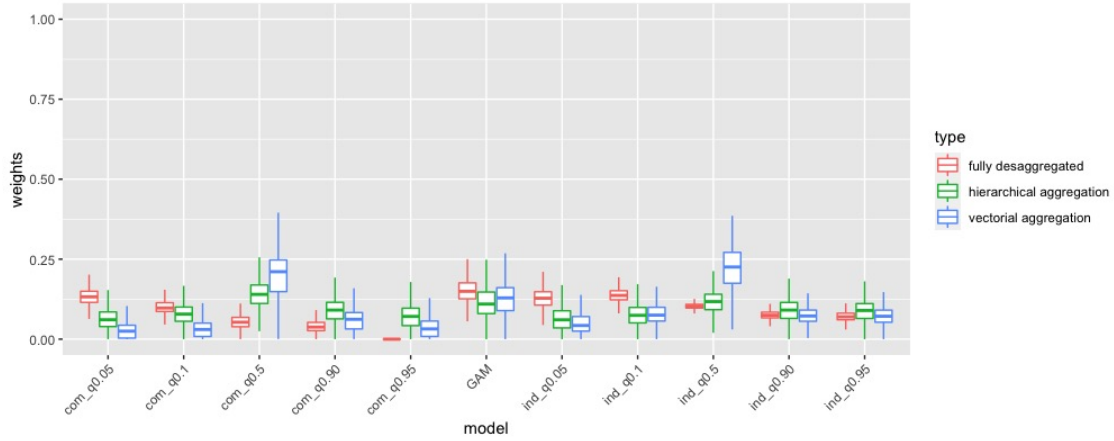
Figure 5: Accumulated Local Effects of the measure of relative frequentation of residential places (left) and workplaces (right) for the RF at national level predicting the quantiles 0.05 (red), 0.5 (black) and 0.95 (blue).

intuitive, given that in the scaled model the aggregation must learn the contribution of the different regions to the national consumption. Looking at the distribution of the weights of the regions in the scaled hierarchical aggregation presented in Figure 6a, we note that the weights do not correspond to the proportion of electricity consumed by the regions (for example, regions with low true weights such as Provence-Alpes-Côte d’Azur may receive more weight in the aggregation than regions with high true weights, such as Île-de-France). Moreover, the weights in the aggregation typically exhibit much more flexibility than the true weights: this phenomenon is all the more striking in the full disaggregated model. The high variability of the weights suggests that some of the models considered are fairly interchangeable. The fact that the scaled hierarchical aggregation outperforms its unscaled counterpart both in the pre-pandemic and in the post-lock-down period suggests that the flexibility provided by the second layer of aggregation used in the scaled model compensates the lack of knowledge of the relative contribution of the different regions.

We see in Figure 6b that all quantiles and GAM experts contribute to the prediction, both in the full disaggregated model and in the hierarchical aggregation. By contrast, the vectorial aggregation gives a predominant weight to the GAM and median stacked RF experts, which appears as the most relevant experts across all regions. Figure 6 shows the weights given by an aggregation predicting the national load using the only the national experts, and the weights



(a) **Weights of the regional and national experts in the prediction at national level.** Red: sum of the weights of the quantile and GAM experts by region, in the aggregation targeting the national load using the full disaggregated approach. Green: weights of the regional experts, and for the national level in the aggregation targeting the national load using the scaled hierarchical approach. Blue: true proportion of the national electricity load consumed by the region.



