

Application-Driven Learning: A Closed-Loop Prediction and Optimization Approach Applied to Dynamic Reserves and Demand Forecasting

Joaquim Dias Garcia

LAMPS, DEE, PUC-Rio & PSR, Rio de Janeiro, Brazil, joaquim@psr-inc.com

Alexandre Street

LAMPS, DEE, PUC-Rio, Rio de Janeiro, Brazil, street@ele.puc-rio.br

Tito Homem-de-Mello

School of Business, UAI, Santiago, Chile, tito.hmello@uai.cl

Francisco D. Muñoz

Facultad de Ingeniería y Ciencias, UAI, Santiago, Chile, fdmunoz@uai.cl

Forecasting and decision-making are generally modeled as two sequential steps with no feedback, following an open-loop approach. In this paper, we present application-driven learning, a new closed-loop framework in which the processes of forecasting and decision-making are merged and co-optimized through a bilevel optimization problem. We present our methodology in a general format and prove that the solution converges to the best estimator in terms of the expected cost of the selected application. Then, we propose two solution methods: an exact method based on the KKT conditions of the second-level problem and a scalable heuristic approach suitable for decomposition methods. The proposed methodology is applied to the relevant problem of defining dynamic reserve requirements and conditional load forecasts, offering an alternative approach to current *ad hoc* procedures implemented in industry practices. We benchmark our methodology with the standard sequential least-squares forecast and dispatch planning process. We apply the proposed methodology to an illustrative system and to a wide range of instances, from dozens of buses to large-scale realistic systems with thousands of buses. Our results show that the proposed methodology is scalable and yields consistently better performance than the standard open-loop approach.

Key words: Application-driven learning, closed-loop prediction and optimization, bilevel optimization, dynamic reserves, forecast, power-systems operation

1. Introduction

The most common approach to make decisions under uncertainty involves three steps. In the first step, one develops a forecast for all uncertainties that affect the decision-making problem based on all information available. In the second step, an action based on the forecast is selected. Finally, in the third step, one implements corrective actions after uncertainties are realized. This three-step procedure constitutes an open-loop forecast-decision process in which the outcomes of the decisions are not considered in the forecasting framework.

In the electricity sector, it is common for system operators to use an open-loop forecast-decision approach. First, loads are forecast based on standard statistical techniques, such as least squares (LS), and reserve requirements are defined by simple rules, based on quantiles, extreme values and standard deviation of forecast errors according to specified reliability standards (Ela et al. 2011). Then, a decision is made to allocate generation resources following an energy and reserve scheduling program (Chen et al. 2013, De Vos et al. 2019). In real-time, reserves are deployed to ensure that power is balanced at every node, compensating for forecast errors.

From the academic perspective, it has been demonstrated that stochastic programming models yield better results than deterministic ones when making decisions under uncertainty because the former takes distributions into consideration. These models provide better results in terms of cost, reliability, and market efficiency compared to deterministic approaches (Wang and Hobbs 2014). Nevertheless, in practical applications, two issues arise: proper modeling of distributions is challenging and tractability imposes small sample sizes for techniques like sample average approximations (SAA). A consequence of this tractability issue is that SAA solutions become sample dependent (Papavasiliou et al. 2014, Papavasiliou and Oren 2013), thereby compromising market transparency and preventing stakeholders' acceptance (Wang and Hobbs 2015). Therefore, most system operators worldwide still rely on deterministic short-term scheduling (economic dispatch or unit commitment) models with exogenous forecasts for loads and reserve requirements (Chen et al. 2013, PJM 2018). Within this context, one alternative to improve the performance of deterministic

scheduling tools is to forecast load and reserve requirements with the goal of minimizing energy and reserve scheduling costs.

There is empirical evidence that system operators rely on *ad hoc* or out-of-market actions—and not just on reserves—to deal with uncertainty in operations. According to the 2019 Annual Report on Market Issues and Performance of the California ISO (CAISO 2020), “...operators regularly take significant out-of-market actions to address the net load uncertainty over a longer multi-hour time horizon (e.g., 2 or 3 hours). These actions include routine upward biasing of the hour-ahead and 15-minute load forecast, and exceptional dispatches to commit and begin to ramp up additional gas-fired units in advance of the evening ramping hours.” Additionally, reserve requirements are, in practice, empirically defined according to further *ad hoc* off-line rules based on off-line analysis (Ela et al. 2011, PJM 2018). These *ad hoc* procedures lack technical formalism and transparency to minimize operating and reliability costs, preventing agents from internalizing it based on scientifically grounded and described methodologies. Consequently, this challenging real-world application requires further research observing the practical issues that need to be addressed to improve the current state-of-the-art of industry practices.

For years, decision-making and forecasting have been treated as two completely separate processes (Bertsimas and Kallus 2019). Many communities, such as Statistics and Operations Research, have studied these problems and developed multiple tools combining probability and optimization. The machine learning community, which combines many ideas from optimization and probability, has also been tackling such tasks and has proposed methods to treat them jointly (Bengio 1997).

Classical forecasting methods do not take the underlying application of the forecast into account. Consequently, hypotheses such as prediction error symmetry in least squares (LS) might not be the best fit for problems with asymmetric outcomes. By acknowledging the asymmetry in particular problems, researchers have attempted to capture it empirically; however, such an approach does not take the application into account directly. Some existing methods do capture asymmetry, such as Quantile Regression (Rockafellar et al. 2008). The interest in exploring asymmetric loss

functions is not new. For instance, Zellner (1986a) and Zellner (1986b) acknowledge that biased estimators can perform even better than those that make accurate predictions of statistical properties of the stochastic variables. The author exemplifies that an overestimation is not as bad as an underestimation for the case of dam construction and attributes a second example about the asymmetry on real estate assessment to Varian (1975).

Within this context, two possible avenues of research are opened to achieve better results: i) a focus on improving the decision-making model (prescriptive framework), which assumes we can change it to consider embedded co-optimized forecasts (Bertsimas and Kallus 2019); or ii) a focus on improving the forecasting model (predictive framework), which assumes we can not change the decision-making process (in our application, defined by system operators' dispatch models), but we can change the forecasts to incorporate, in a closed-loop manner, a given application cost function (Bengio 1997, Elmachoub and Grigas 2021, Muñoz et al. 2022). Therefore, in this paper, we focus on the latter avenue. In Section 2, we provide a literature review on this subject.

1.1. Objective and contribution

The objective of this paper is to present a new closed-loop *application-driven learning* framework to be used in point forecast applications. In the proposed method, the application is characterized through an optimization model, which is then used as part of the estimation problem. In this new framework, both *ex-ante* (planning) and *ex-post* (implementation/assessment) cost-minimization structures of the decision-maker, i.e., the *application schema*, are considered in the prediction process. Therefore, our framework replaces the traditional statistical error minimization objective with a cost-minimization structure of a specific application.

To achieve the objectives described before, we derive the following general technical contributions:

- A new and flexible application-driven learning framework based on a bilevel optimization model. To the best of the authors' knowledge, for the first time in the literature, the relevant case of applications based on linear programming models affected by right-hand-side uncertainty is

addressed with specialized algorithms. Two solutions approaches are presented. The first approach is an exact method based on the KKT conditions of the second-level problem. The second is a scalable heuristic approach suitable for decomposition methods and parallel computing. Although not limited to linear bilevel programs, we show how to design efficient methods tailored for off-the-shelf linear optimization solvers. Additionally, our scalable heuristic method ensures optimal second-level solutions. This is a salient feature of our method in contrast with other methods that rely on surrogates of the second level (Elmachtoub and Grigas 2021) or solve approximations of the KKT conditions with non-linear solvers (Muñoz et al. 2022). In this context, the proposed framework is general and suitable for a wide range of applications relying on the standard structure of the forecast-decision process.

- We provide new asymptotic convergence proofs for both the objective function value and estimated parameters of the proposed application-driven learning method. The convergence proof is completely novel and highlights that the method is asymptotically the best that can be done for a specific application (described by planning and implementation processes) given a forecast functional form (see Corollary 1). It is important to emphasize that Lemma 1 extends the state-of-the-art SAA results from Shapiro et al. (2014) to stationary-ergodic time series, which is also a relevant contribution to the subject. Under the hypothesis of our methodology and based on the aforementioned convergence results, we show that the solution of our method converges to the best estimator in terms of the expected cost of the selected application.

Notwithstanding, our paper also provides relevant contributions to the applied field of power systems operation. In particular, the application of this paper focuses on the problem of demand and reserve requirements forecasting for power system operators. This is a critical problem of the power systems industry, which has not been addressed by previous works on the subject and is currently being tackled by system operators through *ad hoc* out-of-market actions to address the net load uncertainty CAISO (2020). Therefore, as a contribution to this applied field, we propose a new methodology to forecast the load and define the reserve requirements for the power-system

dispatch planning application. The method can be used to either jointly optimize the load forecast and reserve requirements or to optimize only the reserve requirements given an exogenous forecast for the load. In both cases, the optimal solution defines the optimal policy to dynamically allocate reserve requirements so that the expected cost is minimized in the long run. In both cases, agents are provided with a scientifically grounded and comprehensively described methodology that can be used to reduce the number of *ad hoc* procedures currently implemented in practice.

In our specific application, the bilevel formulation can be summarized as follows. The first level accounts for both the predictive model specification (parameters selection) and the cost evaluation metric based on the actual operation of the system for many data points. It is relevant to mention that the methodology is flexible to internalize and address relevant practical issues, such as reserve requirement constraints imposed by regulatory rules, reliability standards (The European Commission 2017, Ela et al. 2011), and risk-aversion metrics (Shapiro et al. 2014). In the second level, the *ex-ante* energy and reserve scheduling process of the system operator is accounted for based on 1) a conditional demand forecast and 2) the definition of adjusted (dynamic) reserve requirements, both defined in the first level as a function of previous data for each point of the training set. Thus, in our bilevel model, we have multiple parallel lower-level problems, each of which represents the one-step-ahead deterministic two-stage scheduling process performed by the system operator for each point of the training dataset. In this context, the second level ensures closed-loop feedback characterizing joint scheduling decisions of energy and reserve allocations without perfect information on the target period data. Such a closed-loop formulation applied to the definition of *conditional load forecasts* and *dynamic reserve requirements* is a salient and original contribution to the subject of power systems operation.

Finally, to empirically corroborate the relevance of our contributions and to demonstrate the applicability, performance, and scalability of our methodology, we benchmark the proposed method with the traditional sequential least squares forecast and energy and reserve scheduling approach. To do that, we analyze the proposed methodology in several case studies using multiple test systems.

First, we present studies with an illustrative single-bus system to explore multiple properties of the methods. Second, we apply the method to various instances based on the IEEE 24-, 118-, and 300-bus test systems to show that the results are of high quality. Third, we consider 10 data sets ranging from 600 to 6000 buses artificially created to combine instances of the 300 bus system to analyze the scalability of the method. Fourth, we consider realistic, very large-scale systems, ranging from 6,468 to 13,659 buses, with infrastructure and conditions that are very close to those of a systems operator performing hour-ahead planning to demonstrate that the methodology can be seriously considered for deployment by operators. Results show that the two proposed application-driven learning approaches (demand and reserves, and only reserves) yield consistently better performance on out-of-sample tests than the benchmark where forecasts and decisions are sequentially carried out. For large and very large systems, where the exact method fails to find solutions within reasonable computational times, the heuristic method exhibits high-quality performance compared to the benchmark for all test systems.

2. Literature Review

We review the literature on 1) forecast models jointly optimized for a given application, hereinafter referred to as application-driven forecast, and 2) uncertainty forecasting and reserve sizing.

2.1. Application-driven forecast models

The ingenious idea of integrating the process of forecasting and optimizing a downstream problem was first proposed in the seminal paper by Bengio (1997). More than twenty years ago, the author emphasized the importance of estimating parameters with the correct goals in mind. In that work, a Neural Network (NN) is trained to forecast stocks with an objective function that describes the portfolio revenue given an allocation based on stocks forecast. Still, in the finance sector, Ghosn and Bengio (1997) proposed new neural network structures and presented extremely promising benchmarks. An attempt to lower the burden of the method was proposed by Garcia and Gençay (2000); the idea is to train multiple prediction models with standard regressions, but choose the

best one in the out-of-sample analysis considering the proper application-driven objective function. The work by Kao et al. (2009) presents another intermediary methodology. The model for estimating forecasts includes both the application objective function and the fitness measure similar to maximum likelihood estimation (MLE). A bi-objective problem is solved with scalarization; the authors look for a good balance between MLE and application value.

Following the key idea of Bengio (1997) closely, the work by Donti et al. (2017) presents a generic algorithm to deal with parameter optimization of forecasting models embedded in stochastic programming problems, that is, parameters estimated considering the loss function of the actual problem. The algorithm is based on the stochastic gradient descent (SGD) method and employs tools for automatically differentiating strongly convex quadratic optimization problems. The method is applied to small prototypical quadratic programming problems and the local solutions obtained are shown to be promising. The idea of extending to combinatorial problems via convexification and regularization (to make it a QP) schemes is presented in Wilder et al. (2019).

More recently, the work Smart “Predict and Optimize” (SPO) (Elmachtoub and Grigas 2021), recognizes the importance of closed-loop estimation. The authors develop an algorithm for the linear programming case that is based on relaxation and convexification of the nonlinear loss function before applying a tailored SGD approach, instead of looking for local solutions with nonlinear methods. To develop the algorithm, the authors limit themselves to linear dependency on features and restrict uncertainty to the objective function. This contrasts with our work that preserves the original planning and implementation functions (detailed in Section 3) and focuses on non-linear methods to obtain global or local solutions. Theoretical improvements were made to the SPO method in El Balghiti et al. (2019), which presents new bounds for the generalization capability of the method. Also following SPO, Mandi et al. (2020) presents an adaptation of SPO for the case of combinatorial problems. SPO is similar to the method proposed by Ryzhov and Powell (2012) to estimate uncertain objective coefficients without considering features in a different context.

While working on this paper, the authors became aware of the work by Muñoz et al. (2022) (developed simultaneously with ours, by a completely different team) that guards many conceptual similarities to the general version of our proposed model. The work of Muñoz et al. (2022)

also focuses on the idea of finding the best forecast for a given application (or context) through a bilevel framework. The framework proposed in Muñoz et al. (2022) is applied to estimate a parameter of the inverse demand curve of a Cournot strategic producer bidding in forward markets. The scalability of their model relies on a nonlinear relaxation of the right-hand side of the complementarity constraints. In our method, we adopt a different approach to overcome the issue of suboptimal lower-level solutions and to tackle very large-scale problems, as will be explained in Section 8.4. Moreover, we prove convergence of our method, whereas Muñoz et al. (2022) only analyze the convergence in a particular example and point to a preprint of our work when discussing the requirement of lower level uniqueness in their bilevel program.

den Boer and Sierag (2021) describes a framework named Decision Based Model Selection, which fits in the general grand scheme of properly combining forecasting and optimizing decisions. The authors propose a two-step procedure. First, they estimate a forecasting model with available data and generate more samples with this model; then, with a new (artificial) data set in hand, the model (forecasting + optimization model + algorithm) with the best decisions is chosen from a discrete list. The framework is general enough to allow forecasts that depend on decisions. Related work in power systems is Bezerra et al. (2016) that proposed to generate multiple forecasting models and then select the best one according to a min-max regret based on the application cost.

