

Adaptive Aggregated Predictions for Renewable Energy Systems

Balázs Csanád Csáji*

András Kovács*

József Váncza*[†]

*Fraunhofer Project Center for Production Management & Informatics
Institute for Computer Science and Control, Hungarian Academy of Sciences

[†]Budapest University of Technology and Economics, Hungary
{balazs.csaji, andras.kovacs, jozsef.vancza}@sztaki.mta.hu

Abstract—The paper addresses the problem of generating forecasts for energy production and consumption processes in a renewable energy system. The forecasts are made for a prototype public lighting microgrid, which includes photovoltaic panels and LED luminaries that regulate their lighting levels, as inputs for a receding horizon controller. Several stochastic models are fitted to historical times-series data and it is argued that side information, such as clear-sky predictions or the typical system behavior, can be used as exogenous inputs to increase their performance. The predictions can be further improved by combining the forecasts of several models using online learning, the framework of prediction with expert advice. The paper suggests an adaptive aggregation method which also takes side information into account, and makes a state-dependent aggregation. Numerical experiments are presented, as well, showing the efficiency of the estimated time-series models and the proposed aggregation approach.

I. INTRODUCTION

Renewable energy systems are vital for sustainability and optimizing their energy flow is a widely studied research area. One of the fundamental problems of building controllers for renewable energy systems is to model energy production and consumption, as they are uncertain and affected by a number of external factors. Such models are essential to make predictions about the future behavior of these systems [10].

Standard approaches, for example, to predict photovoltaic (PV) energy production, include clear-sky and persistence [12], as well as various dynamic models [14], such as autoregressive integrated moving average (ARIMA), artificial neural networks (ANNs), fuzzy and different hybrid models [3].

In this paper we combine several of the aforementioned approaches by fitting stochastic models to time-series data, but also providing the models side information (such as the typical system behavior) as external inputs. We argue that such information increases their performance and by aggregating their predictions online an efficient adaptive forecaster can be constructed. A model predictive controller is also discussed which uses the generated forecasts to calculate a control policy.

We start from the standard framework of prediction with expert advice [4], then refine it to allow situation dependent performance evaluations. The core idea is to use a similarity kernel when calculating the loss of an expert with respect to a given state and weight past costs with their similarity to the current situation. As supported by our simulation experiments, this aggregation approach outperforms standard ones as it can estimate better the actual efficiency of different models.

Previously, the framework of prediction with expert advice was used to predict energy consumption in [8], and online learning with side information was investigated, for example, in [4, 13, 16]. As we will later see in Section V-B, our approach differs from the ones above and one of our main contributions is to use situation dependent losses during prediction which allows the system to adapt to the changing circumstances.

II. E+GRID: ENERGY-POSITIVE PUBLIC LIGHTING

A particular motivation for this research is provided by a work aimed at realizing an *energy-positive public lighting microgrid* (E+grid) system. E+grid uses renewable, solar energy in public lighting services via the appropriate combination of LED luminaries, energy generation and storage, and sensor technologies on the one hand, and novel data processing, prediction, communication, optimization and control methods on the other hand. E+grid balances energy *demand* against *production* and guarantees the required level of street lighting even at times of moderate-duration power outages.

The E+grid system reduces the energy consumption of street lighting by using LED luminaries that regulate their lighting levels according to the actual demand and the surrounding environmental conditions, therefore providing just the required level of lighting at all times.

On the demand side, this is achieved by mounting motion sensors and smart controllers to each light pole, which are then connected by wireless communication. As a vehicle or passerby moves along, the level of lighting is automatically increases before, and, at the same time, dims behind the mover. Smart, autonomous controllers of neighboring luminaries make sure that all this happens in a well-orchestrated way, according to the standards set for local public lighting services.

As for production, energy is generated by PV technology, whereas battery storage provides some limited protection against power outages, as well as opportunity for trading with electricity. Namely, the system has a bi-directional connection to the power grid, hence it can sell and buy electricity at variable energy tariff. Energy flows along all main power lines are monitored by smart meters. A local weather station complementing the system collects weather data and hosts a twilight switch which allows for the lighting periods to comply with the actual environmental conditions. The overall architecture of the E+grid system with its main communication and power lines is presented in Figure 1 (see also [6, 7]).

While the smart LED luminaries are equipped with limited local decision and communication facilities, the heart of the E+grid system is a *central computer* (CC) that monitors and controls its operation. One of the key functions of the CC of particular interest here is the *control of the energy flow*, with regard to the following two criteria:

- warranting island mode operation in case of an eventual power outage for a limited period; and
- minimizing the total energy cost of the public lighting service in the long run.

The E+grid system as a whole is a complex *cyber-physical system* whose elements – the smarter ones even with limited autonomy – are interlinked by lines both of communication and of power. Further on, the system is embedded in a highly uncertain physical and social environment: the energy supply and demand for the lighting service depend not only on the actual location and moment of time, but also on the local weather conditions, while the demand for public lighting depends a lot on the movement of humans and vehicles in the illuminated area. Hence, defining however approximate physical model of the system, together with its interactions with the environment would be far too ambitious. Instead, we rely on data collected by extensive and continuous monitoring activity, analyze their time-series, predict both energy production and consumption and use these forecasts when balancing future, expected consumption and supply. However, so as to keep the reality and the model of the controlled system in as a close correspondence as possible, the predictions and the energy flow control are interleaved: the model is mapped to the (observed) reality time and again, in *model predictive* way of control.

