

Real-Time Risk Analysis with Optimization Proxies

Wenbo Chen, Mathieu Tanneau, Pascal Van Hentenryck
 Georgia Institute of Technology
 wenbo.chen@gatech.edu, {mathieu.tanneau, pascal.vanhentenryck}@isye.gatech.edu

Abstract—The increasing penetration of renewable generation and distributed energy resources requires new operating practices for power systems, wherein risk is explicitly quantified and managed. However, traditional risk-assessment frameworks are not fast enough for real-time operations, because they require numerous simulations, each of which requires solving multiple economic dispatch problems sequentially. The paper addresses this computational challenge by proposing proxy-based risk assessment, wherein optimization proxies are trained to learn the input-to-output mapping of an economic dispatch optimization solver. Once trained, the proxies make predictions in milliseconds, thereby enabling real-time risk assessment. The paper leverages self-supervised learning and end-to-end-feasible architecture to achieve high-quality sequential predictions. Numerical experiments on large systems demonstrate the scalability and accuracy of the proposed approach.

Index Terms—Risk assessment, optimization proxies

I. INTRODUCTION

The growing penetration of distributed energy resources (DERs) and of renewable generation, especially wind and solar generation, is causing a fundamental change of paradigm for power systems operations. Traditional approaches, where uncertainty is low and can be managed through reserve products, are no longer viable, in presence of significant wind and solar generation, and less predictable demand patterns due to DER adoption. Modern power grids thus require new operating practices that explicitly assess operational risk in real time.

Traditional approaches to risk assessment require running large numbers (typically in the thousands) of Monte-Carlo (MC) simulations. Risk and reliability metrics are then computed based on the outputs of these simulations [1]. In power systems, each MC simulation evaluates the behavior of the grid, e.g., generator setpoints and power flows, under a possible realization of load and renewable output, obtained from solving multiple economic dispatch problems sequentially (one per time step). Nevertheless, the computational cost of optimization, combined with the large number of MC simulations needed to obtain an accurate risk estimate, renders this workflow intractable for real-time risk assessment.

The paper proposes a new methodology, Risk Assessment with End-to-End Learning and Repair (RA-E2ELR), to address this computational challenge. RA-E2ELR replaces costly optimization solvers with fast E2ELR optimization proxies, i.e., differentiable programs that are guaranteed to

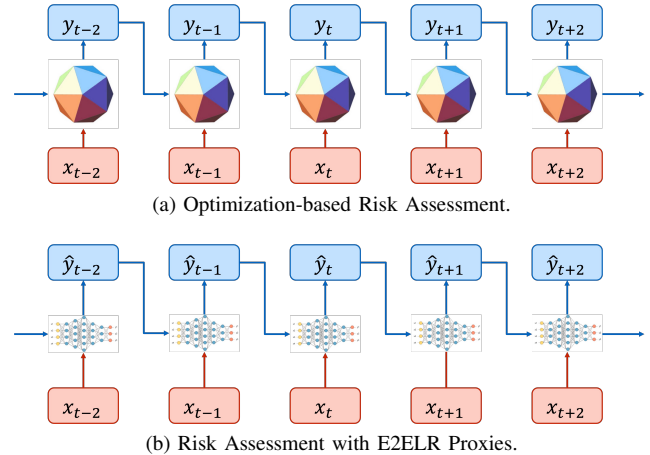


Fig. 1. Illustration of the optimization-based (top) and RA-E2ELR MC simulation pipeline. At each step t , generator setpoints y_t are decided based on current system conditions x_t and previous setpoints y_{t-1} .

produce near-optimal feasible solutions to economic dispatch problems, [2]. RA-E2ELR is sketched in Figure 1 and is capable of assessing risk in electricity markets whose real-time operations are organized around economic dispatch processes that co-optimize energy and reserves. The paper reports numerical experiments on large-scale power grids with thousands of buses, demonstrating the scalability and fidelity of RA-E2ELR. To the best of the authors’ knowledge, RA-E2ELR is the first methodology to real-time risk assessment for the US market-clearing pipeline.

II. LITERATURE REVIEW

A. Risk Assessment in Power Systems

A number of reliability metrics have been proposed to quantify the risk in power systems explicitly such as Lack of ramp probability (LORP) [3], insufficient ramping resource expectation (IRRE) [4], loss of load expectation (LOLE) and expected unserved energy (EUE) [5]. More recently, Stover & al [1] proposed three levels of risk metrics for a quantity of interest: conditional value at risk, reliability, and risk. Computing those risk metrics is based on Monte Carlo (MC) sampling, which requires solving optimization instances for thousands of scenarios. These MC simulations are computationally expensive and intractable in real-time settings. This paper shows how to compute, in real-time, high-fidelity approximations of these risk metrics by leveraging E2ELR optimization proxies. In particular, the paper focuses on the risk that stems from net load uncertainty. This approach

This research was partly supported by NSF award 2112533 and ARPA-E PERFORM award AR0001136.

Model 1 The Economic Dispatch Model [2].

Input: Current load \mathbf{d}_t , previous generation dispatch \mathbf{p}_{t-1}

Output: Generation dispatch \mathbf{p}_t

$$\min_{\mathbf{p}_t} c^\top \mathbf{p}_t + M_{\text{th}} \|\xi_{\text{th}}\|_1 \quad (1a)$$

$$\text{s.t. } \mathbf{e}^\top \mathbf{p}_t = \mathbf{e}^\top \mathbf{d}_t \quad (1b)$$

$$\underline{\mathbf{p}} \leq \mathbf{p}_t \leq \bar{\mathbf{p}} \quad (1c)$$

$$\mathbf{p}_{t-1} - R^\downarrow \leq \mathbf{p}_t \leq \mathbf{p}_{t-1} + R^\uparrow \quad (1d)$$

$$\underline{\mathbf{f}} - \xi_{\text{th}} \leq \Phi(\mathbf{p}_t - \mathbf{d}_t) \leq \bar{\mathbf{f}} + \xi_{\text{th}} \quad (1e)$$

is distinct and complementary to the contingency-based risk perspective associated with security [6], [7], [8], [9].