Sen and Deng (2018) and Deng et al. (2018) describe the so-called Learning Enabled Optimization (LEO). LEO is a framework to combine Statistical Learning (SL), Machine Learning (ML) and (stochastic) Optimization. The idea is to compare a set of predefined SL/ML models with the cost value of the actual application in mind. The main difference here is that the ML/SL are still estimated based on classical methods. More recently, LEO has been extended to Coupled LEO to accommodate decision-dependent uncertainty Liu et al. (2022) solved by derivative-free methods.

Also related is the field of Optimal Learning (Powell and Frazier 2008, Powell and Ryzhov 2012). Optimal Learning is very broad, encompasses many techniques and can be applied to an enormous range of problems, as detailed by the previous references. The idea of estimating unknowns through

experimentation and testing guesses in simulation models is a possible framework for choosing forecast models based on applications. Deeply related to Optimal Learning is Bayesian Optimization (Powell 2010, Frazier 2018), which is a set of methods used to optimize continuous functions, typically low dimensional (possibly noisy) functions that are expensive to evaluate, which might be the case of an application as a cost function.

Not surprisingly, some of the above works formulate the problems as bilevel optimization problems, which is mostly done implicitly and rarely explicitly (Muñoz et al. 2022). For more information on bilevel optimization, the reader is directed to Bard (2013). The work by Dempe (2018) lists hundreds of references related to bilevel optimization, including papers related to parameter optimization. Parameter optimization is frequently modeled as bilevel optimization and has been drawing the attention of many fields, such as ML, control, energy systems, and game theory. This can be thought of as a version of the closed-loop paradigm since these works target the best parameters for algorithms and applications. We refer to (Franceschi et al. 2018) for applications in hyper-parameter tuning. Under a broader ML umbrella, our methodology can be seen as an ML forecast approach where the *application schema* is embedded into the method through its explicit mathematical programming formulation.

2.2. Uncertainty forecasting and reserve sizing

The operation of power systems has been profoundly related to uncertainty handling. The electric load has been among the main challenges for forecasters in power systems for many years. Researchers around the world have proposed the most varied methodologies, ranging from standard linear regressions to Neural Networks (NN); these techniques, along with many others, are reviewed in Hong and Fan (2016) and Van der Meer et al. (2018). The forecast of variable renewable energy sources is probably one of the current big challenges power systems. Although many techniques are already available, forecasting renewable generation such as from wind farms has proven to be significantly harder than load (Van der Meer et al. 2018, Orwig et al. 2014). As shown in reviews (Sweeney et al. 2020, Van der Meer et al. 2018), wind and solar forecasting are divided into two

main trends: i) physical-based methods that rely on topographic models and Numerical Weather Predictions; ii) statistical methods, including Kalman filters, ARMA models, and NN.

Although forecast methods have consistently improved in the last years, the systems must be ready to withstand deviations from predicted values. The widely used approach is to allocate reserves in addition to the power scheduled for each generator to meet demand forecasts. The additional power is scheduled as reserve margins to give the system operator flexibility to handle real-time operations. Many methods have been proposed to account for the variations in loads, contingencies, and variable renewable energy (Holttinen et al. 2012). Furthermore, different reserve sizing rules are applied by different ISOs all over the world (Ela et al. 2011). These rules vary from deterministic *ad hoc* procedures to more statistical-oriented guidelines.

In Ela et al. (2011), examples of real-life *ad hoc* procedures to allocate reserves are presented, most of them relying on static approaches. Although time-varying reserves have been studied in the past, they have re-emerged as dynamic probabilistic reserves (De Vos et al. 2019). In the context of a more sophisticated dynamic probabilistic reserve approach, probabilistic forecasts are frequently used to account for forecast errors. These probabilistic reserves can be sized following a variety of methods with different complexity based on: forecast error standard deviations (Strbac et al. 2007, Holttinen et al. 2012), non-parametric estimation of the forecast error distribution (Bucksteeg et al. 2016), or even machine learning (De Vos et al. 2019). These are all considered stochastic methods and are simple alternatives to capture and incorporate fairly complex dynamics that are challenging for bottom-up approaches.

A prominent alternative to the use of reserves in power systems is the Stochastic Unit commitment. In such applications, many types of reserves can be defined endogenously, targeting cheaper operations on average. However, as described in the review on Unit Commitment by Zheng et al. (2014), there are at least three main barriers toward the wide acceptance of stochastic unit commitment: i) uncertainty modeling, ii) computational performance, and iii) market design. Uncertainty modeling is jointly tackled by statistical modeling of the uncertainty concerning scenario generation and forecasting and by a decision-making framework like risk-averse stochastic optimization,

robust optimization, and so on. Computational performance is the focus of many works like the Lagrangian decomposition (Aravena and Papavasiliou 2020), improved formulations (Knueven et al. 2020), progressive hedging (Gade et al. 2016). Notwithstanding the relevant recent advances in this area, the computational burden and the consequential instability of solutions under small sample sizes still preclude the acceptance of stochastic UC models by ISOs. Finally, the least studied challenge is market design. It requires experimenting and developing rules that are both feasible to be implemented and accepted by stakeholders (Kazempour et al. 2018, Wang and Hobbs 2015). Based on previously reported industry practices and since ISOs currently follow the alternate route and tackle the uncertainty of UC with reserves (Wang and Hobbs 2015), we will also follow this approach to propose a readily practical method.

3. Application-Driven Learning and Forecasting

In this section, we contrast the standard sequential framework, referred to as *open-loop*, and the joint prediction and optimization model, referred to as *closed-loop*. The presentation is in general form to facilitate the description of the solution algorithm, to set notation for the convergence results and to highlight that the method has applications beyond load forecasting and reserve sizing in power systems. We will specialize the bilevel optimization problem for closed-loop load forecasting and reserve sizing in Section 7.

We consider a dataset of historical data $\{y_t, x_t\}_{t \in \mathbb{T}}$, where $\mathbb{T} = \{1, \dots, T\}$. Here y_t are observations of a variable of interest that we want to forecast, while x_t are observations of external variables (covariates or features) that can be used to explain the former. Furthermore, the latter might include lags of y_t as in autoregressive time series models. Additionally, it is worth mentioning that both y_t and x_t can be vector-valued.

The classic forecast-decision approach works as follows. The practitioner *trains* a parametric *forecast* model seeking for the best vector of parameters, θ , such that a loss function, $l(\cdot, \cdot)$, between the conditional forecast for sample t , $\hat{y}_t(\theta, x_t)$, and the actual data, y_t , is minimized, i.e., solving $\min_{\theta} \frac{1}{T} \sum_t l(\hat{y}_t(\theta, x_t), y_t)$. This is frequently done by solving LS optimization problems and finding

$\theta^{LS} \in \arg \min_{\theta} \frac{1}{T} \sum_t \|\hat{y}_t(\theta, x_t) - y_t\|^2$. In the *planning* step, a decision is made by an optimized policy based on the previously obtained forecast, i.e., with $\hat{y}_t^{LS} = \hat{y}_t(\theta^{LS}, x_t)$. This results in a vector $z^*(\hat{y}_t^{LS})$, which in our application comprises the schedule of energy and reserves through generating units. Finally, the actual data y_t is observed, and the decision-maker must adapt to it, for instance, the system operator responds with a balancing re-dispatch, and a *cost*, $G_a(z^*(\hat{y}_t^{LS}), y_t)$, is measured. There is no feedback of the final cost into the forecasting and decision policy, hence, the name *open-loop*.

3.1. The proposed closed-loop application-driven framework

The core of the proposed predictive framework is to explore a feedback structure between the estimated predictive model and the application cost assessment. The general idea is depicted in Figure 1, which also stresses the difference from the open-loop model.

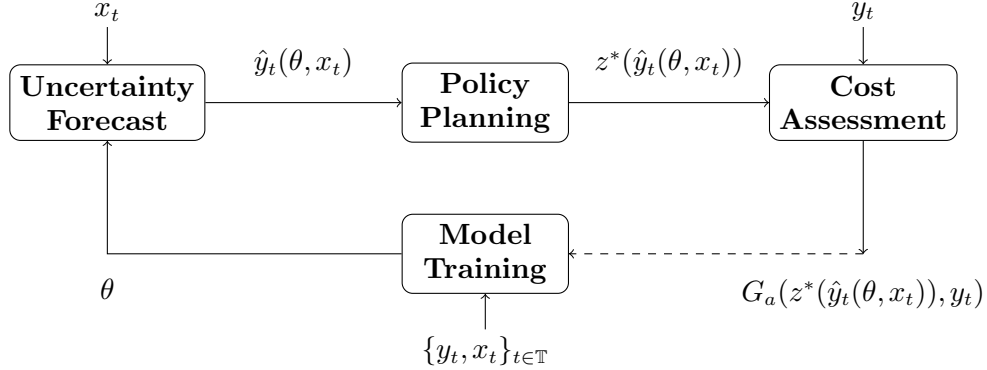


Figure 1 Learning models: considering the dashed line we have the *closed-loop* model, otherwise it represents the *open-loop* model.

The estimation method can be mathematically described through the following bilevel optimization problem (BOP):

$$\theta_T \in \arg \min_{\theta \in \Theta, \hat{y}_t, z_t^*} \frac{1}{T} \sum_{t \in \mathbb{T}} G_a(z_t^*, y_t) \quad (1)$$

$$s.t. \quad \hat{y}_t = \Psi(\theta, x_t) \quad \forall t \in \mathbb{T} \quad (2)$$

$$z_t^* \in \arg \min_{z \in Z} G_p(z, \hat{y}_t) \quad \forall t \in \mathbb{T}, \quad (3)$$

where, for $i \in \{a, p\}$,

$$G_i(z, y) = c_i^\top z + Q_i(z, y) \quad (4)$$

$$Q_i(z, y) = \min_u \{q_i^\top u \mid W_i u \geq b_i - H_i z + F_i y\} \quad (5)$$

Note that the functions in (4) and (5) resemble the formulation of two-stage stochastic programs, in the sense that given a decision z and an observation y , one determines the best corrective action u . In that context, c_i , q_i , W_i , b_i , H_i and F_i ($i \in \{a, p\}$) are parameters defined according to the problem of interest. Note also that the uncertainty y appears only on the right-hand side of the problems defining Q_a and Q_p ; this will be important for our convergence analysis and solutions methods.

In model (1)–(5), $\Psi(\theta, x_t)$ represents a *forecasting* model that depends on both the vector of parameters, θ , and the features vector, x_t , possibly including lags of y_t . The vector \hat{y}_t is the forecast generated for sample (or period) t (comprising load and reserve requirements) conditioned to the vector of features, x_t , as defined in (2). For each t , the forecast \hat{y}_t is used as input in a second-level problem and a decision planning policy, z_t^* , is obtained as a function of \hat{y}_t , i.e., $z_t^*(\hat{y}_t)$. This is done by optimizing the decision-maker *planning* cost function, $G_p(z, \hat{y}_t)$ in (3). Then, the optimized policy z_t^* is evaluated in the first level against the actual realization, y_t , for each t . The evaluation is made under the decision-maker's *assessment* (or *implementation*) cost function, $G_a(z_t^*(\hat{y}_t), y_t)$. Hence, the application is embedded into the estimation process in both the *ex-ante* planning policy and *ex-post* implementation objective (1)–(5). It is worth noticing that the proposed formulation

can be interpreted as an optimization over θ in a back-test, in which for a given θ , the assessment of the forecast performance is completely determined by $G_a(z_t^*(\hat{y}_t), y_t)$. Within this context, the upper level identifies the parameters with the best back-test performance. Furthermore, note that for a fixed θ , there is no coupling between two samples, thereby the model can be decomposed per t . This will be used in one of the proposed solution methods. Finally, there is a slight abuse of notation in (1) because the argmin only retrieves θ_T , a solution for θ with T samples, thus disregarding the rest of the first level decision vectors, \hat{y}_t and z_t^* . Note that θ_T is one among the multiple options in the set of all possible solutions S_T .

One key difference from previous works (Donti et al. 2017, Elmachoub and Grigas 2021, Muñoz et al. 2022) is that G_a and G_p can be different functions. This is extremely useful in the context of power systems operations where planning models might differ from real-time ones. Although model (1)–(5) is fairly general, we specialize to the case of linear programs and right-hand-side uncertainty, (4)–(5), because we will assume polyhedral structure for the set Z . This can be contrasted with previous works that considered strongly quadratic programs (Donti et al. 2017) and objective uncertainty (Elmachoub and Grigas 2021). As mentioned earlier, this specialization will be important for developing our asymptotic convergence results and our solution methods.

4. A Motivating Example

In this interlude, we present a small and illustrative example to showcase how asymmetries can affect open-loop and closed-loop models.

Consider the scheduling process for the next hour of a power system containing a single power plant with a capacity of 4 MW (hence, capable of generating 4 MWh in an hour) and a generation cost of 10 \$/MWh. Consider a penalty of 100 \$/MWh for scheduling the plant below the realized demand and 0 for scheduling above the realized demand. Now, consider a demand for a given hour with the following mass distribution: 0 MWh with a probability of 0.5 and 2 MWh with a probability of 0.5. We consider there are no reserves for the sake of simplicity. Also, suppose there are no external variables, and consider the forecast model: $\Psi(\theta) = \theta$ to obtain a demand

forecast \hat{D} , i.e., $\hat{D} = \theta$. The LS solution for the demand forecast is $\hat{D} = 1$ as it is the minimizer of: $\min_{\theta} 0.5(0 - \theta)^2 + 0.5(2 - \theta)^2$.

Consider the following optimization of the planning function, G_p :

$$g^* \in \arg \min_{g, \delta^{LS}, \delta^{SP}} 10g + 100\delta^{LS} + 0\delta^{SP} \quad (6)$$

$$s.t. \ g + \delta^{LS} - \delta^{SP} = \hat{D} \quad (7)$$

$$0 \leq g \leq 4, \quad 0 \leq \delta^{LS}, \quad 0 \leq \delta^{SP} \quad (8)$$

where g is the generation being scheduled for the related hour, g^* is the schedule decision, δ^{LS} is the value of load shed (missing energy), and δ^{SP} is the load being spilled (excess energy). The optimization of the assessment (or implementation) function, G_a , is given by:

$$\min_{\delta^{LS}, \delta^{SP}} 10g^* + 100\delta^{LS} + 0\delta^{SP} \quad (9)$$

$$s.t. \ \delta^{LS} - \delta^{SP} = D - g^* \quad (10)$$

$$0 \leq \delta^{LS}, \quad 0 \leq \delta^{SP} \quad (11)$$

where g^* is fixed and D is the actual demand that is realized during the implementation.

