
Adversarially Robust Learning for Security-Constrained Optimal Power Flow

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In recent years, the ML community has seen surges of interest in both adversarially
2 robust learning and implicit layers, but connections between these two areas have
3 seldom been explored. In this work, we combine innovations from these areas
4 to tackle the problem of N-k security-constrained optimal power flow (SCOPF).
5 N-k SCOPF is a core problem for the operation of electrical grids, and aims to
6 schedule power generation in a manner that is robust to potentially k simultaneous
7 equipment outages. Inspired by methods in adversarially robust training, we frame
8 N-k SCOPF as a minimax optimization problem – viewing power generation
9 settings as adjustable parameters and equipment outages as (adversarial) attacks
10 – and solve this problem via gradient-based techniques. The loss function of this
11 minimax problem involves resolving implicit equations representing grid physics
12 and operational decisions, which we differentiate through via the implicit function
13 theorem. We demonstrate the efficacy of our framework in solving N-3 SCOPF,
14 which has traditionally been considered as prohibitively expensive to solve given
15 that the problem size depends combinatorially on the number of potential outages.

16 1 Introduction

17 Robust optimization problems are pervasive across many applications and domains – such as electric
18 power systems, supply chain management, and civil engineering – where the goal is to construct
19 some solution that is robust under any allowable instantiation of uncertainty [1, 2]. While the
20 aim is generally that these solutions be *provably* robust, there unfortunately remain many settings
21 where it is either not easy or not possible to construct such solutions. This has often motivated the
22 use of heuristic approaches. For instance, many approaches in adversarially robust deep learning
23 formulate neural network training as a minimax game over neural network parameters and input
24 perturbations, optimizing this problem via gradient-based techniques that do not yield provable
25 robustness guarantees, but are nonetheless effective in practice [3].

26 In this work, we draw inspiration from adversarially robust training to address the problem of N-k
27 security-constrained optimal power flow (SCOPF). SCOPF is a fundamental problem to schedule
28 power generation in a way that is robust to k potential equipment failures (e.g., generator or line
29 outages). Unfortunately, N-k SCOPF is prohibitively expensive to solve at scale, leading grid
30 operators to use rough approximations in practice. To address this challenge, we frame N-k SCOPF
31 as a minimax attacker-defender problem, where the “defender” aims to schedule power generation,
32 and the “attacker” aims to pick adversarial equipment failures. The loss function of this problem
33 requires solving implicit equations representing the physics of the electric grid as well as additional
34 operational decisions that are made after an attack has occurred. As such, we optimize this problem
35 using gradient-based techniques and employ insights from the literature on implicit layers to cheaply
36 compute gradients through the loss function.

37 Our key contributions are:

- 38 • **Formulation for minimax optimization with implicit variables.** To streamline the pre-
39 sentation of concepts, we provide a generic formulation for gradient-based optimization of
40 minimax problems with implicitly-defined variables. While our main focus in this paper
41 is on N-k SCOPF, we believe this generic formulation may also be of broader interest for
42 minimax settings with physics in the loop, as well as for tri-level optimization settings.
- 43 • **Formulation for gradient-based optimization of N-k SCOPF.** We rewrite N-k SCOPF as
44 a continuous minimax optimization problem, and demonstrate how to efficiently compute
45 gradients through relevant implicit components. We also utilize the underlying structure of
46 our optimization solvers to further streamline the outer minimization procedure. Importantly,
47 the per-iteration cost of this approach is agnostic to the number of simultaneous outages k ,
48 despite the combinatorial blowup in the number of associated “contingency scenarios.”
- 49 • **Demonstration on N-1, N-2, and N-3 SCOPF.** We demonstrate the efficacy of our method
50 in addressing SCOPF settings that allow for one, two, or three simultaneous outages on a
51 realistic 4622-node power system with over 38 billion potential N-3 outage scenarios. We
52 find that our method incurs 3-4 \times fewer N-3 feasibility violations than a baseline optimal
53 power flow approach, and requires only 21 minutes to run on a standard laptop.