Smart meters monitoring both the production and the consumption of electricity provide ample input for the CC to *fit stochastic models* to historic data, and to solve the resulting energy flow optimization problem on a rolling horizon. Indeed, an assessment of several different models showed that good predictions of energy production can be obtained by using nonlinear autoregressive exogenous (NARX) models [7], w.r.t. the deviation-normalized root-mean-square error criterion. These models use wavelet-type nonlinearities, where the exogenous components come from a clear-sky model capturing

background knowledge of the domain. Efficient (w.r.t. the same criterion) energy consumption forecasts can also be achieved using Box-Jenkins models, where the exogenous inputs come from averaged historical data. As simulation experiments have shown, with right prediction methods the controller is able to significantly improve the financial balance. In what follows we briefly discuss the applied time-series models, our receding horizon controller and present how the system can make the best out of running several predictive models in parallel.

III. IDENTIFYING QUASI-PERIODIC PV PRODUCTION AND ENERGY CONSUMPTION PROCESSES

One of the main challenges when designing a controller for a renewable energy system is to model or forecast the energy production and consumption signals, as these processes are typically affected by several factors, including the weather (PV production) and human behavior (energy consumption) and are, therefore, typically nonstationary and hard to predict.

A. Standard Approaches

There are several standard forecasting models [3, 10] available, such as *clear-sky models* for PV production [12], which estimate the terrestrial solar radiation under the assumption of a cloudless sky while they typically take the solar elevation angle, site altitude and potentially other (e.g., atmospheric) conditions into account. The arrival of customers (i.e., to the area of the controlled public lighting system) can, for example, be modeled by *Poisson processes* [1] which then can be applied to estimate the future power consumption of the system.

Naturally, there are several refinements of the above mentioned basic approaches, such as *persistence models* or *spatio-temporal forecasts* for PV production. Persistence models assume that the current weather conditions persist and scale the clear-sky predictions for the given horizon with the current deviation from that estimate [12]; while spatio-temporal forecasts usually smooth the forecasts of several monitoring stations, distributed in an area, to generate forecasts for any spatial location within the area covered by the stations [18].

Another approach is to assume an *autoregressive* behavior and apply approximation techniques to fit a sufficiently smooth function to the available historical data. Feed-forward neural networks (e.g., multilayer perceptions or radial basis function networks) are typical choices for such approaches [14].

B. Model Identification and Side Information

Our approach here is to generate (and later aggregate) various time-series models by applying system identification [11] and machine learning [17] techniques while also taking background knowledge, such as clear-sky predictions or the typical system behavior, into account as side information for the models. The time-step of our models was one hour, which we achieved by averaging our observations for each hour.

As the PV production and the energy consumption of the system have a quasi-periodic nature, we treat these signals as the combination of a fully periodic (mean behavior) and an aperiodic part. For the periodic part we simply compute the average value for each hour of the day. Providing such side information as exogenous input to the models may help them

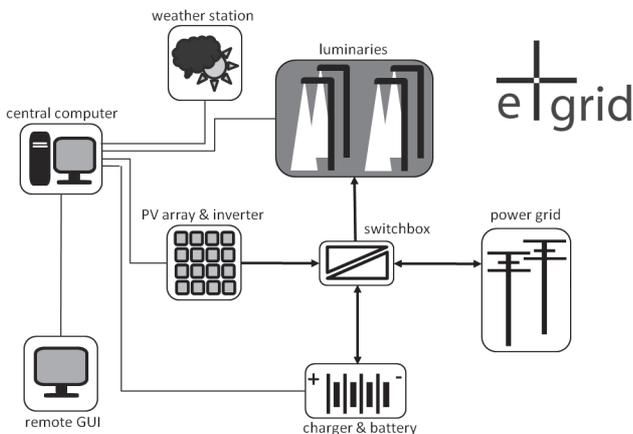


Figure 1. Schematic architecture of the E+grid system (the thick, black arrows and the thin, gray links indicate power and information flow, respectively).

to predict the system, as it is also supported by our experiments presented in Section VI. This is especially the case if the system operates close to its typical behavior, for example, the sky is relatively clear or the consumption is normal.

On the other hand, in cases when the system operates far from its typical states such side information might mislead the models. In these cases it would be good to switch to pure autoregressive models, i.e., without exogenous components. Therefore, our idea is to estimate both types of models and calculate a prediction online based on the recent performances of the models. Before we present how we aggregate the predictions, we briefly overview the relevant approaches from time-series analysis and how the control policy is then computed.

C. Time-Series Analysis

A *time-series* is a data sequence, typically consisting of noisy observations of a dynamical system at discrete time-steps [2]. Estimating models based on time-series data is one of the classical problems of *system identification* [11, 15].