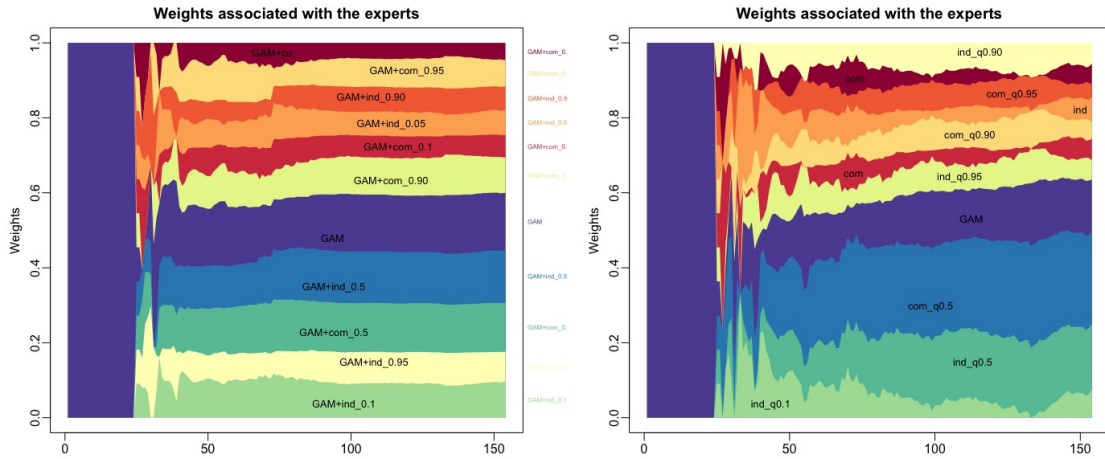
(b) **Weights of the quantile and GAM experts in the prediction at national level.** Weights of the quantile experts and the GAM expert in the aggregation targeting the national load using a full disaggregated approach (red), a hierarchical aggregation approach (green) and a vectorial aggregation approach (blue).

given by the aggregation. Day 25 corresponds to the first day of the pandemic period; before that day only the GAM forecast is available to the the aggregation. We note that the weights in the vectorial aggregation are highly unstable during the beginning of the lock-down. The performance of the vectorial aggregation during this period is worst than that of all other aggregation models, and than that of the stacked RF predicting the median of residuals. This behaviour mirrors the fact that the impact of the pandemic strongly differs from one region to another, as is shown in Figure 7. On the other hand, vectorial aggregation achieves the best performance during the post lock-down period, and appears as a promising approach to predicting consumption under normal circumstances.

5 Conclusions and Future Work

We propose new transfer learning methods designed for forecasting time series observed at different hierarchical scales. We present two different settings and illustrate it with two different usecases:

1. To transfer information from finer scale (an aggregate of smart meters) to wider scale (national) data when the distribution of the data is stable with time, we propose to stack features from GAM obtained at these two scales into random forests.



(a) Weights of the GAM and quantile experts in the first step of hierarchical aggregation, targeting the national load at 7:30 pm.

(b) Evolution of the weights of the GAM and quantile experts in the vectorial aggregation at 7:30 pm.

Figure 6: Evolution of the weights of the quantile stacked RF experts and the GAM expert in the prediction of national load.