54 2 Related work

55 **Adversarial robustness in deep learning.** There has been a growing body of work that aims to
56 parameterize neural networks in a manner that is robust to particular perturbations of their inputs,
57 usually by casting neural network training as an attacker-defender game [3, 4]. While there have
58 been several promising approaches for certifiably robust neural network training [5-7], in general,
59 these approaches do not yet scale to large-sized networks and only address a limited set of threat
60 models. As a result, there has been a lot of research in this area that aims to train robust neural
61 networks using approximate, gradient-based training approaches [8, 9], an approach we adopt in
62 the context of SCOPF. In addition, a key part of this literature has been on constructing strong but
63 cheap-to-compute attacks that can strengthen the outcomes of adversarially robust training, e.g., the
64 fast gradient sign method (FGSM) [8] and projected gradient descent (PGD) attacks [9]. In our
65 experiments, we similarly show how a gradient-based adversarial robustness approach can be used to
66 identify potential grid vulnerabilities, as an input to secure power system optimization.

67 **Implicit layers.** Implicit differentiation techniques have started to be used more widely within deep
68 learning workflows, largely in the context of *implicit layers*, i.e., neural network layers that represent
69 implicit functions [10]. These include differentiable optimization layers [11-17] and physics-based
70 layers [18, 19], among others [20, 21]. Given the large number of parameters within any given deep
71 network, a key aspect of this work has been in finding efficient ways in which to actually compute the
72 relevant derivatives, e.g., by strategically ordering multiplicative operations and reusing the results
73 of previous computations. We similarly employ these kinds of strategies when computing implicit
74 derivatives for our SCOPF optimization approach. We also note that some of the work on implicit
75 differentiation in deep learning and related areas has been in service of solving bi-level optimization
76 problems, e.g., for decision-driven forecasting [22, 23], hyperparameter tuning [24-27], or system
77 identification and control [28]. We similarly consider the use of implicit differentiation for multi-level
78 optimization (in particular, tri-level optimization) in the context of SCOPF.

79 **Security-constrained optimal power flow.** N-k security analysis is widely used to identify which
80 sets of outages (*contingencies*) may be particularly high risk for the electric grid, especially for the
81 case of N-1 [29-31]. However, there have been few attempts at developing methods that actually
82 *solve* for a grid setting that is N-k secure (i.e., secure against any k simultaneous outages) [32, 33].
83 In particular, the computational complexity of N-k SCOPF grows combinatorially with k and the
84 size of the system. Some previous methods to solve N-k SCOPF have included exhaustive methods
85 [33], Bender’s cuts to reduce the number of contingencies analyzed [32], and bi-level optimization
86 frameworks to represent N-k SCOPF [32, 34]. In particular, the work using bi-level optimization [34]
87 developed a systematic attacker-defender approach to address N-3 contingency scenarios, but used a
88 linearized grid model to attain convergence. We extend the bi-level framework by introducing fast
89 gradient calculation methods inspired by the implicit layers literature for solving a realistic non-linear
90 model of the grid, which allows us to scale our approach to a 4622-node system.

91 3 Generic problem formulation

92 Before diving into the details of our SCOPF formulation, we first provide a more generic formulation
 93 for gradient-based minimax optimization over an implicit loss function, which we will later build upon
 94 in the context of SCOPF. In particular, we consider the setting of continuous minimax optimization
 95 problems over “defender” (maximizer) variables $x \in \mathcal{X}$ and “attacker” (minimizer) variables $y \in \mathcal{Y}$;
 96 these are also referred to as first-stage and second-stage decision variables, respectively, in the bi-level
 97 optimization literature. In addition, we allow for “third-stage” decisions $z \in \mathcal{Z}$ that are fully defined
 98 via a set of implicit constraints on x , y , and z .

99 Specifically, we consider problems of the form

$$\begin{aligned} & \underset{x \in \mathcal{X}}{\text{minimize}} \quad \max_{y \in \mathcal{Y}} \ell(x, y, z) \\ & \text{s. t.} \quad g(x, y, z) = 0, \quad z \in \mathcal{Z}, \end{aligned} \tag{1}$$

100 where \mathcal{X} , \mathcal{Y} , and \mathcal{Z} are compact sets; $\ell : \mathcal{X} \times \mathcal{Y} \times \mathcal{Z} \rightarrow \mathbb{R}$ is a standard, continuously differentiable
 101 loss function (e.g., softmax or mean squared error loss); and $g : \mathcal{X} \times \mathcal{Y} \times \mathcal{Z} \rightarrow \mathbb{R}^m$ is defined such
 102 that $g(x, y, z) = 0$ is an implicit function in z with some solution $z \in \mathcal{Z}$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$. We
 103 further restrict ourselves to those functions g that are continuously differentiable with non-singular
 104 Jacobians at their roots, i.e., those functions that are compatible with the implicit function theorem
 105 [3, 35]. We note that this formulation covers a wide range of settings, e.g., many minimax problems
 106 with non-linear physical constraints, or many tri-level optimization problems where z is a solution to
 107 a continuous optimization problem parameterized by x and y (both of which notions we will use in
 108 Section 4 for the setting of N-k SCOPF).