Considering the LS forecast, the operator will run its dispatch *planning* model and say that the generator should be set to $g^* = 1$ (MWh), at the cost of $1 \times 10 = 10$ \$. Later, in the assessment or implementation step, 50% of the time, the total cost will be just 10 \$, as the system was over-prepared, and 50% of the time, the cost will be $1 \times 10 + 1 \times 100 = 110$ \$, thereby leading to an average cost of $0.5 \times 10 + 0.5 \times 110 = 60$ \$.

Now consider a biased forecast of $\hat{D} = 1.1$ (MWh). The generator should be set to $g^* = 1.1$, at cost of $1.1 \times 10 = 11$ \$. Hence, 50% of the time, the total cost will be just 11, as the system was over-prepared, and 50% of the time, the cost will be $1.1 \times 10 + 0.9 \times 100 = 101$ \$, thus leading to an average cost of $0.5 \times 11 + 0.5 \times 101 = 56$ \$. Hence, a biased forecast did perform better than the LS forecast for the problem in question. Notwithstanding, based on our methodology, we can go even further and obtain the best forecast for this application. For this simple case, we can analytically

compute the application-driven cost function given a forecast rule, $\hat{D} = \theta$. For simplicity, we only consider $\hat{D} \in [0, 2]$, which contains the extremes of the support set. Thus, in this case, $g^* = \hat{D}$. So, if $\theta \leq 2$, then the cost is given by $10\theta + 0.5 \times 100(2 - \theta)$. Else, if $\theta \geq 2$, then the average cost is 10θ and the minimizer is clearly $\hat{D} = \theta = 2$, which is clearly different from the one obtained with least squares ($\hat{D} = 1$).

One of the main reasons for the phenomenon described above is that the operator is constrained to a specific decision rule (planning method) that has to consider a given forecast in a predetermined way. While the above example is indeed very stylized to allow for a simple exposition, it carries the fundamental idea that is applied in most power systems: 1) planning models are predetermined and must consider demand forecasts; 2) costs are asymmetric as the cost of not delivering energy is usually more expensive than over-scheduling; 3) some operators have noticed the relevance of the asymmetry and, as described in the introduction, have been implementing *ad hoc* procedures to introduce out-of-market bias on forecasts.

5. Convergence Results

In this section, we discuss some conditions for the convergence of estimators obtained with application-driven joint prediction and optimization. Again, our goal is to obtain the best possible forecast \hat{y}_t , but this is completely defined by the parameters θ since x_t is known. Let S_T be the set of optimal solutions of (1)-(5), so that $\theta_T \in S_T$. We will show that any sequence of θ_T , each in the set S_T , converges to the solution set of the actual expected value formulation of the problem (as opposed to the previously presented sampled version). We will start by describing some assumptions, then we will state and prove the main theorem.

Assumption 1 *There is a unique solution z_t of optimization problem (3) for all possible values of \hat{y}_t .*

In other words, the problem is always feasible and the solution set is a singleton. This is not as restrictive as it seems. The feasibility requirement is similar to the classical assumption of complete

recourse in stochastic programming. The uniqueness requirement is equivalent to the absence of dual degeneracy in a linear program (Borrelli et al. 2003). In this case, the problem in question is dual-degenerate, but it is possible to eliminate this degeneracy by perturbing the objective function—in our case, the vectors c_p and q_p —with small numbers that do not depend on the right-hand-side (RHS) of the problem. Thus, the same perturbation is valid for all possible \hat{y}_t (Megiddo and Chandrasekaran 1989). Another possibility would be resorting to some lexicographic simplex method (Nocedal and Wright 2006). In this setting, we can define the set-valued function:

$$\zeta(y) := \operatorname{argmin}_{z \in Z} G_p(z, y) \quad (12)$$

From Böhm (1975) we know that if $\zeta(y)$ is a compact set for all y then it is a continuous set-valued function. Moreover, since $\zeta(y)$ is a singleton for all possible values of y , then we treat it as a vector-valued function that is continuous and piece-wise affine (Borrelli et al. 2003).

Assumption 2 *The feasibility set Z that appears in (3) is a non-empty and bounded polyhedron.*

Assumption 2 is reasonable since this is the set of implementable solutions of the decision-maker, typically representing physical quantities.

Assumption 3 *The feasibility set of the dual of the problem that defines $Q_a(z, y)$ in (5) is non-empty and bounded.*

Note that this set does not depend on z and y , since they appear in the RHS of the primal problem. Again, this assumption is akin to a relatively complete recourse assumption applied to the problem defining the outer-level function.

We state now our main convergence result.

THEOREM 1. Consider the process given by (1)–(5) and any possible output $\theta_T \in S_T$, for each T . Suppose that (i) Assumptions 1, 2 and 3 hold, (ii) the forecasting function $\Psi(\cdot, \cdot)$ is continuous

in both arguments, (iii) the data process $(X_1, Y_1), \dots, (X_T, Y_T)$ is independent and identically distributed (with (X, Y) denoting a generic element), (iv) the random variable Y is integrable, and (v) the set Θ is compact. Then, with probability 1,

$$\lim_{T \rightarrow \infty} d(\theta_T, S^*) = 0, \quad (13)$$

where d is the Euclidean distance from a point to a set and S^* is defined as

$$S^* = \operatorname{argmin}_{\theta \in \Theta} \mathbb{E} [G_a(\zeta(\Psi(\theta, X)), Y)], \quad (14)$$

with $\zeta(\cdot)$ defined in (12).

Proof: First, notice that $G_i(z, Y)$, $i \in \{a, p\}$, is continuous with respect to its arguments as it is a sum of a linear function and the optimal value of a parametric program (Gal 2010). Recall that ζ is a continuous vector-valued function because of Assumption 1. Hence, $G_a(\zeta(\Psi(\theta, X)), Y)$ is a real-valued continuous function. Next, we show that $G_a(\zeta(\Psi(\theta, X)), Y)$ is integrable. Indeed, since Z is bounded (Assumption 2), it follows that $\zeta(y)$ is bounded for all x by a constant, say K_1 , so that $\|\zeta(y)\| \leq K_1$. By duality, $Q_i(z, y) = \max_{\pi} \{(b_i - H_i z + F_i y)^\top \pi \mid W_i^\top \pi = q_i, \pi \geq 0\}$, but by Assumption 3 the dual variables of $Q_a(z, y)$ are bounded by a constant, say K_2 , so $\|\pi\| \leq K_2$. Thus, by a sequence of applications of Cauchy-Schwarz and triangle inequalities, we have that

$$\begin{aligned} |G_a(\zeta(\Psi(\theta, X)), Y)| &\leq |c_a^\top \zeta(\Psi(\theta, X)) + Q_a(\zeta(\Psi(\theta, X)), Y)| \\ &\leq \|c_a\| \|\zeta(\Psi(\theta, X))\| + \|b_a - H_a \zeta(\Psi(\theta, X)) + F_a Y\| \max_{W_i^\top \pi = q_a, \pi \geq 0} \|\pi\| \\ &\leq K_1 \|c_a\| + K_2 (\|b_a\| + \|H_a \zeta(\Psi(\theta, X))\| + \|F_a Y\|) \\ &\leq K_1 \|c_a\| + K_2 (\|b_a\| + \|H_a\| K_1 + \|F_a\| \|Y\|). \end{aligned}$$

Hence, since Y is integrable (condition (iv) of the Theorem), we have that $G_a(\zeta(\Psi(\theta, X)), Y)$ is integrable.

It follows that the conditions of Theorem 7.53 in Shapiro et al. (2014) are satisfied and we conclude that: (i) the function $\varphi(\theta) := \mathbb{E}[G_a(\zeta(\Psi(\theta, X)), Y)]$ is finite valued and continuous in θ , (ii) by the Strong Law of Large Numbers, for any $\theta \in \Theta$ we have

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T G_a(\zeta(\Psi(\theta, X_t)), Y_t) = \mathbb{E} [G_a(\zeta(\Psi(\theta, X)), Y)] \quad \text{w.p.1}, \quad (15)$$

and (iii) the convergence in (15) is *uniform* in θ . Thus, by Theorem 5.3 in Shapiro et al. (2014), since the set Θ is compact we have that the minimizers (over Θ) of the expression inside the limit on the left-hand side of (15)—i.e., θ_T —converge to the minimizers of the expression on the right-hand side in the sense of (13)–(14). *Q.E.D.*

Remark 1 *Assumption 3 can be replaced by assuming a compact support of Y ; in this case, $G_a(z, y)$ is a continuous function, where both arguments are defined on compact sets, hence it attains a maximum and is trivially integrable.*

While Theorem 1 provides the desired convergence result, condition (iii) of the theorem clearly precludes modeling the situation where the features x_t include (functions of) previous observations y_{t-1}, \dots, y_{t-k} . We now extend that result to the case where the features x_t include only lagged observations of $\{y_t\}$. To do so, suppose that the data process generating $\{Y_t\}_{t=1}^\infty$ is a *stationary ergodic* time series. Stationarity means that the joint distribution of (Y_1, Y_2, \dots, Y_k) is the same as the joint distribution of $(Y_{t+1}, Y_{t+2}, \dots, Y_{t+k})$ for all positive integers t and k . It is a standard assumption in the analysis of time series; see, e.g., Brockwell and Davis (2009) (note that Brockwell and Davis 2009 actually call this notion *strict stationarity*, but elsewhere in the literature it is called just stationarity; see, e.g., White 2014 or Billingsley 1986). On the other hand, an ergodic time series is—roughly speaking—one that exhibits a form of “average asymptotic independence”; a precise definition can be found, for instance, in White (2014). Statistical tests for stationarity and ergodicity of Markovian processes have been developed by Domowitz and El-Gamal (1993).

For reference, we state now a lemma that provides a more general result than Theorem 7.53 in Shapiro et al. (2014):

LEMMA 1. Theorem 7.53 in Shapiro et al. (2014) is still valid if the i.i.d. assumption is replaced with a weaker assumption that the samples form a stationary ergodic process.

Proof: Any measurable function of a stationary ergodic process is also a stationary ergodic process (Billingsley 1986). Moreover, if a process $\{W_t\}_{t=1}^\infty$ is stationary and ergodic, then the classical ergodic theorem (see, e.g., Billingsley (1986)) ensures that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T W_t = \mathbb{E}[W_1] \quad \text{w.p.1.}$$

A closer look at the proof of Theorem 7.53 in Shapiro et al. (2014) shows that the i.i.d. assumption is used only to invoke the Strong Law of Large Numbers, which as shown above, can be replaced by the ergodic theorem in the more general case. *Q.E.D.*

We can now state a more general version of Theorem 1:

THEOREM 2. Theorem 1 is still valid if the assumption that the data process $(X_1, Y_1), \dots, (X_T, Y_T)$ is independent and identically distributed is replaced with the following assumption: X_t is defined as a (measurable) function of Y_1, \dots, Y_t , and the data process generating $\{Y_t\}_{t=1}^\infty$ is a *stationary ergodic* time series.

Proof: Fix $\theta \in \Theta$. Consider the function Φ defined as

$$\Phi(Y_1, \dots, Y_t) := G_a(\zeta(\Psi(\theta, X_t)), Y_t)$$

and process $\{W_t\}_{t=1}^\infty$ defined as $W_t := \Phi(Y_1, \dots, Y_t)$. Under the assumption that $\{Y_t\}_{t=1}^\infty$ is a stationary ergodic time series, it follows that $\{W_t\}_{t=1}^\infty$ is stationary and ergodic, since Φ is measurable. Lemma 1 shows that (15) holds in this case and, therefore, the proof follows the same steps as those of the proof of Theorem 1. *Q.E.D.*

COROLLARY 1. Theorems 1 and 2 imply that if G_a is an assessment function describing the ultimate goal of a given practitioner—e.g., the expected cost incurred when using the forecast function $\Psi(\theta, X)$ within a given application—then, under the conditions of Theorem 1 or Theorem 2, any convergent subsequence of the process $\{\Theta_T\}_{T=1}^\infty$ generated by the estimation process (1)–(5) converges to the (not necessarily unique) best forecast model, $\Psi(\theta^*, X)$, in terms of the related application. That is, $\mathbb{E}[G_a(\zeta(\Psi(\theta^*, X)), Y)] \leq \mathbb{E}[G_a(\zeta(\Psi(\theta, X)), Y)] \quad \forall \theta \in \Theta$. This ensures that our model performs asymptotically better than the classical offline approaches.

Proof: Because of Theorem 1 or Theorem 2, we know that the distance between θ_T and the set S^* of minimizers of the underlying problem (given by (14)) converges to zero as $T \rightarrow \infty$. Thus, any convergent subsequence of the process $\{\Theta_T\}_{T=1}^\infty$ converges to some minimizer $\theta^* \in S^*$. Moreover, since $\{\Theta_T\}_{T=1}^\infty \subseteq \Theta$ and Θ is assumed to be compact, it follows that there exists at least one convergent subsequence. Meanwhile, any other possible choice of θ , obtained by any other estimation method, is merely a feasible solution for the underlying problem, hence, cannot be strictly better than θ^* . Q.E.D.

Corollary 1 highlights that application-driven learning is the best one can do for a fixed triplet of assessment, planning and forecasting functions when the ultimate goal is only minimizing the assessment cost. Although it is possible that other methods lead to the same optimal objective cost, they cannot be better. For other goals, such as minimizing the squared error of forecasts, other methods, such as least squares, will be better since they are inherently aligned with such other goals.

6. Solution Methodology

In this section, we describe solution methods to estimate the forecasting model within the proposed application-driven closed-loop framework described in (1)–(5). First, we present an exact method based on an equivalent single-level mixed integer linear programming (MILP) reformulation of the bilevel optimization problem (1)–(5). This method uses MILP-based linearization techniques to address the Karush Kuhn Tucker (KKT) optimality conditions of the second level and thereby guarantee the global optimality of the solution in exchange for limited scalability. In the sequence, we describe how to use zero-order methods (Conn et al. 2009) that do not require gradients to develop an efficient and scalable heuristic method to achieve high-quality solutions to larger instances. These methods will leverage existing optimization solvers, their current implementations and features.