Discrete-time (observable, causal) stochastic systems with exogenous components (inputs) can be typically written as

$$x_t \triangleq f(z_t, w_t), \quad (1)$$

$$z_t \triangleq (x_{t-1}, x_{t-2}, \dots; u_t, u_{t-1}, \dots), \quad (2)$$

$$w_t \triangleq (n_t, n_{t-1}, \dots), \quad (3)$$

where x_t is the output, u_t is the input and n_t is the noise at time t . Sequence z_t contains the available (cumulative) information at time t , while w_t is the (unobservable) noise up to time t . Note that process $\{u_t\}$ can represent, e.g., a control signal or some available side information. Observe that u_t is included in z_t , but of course x_t is not. Process $\{n_t\}$ is often a white noise or an independent sequence of random variables.

We are usually given a realization of z_t , i.e., $z_t(\omega)$, and want to find a function $\hat{f} \in \mathcal{F}$ from a given model class that optimizes some function of the *prediction errors* defined as

$$\varepsilon_k(\hat{f}) \triangleq x_k(\omega) - \hat{f}(z_k(\omega), 0), \quad (4)$$

where 0 is the sequence of all zeros. An archetypical objective is to minimize $\sum_k \varepsilon_k^2(\hat{f})$ (i.e., least squares or L_2 approach), but it often includes some regularization terms, as well.

Standard stochastic models include general LTI (linear time-invariant) systems, which can be formalized as

$$A(q^{-1})x_t \triangleq \frac{B(q^{-1})}{F(q^{-1})}u_t + \frac{C(q^{-1})}{D(q^{-1})}n_t, \quad (5)$$

where A, B, C, D and F are (finite) polynomials in q^{-1} , the backward *shift* operator (i.e., $q^{-1}x_t = x_{t-1}$), and x_t, y_t, n_t are as previously. Special cases of LTI systems include FIR (finite impulse response), AR (autoregressive), ARX (autoregressive exogenous), ARMAX (autoregressive moving average exogenous) and BJ (Box-Jenkins) type models [11, 15].

LTI models can also take a state space form, that is

$$x_t \triangleq Ax_{t-1} + Bu_t + Cn_t, \quad (6)$$

where A, B and C are given (real or complex) matrices and we assumed that the process $\{x_t\}$ is fully observable.

In some cases linear models are not suitable. Widespread nonlinear models include HW (Hammerstein-Wiener) and NARX (nonlinear autoregressive exogenous) systems [11]. A HW model contains a linear part where the input and the output are transformed by static nonlinearities (i.e., h and g below),

$$x_t \triangleq g\left(\frac{B(q^{-1})}{F(q^{-1})}h(u_t)\right) + n_t, \quad (7)$$

while NARX models take the form of

$$x_t \triangleq g(z_t) + n_t \quad (8)$$

$$= g(x_{t-1}, \dots, x_{t-q}, u_t, \dots, u_{t-s+1}) + n_t, \quad (9)$$

where q, s are the orders of the model and g is a nonlinear function. There are several variants of NARX models depending on the applied nonlinear function g . In our experiments we used (i) wavelets, (ii) multilayer perceptrons (MLPs) and (iii) support vector regression (SVR), to represent g , resulting in three types of models. A NAR model is a special case of a NARX system where there are no inputs. We also used two NAR models in our test, an MLP and an SVR based version.

D. Forecasting with the Identified Models

After having identified a model \hat{f} we may want to use it to generate forecasts. For the aforementioned E+grid project, mean trajectories as well as confidence bounds were needed for the controller. Assuming the input signal $\{u_t\}$ is available in advance, a way to estimate the mean trajectory (i.e., the expected future behavior) is to use zero noise and estimated states for states we do not have information about. That is

$$\hat{x}_i \triangleq \hat{f}(\hat{z}_i(\omega), 0), \quad (10)$$

where 0 is the zero sequence and in \hat{z}_i we recursively use \hat{x}_j , (recall that the process is causal, thus $j < i$) for future states. Of course, for a 1-step prediction we typically do not need such \hat{x}_j variables, they are needed for longer horizons.

In order to get confidence bounds, we need some information about the noise process driving the system. A standard approach is to assume that it is a sequence of zero mean independent and identically distributed (i.i.d.) Gaussian variables. Then, it is enough to estimate the variance of the noise to characterize it. This can be done by [11], e.g.,

$$\hat{\sigma}_n^2(\hat{f}) \triangleq \frac{1}{n-d} \sum_{k=1}^n \varepsilon_k^2(\hat{f}), \quad (11)$$

where d is the dimension of the model class, e.g., if the model class is (linearly) parametrized by $\theta \in \mathbb{R}^d$, then $d = q$.

The choice of Gaussian variables are appropriate if we consider the noise as the composition of several independent effects, in which case the Central Limit Theorem [11] justifies the Gaussian distribution. On the other hand, any other distribution can also be used, including nonparametric ones, e.g., we could even apply the empirical distribution function of the residuals $\{\varepsilon_k(\hat{f})\}$, resulting in a bootstrap-style approach [9].