109 Inspired by the literature on adversarial robustness in deep learning, we propose to solve problem (1)
 110 via gradient-based search on both the inner maximization and outer minimization problems. In
 111 particular, this entails (a) obtaining some (approximately) optimal y for the inner maximization
 112 problem via gradient-based techniques, given some initial value of x , (b) updating x using the
 113 gradient at the optimum of the inner maximization problem, and (c) repeating these steps until
 114 convergence. We now describe steps (a) and (b) in additional detail.

115 3.1 Solving the inner maximization problem

116 Let \bar{x} denote some fixed value for x . The inner maximization problem is then given by

$$\max_{y \in \mathcal{Y}} \ell(\bar{x}, y, z) \quad \text{s. t.} \quad g(\bar{x}, y, z) = 0, \quad z \in \mathcal{Z}. \tag{2}$$

117 We optimize this problem via projected gradient descent. Specifically, let $y = y_0$ denote our initial
 118 guess for the optimal attack, and let \mathcal{P} denote the projection operator. Until convergence (or for some
 119 fixed number of iterations), we then

- 120 (i) Obtain z^* such that $g(\bar{x}, y, z^*) = 0$, $z^* \in \mathcal{Z}$.
- 121 (ii) Update $y \leftarrow \mathcal{P}_{\mathcal{Y}}(y + \gamma \nabla_y \ell(\bar{x}, y, z^*))$ for step size γ .

122 Notably, step (ii) entails obtaining the gradient $\nabla_y \ell(\bar{x}, y, z^*)$. By the chain rule, this involves the
 123 gradient through z^* , which is the solution to a set of implicit equations. Specifically, using the
 124 notation d to denote total derivatives (e.g., gradients) and ∂ to denote partial derivatives, we have

$$\frac{d\ell(\bar{x}, y, z^*)}{dy} = \frac{\partial \ell(\bar{x}, y, z^*)}{\partial y} + \frac{\partial \ell(\bar{x}, y, z^*)}{\partial z^*} \frac{dz^*}{dy}. \tag{3}$$

125 By the implicit function theorem, we can then obtain an expression for dz^*/dy by noting that

$$\frac{dg(\bar{x}, y, z^*)}{dy} = \frac{\partial g(\bar{x}, y, z^*)}{\partial y} + \frac{\partial g(\bar{x}, y, z^*)}{\partial z^*} \frac{dz^*}{dy} = 0 \implies \frac{dz^*}{dy} = - \left(\frac{\partial g(\bar{x}, y, z^*)}{\partial z^*} \right)^{-1} \frac{\partial g(\bar{x}, y, z^*)}{\partial y}, \tag{4}$$

126 which we can plug into Equation (3) to yield our full update.

127 We note that in practice, we seldom want to compute the Jacobian $dz^*/dy \in \mathbb{R}^{\dim(\mathcal{Z}) \times \dim(\mathcal{Y})}$ explicitly
 128 due to the potentially large time and space complexity of doing so; instead, it is often desirable to
 129 compute the left vector-matrix product $(\partial \ell / \partial z^*)(dz^*/dy) \in \mathbb{R}^{\dim(\mathcal{Y})}$ directly. We refer to this strategy
 130 as the “vector-Jacobian product trick.” The details of the relevant computations may vary based on
 131 the particular setting at hand, and we describe how we employ this trick for SCOPF in Section 4.3.

132 3.2 Taking a gradient step in the minimization problem

133 Given some (approximately) optimal y^* and associated z^* from the inner optimization under the
 134 current value of $x = \bar{x}$, the outer optimization problem then becomes

$$\min_{x \in \mathcal{X}} \ell(x, y^*, z^*) \quad \text{s. t.} \quad g(x, y^*, z^*) = 0. \quad (5)$$

135 One option is to then update x via a projected gradient step $x \leftarrow \mathcal{P}_{\mathcal{X}}(x - \beta \nabla_x \ell(x, y^*, z^*))$ for step
 136 size β . To calculate the gradient $\nabla_x \ell(x, y^*, z^*)$, we note that by Danskin’s theorem, we can disregard
 137 the dependence of y^* on x [3] (though we cannot ignore the dependence of z^*). As such, we can
 138 employ a similar process as in Equations (3) and (4), where we treat y^* as constant when computing
 139 gradients with respect to x .¹ We note that while this is one potential process for updating x , we
 140 actually employ a more efficient, domain-specific process for our SCOPF procedure (see Section 4.4).