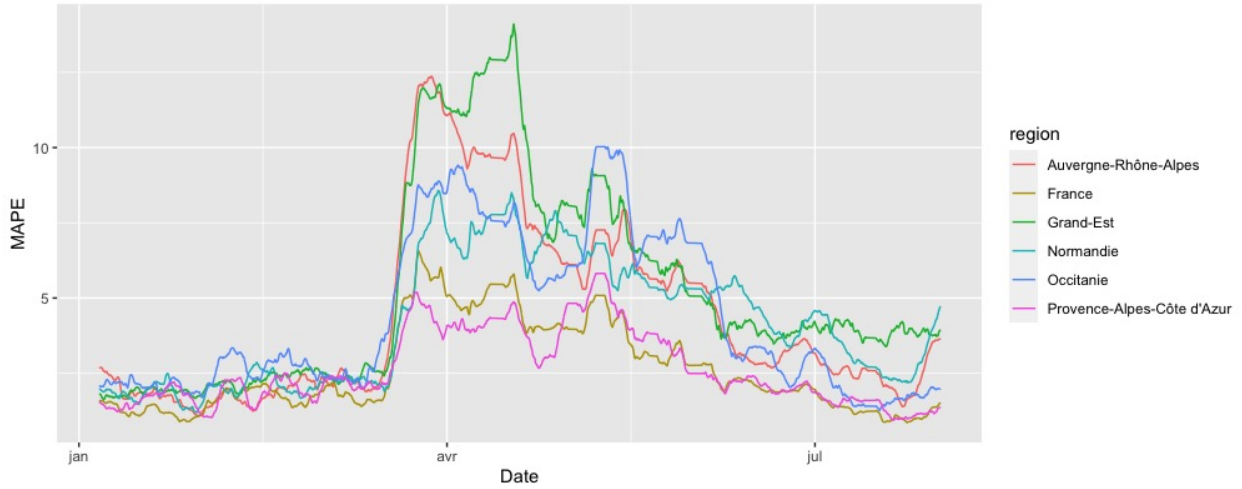


Figure 7: Weekly averaged MAPE of the normalized GAM for the regions Auvergne-Rhône-Alpes, Grand-Est, Normandie, Occitanie, and at the national level.

2. To transfer information from local to global data when the distribution of the data is changing with time we propose hierarchical online aggregation of experts where the experts are generated at a finer scale (regional level) using quantile stacked random forest.

We demonstrate the interest of our proposed approach in both cases. In both cases transfer learning by RF stacking at a single scale improve significantly the forecasting performance of single GAM or RF model: 14% of improvement over GAM and 9% over the RF for case 1, 38% over GAM for case 2. It demonstrates our original intuition that stacked RF gather both the ability of GAM to extrapolate and RF to model automatically interaction between covariates.

Regarding multi-scale transfer performances, we also obtained convincing results. In case 1, we improved day-ahead forecasting performance of the wider scale stacked RF of about 1.5% with our multi-scale transfer algorithm.

For case 2, the best hierarchical aggregation algorithm improve about 10% the stacked RF at a wider scale. Our relatively simple strategy of re-scaling plus aggregation behaves well in this 2 scale hierarchy. We also saw that introducing strong constrains in the aggregation weight (vectorial aggregation) can be an interesting transfer strategy when the experts behave similarly at the different scales of the hierarchy. This is true during the post covid period but not during the hard lockdown in France of march-april 2020 (we suspect that the effect of COVID over the electricity load impact the different regions in a desynchronised way).

The main learner used in the paper for the final forecasting is based on stacked RF. We could have chosen other machine learning methods such as tree based gradient boosting or neural networks, this can be tested in future work. We showed that automatic variable selection when forecasting didn't show any improvement. However, in a high dimensional setting with a large number of features generated when learning the source, we believe that a possible approach that is worth exploring, is to use for forecasting a regression-reinforced random forest (RFRF) approach that may have better prediction performance than RFs. The idea behind RFRF is to borrow the strength of penalized parametric regression to improve RF. For example, for RFRFs, we may run a SCAD (or LASSO) (see Fan and Li (2001)) based selection before RF, then construct a RF on the residuals from the SCAD (or LASSO) penalized fit. Preliminary simulation results show that RFRFs can capitalize on the strength of both parametric and nonparametric methods and may give reliable predictions in high-dimensional extrapolation problems such as those encountered in transfer learning.

For case 1, we didn't investigate clustering of smart meter data to generate diverse GAM feature but this is clearly a possible improvement. Introducing hierarchical constrains in the weights as proposed in Brégère and Huard (2021) is also a potential perspective for case 2. Also, finding the good warping of weights constrain for vectorial aggregation could be a way to improve the performance of this method on desynchronized data.

Acknowledgement

The authors acknowledge the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-20-CE40-0025-01 (T-REX project).