After the distribution of the noise was estimated, Monte Carlo experiments can be carried out, using the last values of our observations as initial states and randomly generated noise according to the identified noise distribution, to generate simulated trajectories. Then, approximate upper [lower] confidence

bounds can be calculated from the simulated trajectories by finding the smallest [largest] sequence that is larger [smaller] than a given percentage, for example 95 %, of the trajectories.

IV. CONTROLLING THE ENERGY FLOW

In the E+grid project we use the generated energy production and consumption forecasts, i.e., the mean and confidence bound trajectories, to generate a policy to control the energy flow in the lighting system. Particularly, we apply a receding horizon controller, namely, in each step we compute an open-loop control sequence for a given horizon T and the environmental feedback is incorporated by recalculating the control sequence, taking new forecasts into account, after each iteration. The policy thus can be seen as rollout type, while the method is a variant of model predictive control (MPC).

Now, we discuss one iteration of the controller, namely computing a finite-horizon open-loop control sequence. In each step, the open-loop policy is the solution of an optimization problem. The input contains the expected future energy production, $\{C_t^+\}$ and consumption $\{C_t^-\}$, as well as stochastically guaranteed lower confidence bounds on production $\{\underline{C}_t^+\}$ and upper confidence bounds on consumption $\{\overline{C}_t^-\}$. The control policy must be *robust* in the sense that it must guarantee island-mode operation for a given amount of time for a single power cut arising at any point in time, even in the worst-case scenario defined by $\{\underline{C}_t^+\}$ and $\{\overline{C}_t^-\}$. This requirement can be fulfilled by maintaining the appropriate state of charge, $\{B_t\}$, in the battery. The battery is characterized by its capacity \overline{B} , maximum charge and discharge rates R^+ and R^- , the initial state of charge b_0 and the efficiency of charging β . A method for computing $\{B_t\}$ from $\{C_t^+\}$ and $\{\overline{C}_t^-\}$, together with the detailed assumptions, is presented in [7].

Given the above input data, an open-loop control sequence defining the optimal electricity purchase rate x_t^+ , grid feed-in rate x_t^- , as well as a battery charge rate r_t^+ and discharge rate r_t^- is sought for each time period t that minimize the total energy cost in the system subject to time-varying electricity purchase and feed-in prices Q_t^+ and Q_t^- . A linear programming formulation of this problem is presented below.

$$\text{minimize } \sum_{t=1}^T (Q_t^+ x_t^+ - Q_t^- x_t^-) \quad (12)$$

subject to

$$C_t^+ - C_t^- + x_t^+ - x_t^- = r_t^+ - r_t^- \quad \forall t \quad (13)$$

$$\beta r_t^+ - r_t^- = b_t - b_{t-1} \quad \forall t \quad (14)$$

$$\underline{B}_t \leq b_t \leq \overline{B} \quad \forall t \quad (15)$$

$$0 \leq r_t^+ \leq R^+ \quad \forall t \quad (16)$$

$$0 \leq r_t^- \leq R^- \quad \forall t \quad (17)$$

$$0 \leq x_t^+, x_t^- \quad \forall t \quad (18)$$

The objective (12) encodes minimizing the total cost of energy, i.e., the difference of the price of energy purchased and sold. Constraint (13) ensures that the energy balance in the system is maintained. Equality (14) defines the state of charge in the battery based on the charge and discharge rates. Finally, box constraints (15-18) define the range of the variables.

V. ADAPTIVE FORECAST AGGREGATION

Now, we turn our attention to aggregating predictions of various time-series models, in order to increase the performance and achieve adaptive behavior. This approach has a number of benefits, e.g., it is sometimes hard to select just one model in advance, as several models may have similar fit values, while they behave differently in various parts of the state space. Then, we do not have to select just one of them, but we can select online in each time step which model to use. As evaluating a model is typically computationally cheap, this does not come with heavy computational burden. Moreover, later we may purposefully estimate models for different situations, to specifically have them focused on different parts of the state space. For example, we may train PV production models for different weather conditions or may estimate energy consumption for different typical scenarios (e.g., rush hour, big event, etc.). Then, we do not have to detect which situation we are in, as the aggregation mechanism automatically selects the best model to use based on their recent performances.

Before describing our aggregation mechanism, which takes state information into account, first, we overview the standard framework of sequential prediction with expert advice [4].

A. Prediction with Expert Advice

In the standard framework, we sequentially face the problem of predicting a process, called the *environment*, based on the predictions and past performances (measured by their amount of mispredictions) of a pool of *experts*. We refer to the entity making the aggregated predictions as the *learner*.

We denote the state of the process, the prediction of expert i and the aggregated prediction of the learner, at time t , by x_t , $\hat{x}_{i,t}$ and \hat{p}_t , respectively. The *loss* of expert i and the learner, at time t , is defined as their cumulative cost up to t , that is

$$L_{i,t} \triangleq \sum_{k=1}^t \ell(\hat{x}_{i,k}, x_k), \quad (19)$$

$$\hat{L}_t \triangleq \sum_{k=1}^t \ell(\hat{p}_k, x_k), \quad (20)$$

where $\ell(\cdot)$ is a cost function. A typical choice is, for example, $\ell(x, y) = \|x - y\|_p^q$ for some p -norm and $q > 0$.