6.1. MILP-based exact method

Our first approach consists of solving the bilevel problem (1)–(3) with standard techniques based on the KKT conditions of the second-level problem (Fortuny-Amat and McCarl 1981). Thus, the

resulting single-level nonlinear equivalent formulation can be reformulated as a MILP and solved by standard commercial solvers. The conversion between the KKT form to the MIP form can be done by numerous techniques (Siddiqui and Gabriel 2013, Pereira et al. 2005, Fortuny-Amat and McCarl 1981), all of which have pros and cons. These techniques are implemented and automatically selected by the open-source package BilevelJuMP.jl (Dias Garcia et al. 2022). This new package was conceived to allow users to formulate bilevel problems in JuMP (Dunning et al. 2017) and solve them with multiple off-the-shelf optimization solvers.

For the sake of completeness, we write the single-level nonlinear reformulation of the bilevel problem (1)–(5) in (16)–(20). For simplicity, in this model, we assume that $Z = \{z | Ax \geq h\}$ and that Θ is polyhedral.

$$\min_{\theta \in \Theta, \hat{y}_t, z_t^*, u_t, \pi_t} \quad \frac{1}{T} \sum_{t \in \mathbb{T}} [c_a^\top z_t^* + Q_a(z_t^*, y_t)] \quad (16)$$

$$s.t. \quad \forall t \in \mathbb{T} :$$

$$\hat{y}_t = \Psi(\theta, x_t) \quad (17)$$

$$W_p y_t + H_p z_t^* \geq b_p + F_p \hat{y}_t ; \quad A z_t^* \geq h \quad (18)$$

$$W_p^\top \pi_t = q_p ; \quad H_p^\top \pi_t + A^\top \mu_t = c_p ; \quad \pi_t, \mu_t \geq 0 \quad (19)$$

$$\pi_t \perp W_p u_t + H_p z_t^* - b_p - F_p \hat{y}_t ; \quad \mu_t \perp A z_t^* - h \quad (20)$$

Equations (16) and (17) are the same as (1) and (2). (3) was replaced by (18)–(20). (18) are the primal feasibility constraint, (19) are the dual feasibility constraints, and (20) represents the complementarity constraints.

6.2. Scalable heuristic method

The proposed class of methods will make extensive use of the way of thinking described in Figure 1. In other words, the core algorithm decomposes the problem as follows: We call this method a pseudo-algorithm because a few steps are not well specified, namely *Initialization*, *Update*, and *Convergence* check, allowing for a wide range of possible specifications. *Initialization* can be as

Algorithm 1: Pseudo algorithm

Result: Optimized θ Initialize θ ;**while** *Not converged* **do** Update θ ; **for** $t \in \mathbb{T}$ **do** Forecast: $\hat{y}_t \leftarrow \Psi(\theta, x_t)$; Plan Policy: $z_t^* \leftarrow \arg \min_{z \in Z} G_p(z, \hat{y}_t)$; Cost Assessment: $cost_t \leftarrow G_a(z_t^*, y_t)$ **end** Compute cost: $cost(\theta) \leftarrow \sum_{t \in \mathbb{T}} (cost_t)$ **end**

simple as θ receiving a vector of zeros, which might not be good if the actual algorithm is a local search. One alternative that will be applied in the case study section is the usage of traditional models as starting points, for instance, the ordinary least squares. In the case study, we will initialize the algorithm with the LS estimate, this guarantees that the algorithm will return at most the same cost as the open-loop framework in the training sample. There are many possibilities for the *convergence* test. For instance, iteration limit, time, the variation of the objective function value, and other algorithm-specific tests. Finally, the *update* step depends on the selected concrete algorithm that is ultimately minimizing the non-trivial $cost(\theta)$ function.

We will focus on a derivative-free local search algorithm named Nelder-Mead (Conn et al. 2009). Notwithstanding, it is relevant to highlight the generality of the proposed pseudo-algorithm. For instance, gradient-based algorithms could also be developed based on numerical differentiation and automatic differentiation (Nocedal and Wright 2006). In this context, gradient calculation would enable the usage of Gradient Descent and BFGS-like algorithms (Nocedal and Wright 2006).

The main features of the above-proposed pseudo-algorithm are: 1) it is suitable for parallel computing (the loop in the sample \mathbb{T} is intrinsically decoupled); 2) each step is based on a deterministic

LP defining the second-level variables in (3), suitable for off-the-shelf commercial solvers that find globally optimal solutions in polynomial time; 3) each inner step can significantly benefit from warm-start processes developed in linear programming solvers (e.g., the dual simplex warm-start is extremely powerful, and many times only a handful of iterations will be needed in comparison to possibly thousands of iterations if there were no warm-start, cf. Nocedal and Wright (2006)). It is worth emphasizing that the aforementioned feature 2) allows for an exact (always optimal) description of the second-level problem. In our approach, we keep the second level exact and face the challenge of optimizing a nonlinear problem on the upper level. In contrast, Muñoz et al. (2022) choose to relax the complementarity constraints and deal with a nonlinear program lacking the benefits of the above-mentioned features 1) to 3). As will be illustrated in the case study, our choice is supported by empirical evidence about the shape of the nonlinear function faced in the objective function. Additionally, it is usual in more complex estimation processes (like maximum likelihood-based methods) to rely on nonlinear optimization methods to select the best parameters (Henningsen and Toomet 2011). Moreover, although not convex, as explored in our case study, the objective function exhibits relevant properties that facilitate the search within the parameters domain. Finally, note that this heuristic approach allows for a wider set of forecast models, such as NN and other machine learning models, as it only requires that a forecast can be pointwise obtained and its performance evaluated by the cost function for a given trial solution (parameters).

One caveat is that variations on θ can lead to possibly infeasible results for the *Policy Planning* and *Cost Assessment* optimization problems. Consequently, we require complete recourse for such problems. In cases where this property does not hold, it is always possible to add artificial (slack) variables with high penalty costs in the objective function to keep the problem feasible. In the energy and reserve dispatch problem, this requirement is addressed by imbalance variables (load and renewable curtailment decisions).

7. Application-Driven Load Forecasting and Reserve Sizing

In this work, we focus on the energy and reserve scheduling problem of power systems (Chen et al. 2013, Kirschen and Strbac 2018). In this problem, we aim to obtain the best joint conditional

point-forecast for the vector of nodal demands, \hat{D}_t , and vectors of up and down zonal or nodal reserve requirements, $\hat{R}_t^{(up)}$ and $\hat{R}_t^{(dn)}$. While the forecast vector of nodal loads represents, e.g., the next hour operating point target that system operators and agents should comply with, up- and down-reserve requirements represent a forecast of the system's resource availability (or security margins), defined per zone or node, allowing the system to withstand load deviations. Note that we can think of loads as a general net load that corresponds to load minus non-dispatchable (e.g., renewable) generation.

The inputs of the problem are: vectors of historical data of dependent and explanatory variables, $\{y_t, x_t\}_{t \in \mathbb{T}}$, including lags of demand, D_{t-1}, \dots, D_{t-k} , and possibly other exogenous covariates such as climate and weather indices (or forecasts), dummy variables, and all sort of non-linear machine-learning-based forecasts for the dependent variables; vectors of data associated with generating units; maximum generation capacity, G , dispatch costs or offers, c ; maximum up- and down-reserve capacity, $\bar{r}^{(up)}$ and $\bar{r}^{(dn)}$; up- and down-reserves costs, $p^{(up)}$ and $p^{(dn)}$; load-shed and spillage penalty costs, λ^{LS} and λ^{SP} ; network data comprising the vector of transmission line capacities F ; and network sensitivity matrix, B , describing the network topology and physical laws of electric circuits. Additionally, it is important to mention that the input data describing the system characteristics can be provided under two perspectives: 1) under the perspective of the actual *ex-post* (or assessed/implemented) operation, i.e., based on the observed demand data and most accurate system's description for optimizing the function G_a defined in (1); and 2) under the *ex-ante* planning perspective, G_p , which is accounted for in (3) based on observed features, such as previous information, and system operator's description of the system considered in the dispatch tool. While the former has already been listed at the beginning of this paragraph, the latter uses the same symbols but with a tilde above, i.e., $\tilde{c}, \tilde{p}^{(up)}, \tilde{p}^{(dn)}, \tilde{G}, \tilde{B}$, etc. For a simple matrix representation of the problem, we define e to be a vector with one in all entries and an appropriate dimension. M is an incidence matrix with buses in rows and generators in columns that is one when the generator lies in that bus and zero otherwise. Similarly, N is an incidence matrix with generators in columns and

reserve zones in rows, which is one if the generator lies in that area. Thus, we study the following particularization of the closed-loop framework proposed in (1)–(5):

$$\min_{\substack{\theta_D, \theta_{R^{up}}, \theta_{R^{dn}}, \\ \hat{D}_t, \hat{R}_t^{(up)}, \hat{R}_t^{(dn)}, g_t, \delta_t^{LS}, \delta_t^{SP}, \\ g_t^*, r_t^{(up)*}, r_t^{(dn)*}}} \frac{1}{T} \sum_{t \in \mathbb{T}} [c^\top g_t^* + p^{(up)\top} \hat{r}_t^{(up)*} + p^{(dn)\top} \hat{r}_t^{(dn)*} + \lambda^{LS} \delta_t^{LS} + \lambda^{SP} \delta_t^{SP}] \quad (21)$$

$$s.t. \quad \forall t \in \mathbb{T} :$$

$$\hat{D}_t = \Psi_D(\theta_D, x_t) \quad (22)$$

$$\hat{R}_t^{(up)} = \Psi_{R^{(up)}}(\theta_{R^{(up)}}, x_t) \quad (23)$$

$$\hat{R}_t^{(dn)} = \Psi_{R^{(dn)}}(\theta_{R^{(dn)}}, x_t) \quad (24)$$

$$e^\top (Mg_t - \delta_t^{SP}) = e^\top (D_t - \delta_t^{LS}) \quad (25)$$

$$-F \leq B(Mg_t + \delta_t^{LS} - D_t - \delta_t^{SP}) \leq F \quad (26)$$

$$g_t^* - r_t^{(dn)*} \leq g_t \leq g_t^* + r_t^{(up)*} \quad (27)$$

$$\delta_t^{LS}, \delta_t^{SP}, \hat{R}_t^{(up)}, \hat{R}_t^{(dn)}, g_t \geq 0 \quad (28)$$

$$(g_t^*, r_t^{(up)*}, r_t^{(dn)*}) \in \arg \min_{\substack{\hat{g}_t, \hat{\delta}_t^{LS}, \hat{\delta}_t^{SP}, \\ \hat{r}_t^{(up)}, \hat{r}_t^{(dn)}}} [\tilde{c}^\top \hat{g}_t + \tilde{p}^{(up)\top} \hat{r}_t^{(up)} + \tilde{p}^{(dn)\top} \hat{r}_t^{(dn)} + \tilde{\lambda}^{LS} \hat{\delta}_t^{LS} + \tilde{\lambda}^{SP} \hat{\delta}_t^{SP}] \quad (29)$$

$$s.t. \quad e^\top (M\hat{g}_t - \hat{\delta}_t^{SP}) = e^\top (\hat{D}_t - \hat{\delta}_t^{LS}) \quad (30)$$

$$-\tilde{F} \leq \tilde{B}(M\hat{g}_t + \hat{\delta}_t^{LS} - \hat{D}_t - \hat{\delta}_t^{SP}) \leq \tilde{F} \quad (31)$$

$$N\hat{r}_t^{(up)} = \hat{R}_t^{(up)} \quad (32)$$

$$N\hat{r}_t^{(dn)} = \hat{R}_t^{(dn)} \quad (33)$$

$$\hat{g}_t + \hat{r}_t^{(up)} \leq \tilde{G} \quad (34)$$

$$\hat{g}_t - \hat{r}_t^{(dn)} \geq 0 \quad (35)$$

$$\hat{r}_t^{(up)} \leq \bar{r}^{(up)} \quad (36)$$

$$\hat{r}_t^{(dn)} \leq \bar{r}^{(dn)} \quad (37)$$

$$\hat{g}_t, \hat{r}_t^{(up)}, \hat{r}_t^{(dn)}, \hat{\delta}_t^{LS}, \hat{\delta}_t^{SP} \geq 0. \quad (38)$$

In (21)–(38), the objective function of the upper level problem (21) comprises the sum of the actual operating cost, the cost of scheduled reserves, and the implemented load-shed and renewable spillage costs for all periods within the dataset, i.e., for $t \in \mathbb{T}$. In the upper level, constraints (22)–(24) define the forecast model. Note that all periods are coupled by the vector of parameters $\theta = [\theta_D^\top, \theta_{R^{(up)}}^\top, \theta_{R^{(dn)}}^\top]^\top$, which do not depend on t . These parameters define the forecast model that will be applied to each t for demand, as per \hat{D}_t in (22), for up reserve requirements, as per $\hat{R}_t^{(up)}$ in (23), and for down reserve requirements, as per $\hat{R}_t^{(dn)}$ in (24). The forecast models are defined by functions Ψ_D , $\Psi_{R^{(up)}}$, and $\Psi_{R^{(dn)}}$ that transform parameters and the historical data on load and reserve requirement forecasts. For the sake of simplicity and didactic purposes, in this work, we assume affine regression models. The reserves are parts of the forecast vector \hat{y}_t because the method optimizes a model for them. However, reserves historical data need not be in y_t , since, typically, the choice of reserves in each period is not based on past values of reserves.

Constraints (25)–(28) together with the objective function (21) particularize G_a from (1). They assess the *ex-post* operating cost (first term, $c_t^T g_t$, and the last two terms, $\lambda_t^{LS} \delta_t^{LS} + \lambda_t^{SP} \delta_t^{SP}$) of the actual dispatch given the *ex-ante* planned generation, g_t^* , and allocated up and down reserves, $r_t^{(up)*}$ and $r_t^{(dn)*}$, defined by the second level (29)–(38). Constraint (25) accounts for the *ex-post* energy balance constraint, where total generation meets total observed load data. The left-hand side of the constraint is the sum of generated energy in all buses, with Mg_t resulting in the nodal generation injection vector (total generation per bus). δ_t^{SP} represents the nodal generation spilled per bus (positive load imbalance decision). The right-hand side of the constraint accounts for the net-nodal load vector (observed net-demand vector D_t). δ_t^{LS} represents the decision vector of nodal load-shed (negative load imbalance decision). Finally, constraint (26) limits the flow of energy through each transmission line to pre-defined bounds and (27) limits the *ex-post* generation to respect the operation range defined by the *ex-ante* planned generation, g_t^* , and allocated up and down reserves, $r_t^{(up)*}$, $r_t^{(dn)*}$. Constraint (28) ensures the positiveness of slack, generation, and reserve requirement variables.