B. Optimization Proxies for Sequential OPF

Some recent papers [10], [11], [12], [13], [14] focus on developing optimization proxies for multi-period ACOPF problems. They model the proxies with deep neural networks, typically combined with repair steps for feasibility. The proxies are trained under the framework of Reinforcement Learning (RL) with different algorithms. For instance, in [10], proxies are pretrained with imitation learning and then improved using proximal policy optimization. The results are reported for systems with up to 200 buses. In [11], the authors train proxies using deep deterministic policy gradient. The reward function is augmented with constraint violations to improve the feasibility. The results are reported on the 118-bus system. *Those methodologies cannot guarantee feasibility and thus cannot be used for risk assessments.*

Subsequent results addressed feasibility issues by designing complex optimization-based feasibility layers such as power flow solvers and convex-concave procedures [12], or holomorphic embedding [13]; they also use more advanced RL algorithms such as the soft actor-critic framework. Because of the complexity of their optimization-based repair steps, *the improvements due to the proxies compared to the optimization algorithms are not sufficient for real-time risk assessment.* The results are only reported up to systems with 300 buses.

III. OPERATIONAL RISK ASSESSMENT IN POWER GRIDS

Consider a power grid with buses $\mathcal{N} = \{1, \dots, N\}$, and branches $\mathcal{E} = \{1, \dots, E\}$. Assume, without loss of generality, that the costs are linear and that exactly one generator and one load are attached to each bus.

A. The Economic Dispatch Formulation

Model 1 presents the Economic Dispatch (ED) model to be solved at time step $t \in \mathbb{Z}$. For simplicity, the formulation omits the reserve constraints but RA-E2ELR has no difficulty in handling them [2]. The duration of each time step is assumed to be 1. In the US, real-time electricity markets are typically cleared every 5 minutes by solving a security-constrained economic dispatch (SCED) problem, which extends Model 1 to ensure N-1 security. The ED problem at time t takes, as inputs, the current demand \mathbf{d}_t and the prior generation

Algorithm 1 Monte-Carlo Simulation

Require: Demand scenario $(\mathbf{d}_1^s, \dots, \mathbf{d}_T^s)$

1: **for** $t = 1 \dots T$ **do**

2: $\bar{\mathbf{p}}_t^s, \underline{\mathbf{p}}_t^s \leftarrow \max(\underline{\mathbf{p}}, \mathbf{p}_{t-1}^s - R^\downarrow), \min(\bar{\mathbf{p}}, \mathbf{p}_{t-1}^s + R^\uparrow)$

3: $\mathbf{p}_t^s \leftarrow \text{ED}(\mathbf{d}_t^s, \bar{\mathbf{p}}_t^s, \underline{\mathbf{p}}_t^s)$

4: **end for**

5: **Return** $(\mathbf{p}_1^s, \dots, \mathbf{p}_T^s)$

dispatch \mathbf{p}_{t-1} , and outputs the generation dispatch \mathbf{p}_t . For generators, the minimum output, maximum output, upward and downward ramping capacity are denoted by $\underline{\mathbf{p}}$, $\bar{\mathbf{p}}$, R^\uparrow and R^\downarrow , respectively. The Power Transfer Distribution Factor (PTDF) matrix is denoted by $\Phi \in \mathbb{R}^{E \times N}$, and the lower and upper thermal limits are denoted by $\underline{\mathbf{f}}$ and $\bar{\mathbf{f}}$.

The ED objective (1a) minimizes the sum of total generation costs and thermal violation penalties. Constraint (1b) enforces power balance in the system. Constraints (1c) enforce generation minimum and maximum limits, and constraints (1d) enforce ramping constraints. Constraints (1c)–(1d) can be combined into

$$\underline{\mathbf{p}}_t \leq \mathbf{p}_t \leq \bar{\mathbf{p}}_t, \quad (2)$$

where $\underline{\mathbf{p}}_t = \max(\underline{\mathbf{p}}, \mathbf{p}_{t-1} - R^\downarrow)$ and $\bar{\mathbf{p}}_t = \min(\bar{\mathbf{p}}, \mathbf{p}_{t-1} + R^\uparrow)$. Finally, constraints (1e) express thermal constraints using a PTDF formulation. Following the practice of US system operators [15], thermal constraints are treated as soft, i.e., thermal violations ξ_{th} are penalized using a high price M_{th} .

B. The Operational Risk Assessment Framework

Principled risk assessment is based on executing large numbers of MC simulations [1]. Each individual MC simulation emulates the system behavior under given (net) load conditions. The results of multiple MC simulations are then aggregated into risk and reliability metrics.

Formally, consider a simulation horizon $\mathcal{T} = \{1, \dots, T\}$. In a day-ahead setting, \mathcal{T} may correspond to each hour of the operating day; in a real-time setting, \mathcal{T} may capture the subsequent few hours at a 5-minute granularity. A set of scenarios $\mathcal{S} = \{1, \dots, S\}$, typically generated by a probabilistic forecasting model, is available. Algorithm 1 formalizes the simulation algorithm, which is also illustrated in Figure 1a. Each simulation s takes as input a sequence of load $(\mathbf{d}_1^s, \dots, \mathbf{d}_T^s)$, where \mathbf{d}_t^s denotes the vector of loads at time t and for scenario s . At each time t , the minimum and maximum limits $\underline{\mathbf{p}}_t^s, \bar{\mathbf{p}}_t^s$ are updated. Then, current generation dispatches \mathbf{p}_t^s are computed by solving the ED problem in Model 1. Finally, the simulation returns the sequence of generation dispatches $(\mathbf{p}_1^s, \dots, \mathbf{p}_T^s)$.

Risk metrics are computed by aggregating the outputs of multiple MC simulations, typically hundreds or thousands. This work considers the three levels of risk metrics proposed in [1]: conditional value at risk, probability of failure, and risk.

Let \mathcal{Q} be a so-called *quantity of interest* (QoI), whose value depends on the state of the system. Namely, denote by

$$Q_t^s = \mathcal{Q}(\mathbf{d}_t^s, \mathbf{p}_t^s) \quad (3)$$

the QoI value at time t in scenario s . For instance, if \mathcal{Q} is the system power imbalance, then $Q_t^s = \mathbf{e}^\top (\mathbf{p}_t^s - \mathbf{d}_t^s)$.