The *regret* of the learner at time t is defined as

$$\hat{R}_t \triangleq \hat{L}_t - \inf_i L_{i,t}, \quad (21)$$

which therefore measures the (relative) performance of the learner with respect to the best (in terms of loss) expert.

The learner aims at minimizing his regret while repeatedly addressing the problem. The general protocol of prediction with expert advice is summarized in Table I.

One of the desirable properties an aggregation rule can have is vanishing per-round regret, called *Hannan consistency*, i.e.,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \hat{R}_t \leq 0, \quad (22)$$

where the convergence is uniform over all possible outcome sequences and expert advice sequences (nonstochastic setting), or the convergence is almost sure (stochastic setting).

The archetypical aggregation rule to compute the learner's prediction for time t is to form a convex combination. This approach is called the *weighted average forecaster*,

$$\hat{p}_t \triangleq \frac{\sum_{i=1}^n w_{i,t-1} \hat{x}_{i,t}}{\sum_{i=1}^n w_{i,t-1}}, \quad (23)$$

assuming n experts, where the weights are often defined as $w_{i,t-1} = \nabla \Phi(R_{t-1})_i$, for some *potential* function

$$\Phi(R_{t-1}) \triangleq \psi \left(\sum_{i=1}^n \phi(R_{i,t-1}) \right), \quad (24)$$

where ϕ, ψ are nonnegative, twice differentiable functions as well as ϕ is increasing and ψ is strictly increasing and concave; and where R_{t-1} is the *regret vector* for time $t-1$, that is

$$R_{t-1} \triangleq \left(\hat{L}_{t-1} - L_{1,t-1}, \dots, \hat{L}_{t-1} - L_{n,t-1} \right)^T. \quad (25)$$

A celebrated potential is the *exponential potential*, i.e.,

$$\Phi_\eta(r) \triangleq \frac{1}{\eta} \left(\sum_{i=1}^n \exp(\eta r_i) \right), \quad (26)$$

where $\eta > 0$ is a user-chosen design-parameter and r is a regret vector. This approach leads to the *exponentially weighted average forecaster*, which can be simplified to

$$\hat{p}_t = \frac{\sum_{i=1}^n \exp(-\eta L_{i,t-1}) \hat{x}_{i,t}}{\sum_{i=1}^n \exp(-\eta L_{i,t-1})}. \quad (27)$$

It is known that if we apply a time-dependent η with

$$\eta_t \triangleq \sqrt{8 \ln(n)/t}, \quad (28)$$

then the exponentially weighted average forecaster is Hannan consistent [4], more precisely its regret is bounded by

$$\hat{R}_t \leq 2\sqrt{\frac{t}{2} \ln(n)} + \sqrt{\frac{\ln(n)}{8}}, \quad (29)$$

if $\ell(\cdot, \cdot) \in [0, 1]$ and it is convex in its first argument.

Table I. STANDARD SCHEME OF PREDICTION WITH EXPERT ADVICE

PROTOCOL: PREDICTION WITH EXPERT ADVICE	
1.	Losses of the learner and the experts are set to zero, $L_0 := 0$ and $\{L_{i,0} := 0\}_i$;
2.	For each round $t = 1, 2, \dots$ do
3.	Experts announce their forecasts $\{\hat{x}_{i,t}\}_i$ for time t ;
4.	Learner announces his forecast \hat{p}_t for time t ;
5.	Environment announces the "true" outcome, x_t ;
6.	Learner and the experts incur costs, i.e., $\ell(\hat{p}_t, x_t)$ and $\{\ell(\hat{x}_{i,t}, x_t)\}_i$, respectively;
7.	Losses get updated, $L_t := L_{t-1} + \ell(\hat{p}_t, x_t)$, and similarly, $\{L_{i,t} := L_{i,t-1} + \ell(\hat{x}_{i,t}, x_t)\}_i$;
8.	Repeat

An improved bound can be achieved, if some a priori information is available on the loss of the best expert, by using

$$\eta_t \triangleq \ln \left(1 + \sqrt{2 \ln(n)/L_t^*} \right), \quad (30)$$

where $L_t^* > 0$ is the loss of the best expert at time t . Since L_t^* is not available in advance, only after the t th round, it can be estimated by the loss of the currently best expert [4].

B. State-Dependent Aggregation

This section presents our state-dependent learner, which tries to take into account some side information to decide which experts (time-series models in our case) may have better performances in the current situation. We note that the use of side information in the framework of prediction with expert advice was previously addressed, e.g., in [4, 13, 16]. However, these approaches differ from ours, as they assume white box models, i.e., that they know how the experts make their predictions, and use a state-independent weighting. The experts are typically assumed to have linearly parametrized models, either simple linear regressions or some elements of a finite dimensional linear space of nonlinear functions (e.g., Reproducing Kernel Hilbert Space) and the learner aims at iteratively approximating the best weights of this regression. On the other hand, in our approach we treat the experts as black boxes (we do not assume the knowledge about how they make their predictions) and compute a weighting that depends on the available side information through state-dependent losses.