The planning policy defines variables g_t^* , $r_t^{(up)*}$, and $r_t^{(dn)*}$ under the conditional information available in vector x_t . These variables should respect the optimality of the market's or system operator's *ex-ante* scheduling (or planning policy), as per (29), based on the vector of load forecasts, \hat{D}_t , and vectors of reserve requirements, $\hat{R}_t^{(up)}$ and $\hat{R}_t^{(dn)}$, for the next period (e.g., hour). Again, it is relevant to note that this planning policy may differ from the actual *ex-post* implemented policy (based on the observed load data D_t), resulting in the actual operating cost considered in the objective function of the first level (21). Within this context, constraints (29)–(38) detail the second-level problem, which represents the optimization that takes place in the planning phase at each period to define a generation and reserve schedule for the next period. Thus, these constraints are particularizing the general model, G_p in (3). In the proposed closed-loop framework, (29) is key. It allows us to define, within a problem that seeks the best forecast model aiming to minimize the *ex-post* operation cost, the objective of an *ex-ante* scheduling problem minimizing energy and reserve costs for a conditioned load forecast (\hat{D}_t) and reserve requirements ($\hat{R}_t^{(up)}$ and $\hat{R}_t^{(dn)}$). Constraints (30) and (31) are similar to (25) and (26). Expressions (32) and (33) ensure that the total reserve requirements ($\hat{R}_t^{(up)}$ and $\hat{R}_t^{(dn)}$, which are considered as parameters for the lower-level problem) must be allocated among generators in the form of up and down reserves ($\hat{r}_t^{(up)}$ and $\hat{r}_t^{(dn)}$ – second-level decision vectors). Constraints (34) and (35) limit the scheduled generation and reserves range (up and down) to generators' physical generation limits. Constraints (36) and (37) limit the maximum amount of reserves that can be allocated in each generating unit, and (38) ensures positiveness of the generation and reserve decision vectors of the second level. We highlight that G_a and G_p are cost functions represented by linear programs. Although their objective functions have the same structure, their above-described constraints are different.

The proposed application-driven framework provides system operators with the flexibility to either jointly optimize the load forecast and the reserve requirements or only one of them. Additionally, this model is powerful because the load forecasts and both the size and location of reserve requirements are defined in the best way possible to maximize the *ex-post* performance of the

operation. While the lower level, (29)–(38), ensures an *ex-ante* (planning) generation schedule and reserve allocation compatible with the system operator’s information level (best schedule given the previous hour conditional forecast for t), the upper level selects the parameters of the forecast model aiming to minimize the average *ex-post* (assessment/implementation) operating cost for a large dataset. In this sense, depending on the network details considered in the assessment part of the model, it also helps in mitigating reserve deliverability issues associated with *ad hoc* procedures used in industry practices.

In the next section, we will compare a few variants of this problem by optimizing all or some of the parameters θ_D , $\theta_{R(up)}$, and $\theta_{R(dn)}$. As described in the contributions section, two novel methods are the one that optimizes all parameters and the one that gets a fixed θ_D and only optimizes reserves ($\theta_{R(up)}$, and $\theta_{R(dn)}$). The case of fixed reserves and optimized θ_D will also be presented but is not particularly meaningful in practice.

8. Case Studies

This section presents case studies to demonstrate the methodology’s applicability and how the closed-loop framework can outperform the classic open-loop scheme in multiple variants of the load forecasting and reserve sizing problem defined in Section 7. First, we show that the Heuristic method of Section 6.2 can achieve close to optimal solutions in a fraction of the time required by the Exact method of Section 6.1. Second, we study the estimated parameters’ and forecasts’ empirical properties and contrast them with the classical least squares (LS) estimators. After that, we briefly explore how the method can estimate a reserves model with features. Moreover, we apply the method to highlight the relationship between load-shed cost and the estimated parameters. Finally, we show that the heuristic algorithm finds good quality local-optimal parameters systematically outperforming the LS open-loop benchmark for instances far larger than those solved in previously reported works tackling closed-loop bilevel frameworks. We used the same Dell Notebook (Intel i7 8th Gen with 4 cores at 1.99Ghz, 16Gb RAM) for all studies except when otherwise noted.

8.1. Power systems cases and datasets

We consider multiple power system cases throughout this section. The first is a single bus system defined by us, with 1 zone, 1 load (with long-term average of 6) and 4 generators (with capacities 5, 5, 2.5, 2.5 and costs 1, 2, 4, 8). The other three are typical test systems based on realistic networks used by the power system community, namely, "IEEE 24bus rts" (38 lines, 33 generators, 17 loads, 4 zones), "IEEE 118 bus" (186 lines, 54 generators, 99 loads, 7 zones) and "IEEE 300 bus" (411 lines, 69 generators, 191 loads, 10 zones). The base datasets were obtained from PG-LIB-OPF (Babaeinejadsarookolaee et al. 2019). The zone definition is standard for the 24 buses case. For the 118 and 300 bus cases, we used the zones defined by Li et al. (2015). We also considered very large cases, with thousands of buses, by connecting copies of the 300-bus system. Finally, we consider realistic power systems with more than 6000 buses.

Henceforth, we will refer to the test systems by their number of buses. We only considered as loads the buses with positive demand in the original files. We performed some modifications in the case data: all flow limits were set to 75% their rate; we modified the demand values of the cases 24, 118 and 300 by the factors: 0.9, 1.3 and 0.9, respectively. These modifications were made to stress the systems. Deficit and generation curtailment costs were defined, respectively, as 8 and 3 times the most expensive generator cost. All generators were allowed to have up to 30% their capacity allocated to reserves, and their reserve allocation costs were set to 30% their nominal costs. We only considered the linear component of the generators' costs in all instances. In all datasets, we used demand values as the long-term average of AR(1) processes for each bus. The AR(1) coefficients were set to 0.9 and the AR(0) coefficients were set so that we get the desired long-term averages. For the sake of simplicity, load profiles were generated independently. The coefficient of variation of all simulated load stochastic processes was 0.4. Negative demands were truncated to zero, although they could represent an excess renewable generation.

8.2. Studied models and notation

In most of the following sections, we will consider simple forecast models so that we can detail experiment results clearly. Therefore, unless otherwise mentioned, the model used to forecast loads in each node is the following:

$$\hat{D}_t = \Psi_D(\theta_D, x_t) = \theta_D(0) + \theta_D(1)D_{t-1}, \quad (39)$$

For the single bus case, we set the “real”, or population, values as $\theta_D(0) = 0.6$ and $\theta_D(1) = 0.9$, resulting in the long-term average $\theta_D(0)/(1 - \theta_D(1)) = 6$, defined in section 8.1. Note that such a choice of coefficients ensures that the input process is stationary and ergodic (see, e.g., Billingsley 1986, p.495), thereby satisfying the condition in Theorem 2. Because the stochastic model for loads is homoscedastic, we set the reserve models to $AR(0)$ – i.e., a number that does not depend on previous values of reserves – since it is customary to set the reserves just in terms of variability of loads:

$$\hat{R}_t^{(up)} = \Psi_{R^{(up)}}(\theta_{R^{(up)}}, x_t) = \theta_{R^{(up)}}(0), \quad (40)$$

$$\hat{R}_t^{(dn)} = \Psi_{R^{(dn)}}(\theta_{R^{(dn)}}, x_t) = \theta_{R^{(dn)}}(0). \quad (41)$$

8.3. Exact vs heuristic method comparison

In this first experiment, we aim to compare the exact and the heuristic methods to check the quality of the latter for an instance that the exact method is capable of reaching global optimal solutions. To that end, we consider the single-bus test system.

We started by solving 10 instances for each $T \in \{15, 25, 50, 75\}$. All instances solved with the exact method converged within a gap lower than 0.1% using the Gurobi solver or stopped after two hours. The heuristic method was terminated when the objective function presented a decrease lower than 10^{-7} between consecutive iterations. We used a Nelder-Mead implementation found in Mogensen and Riseth (2018). To compare the results, we plotted the ratio of objective values in Figure 2(a) and the time ratio in Figure 2(b). We can observe that the heuristic method achieves high-quality

solutions for almost all instances. Although the exact method is competitive for $T \in \{15, 25\}$, the heuristic method is much faster with an average solve time of 4.4s, for $T = 50$, and 5.9s, for $T = 75$, compared to 1200s and 6670s for the exact method. Four instances with $T = 75$ did not converge with the exact method after two hours.

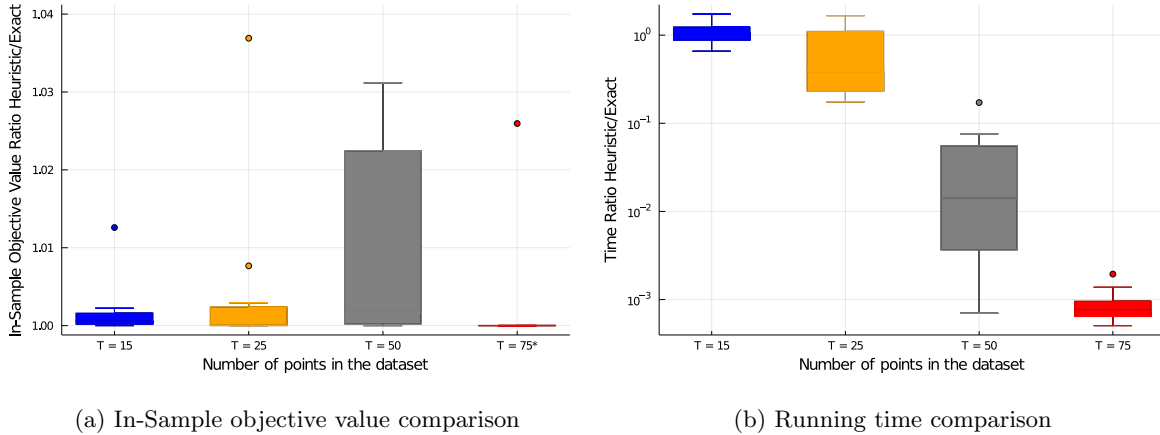


Figure 2 (a) Objective of Heuristic method divided by the objective of Exact method for the same datasets.

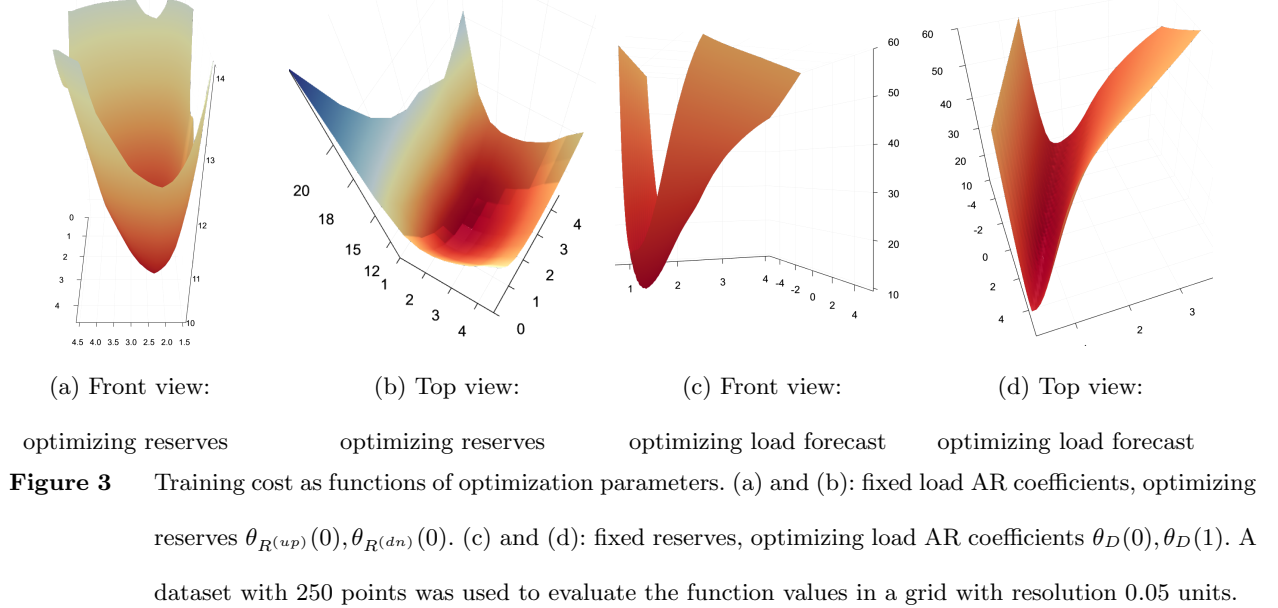
*Four problems not considered for $T = 75$: the exact method found no solution in the given time.

(b) Time to solve the same problem (in log scale): Heuristic method divided by Exact method.

Next, we further analyze the shape of the objective cost, G_a in (1), as a function of the parameters, θ , to better understand how good the heuristic solutions can be. Given one dataset with 250 points, we fixed the demand autoregressive parameters to the LS estimation and plot in \mathbb{R}^3 the cost as a function of the reserve requirement parameters. Two views of this function are presented in Figure 3 (a) and (b). We also plot the cost as a function of the $AR(0)$ and $AR(1)$ coefficients of the demand forecasting model, in this case, the reserve requirement parameters were fixed to the exogenous values of ± 1.96 standard deviations of the LS estimation of load forecast. This is presented in two views in Figure 3 (c) and (d).

We can note that both functions are reasonably well-behaved and suited for local search algorithms, even though the functions are non-convex. This is a relevant feature supporting the choice

for our heuristic approach as previously described at the end of Section 6.2. We also note a smoothing effect due to the average in the objective function (Shapiro et al. 2014). Therefore, it is expected to have more well-behaved functions as the sample size grows.



8.4. Asymptotic behavior and biased estimation

Now we focus only on the heuristic method to analyze how the estimates behave with respect to the dataset size variation. We will see that they actually converge in our experiments. Moreover, we empirically show through out-of-sample studies that using the closed-loop model is strictly better than the open-loop one provided we have a reasonable dataset size in the training step. It will be possible to see that a method with too many parameters might overfit the model for reduced dataset sizes and not generalize well enough. From now on, we will use the following nomenclature and color code to refer to the different models:

- LS-Ex (red): This is the benchmark model representing the classical open-loop approach. It uses LS to estimate demand and an exogenous reserve requirement.
- LS-Opt (blue): This is a partially optimized model, where least squares are used to estimate demand and only reserve requirements are optimized with the application-driven framework.

- Opt-Ex (yellow): This is also a partially optimized model, where demand is optimized, whereas reserve requirements are still exogenously defined. This model is not particularly meaningful in practice. We show it in some studies for completeness.

- Opt-Opt (green): This is the fully optimized model, where both demand forecast and reserve requirements are jointly optimized.

Both (closed-loop) methods LS-Opt and Opt-Opt are novel contributions proposed in this work, to be contrasted with the benchmark (open-loop) method, LS-Ex. For didactic purposes, in all cases tested in this section, up and down reserve requirements were defined as ± 1.96 standard deviations, respectively, of the estimated residuals from the LS demand forecast. Note, however, that other exogenous *ad hoc* rule could be used (Ela et al. 2011).