1) *Conditional Value at Risk (CVAR)*: Let $\alpha \in [0, 1]$ be a given significance level, typically chosen by a system operator. The CVAR metric for time t is

$$\text{CE}(\mathcal{Q}_t) = \mathbb{E}[\mathcal{Q}_t | \mathcal{Q}_t \geq Q_t(\alpha)] \approx \frac{\sum_{s \in \mathcal{S}} Q_t^s \cdot \mathbb{1}_{Q_t^s \geq Q_t(\alpha)}}{\sum_{s \in \mathcal{S}} \mathbb{1}_{Q_t^s \geq Q_t(\alpha)}}, \quad (4)$$

where $\mathbb{1}$ denotes the indicator function and $Q_t(\alpha)$ denote the α quantile of (Q_t^1, \dots, Q_t^S) . Note that the CVAR metric can be adapted to capture the left tail of the distribution [1].

2) *Probability of Failure*: To assess system reliability, the second metric considers the probability of failure. Namely, let \bar{Q} be a desired reliability threshold, i.e., the system is considered to fail if $\mathcal{Q} \geq \bar{Q}$. Then, the probability of failure at time t is given by

$$\mathbb{P}(\mathcal{Q}_t \geq \bar{Q}) \approx \frac{1}{S} \sum_{s \in \mathcal{S}} \mathbb{1}_{Q_t^s \geq \bar{Q}}. \quad (5)$$

Note that reliability can be measured for the entire system, or for individual components.

3) *Risk*: Risk quantifies the monetary consequences of failure. The risk at time t is given by

$$\text{Risk}(\mathcal{Q}_t) = \mathbb{E}[c(\mathcal{Q}_t)] \approx \frac{1}{S} \sum_{s \in \mathcal{S}} c(Q_t^s), \quad (6)$$

where c denotes a cost function. For instance, if the QoI \mathcal{Q} denotes system power imbalance, a natural choice for c could be $c = \text{VOLL} \times \mathcal{Q}$, where VOLL is the Value of Lost Load. Compared to the previous two metrics, risk captures the magnitude of adverse events and their financial consequences.

IV. RISK ASSESSMENT WITH RA-E2ELR

The main limitation of existing risk assessment frameworks is their computational cost. Indeed, running thousands of simulations requires solving tens to hundreds of thousands of ED problems. This approach is not tractable, especially as more resources are integrated into the grid and the complexity of ED problems increases in the future.

A. Proxy-Based Risk Assessment

To enable fast risk assessment, the paper replaces costly optimization solvers with fast optimization proxies, i.e., *differentiable programs that approximate the input-output mapping of optimization solvers*. Figure 1b illustrates the proposed risk assessment framework, which replaces the ED optimization in step 3 of Algorithm 1 with an optimization proxy. At each time step t , the ED proxy takes as input the ED problem data \mathbf{x}_t , and outputs a near-optimal solution $\hat{\mathbf{y}}_t$ in milliseconds. It is important to note that, in the present setting, the proxy

prediction at time t is part of the input of the subsequent prediction. The sequential nature of the prediction task is in stark contrast with existing literature on optimization proxies, which consider individual ED instances.

The proxy-based risk assessment task can therefore be stated as follows. Given scenario s with load $(\mathbf{d}_1^s, \dots, \mathbf{d}_T^s)$, the model should output a sequence $(\hat{\mathbf{p}}_1^s, \dots, \hat{\mathbf{p}}_T^s)$ of generator dispatches that is as close as possible to the sequence $(\mathbf{p}_1^s, \dots, \mathbf{p}_T^s)$ produced by an optimization solver.

Risk metrics can then be evaluated based on the predicted sequence $(\hat{\mathbf{p}}_1^s, \dots, \hat{\mathbf{p}}_T^s)$. Given a QoI \mathcal{Q} , let \hat{Q}_t^s denote the value of the QoI in scenario s at time t given the predicted generation dispatches. Eqs. (4), (5), (6) then become

$$\text{CE}(\mathcal{Q}_t) \approx \frac{\sum_{s \in \mathcal{S}} \hat{Q}_t^s \cdot \mathbb{1}_{\hat{Q}_t^s \geq \hat{Q}_t(\alpha)}}{\sum_{s \in \mathcal{S}} \mathbb{1}_{\hat{Q}_t^s \geq \hat{Q}_t(\alpha)}}, \quad (7)$$

$$\mathbb{P}(\mathcal{Q}_t \geq \bar{Q}) \approx \frac{1}{S} \sum_{s \in \mathcal{S}} \mathbb{1}_{\hat{Q}_t^s \geq \bar{Q}}, \quad (8)$$

$$\text{Risk}(\mathcal{Q}_t) = \mathbb{E}[c(\mathcal{Q}_t)] \approx \frac{1}{S} \sum_{s \in \mathcal{S}} c(\hat{Q}_t^s), \quad (9)$$

where $\hat{Q}_t(\alpha)$ denotes the α quantile of $(\hat{Q}_t^1, \dots, \hat{Q}_t^S)$.

Observe that the proposed methodology predicts *individual generation dispatches*, resulting in a risk assessment for *individual components* compared to existing coarser-grained proposals, e.g., the zonal approach proposed in [16].

The proxies considered in this paper all extend the baseline DNN architecture shown in Figure 2. They differ in how they handle constraint violations. The baseline DNN consists of encoders, hidden layers, and a decoder. The DNN takes as input the load vector \mathbf{d} and minimum/maximum limits $\underline{\mathbf{p}}, \bar{\mathbf{p}}$, which are each fed to an encoder. The encoders are small multi-layer perceptron (MLP) models that use fully-connected layers with Rectified Linear Unit (ReLU) activation

$$\text{ReLU}(x) = \max(0, x).$$

The output of the encoders, referred to as embeddings, are passed to hidden layers, which are also fully connected layers with ReLU activations. A sigmoid activation is applied to the output of the last hidden layer, which yields a vector $\mathbf{z} \in [0, 1]^N$. The sigmoid function is defined as

$$\sigma(x) = \frac{1}{1 + e^{-x}} \in [0, 1]. \quad (10)$$

The predicted dispatch $\hat{\mathbf{p}}$ is then obtained as

$$\hat{\mathbf{p}} = \mathbf{z} \cdot \underline{\mathbf{p}} + (1 - \mathbf{z}) \cdot \bar{\mathbf{p}}, \quad (11)$$

where all operations are element-wise. This approach ensures that $\underline{\mathbf{p}} \leq \hat{\mathbf{p}} \leq \bar{\mathbf{p}}$, which also ensures the feasibility of the ramping constraints due to the structure of Algorithm 7. However, the predicted dispatch may not satisfy power balance, i.e., $\mathbf{e}^\top \hat{\mathbf{p}} = \mathbf{e}^\top \mathbf{d}$ may not hold.