141 4 Addressing N-k SCOPF via adversarially robust optimization

142 Having presented this generic formulation, we now introduce our approach, CAN ∂ Y, for addressing
 143 security-constrained optimal power flow. In particular, we consider the problem of N-k SCOPF, where
 144 power generation must be scheduled so as to be feasible and low-cost in the *absence* of equipment
 145 outages (“base case”) as well as to be robust to any k simultaneous outages of power generators or
 146 lines that may occur (“contingency cases”). We note that the set of *contingencies* – i.e., allowable
 147 combinations of outages – is combinatorial in the number of potential outages, making the SCOPF
 148 problem extremely computationally expensive. For instance, a realistic 4622 node system with 6137
 149 potential single outages has ~ 38.5 billion contingency scenarios to consider under the N-3 setting.

150 In the rest of this section, we first more formally define the N-k SCOPF problem. We then show how
 151 we rewrite N-k SCOPF as a minimax problem of the form (1), in particular by forming a compact
 152 outer approximation to the contingency space. Finally, we describe how we solve this problem using
 153 a combination of gradient-based techniques and domain-specific enhancements.

154 4.1 Defining N-k SCOPF

155 Let x denote the *dispatch* – i.e., setpoints of *real power*² and voltage magnitude – at all power
 156 generators on the electricity system, and let \mathcal{X} represent generator-wise box constraints on the
 157 dispatch. Let \mathcal{C} denote the set of potential contingencies, i.e., all sets of exactly k potential outages.
 158 Finally, let $z^{(i)} \in \mathcal{Z}_i(x, c^{(i)})$ represent slightly adjusted settings of real power and voltage magnitude
 159 that the power system operator can create after scheduling x and then observing some contingency
 160 $c^{(i)} \in \mathcal{C}$, where the \mathcal{Z}_i represent box constraints. Then, the N-k SCOPF problem can be expressed as

$$\begin{aligned} & \underset{x \in \mathcal{X}}{\text{minimize}} \quad f_{\text{base}}(x) + \sum_{(z^{(i)}, c^{(i)})} f_{\text{cont}}(z^{(i)}, c^{(i)}) \\ & \text{subject to} \quad g_{\text{flow,base}}(x, w_{\text{base}}) = 0, \quad w_{\text{base}} \in \mathcal{W}_{\text{base}} \end{aligned} \quad (6)$$

$$\begin{aligned} & z^{(i)} \in \underset{z^{(i)} \in \mathcal{Z}_i(x, c^{(i)})}{\text{argmin}} f_{\text{cont}}(z^{(i)}, c^{(i)}) \quad \forall c^{(i)} \in \mathcal{C}, \\ & \text{s. t.} \quad g_{\text{flow,cont}}(z^{(i)}, w^{(i)}, x) = 0, \quad w^{(i)} \in \mathcal{W}_i(x, c^{(i)}) \end{aligned}$$

161 where $f_{\text{base}} : \mathcal{X} \rightarrow \mathbb{R}$ represents base case power production costs; $g_{\text{flow,base}} : \mathcal{X} \times \mathcal{W}_{\text{base}} \rightarrow \mathbb{R}^{n_{\text{bus}}}$
 162 represents the non-linear power flow equations in the base case, with n_{bus} being the number of power
 163 system nodes; w_{base} represents electrical quantities that result from solving the base case power flow
 164 equations (e.g., *reactive powers* and voltage angles), with box constraints (device limits) represented
 165 by $\mathcal{W}_{\text{base}}$; and $f_{\text{cont}} : \mathcal{Z}_i \times \mathcal{C} \rightarrow \mathbb{R}$, $g_{\text{flow,cont}} : \mathcal{Z}_i \times \mathcal{W}_i \times \mathcal{X} \rightarrow \mathbb{R}^n$, and $w^{(i)} \in \mathcal{W}_i(x, c^{(i)})$ represent
 166 their respective contingency-case counterparts. (See Appendix A for a more explicit formulation.)