We empirically compare and analyze the convergence of the four demand and reserve requirement forecast models mentioned above. We varied the dataset size used in the estimation process from 50 to 1000 observations. For each dataset size, we performed 100 trial estimations, with different datasets generated from the same process, to study the convergence. To evaluate the out-of-sample performance of each one of the 100 estimates for each dataset size, we compute the objective function, G_a in (1), for a single fixed dataset with 10,000 new observations (generated with the same underlying process but different from all other data used in the estimation processes). In the following plots, lines represent mean values among the 100 estimated costs with the in-sample or out-of-sample data, and shaded areas represent the respective 10% and 90% quantiles.

The average operation cost for each dataset size is presented in Figure 4. The vertical axis shows in-sample costs, while the horizontal axis shows the dataset size used for the estimation procedure. It is possible to see that the method that co-optimizes reserves requirements and demand forecasts finds lower costs than the others. This is expected because this method has more degrees of freedom (it is a relaxed version of the others) on the parameter estimation and this is the objective function being minimized. Also, as expected, the LS plus exogenous reserves requirement model finds higher costs than the others. For the same reason, it does not allow for improvements

by the local optimization method as it can be seen as a constrained version of the others. The other two methods are always in between and *Opt-Ex* is always below *LS-Opt*, which shows that demand forecasting might have a larger effect than reserve allocation in this test system.

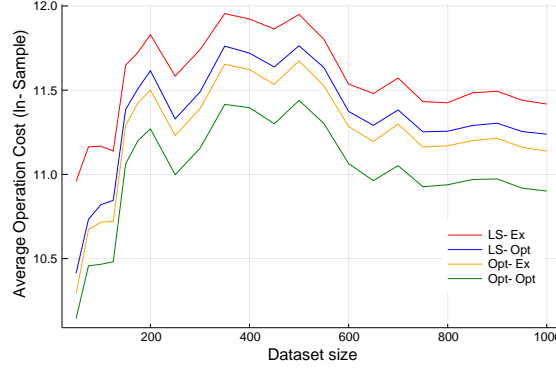


Figure 4 Average operation cost in training sample (in-sample) versus dataset size.

Figures 5 (a) and (b) depict the same costs but in the out-of-sample data. Hence, they measure how well the models generalize to data it has never seen before. We can see that the models allowing more parameters to be endogenously optimized perform much better than models with exogenously defined forecasts. Thus, we see that the application-driven learning framework works successfully on out-of-sample data when estimated with datasets larger than 150. However, we note that these steady improvements require more data than the classic exogenous models, as shown in Figure 5 (b). Between 50 and 120 points, the model with more optimization flexibility, Opt-Opt, exhibits a more significant cost variance. This is due to excessive optimization in a small dataset that led to overfitting and poor generalization. Note that, in this work, we did not consider any regularization procedure to avoid this issue. However, our optimization-based framework is suitable for well-known shrinkage operators (Tibshirani 2011) that can be readily added in the objective function (1).

Figure 6 shows how the estimated parameters behave as functions of the estimation dataset size. In Figure 6 (a) and (b) we can see that the load model parameters are indeed converging to long-run values. It is also clear to see the bias in those parameters. The constant term is greatly

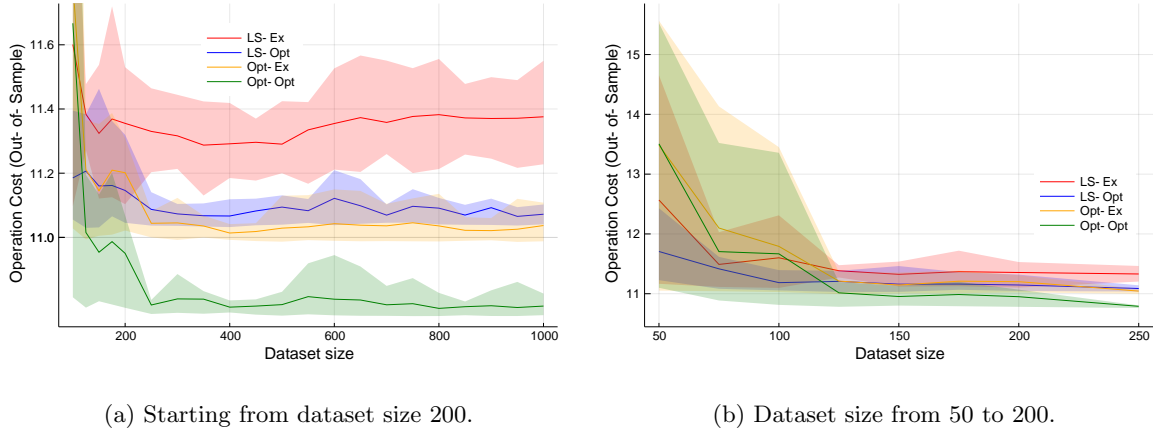


Figure 5 Out-of-sample average operation cost versus (in-sample) dataset size. Lines represent the average of the 100 estimation trials. Shaded areas represent the 10% and 90% quantiles. All trials are evaluated on a single out-of-sample dataset with size 10,000 observations.

increased while the autoregressive coefficient is slightly reduced. Ultimately this leads to a larger forecast value, which can be interpreted as the application risk adjustment due to the asymmetric imbalance penalization costs (load-shed is much higher than the spillage cost). Thus, the Opt-Opt model will do the best possible to balance these costs, thereby prioritizing the load-shed by increasing the forecast level. The fixed reserves model (Opt-Ex) is less biased because the fixed reserves constrain how much the load model can bias due to the risk of not having enough reserves to address lower demand realizations. Note that the red (LS-Ex) is on top of the purple (LS-Opt) since both use the same LS estimates for demand, which exhibits the lowest variance.

In Figure 6 (c) and (d), we see that the Opt-Opt model greatly increases the downward reserve and decreases the upward reserve, both consistent with the change in the demand forecast parameters. Closed-loop estimation of only reserves led to increased up reserves that are the most expensive to violate, while downward reserves are mostly unaffected, this might be an artifact of the estimation model that uses the open-loop estimation as a starting point. The Opt-Opt model is limited to 3 because that is the maximum reserve that can be allocated (30% of the generators' capacity).

To highlight the bias on load forecast we present, in Figure 7 (a), a histogram of deviations: $error := realization - forecast$. Negative values mean that the forecast value was above the realization. The LS estimation leads to an unbiased estimator, seen in the red histogram centered on

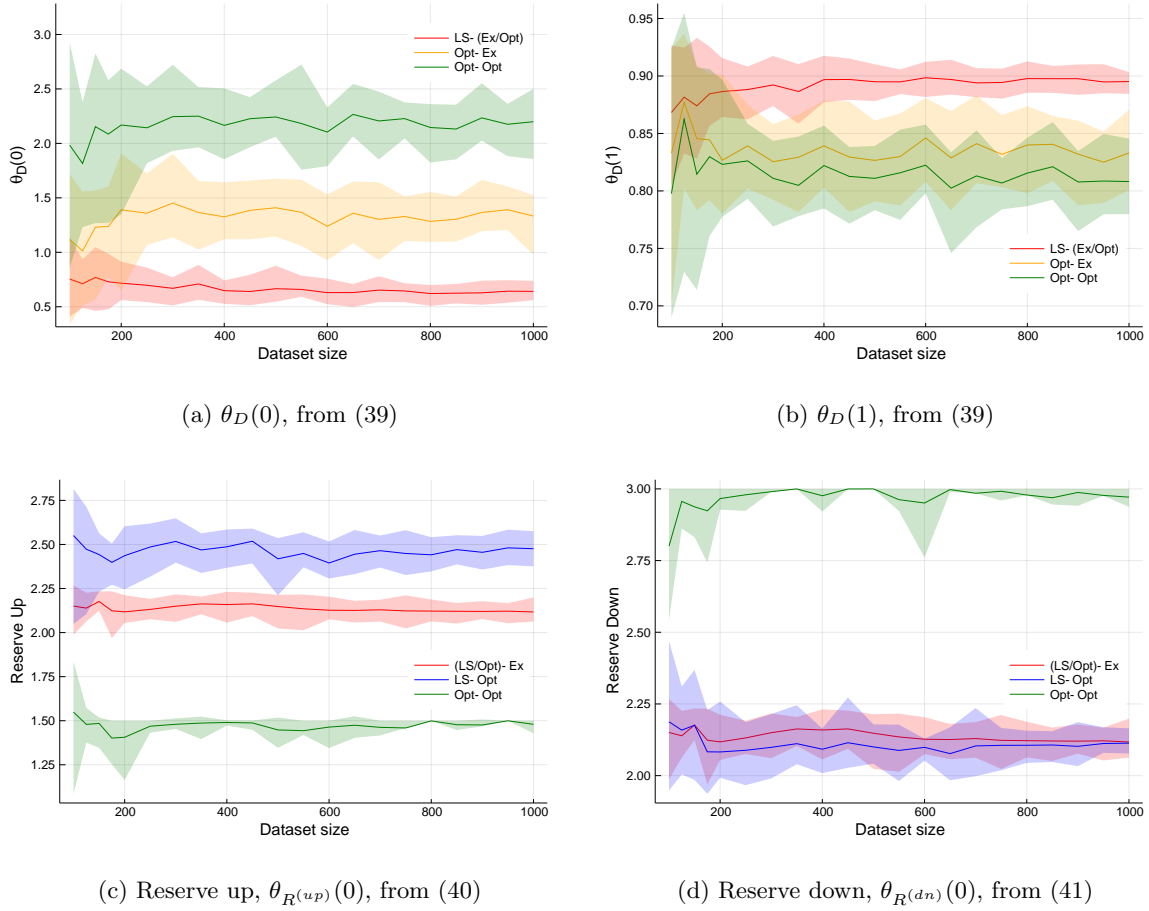
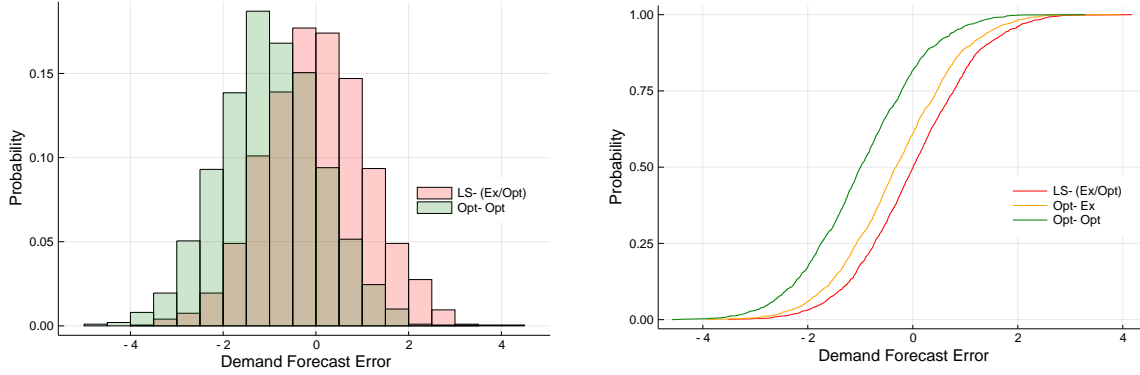


Figure 6 Estimated parameters versus dataset size. Lines represent the average of the 100 estimation trials. Shaded areas represent the 10% and 90% quantiles. (a) and (b) Load coefficients, the models LS-Ex and LS-Opt coincide, thereby are presented as LS-(Ex/Opt). (c) and (d) Reserve coefficients The models LS-Ex and Opt-Ex coincide, thereby are presented as (LS/Opt)-Ex.

zero. On the other hand, the forecast from the fully endogenous model is clearly biased, as it consistently forecasts higher values than the realizations. This fact is corroborated by the cumulative distribution functions displayed in Figure 7 (b).

8.5. A reserves model with features

This experiment aims to show that it is also possible to consider a reserves model with features within the proposed scheme, i.e., conditioned to external information being dynamically revealed to the system operator. There are examples of works in the literature that considered the load to be heteroscedastic (Van der Meer et al. 2018). Hence, we will study here a simple formulation



(a) Two histograms are shown, the third color is their intersection (LS-Ex and LS-Opt coincide). (b) Accumulated Probability – out of the four models, three are shown here (LS-Ex and LS-Opt coincide).

Figure 7 Forecast error (observation – forecast) in a histogram, comparing the fully optimized model with least squares estimation. Negative values mean forecast was larger than actual realization.

of demand time series with time-varying variance: We considered an exogenous variable E_t that follows an autoregressive process of order one, for which the noise term, ε_t , has zero mean and variance σ_E^2 , (42). The variance of the demand process is the square of E_t , as shown in (43).

$$E_t = \phi(0) + \phi(1)E_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma_E^2), \quad (42)$$

$$D_t = \theta_D(0) + \theta_D(1)D_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, E_t^2), \quad (43)$$

Formulations in (42) and (43) were only used to generate synthetic input data for this case study. In this study, we did not modify the demand forecast model (39) that is used to specify (22). On the other hand, we allowed for features in the reserve sizing, that is, the reserve will vary with external information. This time dependency will be considered through contextual information (features, explanatory variables). As a driver for demand variance changes, E_t is reasonable contextual information for estimating dynamic reserve margins. Thus, we replaced (40) and (41) that specify (23) and (24) by the following models for the reserve requirements, (44) and (45):

$$\hat{R}_t^{(up)} = \Psi_{R^{(up)}}(\theta_{R^{(up)}}, x_t) = \theta_{R^{(up)}}(0) + \theta_{R^{(up)}}(1)E_t, \quad (44)$$

$$\hat{R}_t^{(dn)} = \Psi_{R^{(dn)}}(\theta_{R^{(dn)}}, x_t) = \theta_{R^{(dn)}}(0) + \theta_{R^{(dn)}}(1)E_t. \quad (45)$$

The results of this experiment are depicted in Figure 8. We refer to the model from the previous sections as *simple reserves model*, Figure 8 (a), and the model defined here as *reserves model with features*, Figure 8 (b). The system cost was clearly reduced by considering features in the reserves for both Opt-Opt and LS-Opt models. In the case with features, the Opt-Opt model gain over the LS-Opt is significantly reduced because the reserve sizing method given by (44) and (45) is able to capture asymmetries by incorporating the variable E_t , therefore this reduces the pressure for an increased bias on demand forecast. Moreover, the demand forecast models do not depend explicitly on E_t for both Opt-Opt and LS-Opt. Hence, in these instances, the reserve sizing, dependent on E_t , is able to optimize up- and down-reserve requirements that incorporate most of the asymmetries on the cost function even with a demand forecast with zero bias. This is a relevant insight that the proposed framework provides.