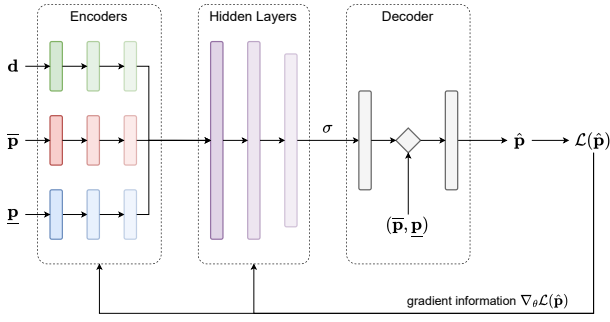


Fig. 2. The Baseline DNN architecture. The hidden layers output $\mathbf{z} \in [0, 1]^N$ using sigmoid activation σ , which is combined with min/max limits $\underline{\mathbf{p}}, \bar{\mathbf{p}}$ to ensure $\underline{\mathbf{p}} \leq \hat{\mathbf{p}} \leq \bar{\mathbf{p}}$. The DeepOPF, DNN and E2ELR architectures extend the DNN architecture by adding a feasibility mechanism after the decoder.

B. The E2ELR Optimization Proxy

The E2ELR architecture [2] enforces min/max limits, power balance constraints, and reserve constraints via dedicated, differentiable feasibility layers. For the purpose of this paper, it is sufficient to present the repair layer for power balance, which is inserted between the approximation $\hat{\mathbf{p}}$ and the computation of the loss function. This repair layer (see [2]) takes $\hat{\mathbf{p}}$ as input and outputs

$$\tilde{\mathbf{p}} = \begin{cases} (1 - \eta^\uparrow)\hat{\mathbf{p}} + \eta^\uparrow\bar{\mathbf{p}} & \text{if } \mathbf{e}^\top \hat{\mathbf{p}} < \mathbf{e}^\top \mathbf{d} \\ (1 - \eta^\downarrow)\hat{\mathbf{p}} + \eta^\downarrow\underline{\mathbf{p}} & \text{if } \mathbf{e}^\top \hat{\mathbf{p}} \geq \mathbf{e}^\top \mathbf{d} \end{cases} \quad (12)$$

where $\underline{\mathbf{p}} \leq \hat{\mathbf{p}} \leq \bar{\mathbf{p}}$, and $\eta^\uparrow, \eta^\downarrow$ are defined as follows:

$$\eta^\uparrow = \frac{\mathbf{e}^\top \mathbf{d} - \mathbf{e}^\top \hat{\mathbf{p}}}{\mathbf{e}^\top \bar{\mathbf{p}} - \mathbf{e}^\top \hat{\mathbf{p}}}, \quad \eta^\downarrow = \frac{\mathbf{e}^\top \hat{\mathbf{p}} - \mathbf{e}^\top \mathbf{d}}{\mathbf{e}^\top \hat{\mathbf{p}} - \mathbf{e}^\top \underline{\mathbf{p}}}. \quad (13)$$

The repaired dispatch $\tilde{\mathbf{p}}$ is feasible if and only if ED is feasible (see Theorem 1 in [2]).

C. The Training Methodology

The training of E2ELR is based on self-supervised learning [17], [2], which directly minimizes the objective function (1a). Consider a dataset with M data points

$$\mathcal{D} = \{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^M\}, \quad (14)$$

where $\mathbf{x}^t = \{\mathbf{d}^t, \bar{\mathbf{p}}^t, \underline{\mathbf{p}}^t\}$. Let f_Θ be the machine learning model with trainable parameters Θ . The training of f_Θ amounts to solving the optimization problem

$$\Theta^* = \arg\min_{\Theta} \frac{1}{M} \sum_{t=1}^M \mathcal{L}(f_\Theta(\mathbf{x}^t)), \quad (15)$$

In E2ELR, the loss function is the objective function of ED:

$$\mathcal{L}(\hat{\mathbf{p}}) = c(\hat{\mathbf{p}}) + M_{\text{th}} \|\xi_{\text{th}}(\hat{\mathbf{p}})\|_1. \quad (16)$$

Observe that E2ELR only uses a set of inputs for training: it does not require labeled data.

V. EXPERIMENTAL RESULTS

A. The Machine Learning Architectures

A key advantage of RA-E2ELR is its ability to produce ED-feasible solutions quickly at each stage of the MC simulations. To demonstrate these benefits, the paper compares RA-E2ELR to other ML architectures, which offer no feasibility guarantees. To train these infrastructures, the loss function should be generalized to $\mathcal{L}(\hat{\mathbf{p}}) = \varphi(\hat{\mathbf{p}}) + \lambda\psi(\hat{\mathbf{p}})$, where the first term $\varphi(\hat{\mathbf{p}})$ is the objective value of ED

$$\varphi(\hat{\mathbf{p}}) = c(\hat{\mathbf{p}}) + M_{\text{th}} \|\xi_{\text{th}}(\hat{\mathbf{p}})\|_1 \quad (17)$$

and the second term $\psi(\hat{\mathbf{p}})$ penalizes constraint violations. For instance, the power balance violations can be formalized as

$$\psi(\hat{\mathbf{p}}) = M_{\text{pb}} |\mathbf{e}^\top \mathbf{d} - \mathbf{e}^\top \hat{\mathbf{p}}|, \quad (18)$$

where M_{pb} are penalty coefficients.

The evaluation considers three ML architectures to show the benefits of RA-E2ELR: the baseline DNN architecture presented earlier, Deep-OPF [18], and DC3 [19].

1) *DeepOPF*: The DeepOPF architecture [18] extends the DNN model by adding a equality completion step that enforces system power balance (1b). Let $\hat{\mathbf{p}}$ denote the output of DNN and assume bus 1 is the slack bus without loss of generality. The equality completion layer outputs $\tilde{\mathbf{p}} \in \mathbb{R}^N$, where

$$\tilde{\mathbf{p}}_1 = \mathbf{e}^\top \mathbf{d} - \sum_{i>1} \hat{\mathbf{p}}_i, \quad \tilde{\mathbf{p}}_i = \hat{\mathbf{p}}_i, \quad \forall i > 1, \quad (19)$$

thus ensuring the satisfaction of the power balance constraint (1b). However, the dispatch of the slack generator $\tilde{\mathbf{p}}_1$ is not guaranteed to respect its minimum and maximum limits.