We also start with a pool of experts (time-series models):

$$\mathcal{E} \triangleq \{f_1, \dots, f_n\}, \quad (31)$$

which may take some state information into account (e.g., the past of the process and the exogenous inputs in our case). Let s_t denote the side information available at time t , then the prediction of the expert i at time t can be written as

$$\hat{x}_{i,t} \triangleq f_i(s_t). \quad (32)$$

We assume that we have a similarity kernel, $k(\cdot, \cdot) \in [0, 1]$, available which measures how similar two states are. If $k(s, r)$ is close to one, it shows that states s and r are very similar, while a small $k(s, r)$ indicates dissimilar states.

Designing a similarity kernel may need domain specific knowledge. In our case, $s_t = z_t$, see equation (2), and during the experiments, presented in Section VI, we simply used

$$k(z_t, z_k) \triangleq 1 - (\ell(x_{t-1}, u_{t-1}) - \ell(x_{k-1}, u_{k-1}))^2, \quad (33)$$

where $0 \leq x_t, x_k, u_t, u_k \leq 1$, for all t, k . Therefore, the side information indicates how far the system is operating from its typical behavior, and two situations are similar if they have similar distances from the typical behavior.

Assuming we have a similarity kernel k , we can define the discounted state-dependent loss of the expert i at time t by

$$K_{i,t}(s) \triangleq \sum_{k=1}^t \gamma^{t-k} k(s, s_k) \ell(\hat{x}_{i,k}, x_k) \quad (34)$$

$$= \sum_{k=1}^t \gamma^{t-k} k(s, s_k) \ell(f_i(s_k), x_k), \quad (35)$$

where $\gamma \in [0, 1]$ is a discount factor, to decrease the relevance of past losses, and we take into account the similarity of the past states to situation z when calculating the loss for z .

Discounting helps to focus on the recent events, as it renders less weights to costs incurred long time ago. This makes sense even if we do not use state-dependent aggregation. Thus, during our numerical experiments, we also used discounted losses when we calculated the predictions of the classical exponentially weighted average forecaster, as it increased its performance and made the comparison fairer.

After the state dependent losses were calculated, we may even use the standard aggregation approaches, e.g., weighted average, with $\{L_{i,t-1}\}$ having replaced with $\{K_{i,t-1}(s_t)\}$. The exponentially weighted average forecaster, e.g., becomes

$$\hat{p}_t(s_t) = \frac{\sum_{i=1}^n \exp(-\eta K_{i,t-1}(s_t)) \hat{x}_{i,t}}{\sum_{i=1}^n \exp(-\eta K_{i,t-1}(s_t))}. \quad (36)$$

In the section about our experiments below, we referred the above modification of the EWAF (exponentially weighted average forecaster) as SDAF (state-dependent average forecaster).

Table II summarizes the protocol of our state-dependent aggregation. It can be seen as a generalization of the classical approach, as having a kernel which is constant 1 and discount factor $\gamma = 1$ returns us to the standard definition of loss.

VI. EXPERIMENTAL RESULTS

In this section we present numerical experiments about fitting stochastic models to time-series data coming from a renewable energy system. Using side information during prediction will also be studied, both in the cases of individual models and the aggregated learner, i.e., state-dependent aggregation. Regarding the applied MPC controller, some preliminary results, without aggregation, can be found in [7].

We estimated 12 time-series models, which were briefly discussed in Section III-C, and tested their abilities to predict PV production and energy consumption data. For the energy production case the measured quantities were the PV current (A) and PV voltage (V), from which the PV power was calculated (voltage \times current). The PV production data were preprocessed by removing outliers, they were normalized and

Table II. SCHEME OF PREDICTION WITH SIDE INFORMATION

PROTOCOL: STATE-DEPENDENT AGGREGATION	
1.	For each round $t = 1, 2, \dots$ do
2.	Environment announces side information s_t for t ;
3.	Similar losses are calculated $\{K_{i,t-1}(s_t)\}_i$;
4.	Experts announce their forecasts $\{\hat{x}_{i,t} = f_i(s_t)\}_i$;
5.	Learner announces forecast, $\hat{p}_t(s_t)$;
6.	Environment announces the “true” outcome, x_t ;
7.	Learner and the experts incur costs, i.e., $\ell(\hat{p}_t, x_t)$ and $\{\ell(\hat{x}_{i,t}, x_t)\}_i$, respectively;
8.	Repeat

Table III. PERFORMANCE EVALUATION

Time-Series Model		Production: Loss		Consumption: Loss	
Name	Side Info	Estimation	Validation	Estimation	Validation
FIR	+	16.92	13.94	27.15	30.18
AR	-	6.27	7.78	13.67	21.38
ARX	+	5.51	7.07	10.16	25.76
ARMA	-	5.68	7.81	16.53	22.09
BJ	+	5.15	7.06	9.17	18.31
STATE	+	5.21	6.96	9.28	26.28
HW	+	10.68	14.46	23.28	30.15
WAVE	+	4.21	9.29	6.95	20.07
MLP	-	5.24	9.83	13.64	25.04
MLPX	+	4.02	8.58	9.61	19.84
SVR	-	6.45	7.38	11.23	20.15
SVRX	+	6.37	7.18	5.37	16.43
Aggregation Method		Loss (Regret)		Loss (Regret)	
EWAF	-	4.93 (0.91)	6.89 (-0.07)	8.08 (2.71)	14.91 (-1.52)
SDAF	+	4.32 (0.30)	6.75 (-0.21)	5.43 (0.06)	14.59 (-1.84)

averaged in windows of one hour length, as this is the resolution of the applied MPC controller, namely, the time the forecasts and the control policy are recalculated. Averaging helped to achieve a better signal-to-noise ratio, too. The energy consumption data were also preprocessed similarly to the production data: they were cleaned from corrupted measurements, normalized and averaged in one hour wide windows.