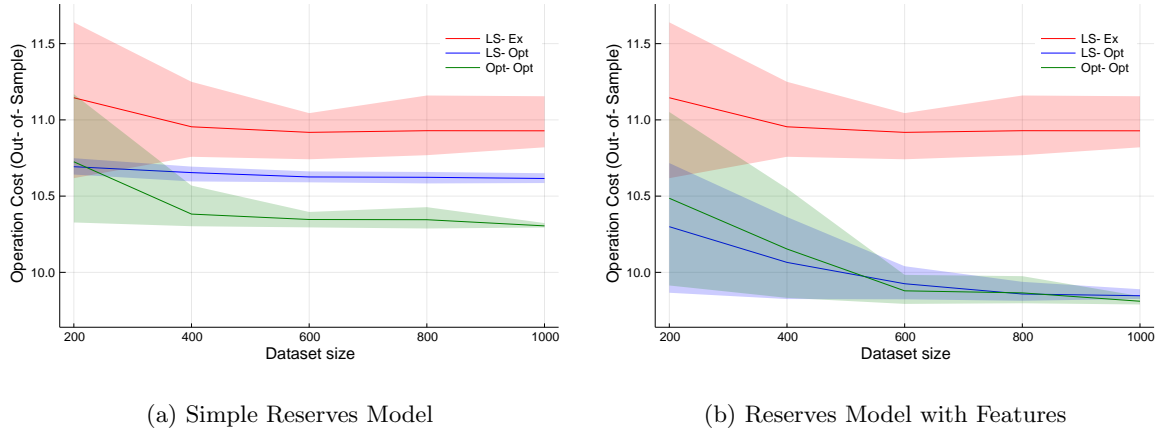


Figure 8 Variance driven by exogenous variable. Lines represent the average of the 100 estimation trials. Shaded areas represent the 10% and 90% quantiles.

8.6. Reserves and demand forecast as a function of the load-shed cost

This section aims to spotlight the dependency of the estimated parameters on the deficit cost, which is the largest violation penalty in this problem. We varied the deficit cost between 15 and 100 and estimated parameters with one single dataset with 1,000 observations.

Figure 9 depicts the up and down reserves and steady-state demand ($\theta_D(0)/(1 - \theta_D(1))$). We added the solid line in red, *Demand LS*, to represent the steady-state demand and dashed red, *Reserve Ex*, to show the obtained reserve requirement from the residue of the LS estimation model. So, as these two values are exogenously calculated, they do not vary with the load-shed cost. The solid green line, *Demand Opt-Opt*, shows an increasing bias as the load-shed cost grows, corroborating the expected behavior. The dashed green line also shows that the best response is to increase reserve levels as the load-shed cost grows. The dashed blue lines are the reserve margins around the solid LS-based demand (red line) and they are also clearly affected by the load-shed cost.

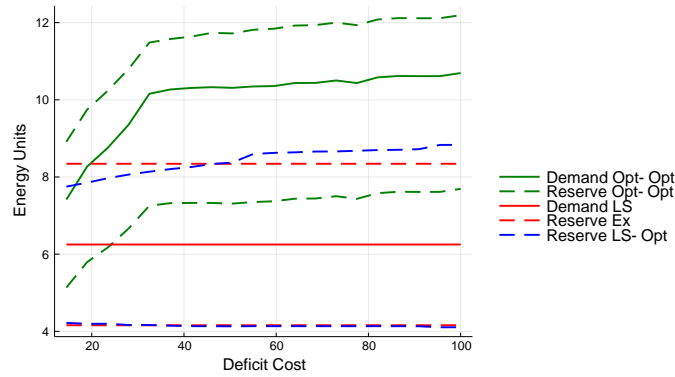


Figure 9 Long-run averages and reserve margins, in units of energy, as functions of deficit cost

8.7. Large-scale optimization

In this experiment, we apply the closed-loop estimation methodology to larger power system networks. As in the previous section, we focus our attention on the three models: LS-Ex, LS-Opt and Opt-Opt. Our primary goal in this case study is to demonstrate that the methodology is capable of consistently obtaining high-quality solutions that significantly improve the standard procedure in larger systems. For each of the 24, 118 and 300 bus systems, we considered 100 datasets, each of which with 1,000 observations for the training stage. Each instance considered a time limit of 15 minutes for training, i.e., for solving (21)–(38) with $T = 1,000$. For the testing stage, we evaluated the out-of-sample cost using a common set of 10,000 observations for each of the 100 obtained

solutions. In Table 1, we present the mean and standard deviation of the 100 estimates of the operation costs for both the test and training stages.

As can be seen in Table 1, some training times are slightly above 15 minutes because of problem-building time. We can see improvements of 5% in the 24 bus system and 3% in the 118 system. In the 300 bus test case, we have a 4.8% improvement by endogenously sizing reserves and an 11.4% improvement by jointly estimating the closed-loop demand forecast and reserve size. The latter case hit the time limit frequently, hence results could be even better with more time.

System	Model	Test Cost(\$)		Train Cost (\$)		Train Time (s)	
		Mean	Std	Mean	Std	Mean	Std
24	LS-Ex	414.70	1.14	397.01	4.39	0.00	0.00
	LS-Opt	398.48	0.69	378.49	93.00	450.81	90.35
	Opt-Opt	398.20	1.04	376.35	14.04	643.89	11.16
118	LS-Ex	2956.01	11.33	3163.74	3.78	0.00	0.00
	LS-Opt	2829.86	4.20	3041.28	4.86	639.87	4.97
	Opt-Opt	2815.85	4.57	3029.93	13.20	911.43	12.37
300	LS-Ex	7697.25	40.78	7646.84	26.09	0.00	0.00
	LS-Opt	7329.44	36.62	7278.88	18.75	803.15	34.08
	Opt-Opt	6820.47	37.67	6787.65	294.53	923.38	25.12

Table 1 Results for the 24, 118 and 300-bus systems

8.8. Scalability of the proposed method

Although many systems are frequently represented with networks with a few hundred buses—for instance, COES (2021) considers 255 buses in Peru and CND (2021) considers 140 buses in Panama—we further analyze the algorithm’s scalability and its consequent applicability to an even broader range of very large power system networks. We created instances with the number of buses ranging from 600 to 6000. These instances were created by connecting multiple copies of the 300 bus case. The optimization was performed on an Intel Xeon E5-2680 with 12 cores at 2.50GHz, 128Gb RAM. We generated a single training dataset with 1,000 observations for each instance and optimized each problem considering four- and twelve-hour time limits. Then, we evaluated

the solutions obtained with each method and computational time limit with a common dataset of 10,000 out-of-sample observations.

The results, in terms of Test and Train cost, are depicted in Table 2. It is clear that the proposed methods (LS-Opt and Opt-Opt) consistently outperformed the least-squares benchmark (LS-Ex) despite reaching the time limit for training. Moreover, the co-optimization scheme (Opt-Opt) was again better than the other two methods (LS-Ex and LS-Opt) in both in-sample and out-of-sample evaluation. Also, the method performed well, generalizing from training to testing. For the larger systems, we see smaller relative improvements in the cost functions. This is due to (i) the increased dimension of linear programs solved in each iteration, and (ii) the number of parameters, which increases with the system size. Reasons (i) and (ii) imply that, for a given computational budget, fewer iterations are run in the training stage, possibly yielding sub-optimal solutions. Thus, it is conceivable that much better results could be obtained with more time (or processing capacity), as the 12-hour runs lead to improvements that are more than three times those obtained with 4-hour runs. Although we were limited to 12 cores, one could use one core (or more) per observation in the training dataset, leading to massive speed-ups. It is worth mentioning that, in all cases, the LS estimation found coefficients very close to the true ones. This fact, together with the gains shown in Tables 1 and 2 for many instances of different sizes, indicate that the forecast bias introduced by our methodology is consistent in promoting an improved operation. Notably, relevant gains were found even in cases of very large-scale power systems. For instance, for the 3000-bus system, in this simulation run the gain was 13% with the 12-hour time limit. Therefore, the results provide strong evidence that our method is capable of producing meaningful gains in practice.

9. Tests on very large systems

In this section, we emulate a realistic usage of the method, by testing it in realistic large-scale instances. The ones marked with “*rte*” represent the French grid in 2013, using large-scale instances based on real systems (Josz et al. 2016). The ones marked with “*pegase*” are synthetic but based on the European grid (Fliscounakis et al. 2013). All these instances are available in Babaeinejad-sarookolae et al. (2019). In this section, all system data in the test cases are considered without

Buses	Time (h)	Test					Train				
		Opt-Opt		LS-Opt		LS-Ex	Opt-Opt		LS-Opt		LS-Ex
		(\$)	(%)	(\$)	(%)	(\$)	(\$)	(%)	(\$)	(%)	(\$)
600	4	15256	21.51	16741	13.87	19437	14898	22.73	16552	14.16	19281
600	12	15226	21.66	16739	13.88	19437	14873	22.86	16563	14.10	19281
1200	4	35261	26.38	40103	16.28	47899	34985	26.87	39994	16.41	47843
1200	12	35158	26.60	38843	18.91	47899	34842	27.17	38787	18.93	47843
1800	4	60043	18.61	69355	5.98	73770	61033	17.97	69915	6.03	74405
1800	12	53378	27.64	62786	14.89	73770	54294	27.03	63224	15.03	74405
2400	4	81018	20.23	94675	6.78	101561	78686	19.57	91278	6.70	97832
2400	12	80330	20.91	93739	7.70	101561	77969	20.30	90413	7.58	97832
3000	4	120389	5.18	125465	1.19	126972	119214	5.27	124328	1.20	125841
3000	12	110302	13.13	122560	3.47	126972	109231	13.20	121382	3.54	125841
3600	4	149141	3.51	153484	0.69	154558	147110	3.50	151404	0.68	152439
3600	12	136479	11.70	150303	2.75	154558	134664	11.66	148332	2.69	152439
4200	4	177451	1.72	179785	0.43	180555	174395	1.70	176651	0.43	177406
4200	12	165963	8.08	177444	1.72	180555	162820	8.22	174309	1.75	177406
4800	4	206358	1.22	208300	0.29	208910	209816	1.28	211927	0.29	212546
4800	12	197707	5.36	206539	1.13	208910	201294	5.29	210197	1.11	212546
5400	4	232071	1.05	233992	0.23	234524	228660	1.23	231018	0.21	231514
5400	12	225203	3.97	232658	0.80	234524	222018	4.10	229769	0.75	231514
6000	4	260427	0.96	262493	0.18	262963	262049	0.99	264198	0.18	264668
6000	12	255384	2.88	261500	0.56	262963	257062	2.87	263118	0.59	264668

Table 2 Results for very large-scale systems. LS-Ex is the reference, only costs in \$ are shown. For both Opt-Opt and LS-Opt costs are shown in \$ and their improvement compared to LS-Ex is shown in %.

any modifications. Moreover, we limited the computing time to 30 minutes so that the method can be used for dynamic reserve sizing and load forecast in hour-ahead dispatch. Note that using the method in other time scales, such as day-ahead dispatch, would also have great potential to benefit real systems around the globe. Even systems with slower dynamics running week-ahead dispatch could be benefited. The number of observations used for training was 600, and the number of out-of-sample observations for testing was 10000. We used larger servers with 64 processes in parallel and 1024 Gb of RAM to perform the computations, all instances are available in Amazon Web Services, AWS, (Amazon 2022).

Dataset	Test					Train				
	Opt-Opt		LS-Opt		LS-Ex	Opt-Opt		LS-Opt		LS-Ex
	(\$)	(%)	(\$)	(%)	(\$)	(\$)	(%)	(\$)	(%)	(\$)
6468-rte	19870	8.56	20026	7.85	21731	20278	8.50	20462	7.67	22163
6470-rte	25579	9.21	25849	8.25	28174	24004	10.38	24235	9.52	26785
6495-rte	28726	9.77	28988	8.94	31835	28400	10.52	28666	9.68	31739
6515-rte	34397	2.01	34557	1.56	35103	33757	2.01	33874	1.55	34354
9241-pegase	76438	12.89	76662	12.62	87750	75315	10.31	75509	10.08	83973
13659-pegase	82608	7.97	82908	7.63	89760	83930	7.50	84196	7.21	90737

Table 3 Results for very large-scale systems. LS-Ex is the reference, only costs in \$ are shown. For both Opt-Opt and LS-Opt costs are shown in \$ and their improvement compared to LS-Ex is shown in %.

Results are shown in Table 3. The number in the dataset name represents the number of buses in the system. These results provide even stronger evidence that our method is capable of producing meaningful gains in practice.

10. Conclusions

A mathematical framework for application-driven learning is proposed building upon the ideas of bilevel optimization. Asymptotic convergence is demonstrated, and we show that the proposed solution converges to the best forecast model on average for the selected application. Two solution techniques are presented: one exact, which is based on the KKT conditions of the second-level problem, and one heuristic, which scales well for very large instances while showcasing relevant gains in comparison to the benchmark. We apply the general application-driven framework to jointly estimate the parameters of a load and dynamic reserve requirements forecast model in a closed-loop fashion. Following the same framework, we also presented a novel (closed-loop) method to estimate dynamic reserves if the demand forecast is fixed. The proposed methods are compared with the classical sequential open-loop procedure (benchmark), where the forecast models are estimated based on least squares and used in the decision-making process. Regarding our application case, the proposed framework finds support in current industry practices, where *ad hoc* procedures are implemented to bias load forecasts aiming to reduce risks empirically. The application of our model provides not only a theoretically-grounded understanding of such procedures but also a flexible

computational tool for testing current practices and jointly determining the optimal bias and reserve requirements.

The reported numerical experience allows highlighting the following main empirical results and insights: 1) There exists an optimal bias in the load forecast maximizing the performance of a system or market operator in the long run. Moreover, the optimal bias in the load forecast is not disconnected from the optimal reserve requirements (reserve sizing problem). 2) Reserve requirements sizing and location are intrinsically dependent on the load-shed cost and system’s characteristics, and a policy to dynamically allocate them can be optimally determined by our framework even in the case where we impose an exogenous estimate for the demand forecast (e.g., least squares or other methods). 3) Our model can endogenously define the optimal reserve sizing across the network by defining zonal reserve requirements that will best perform given the system operator’s description of the network. 4) We show for realistic test systems, e.g., IEEE 300-bus system, that both models (optimizing only reserves and co-optimizing load forecast and reserves) were capable of significantly improving the long-run operation cost. Moreover, we demonstrated that the method scales to even larger power systems with thousands of buses and leads to consistent improvements in the long-run operation costs even with limited computational resources. Finally, we empirically proved that the method scales well and is able to handle large-scale realistic power systems, ranging from 6,468 to 13,659 buses, with a limited training time of 30 minutes resorting to cloud computing. Consequently, the method might be well fit for realistic hour-ahead planning besides day- and week-ahead. 5) We show that the proposed heuristic solution method can provide high-quality solutions in reasonable computational time. This is mostly due to the selected approach, which i) initializes the search method with the traditional least-square estimates, thereby only moving to another point if an improvement in performance is found, and ii) allows the estimation problem to be decomposed per observation and the second-level problem to be solved till global optimality in polynomial time. Additionally, it also leverages mature linear programming-based warm-start technologies and algorithms to scale up the performance in larger instances. This pattern is consistently observed in all test systems, corroborating the proposed framework’s effectiveness in finding improved estimates for both load and reserve requirements.