2) *Deep Constraint Completion and Correction (DC3)*: The DC3 architecture [19] uses the same equality completion mechanism as DeepOPF, with an additional inequality correction step. The latter mitigates potential violations of the min/max limits for the slack generator, and uses unrolled gradient descent mechanism that minimizes the violation of inequality constraints. The number of gradient steps in the correction mechanism is set to a fixed value K . If K is large enough, then DC3 returns feasible solutions, albeit at the cost of longer inference times. A smaller K improves inference time, but may yield constraint violations.

B. Data Generation

Experiments are carried on the IEEE 300-bus system (ieee300) [20] and the Pegase 1354-bus system (pegase1k) [21], using data from the UnitCommitment (UC) package [22]. The experiments use the network topology and cost information in the original test cases. Load time series and the minimum/maximum limits and ramping rates of the generators are obtained from [22]. For ease of presentation, all generators are assumed to be online; the proposed framework naturally extends to varying commitment decisions. For each system, [22] provides 365 benchmark instances, one per day in 2017, with each instance spanning a horizon of $T=36$ hours. This initial dataset of 365 load profiles is augmented

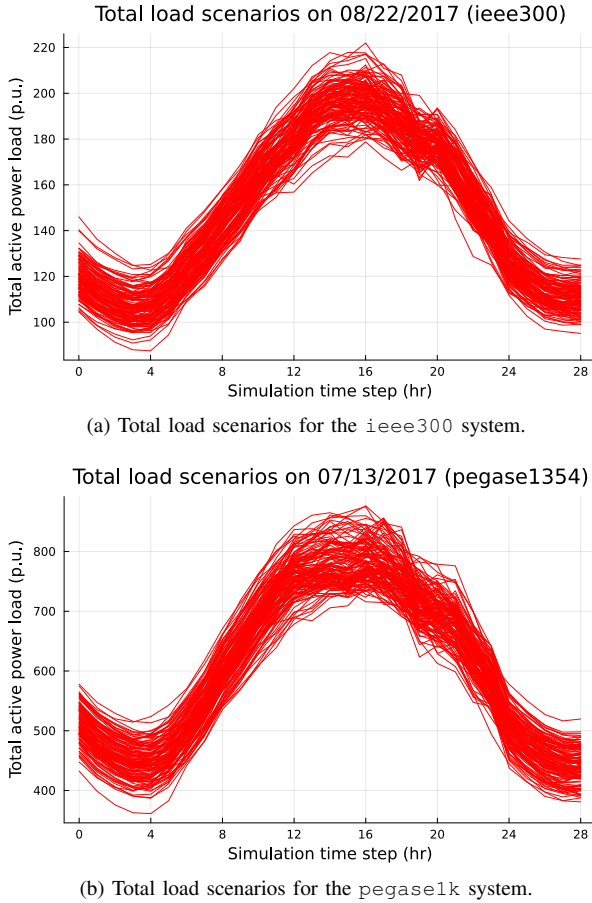


Fig. 3. Total load scenarios used in the experiments. 128 scenarios are depicted for a 28-hour simulation window starting at midnight of the first day (time step 0), and ending at 4am of the subsequent day (time step 28).

by generating addition load samples, using the methodology proposed in [23].

For a given system and a given day, bus-level load scenarios are of the form $\mathbf{d}_{i,t} = \xi_{i,t} \times (\gamma_t \times \bar{P}^d)$, where \bar{P}^d denotes the maximum total load in the reference UC instance, $\gamma \in \mathbb{R}^T$ captures the distribution of load-to-peak ratio, and $\xi_{i,t}$ denotes bus-level scaling factors. The distribution of γ is estimated from the original 365 instances using a Gaussian copula model with Beta marginal distributions. Note that, in the instances provided by [22], each bus i always accounts for a fixed proportion ξ_i^{ref} of total demand. Therefore, nodal disaggregation factors ξ are of the form $\xi_{i,t} = \xi_i^{\text{ref}} \times \eta_{i,t}$, where $\eta_{i,t}$ is uncorrelated white noise with log-normal distribution of mean 1 and 5% standard deviation.

Figure 3 depicts daily total load scenarios for the `ieee300` and `pegaselk` systems. The paper selected August 22nd, 2017 for `ieee300`, which was found to encounter capacity shortages, and July 13th, 2017 for `pegaselk`, which is the day that experienced the highest load over the year. A total of 1024 and 2048 scenarios are generated for the 300 and 1354-bus system, respectively. Each dataset is then split between training (80%), validation (10%) and testing (10%).

C. Implementation Details

The optimization problems are formulated in Julia using JuMP [24], and solved with Gurobi 9.5 [25], with a single CPU thread and its default parameters. All deep learning models are implemented using PyTorch [26] and trained using the Adam optimizer [27]. All models are hyperparameter tuned using a grid search, with the learning rate from $\{10^{-2}, 10^{-3}\}$, hidden dimension from $\{128, 256\}$ and penalty term of constraint violation from $\{0.1, 1\}$. For each system, the model with the lowest validation loss is selected and the performances on the test set are reported. During training, the learning rate is reduced by a factor 10 if the validation loss shows no improvement for a consecutive sequence of 10 epochs. In addition, training is stopped if the validation loss does not improve for consecutive 20 epochs. Experiments are conducted on dual Intel Xeon 6226@2.7GHz machines running Linux, on the PACE Phoenix cluster [28]. The ML models are trained on Tesla V100-PCIE GPUs with 16GBs HBM2 RAM.


D. Comparison Framework

The risk simulations follow Algorithm 1. The first 4 hours of each simulation are used as warm-up, and results are reported for a 24-hour window starting at 4am and ending at 4am the following day. Indeed, while all simulations consider the same initial generation setpoint (as specified in the original UC instance), the initial total load varies across individual scenarios. For instance, as shown in Figure 3, in the `ieee300` system, total load at midnight ranges from 100p.u. to 140p.u. This, in turn, creates artificially high ramping requirements during the first hours of the simulation, as generators' dispatches are adjusted from their (fixed) initial setpoint to match total demand. The duration of the warmup period was set to 4 hours based on preliminary results, which showed that such effects disappeared after a few hours.


Finally, to ensure fair comparison between all ML models, the paper implements the following mechanism. All generation dispatches are clipped, if necessary, to ensure that minimum/maximum limits (2) are satisfied. This mechanism is executed at each time step of every simulation, and captures the fact that generators cannot adjust their output beyond their physical capabilities. While this does not affect the output of DNN and E2ELR, it may modify the output of DeepOPF and DC3, thus resulting in power balance violations instead of violations of the slack bus constraints.