Linear and nonlinear models were fitted (the first six rows of Table III show the linear ones) and the effects of providing side information to the time-series models (see Section III-B) were also studied. In this case, side information means the typical (periodic average) value for that hour (from historical data), or the clear-sky prediction for that hour. The MLP and SVR based models were tested with (MLPX and SVRX) and without (MLP and SVR) side information as exogenous input.

The orders of the models were either selected automatically by the estimation method or several variants were tested and the best one was selected. Due to space limitations we cannot discuss selecting the orders of the models in detail, but for most of the applied time-series models the orders with respect to the autoregressive parts (i.e., the past of the process) were between 4 and 7, while with respect to the exogenous inputs (i.e., the side information) they were between 2 and 5.

Table III shows the total losses, as defined by equation (19) with $\ell(x, y) = \|x - y\|_2^2$, of the time-series models for both the energy production and consumption processes. The losses were calculated on a normalized data and the samples contained 1000 measurements in all cases. The loss values are given for both the estimation (learning) data and the validation (test) data for each model. The results support that providing side information as exogenous inputs to the models increases their performance, as it can be observed, for example, by the AR/ARX, MLP/MLPX and SVR/SVRX rows, where “X” indicates that the models have exogenous components.

The results also indicate that such models can effectively predict energy production and consumption processes on a short horizon. For the PV production case, linear, BJ (Box-Jenkins) and STATE (state space), and nonlinear autoregres-

sive, WAVE (wavelet) and MLP (multilayer perceptron) based models provided the smallest loss values. For the consumption case, BJ and SVR (support vector regression) based models were the best with respect to the total loss criterion.

The last two rows of Table III illustrate the performance of two aggregation methods: the classical exponentially weighted average forecaster (EWA) and our state-dependent (exponentially weighted) average forecaster (SDAF), presented in Section V-B. Parameter η was time-dependent and set according to formula (30), but for SDAF the losses were state-dependent, i.e., kernelized. We used discounted costs with $\gamma = 0.95$ in both cases (EWA, SDAF) to set η , in order to put more emphasis on the recent events when combining the predictions. Table III on the other hand shows total (undiscounted) costs, as we are interested in the uniform performance of the learners.

The resulting regrets show that combining predictions with online learning leads to good performance. The regret was often negative indicating that the aggregated predictor outperformed the best time-series model. It can also be observed that even if all the models were trained on the same (estimation) data, SDAF could still outperform EWA. A difference increase is expected, in favor of SDAF, when using models specially fitted to various scenarios (for example, different weather conditions or consumption schemes).

Note that for the aggregation methods there is no learning or estimation phase (as for the time-series models), thus there is no reason to expect smaller regret on the estimation data. The time-series models are expected to achieve smaller losses on the estimation (learning) data, but the regret is measured w.r.t. the performance of the best time-series model, thus it might even be harder to achieve smaller regret on the estimation data. This helps explaining why the learners achieved smaller regrets on the validation (test) data during our experiments.

Figures 2 and 3 illustrate a situation when SDAF could show up better tracking capabilities than the classical EWA. The PV production was much lower than expected, due to the weather conditions, and SDAF realized faster which models are better in such cases. Figure 2 shows the real (normalized) PV production of the environment, the prediction of the best time-series model, EWA and SDAF. Figure 3 shows the discounted regret during the same period, with discount factor $\gamma = 0.95$. As the environment does not have any regret, Figure 3 plots the average regret of the 12 time-series models instead.

VII. CONCLUSION

The paper investigated the problem of generating forecasts for energy production and consumption processes which are used for model predictive control. Several time-series models were fitted to historical data and it was demonstrated that providing side information increases their performance. An adaptive state-dependent aggregation approach was proposed which uses situation dependent performance evaluation of the experts. Several experiments were presented which demonstrate the effectiveness of the applied stochastic models and the suggested state-dependent aggregation approach.

Our motivation was the E+grid system, which is an energy positive public lighting microgrid. The physical E+grid system

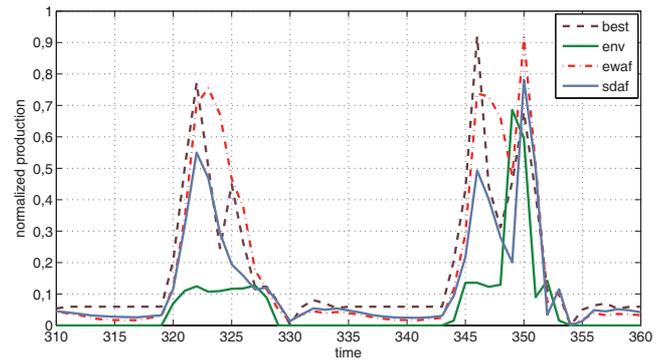


Figure 2. Predicting energy production with aggregated forecasters.