References

- AECOM (2018). Energy demand research project: Early smart meter trials, 2007-2010. Technical report, UK Data Service.
- Amato, U., Antoniadis, A., De Feis, I., Goude, Y., and Lagache, A. (2021). Forecasting high resolution electricity demand data with additive models including smooth and jagged components. *International Journal of Forecasting*, 37(1):171–185.
- Balestriero, R., Pesenti, J., and LeCun, Y. (2021). Learning in high dimension always amounts to extrapolation. Technical report, arXiv:2110.09485.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.
- Breiman, L., Friedman, J., Olshen, R., and Stone, C. (1984). *Classification and regression trees*. Chapman & Hall/CRC.
- Brégère, M. and Huard, M. (2021). Online hierarchical forecasting for power consumption data. *International Journal of Forecasting*, to appear.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA.
- de Vilmarest, J. and Goude, Y. (2021). State-space models win the ieeedataport competition on post-covid day-ahead electricity load forecasting. Technical report, arXiv:2110.00334.
- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.*, 96:1348–1360.
- Fan, S. and Hyndman, R. J. (2012). Short-term load forecasting based on a semi-parametric additive model. *IEEE Transactions on Power Systems*, 27(1):134–141.
- Farrokhbadi, M., Browell, J., Wang, Y., Makonin, W., and Zareipour, H. (2021). Day-ahead electricity demand forecasting competition: Post-covid paradigm. Technical report.

- Gaillard, P. and Goude, Y. (2016). opera: Online prediction by expert aggregation. URL: <https://CRAN.R-project.org/package=opera>. r package version, 1.
- Gaillard, P., Stoltz, G., and Van Erven, T. (2014). A second-order bound with excess losses. In *Conference on Learning Theory*, pages 176–196.
- Genuer, R., Poggi, J.-M., and Tuleau-Malot, C. (2015). VSURF: An R Package for Variable Selection Using Random Forests. *The R Journal*, 7(2):19–33.
- Goehry, B., Goude, Y., Massart, P., and Poggi, J.-M. (2019). Aggregation of multi-scale experts for bottom-up load forecasting. *IEEE Transactions on Smart Grid*, 11(3):1895–1904.
- Goude, Y., Nedellec, R., and Kong, N. (2013). Local Short and Middle term Electricity Load Forecasting with semi-parametric additive models. *IEEE transactions on smart grid*, 5(1):440 – 446.
- Kursa, M. B. and Rudnicki, W. R. (2010). Feature selection with the Boruta package. *Journal of Statistical Software*, 36(11):1–13.
- Obst, D., de Vilmarest, J., and Goude, Y. (2021a). Adaptive methods for short-term electricity load forecasting during covid-19 lockdown in france. *IEEE Transactions on Power Systems*, 36(5):4754–4763.
- Obst, D., Ghattas, B., Cugliari, J., Oppenheim, G., Claudel, S., and Goude, Y. (2021b). Transfer learning for linear regression: a statistical test of gain. Technical report, arXiv:2102.09504.
- Olivas, E. S., Guerrero, J. D. M., Sober, M. M., Benedito, J. R. M., and Lopez, A. J. S. (2009). *Handbook Of Research On Machine Learning Applications and Trends: Algorithms, Methods and Techniques - 2 Volumes*. Information Science Reference - Imprint of: IGI Publishing, Hershey, PA.
- Pan, S. J. and Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.
- Schellong, W. (2011). *Energy demand analysis and forecast*. BoD–Books on Demand.
- Wang, Y., Chen, Q., Hong, T., and Kang, C. (2019). Review of smart meter data analytics: Applications, methodologies, and challenges. *IEEE Transactions on Smart Grid*, 10(3):3125–3148.
- Wood, S. (2017). *Generalized Additive Models: An Introduction with R*. Chapman and Hall/CRC, 2 edition.
- Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., and He, Q. (2020). A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76.
- Ziel, F. (2021). Smoothed bernstein online aggregation for day-ahead electricity demand forecasting. Technical report, arXiv:2107.06268.