E. Global Risk Analysis

Figures 4 and 5 compare ground truth and proxy-based risk profiles for the `ieee300` and `pegaselk` systems, respectively. The ground truth risk profiles are obtained by executing Algorithm 1 with an optimization solver (GRB). The proxy-based risk profiles are obtained by replacing the optimization solver with an optimization proxy. Each figure reports, for each hour: the 90%-CVaR (level 1) for system-wide power imbalance and total thermal violations, the probability of system-wide power imbalance and thermal violations, and the risk of power imbalance and total thermal violations.




(a) Risk analysis for `ieee300` and system-wise power imbalance. Left: CVaR of system-wide power imbalance (Level 1). Middle: Probability of system-wide power imbalance (Level 2). Right: Risk of system-wide power imbalance (Level 3).




(b) Risk analysis for `ieee300` and thermal violations. Left: CVaR of total thermal violations (Level 1). Middle: Probability of non-zero total thermal violations (Level 2). Right: Risk of total thermal violations (Level 3).

Fig. 4. Risk assessment results on the `ieee300` system.



(a) Risk analysis for `pegase1k` and system-wise power imbalance. Left: CVaR of system-wide power imbalance (Level 1). Middle: Probability of system-wide power imbalance (Level 2). Right: Risk of system-wide power imbalance (Level 3).



(b) Risk analysis for `pegase1k` and thermal violations. Left: CVaR of total thermal violations (Level 1). Middle: Probability of non-zero total thermal violations (Level 2). Right: Risk of total thermal violations (Level 3).

Fig. 5. Risk assessment results on the `pegase1k` system.

First observe that RA-DNN, RA-DC3 and RA-DeepOPF consistently over-estimates power imbalances. This is most striking for RA-DNN, which almost always predicts non-zero power imbalances, while RA-DeepOPF and RA-DC3 wrongly predict system imbalances during the afternoon hours on `ieee300` and throughout the day on `pegase1k`. This is because the DNN, DeepOPF, and DC3 are not guaranteed to

output solutions that satisfy all ED constraints. In contrast, the E2ELR architecture, which always outputs ED-feasible solutions, perfectly matches the ground truth risk assessment for power imbalances.

Second, all proxies tend to over-estimate thermal violations, which is more flagrant on the larger `pegase1k` system (Figure 5b). RA-DNN, RA-DeepOPF and, to a lesser extent, RA-

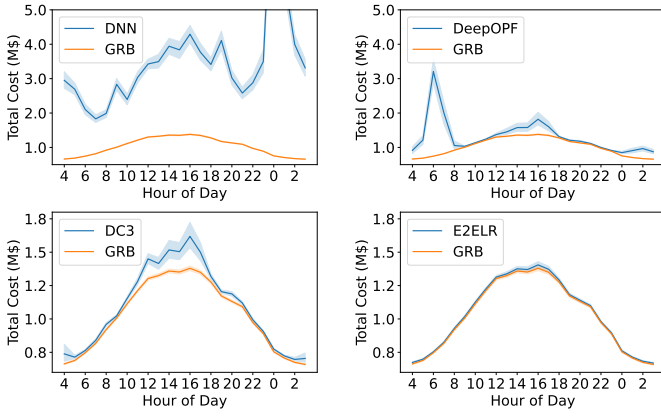


Fig. 6. Hourly distributions of total cost on `pegase1k`: comparison between ground truth (GRB) and proxies. Top: RA-DNN (left) and RA-DeepOPF (right); Bottom: RA-DC3 (left) and RA-E2ELR (right). All values in M\$.

TABLE I
IMPACT OF E2ELR REPAIR LAYER ON SOLUTION COSTS
(`PEGASE1K` SYSTEM ON 07/13/2017)

Costs	E2ELR		Gurobi
	Unrepaired	Repaired	
Generation Dispatch	1.199	1.213	1.201
Power Imbalances	1.200	0.000	0.000
Thermal Violations	0.098	0.109	0.092
Total	2.493	1.323	1.293

All values are averaged over are in M\$

DC3, incorrectly predict thermal violations during the morning hours (5am–9am) and during the night (10pm–2am). While the E2ELR-based risk profile over-estimates thermal violations during the afternoon (12pm–5pm), Figure 5b shows that RA-DC3 and RA-E2ELR correctly identify those hours during which congestion occurs. It is important to note that inaccurate risk predictions may result in unnecessary preventive actions, which result in higher economic cost for the system.

Finally, the over-estimation of thermal violations by RA-E2ELR can be explained by two factors. On the one hand, the training process may only converge to a local optimum, yielding sub-optimal solutions at inference. On the other hand, the E2ELR power balance repair layer, while successful at eliminating power balance violations, tends to yield higher thermal violations. This is because the repair layer adjusts the output of all generators to satisfy power balance (see Eqs. (12) and (13)), without taking into account power flows and potential thermal violations. To support this latter point, Table I compares the averaged itemized costs of the E2ELR predictions on the `pegase1k` system, before and after the repair layer. While generation dispatch costs remain within 1% of the ground truth, the repair layer yields to a complete elimination of power imbalances, but a 10% increase in thermal violation costs.

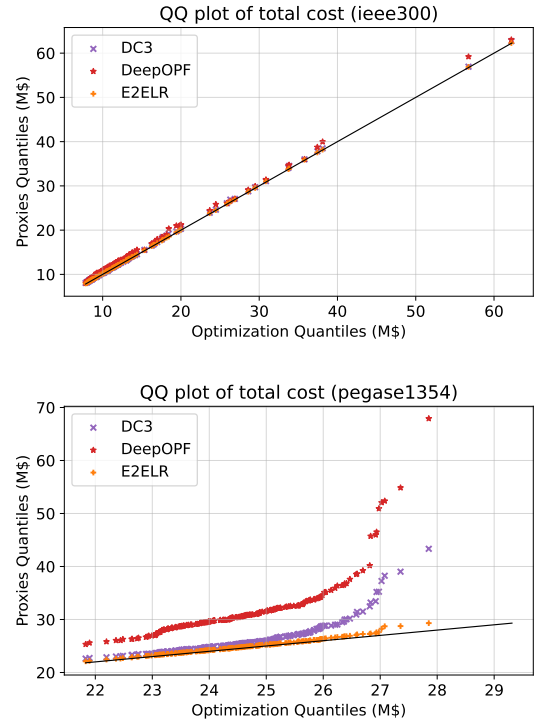


Fig. 7. QQ plot of the total cost on 300 IEEE (top) and 1354 Pegase system (bottom). RA-DNN is omitted due to its poor performance.