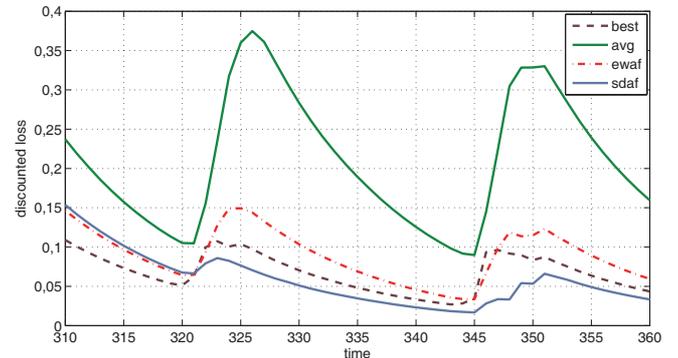


Figure 3. Discounted loss during the same period as in Figure 2.

is currently under construction in cooperation between GE Hungary Ltd, the Budapest University of Technology and Economics, and two institutes of the Hungarian Academy of Sciences: the Institute for Technical Physics and Materials Science and the Institute for Computer Science and Control.

There are several future research directions, including investigating the consistency of state-dependent aggregated forecasts as well as analyzing and designing similarity kernels.

ACKNOWLEDGMENTS

This research project has been partially supported by the grants of the National Development Agency (NFÜ), Hungary, under contract numbers KTIA KMR 12-1-2012-0031 and NFÜ ED-13-2-2013-0002 and by the Hungarian Scientific Research Fund (OTKA) under contract number 111797. PV production data were provided by the Department of Energy Engineering of the Budapest University of Technology and Economics. The SVR related experiments were performed with the help of the LibSVM library [5]. B. Cs. Csáji acknowledges the support of the János Bolyai Research Fellowship of the Hungarian Academy of Sciences, no. BO/00683/12/6.

REFERENCES

- [1] Søren Asmussen. *Applied Probability and Queues*. Springer, 2nd edition, 2003.
- [2] George E.P. Box, Gwilym M. Jenkins, and Gregory C. Reinsel. *Time Series Analysis*. Prentice-Hall, 1994.

- [3] João Paulo da Silva Catalão, editor. *Electric Power Systems: Advanced Forecasting Techniques and Optimal Generation Scheduling*. Taylor & Francis, 2012.
- [4] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [5] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:1–27, 2011.
- [6] Balázs Cs. Csáji, Borbála Háty, András Kovács, Gianfranco Pedone, Tibor Révész, and József Váncza. Monitoring and controlling energy-positive public lighting: The E+grid system. *ERCIM News*, 98:36–37, 2014.
- [7] Balázs Cs. Csáji, András Kovács, and József Váncza. Prediction and robust control of energy flow in renewable energy systems. In *19th World Congress of the International Federation of Automatic Control, Cape Town, South Africa*, pages 3663–3669, 2014.
- [8] Marie Devaine, Pierre Gaillard, Yannig Goude, and Gilles Stoltz. Forecasting electricity consumption by aggregating specialized experts. *Mach. Learn.*, 90:231–260, 2013.
- [9] B. Efron and R.J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall, New York, 1993.
- [10] Jan Kleissl. *Solar Energy Forecasting and Resource Assessment*. Elsevier Science, 2013.
- [11] Lennart Ljung. *System Identification: Theory for the User*. Prentice-Hall, 2nd edition, 1999.
- [12] Daryl R. Myers. *Solar Radiation: Practical Modeling for Renewable Energy Applications*. CRC Press, 2013.
- [13] György Ottucsák and László Györfi. Sequential prediction of binary sequence with side information only. In *Information Theory, 2007. ISIT 2007. IEEE International Symposium on*, pages 2351–2355, June 2007.
- [14] Christophe Paoli, Cyril Voyant, Marc Muselli, and Marie-Laure Nivet. Forecasting of preprocessed daily solar radiation time series using neural networks. *Solar Energy*, 84(12):2146–2160, 2010.
- [15] Torsten Söderström and Petre Stoica. *System Identification*. Prentice Hall, 1989.
- [16] Vladimir Vovk and Fedor Zhdanov. Prediction with expert advice for the Brier game. *Journal of Machine Learning Research*, 10:2445–2471, December 2009.
- [17] Louis A Wehenkel. *Automatic learning techniques in power systems*. Number 429 in Power Electronics and Power Systems. Springer, 1998.
- [18] Dazhi Yang, Chaojun Gu, Zibo Dong, Panida Jirutitijaroen, Nan Chen, and Wilfred M Walsh. Solar irradiance forecasting using spatial-temporal covariance structures and time-forward kriging. *Renewable Energy*, 60:235–245, 2013.