The proposed framework is fairly general, but we focused on right-hand side uncertainty and linear models as this was the relevant setting for the reserve requirement forecasting problem. Possible extensions of this work could aim at generalizing this setting. Indeed, the exact method would work for objective-function uncertainty and for non-linear models with strong duality like many conic programs. The heuristic method is even more flexible, it could accommodate uncertainty anywhere in the models, and it could be used in many non-linear (even integer) problems. Also, the heuristic method can be benefited from online optimization approaches and warm starts based on the solution of a previous execution of the algorithm with similar data as would occur in a real-time application. This could bring relevant gains to the optimization of sequential hour-ahead operation. Other aspects of estimation procedures to be explored include regularization with shrinkage operators (Tibshirani 2011), or the addition of LS estimates in the objective (Kao et al. 2009) with a penalizing coefficient to provide a balance between classical and application-driven forecasts. Moreover, it would be interesting to conduct experiments to verify the empirical performance of various statistical models (ψ), such as vector autoregressive models. Proving convergence conditions for many of the aforementioned changes in the model would provide nice contributions to the literature.

Acknowledgments

The first author was partially supported by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001. The work of the second author was also partially supported by FAPERJ and CNPq. The third author acknowledges the support of grants FONDECYT 1221770 and ACT-192094 from ANID, Chile. The authors also thank PSR for making PSRCloud available for the experiments in Section 9.

References

- Amazon (2022) EC2 amazon web services, aws, 2022. URL <https://aws.amazon.com/ec2/>.
- Aravena I, Papavasiliou A (2020) Asynchronous lagrangian scenario decomposition. *Mathematical Programming Computation* 1–50.

- Babaeinejadsarookolae S, Birchfield A, Christie RD, Coffrin C, DeMarco C, Diao R, Ferris M, Fliscounakis S, Greene S, Huang R, et al. (2019) The power grid library for benchmarking ac optimal power flow algorithms. *arXiv preprint arXiv:1908.02788* .
- Bard JF (2013) *Practical bilevel optimization: algorithms and applications*, volume 30 (Springer Science & Business Media).
- Bengio Y (1997) Using a financial training criterion rather than a prediction criterion. *International Journal of Neural Systems* 8(04):433–443.
- Bertsimas D, Kallus N (2019) From predictive to prescriptive analytics. *Management Science* .
- Bezerra B, Veiga Á, Barroso LA, Pereira M (2016) Stochastic long-term hydrothermal scheduling with parameter uncertainty in autoregressive streamflow models. *IEEE Transactions on Power Systems* 32(2):999–1006.
- Billingsley P (1986) *Probability and Measure*. Wiley Series in Probability and Statistics (Wiley), ISBN 9780471804789.
- Böhm V (1975) On the continuity of the optimal policy set for linear programs. *SIAM Journal on Applied Mathematics* 28(2):303–306.
- Borrelli F, Bemporad A, Morari M (2003) Geometric algorithm for multiparametric linear programming. *Journal of optimization theory and applications* 118(3):515–540.
- Brockwell PJ, Davis RA (2009) *Time series: theory and methods* (Springer science & business media).
- Bucksteeg M, Niesen L, Weber C (2016) Impacts of dynamic probabilistic reserve sizing techniques on reserve requirements and system costs. *IEEE Transactions on Sustainable Energy* 7(4):1408–1420.
- CAISO (2020) 2019 Annual Report on Market Issues and Performance. Department of Market Monitoring, California Independent System Operator. URL <http://www.caiso.com/Documents/2019AnnualReportonMarketIssuesandPerformance.pdf>.
- Chen Y, Gribik P, Gardner J (2013) Incorporating post zonal reserve deployment transmission constraints into energy and ancillary service co-optimization. *IEEE Transactions on Power Systems* 29(2):537–549.
- CND (2021) Centro nacional de despacho, informe de operaciones, 2021. URL <https://www.cnd.com.pa/index.php/informes/categoria/informes-de-operaciones?tipo=68&anio=2021&semana=1>.

- COES (2021) Programa mediano plazo operación, 2021, mayo, informes. URL <https://www.coes.org.pe/Portal/Operacion/ProgOperacion/ProgMedianoPO>.
- Conn AR, Scheinberg K, Vicente LN (2009) *Introduction to derivative-free optimization* (SIAM).
- De Vos K, Stevens N, Devolder O, Papavasiliou A, Hebb B, Matthys-Donnadieu J (2019) Dynamic dimensioning approach for operating reserves: Proof of concept in belgium. *Energy policy* 124:272–285.
- Dempe S (2018) *Bilevel optimization: theory, algorithms and applications* (TU Bergakademie Freiberg, Fakultät für Mathematik und Informatik).
- den Boer AV, Sierag DD (2021) Decision-based model selection. *European Journal of Operational Research* 290(2):671–686.
- Deng Y, Liu J, Sen S (2018) Coalescing data and decision sciences for analytics. *INFORMS TutORials in Operations Research*, 20–49 (INFORMS).
- Dias Garcia J, Bodin G, Street A (2022) BilevelJuMP.jl: Modeling and Solving Bilevel Optimization in Julia. *arXiv preprint arXiv:2205.02307* .
- Domowitz I, El-Gamal MA (1993) A consistent test of stationary-ergodicity. *Econometric Theory* 9(4):589–601.
- Donti P, Amos B, Kolter JZ (2017) Task-based end-to-end model learning in stochastic optimization. *Advances in Neural Information Processing Systems*, 5484–5494.
- Dunning I, Huchette J, Lubin M (2017) Jump: A modeling language for mathematical optimization. *SIAM Review* 59(2):295–320.
- El Balghiti O, Elmachtoub A, Grigas P, Tewari A (2019) Generalization bounds in the predict-then-optimize framework. *Advances in Neural Information Processing Systems*, 14389–14398.
- Ela E, Milligan M, Kirby B (2011) Operating reserves and variable generation. Technical Report NREL/TP-5500-51978, National Renewable Energy Lab. (NREL), Golden, CO (United States).
- Elmachtoub AN, Grigas P (2021) Smart “predict, then optimize”. *Management Science* URL <https://doi.org/10.1287/mnsc.2020.3922>.

- Fliscounakis S, Panciatici P, Capitanescu F, Wehenkel L (2013) Contingency ranking with respect to overloads in very large power systems taking into account uncertainty, preventive, and corrective actions. *IEEE Transactions on Power Systems* 28(4):4909–4917.
- Fortuny-Amat J, McCarl B (1981) A representation and economic interpretation of a two-level programming problem. *Journal of the operational Research Society* 32(9):783–792.
- Franceschi L, Frasconi P, Salzo S, Grazzi R, Pontil M (2018) Bilevel programming for hyperparameter optimization and meta-learning. *International Conference on Machine Learning*, 1568–1577 (PMLR).
- Frazier PI (2018) Bayesian optimization. *INFORMS TutORials in Operations Research*, 255–278 (Informs).
- Gade D, Hackebeil G, Ryan SM, Watson JP, Wets RJB, Woodruff DL (2016) Obtaining lower bounds from the progressive hedging algorithm for stochastic mixed-integer programs. *Mathematical Programming* 157(1):47–67.
- Gal T (2010) *Postoptimal Analyses, Parametric Programming, and Related Topics: degeneracy, multicriteria decision making, redundancy* (Walter de Gruyter).
- Garcia R, Gençay R (2000) Pricing and hedging derivative securities with neural networks and a homogeneity hint. *Journal of Econometrics* 94(1-2):93–115.
- Ghosh J, Bengio Y (1997) Multi-task learning for stock selection. *Advances in neural information processing systems*, 946–952.
- Henningsen A, Toomet O (2011) maxlik: A package for maximum likelihood estimation in r. *Computational Statistics* 26(3):443–458.
- Holttinen H, Milligan M, Ela E, Menemenlis N, Dobschinski J, Rawn B, Bessa RJ, Flynn D, Gomez-Lazaro E, Detlefsen NK (2012) Methodologies to determine operating reserves due to increased wind power. *IEEE Transactions on Sustainable Energy* 3(4):713–723.
- Hong T, Fan S (2016) Probabilistic electric load forecasting: A tutorial review. *International Journal of Forecasting* 32(3):914–938.
- Josz C, Fliscounakis S, Maeght J, Panciatici P (2016) Ac power flow data in matpower and qcqp format: itesla, rte snapshots, and pegase. *arXiv preprint arXiv:1603.01533* .

- Kao Yh, Roy BV, Yan X (2009) Directed regression. *Advances in Neural Information Processing Systems*, 889–897.
- Kazempour J, Pinson P, Hobbs BF (2018) A stochastic market design with revenue adequacy and cost recovery by scenario: Benefits and costs. *IEEE Transactions on Power Systems* 33(4):3531–3545.
- Kirschen DS, Strbac G (2018) *Fundamentals of power system economics* (John Wiley & Sons).
- Knueven B, Ostrowski J, Watson JP (2020) On mixed-integer programming formulations for the unit commitment problem. *INFORMS Journal on Computing* 32(4):857–876.
- Li H, Bose A, Venkatasubramanian VM (2015) Wide-area voltage monitoring and optimization. *IEEE Transactions on Smart Grid* 7(2):785–793.
- Liu J, Li G, Sen S (2022) Coupled learning enabled stochastic programming with endogenous uncertainty. *Mathematics of Operations Research* 47(2):1681–1705.
- Mandi J, Stuckey PJ, Guns T, et al. (2020) Smart predict-and-optimize for hard combinatorial optimization problems. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 1603–1610.
- Megiddo N, Chandrasekaran R (1989) On the ε -perturbation method for avoiding degeneracy. *Operations Research Letters* 8(6):305–308.
- Mogensen PK, Riseth AN (2018) Optim: A mathematical optimization package for julia. *Journal of Open Source Software* 3(24).
- Muñoz MA, Pineda S, Morales JM (2022) A bilevel framework for decision-making under uncertainty with contextual information. *Omega* 108:102575.
- Nocedal J, Wright S (2006) *Numerical optimization* (Springer Science & Business Media).
- Orwig KD, Ahlstrom ML, Banunarayanan V, Sharp J, Wilczak JM, Freedman J, Haupt SE, Cline J, Bartholomy O, Hamann HF, et al. (2014) Recent trends in variable generation forecasting and its value to the power system. *IEEE Transactions on Sustainable Energy* 6(3):924–933.
- Papavasiliou A, Oren SS (2013) Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network. *Operations Research* 61(3):578–592.

- Papavasiliou A, Oren SS, Rountree B (2014) Applying high performance computing to transmission-constrained stochastic unit commitment for renewable energy integration. *IEEE Transactions on Power Systems* 30(3):1109–1120.
- Pereira MV, Granville S, Fampa MH, Dix R, Barroso LA (2005) Strategic bidding under uncertainty: a binary expansion approach. *IEEE Transactions on Power Systems* 20(1):180–188.
- PJM (2018) PJM Manual 11 : Energy and Ancillary Services Market Operations (97):200, URL <https://pjm.com/-/media/documents/manuals/archive/m11/m11v97-energy-and-ancillary-services-market-operations-07-26-2018.ashx>.
- Powell WB (2010) The knowledge gradient for optimal learning. *Wiley Encyclopedia of Operations Research and Management Science* .
- Powell WB, Frazier P (2008) Optimal learning. *INFORMS TutORials in Operations Research*, 213–246 (InformS).
- Powell WB, Ryzhov IO (2012) *Optimal learning*, volume 841 (John Wiley & Sons).
- Rockafellar RT, Uryasev S, Zabarankin M (2008) Risk tuning with generalized linear regression. *Mathematics of Operations Research* 33(3):712–729.
- Ryzhov IO, Powell WB (2012) Information collection for linear programs with uncertain objective coefficients. *SIAM Journal on Optimization* 22(4):1344–1368.
- Sen S, Deng Y (2018) Learning enabled optimization: Towards a fusion of statistical learning and stochastic programming. *Optimization Online* .
- Shapiro A, Dentcheva D, Ruszczyński A (2014) *Lectures on stochastic programming : modeling and theory* (SIAM), 2nd edition.
- Siddiqui S, Gabriel SA (2013) An sos1-based approach for solving mpecs with a natural gas market application. *Networks and Spatial Economics* 13(2):205–227.
- Strbac G, Shakoor A, Black M, Pudjianto D, Bopp T (2007) Impact of wind generation on the operation and development of the uk electricity systems. *Electric Power Systems Research* 77(9):1214–1227.
- Sweeney C, Bessa RJ, Browell J, Pinson P (2020) The future of forecasting for renewable energy. *Wiley Interdisciplinary Reviews: Energy and Environment* 9(2):e365.

- The European Commission (2017) Commission regulation (eu) 2017/1485 - establishing a guideline on electricity transmission system operation.
- <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32017R1485&from=EN>.
- Tibshirani R (2011) Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73(3):273–282.
- Van der Meer DW, Widén J, Munkhammar J (2018) Review on probabilistic forecasting of photovoltaic power production and electricity consumption. *Renewable and Sustainable Energy Reviews* 81:1484–1512.
- Varian H (1975) A bayesian approach to real estate assessment. Fienberg SE, Zellner A, eds., *Studies in Bayesian Econometrics and Statistics in Honor of Leonard J. Savage* (Amsterdam: North-Holland).
- Wang B, Hobbs BF (2014) A flexible ramping product: Can it help real-time dispatch markets approach the stochastic dispatch ideal? *Electric Power Systems Research* 109:128–140.
- Wang B, Hobbs BF (2015) Real-time markets for flexiramp: A stochastic unit commitment-based analysis. *IEEE Transactions on Power Systems* 31(2):846–860.
- White H (2014) *Asymptotic theory for econometricians* (Academic press).
- Wilder B, Dilkina B, Tambe M (2019) Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1658–1665.
- Zellner A (1986a) Bayesian estimation and prediction using asymmetric loss functions. *Journal of the American Statistical Association* 81(394):446–451.
- Zellner A (1986b) Biased predictors, rationality and the evaluation of forecasts. *Economics Letters* 21(1):45–48.
- Zheng QP, Wang J, Liu AL (2014) Stochastic optimization for unit commitment—a review. *IEEE Transactions on Power Systems* 30(4):1913–1924.