F. Global Cost Analysis

In addition to operational risk, grid operators rely on accurate cost estimates to operate the grid as economically as possible. Therefore, Figure 6 depicts, for each proxy architecture, the hourly distribution of total costs for the `pegase1k` system. Here, total costs include generation costs, thermal violation costs, and power balance violation costs. Each figure displays the 2.5%, mean and 97.5% quantiles of ground truth cost distribution (in orange), and the cost distribution obtained by the proxy (in blue).

First, throughout the entire day, RA-DNN significantly mis-estimates total costs. Similarly, RA-DeepOPF over-estimates costs costs, especially in the early morning (4am–8am), afternoon (2pm–5pm) and at night (12am–3am). Both of these are caused by incorrect predictions of system imbalance and thermal violations (see Figure 5). While RA-DC3 is more accurate, it still over-estimates total costs between 10am and 7pm, also mainly because of errors in power balance violations. The RA-E2ELR architecture achieves the best results, and almost perfectly matches the ground truth cost distribution.

To complement the above analysis, Figure 7 presents quantile-quantile (QQ) plots that compare, for both systems, the distribution of total costs (over the entire day) obtained by the ground truth optimization-based and proxy-based simulations. RA-DNN is omitted from the figure due to its poor performance. For `ieeee300`, all proxies accurately capture the distribution of total costs. However, for `pegase1k`, RA-DeepOPF and RA-DC3 significantly deviate from the ground

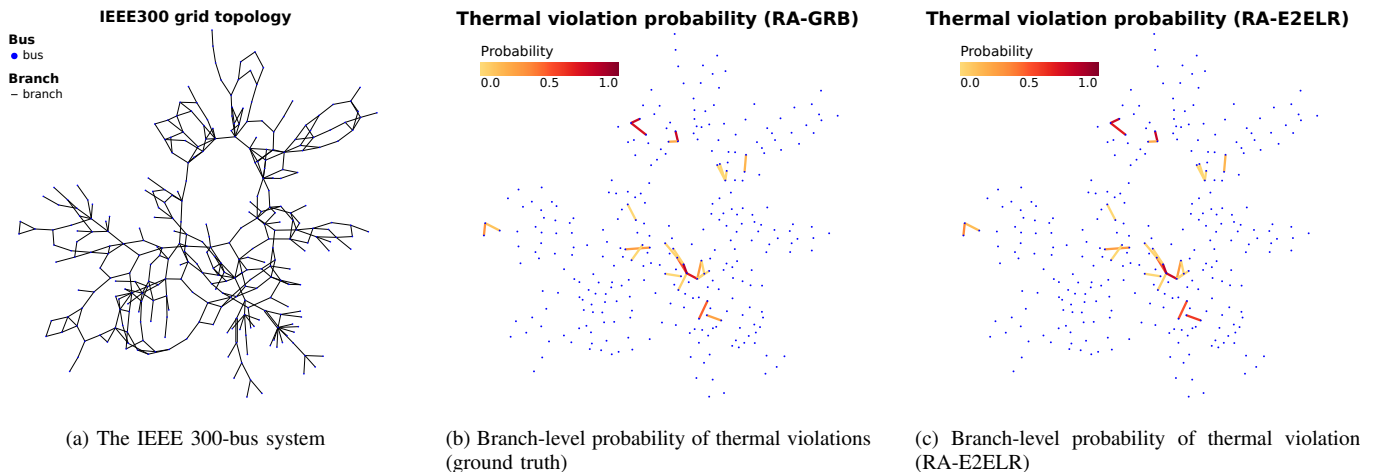


Fig. 8. The IEEE300 grid topology, and branch-level probability of thermal violations for the `ieee300` system at 12pm.

truth distribution: both architectures significantly over-estimate total costs for costlier (and riskier) scenarios. This is mostly due to incorrect predictions of system imbalances, which incur high penalty costs. In contrast, RA-E2ELR matches the ground truth distribution almost perfectly.

G. Component-level Risk Analysis

One of the advantages of the proposed proxy-based risk assessment frameworks is its granularity, i.e., it can produce risk profiles down to individual components. To demonstrate this capability, Figure 8 compares ground truth and proxy-based probability of thermal violations on each individual branch in `ieee300` at 12pm, which is the hour with the most violations. Figure 8a displays the original grid, Figure 8b displays the ground truth result, and Figure 8c displays the proxy-based results obtained with RA-E2ELR. The proxies correctly identify all branches with non-zero probability of thermal violations.

H. Computing Times

Table II compares computing times for proxy-based and optimization-based (GRB) risk assessment methods. Computing times for GRB scale roughly linearly with the size of the system, taking an average 0.5 and 2.2s per simulation per CPU core. Note that congested scenarios typically lead to increased computing times. In contrast, the simulation time of the proxies is consistent across different systems, typically taking less than 100 milliseconds to run 100 simulations. Specifically, RA-E2ELR achieves 30.3x speed up than GRB, even under the assumption that hundreds of ED instances can be solved in parallel. Among the proxies, RA-DC3 is the slowest because every inference requires 200 gradient steps.

The performance of Gurobi suggests a simple parallelization approach, wherein each simulation is executed on a single CPU core. It is important to note that, because such simulations must be executed locally to comply with confidentiality restrictions, this approach would incur substantial

TABLE II
RISK SIMULATION COMPUTING TIMES

Systems	E2ELR	DC3 [†]	DeepOPF	DNN	GRB*
<code>ieee300</code>	71.5	192.0	77.6	50.2	488.0
<code>pegase1k</code>	72.0	211.4	59.0	49.6	2188.0

All times in milliseconds. ML inference times are per batch of 100 scenarios on 1 GPU. *Average time per scenario using 1 CPU. [†]with 200 gradient steps.

hardware and energy costs. Furthermore, real-time market-clearing problems are becoming harder to solve, mainly due to a growing number of resources, additional reserve products, the incorporation of multiple time periods, and the potential transition to stochastic formulations and nonlinear AC-based formulations. Since each simulation requires the *sequential* solving of multiple ED problems, it is clear that optimization technology alone will not be tractable enough to execute numerous simulations in real-time. In contrast, learning-based approaches shift the computational effort offline during training, and offers fast, online inference.

VI. CONCLUSION

The paper has presented the first risk-assessment framework for power systems, RA-E2ELR, that uses end-to-end feasible optimization proxies. The proposed framework leverages the speed of optimization proxies and the feasibility guarantees of the E2ELR architecture. Numerical experiments demonstrated that RA-E2ELR yields high-quality risk estimates, and a 30x speedup over optimization-based simulations. This contrasts with previously-proposed architectures like DeepOPF and DC3, which offer no significant speed advantage over E2ELR but whose lack of feasibility guarantees results in over-predicting risks and costs.

While the paper has laid the foundation for learning-based risk assessment, several avenues of future work are important to explore. This includes the extension of E2ELR architectures to nonlinear AC formulations and to security-constrained

formulations. Another important direction is the evaluation of N-1 security, which would require the evaluation of post-contingency states and emergency actions, see, e.g., [7], [8].

REFERENCES

- [1] O. Stover, P. Karve, and S. Mahadevan, "Reliability and risk metrics to assess operational adequacy and flexibility of power grids," *Reliability Engineering & System Safety*, vol. 231, p. 109018, 2023.
- [2] W. Chen, M. Tanneau, and P. V. Hentenryck, "End-to-End Feasible Optimization Proxies for Large-Scale Economic Dispatch," *IEEE Transactions on Power Systems*, pp. 1–12, 2023.
- [3] A. A. Thatte and L. Xie, "A metric and market construct of inter-temporal flexibility in time-coupled economic dispatch," *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3437–3446, 2015.
- [4] E. Lannoye, D. Flynn, and M. O'Malley, "Evaluation of power system flexibility," *IEEE Transactions on Power Systems*, vol. 27, no. 2, pp. 922–931, 2012.
- [5] "North american electric reliability corporation. probabilistic adequacy and measures," *Technical reference report final.*, 2018.
- [6] F. Bouffard, F. D. Galiana, and A. J. Conejo, "Market-clearing with stochastic security-part i: formulation," *IEEE Transactions on Power Systems*, vol. 20, no. 4, pp. 1818–1826, 2005.
- [7] Q. Wang, J. D. McCalley, T. Zheng, and E. Litvinov, "A computational strategy to solve preventive risk-based security-constrained opf," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1666–1675, 2013.
- [8] —, "Solving corrective risk-based security-constrained optimal power flow with lagrangian relaxation and benders decomposition," *International Journal of Electrical Power & Energy Systems*, vol. 75, pp. 255–264, 2016.
- [9] F. Geth, H. Ergun, D. van Hertem, E. Heylen, and I. B. Sperstad, "Garpur d7. 3-a broader comparison of different reliability criteria through the garpur quantification platform," 2017.
- [10] Y. Zhou, B. Zhang, C. Xu, T. Lan, R. Diao, D. Shi, Z. Wang, and W.-J. Lee, "A Data-Driven Method for fast AC Optimal Power Solutions via Deep Reinforcement Learning," *Journal of Modern Power Systems and Clean Energy*, vol. 8, no. 6, pp. 1128–1139, 2020.
- [11] Z. Yan and Y. Xu, "Real-time optimal power flow: A lagrangian based deep reinforcement learning approach," *IEEE Transactions on Power Systems*, vol. 35, no. 4, pp. 3270–3273, 2020.
- [12] A. R. Sayed, C. Wang, H. Anis, and T. Bi, "Feasibility constrained online calculation for real-time optimal power flow: A convex constrained deep reinforcement learning approach," *IEEE Transactions on Power Systems*, 2022.
- [13] A. R. Sayed, X. Zhang, G. Wang, C. Wang, and J. Qiu, "Optimal operable power flow: Sample-efficient holomorphic embedding-based reinforcement learning," *IEEE Transactions on Power Systems*, 2023.
- [14] L. Zeng, M. Sun, X. Wan, Z. Zhang, R. Deng, and Y. Xu, "Physics-constrained vulnerability assessment of deep reinforcement learning-based scopf," *IEEE Transactions on Power Systems*, 2022.
- [15] MISO, "Energy and operating reserve markets," 2022, Business Practices Manual Energy and Operating Reserve Markets.
- [16] O. Stover, P. Karve, S. Mahadevan, W. Chen, H. Zhao, M. Tanneau, and P. Van Hentenryck, "Just-in-time learning for operational risk assessment in power grids," *arXiv preprint arXiv:2209.12762*, 2022.
- [17] S. Park and P. Van Hentenryck, "Self-supervised primal-dual learning for constrained optimization," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 4, 2023, pp. 4052–4060.
- [18] X. Pan, T. Zhao, M. Chen, and S. Zhang, "Deepopf: A deep neural network approach for security-constrained dc optimal power flow," *IEEE Transactions on Power Systems*, vol. 36, no. 3, pp. 1725–1735, 2020.
- [19] P. L. Donti, D. Rolnick, and J. Z. Kolter, "DC3: A learning method for optimization with hard constraints," *preprint arXiv:2104.12225*, 2021.
- [20] University of Washington, Dept. of Electrical Engineering, "Power Systems Test Case Archive." [Online]. Available: <http://www.ee.washington.edu/research/pstca/>
- [21] C. Jozs, S. Fliscounakis, J. Maeght, and P. Panciatici, "AC Power Flow Data in MATPOWER and QCQP Format: iTesla, RTE Snapshots, and PEGASE," 2016.
- [22] A. S. Xavier, A. M. Kazachkov, O. Yurdakul, and F. Qiu, "UnitCommitment.jl: A Julia/JuMP Optimization Package for Security-Constrained Unit Commitment." Jul. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6857290>
- [23] A. S. Xavier, F. Qiu, and S. Ahmed, "Learning to Solve Large-Scale Security-Constrained Unit Commitment Problems," *INFORMS Journal on Computing*, vol. 33, no. 2, pp. 739–756, 2021.
- [24] I. Dunning, J. Huchette, and M. Lubin, "JuMP: A Modeling Language for Mathematical Optimization," *SIAM Review*, vol. 59, no. 2, pp. 295–320, 2017.
- [25] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023. [Online]. Available: <https://www.gurobi.com>
- [26] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," in *NIPS-W*, 2017.
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [28] PACE, *Partnership for an Advanced Computing Environment (PACE)*, 2017. [Online]. Available: <http://www.pace.gatech.edu>