¹Technically, Danskin’s theorem only holds when y^* is a unique optimum of the inner maximization problem. However, in the adversarially robust training literature, the conditions of Danskin’s theorem do not necessarily hold – in particular, the inner maximization problem often does not have a unique optimum, and many implementations tend to generate approximate (rather than exact) optima [8, 9] – but this method of computing gradients is used in practice regardless [3].

²Modern electric power systems are generally alternating current (AC) systems, in which all electrical quantities – e.g., powers and voltages – are considered to be complex-valued. In particular, the terms *real power* and *reactive power* refer to the real and imaginary components, respectively, of *complex power*.

167 4.2 Rewriting N-k SCOPF as a minimax problem

168 We reformulate the SCOPF problem (6) as an attacker-defender game, where the defender must
 169 choose a dispatch that is robust to potential “worst-case” contingencies chosen by an attacker. In
 170 particular, since the contingency set \mathcal{C} is discrete, we create a continuous outer approximation to
 171 this set in order to enable the use of gradient-based techniques. Specifically, let n_o be the number of
 172 generators or power lines that can potentially experience an outage. Then, for any $y \in [0, 1]^{n_o}$, we
 173 define the i th entry as follows:

$$y_i = \begin{cases} 1 & \text{iff outage } i \text{ is fully active,} \\ 0 & \text{iff outage } i \text{ is not active,} \\ \alpha_i \in (0, 1) & \text{iff outage } i \text{ is partially active with fraction } \alpha_i. \end{cases} \quad (7)$$

174 The first two notions presented in Equation (7) are standard in power systems: the generator or line
 175 pertinent to outage i is either fully operational or out of service. We newly define the notion of a
 176 partial outage with fraction α_i as one in which the power flowing through the outage device during
 177 normal operation has been reduced by a factor of α_i . For instance, we model a partial contingency on
 178 a transmission line or transformer device as reducing its admittance (i.e., ability to conduct current)
 179 by a factor of α_i . Similarly, we restrict the power produced by a generator undergoing a partial
 180 contingency by multiplying its power output by α_i .

181 Given these notions, we define our “threat model” for the N-k SCOPF setting to contain all vectors
 182 y with an L1-norm of at most k , i.e., $\mathcal{Y} := \{y : y \in [0, 1]^{n_o}, \|y\|_1 \leq k\}$. Notably, the original
 183 contingency set \mathcal{C} is fully represented within \mathcal{Y} , and in fact, all scenarios with *up to* k simultaneous
 184 potential outages are also represented. As such, \mathcal{Y} represents a much broader set of potential
 185 contingencies than specified in the original problem. (Relevantly for projected gradient descent, this
 186 is also a convex set.) Using this set, we can then write our reformulation of the SCOPF problem as

$$\text{minimize}_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f_{\text{base}}(x) + f_{\text{cont}}(z, y) + \frac{1}{2} \|s\|_2^2 \quad (8a)$$

$$\text{subject to } g_{\text{flow,base}}(x, w_{\text{base}}) = 0, \quad w_{\text{base}} \in \mathcal{W}_{\text{base}} \quad (8b)$$

$$z, s \in \underset{z \in \mathcal{Z}(x, y), s \in \mathbb{R}^{n_{\text{bus}}}}{\text{argmin}} f_{\text{cont}}(z, y) + \frac{1}{2} \|s\|_2^2 \quad (8c)$$

$$\text{s. t. } g_{\text{flow,cont}}(z, w_{\text{cont}}, x) + s = 0, \quad w_{\text{cont}} \in \mathcal{W}_{\text{cont}}(x, y),$$

187 where $s \in \mathcal{S} := \mathbb{R}^{n_{\text{bus}}}$ are slack variables representing potential infeasibilities in the third-stage
 188 optimization problem, as necessitated by the expanded contingency set. In particular, the goal
 189 of the attacker is to now to find a set of partial outages that not only increase the cost of power
 190 generation, but also create instabilities in the grid, as captured by s . As we hinted at in Section 3 this
 191 is a minimax optimization problem with implicit constraints over the third-stage variables z , w_{base} ,
 192 w_{cont} , and s , incorporating both non-linear equality constraints (8b) as well as an optimization-based
 193 constraint (8c).

194 4.3 Obtaining attack gradients

195 As described in Section 3.1, we aim to find the worst-case attack via projected gradient descent. In
 196 particular, we must compute the gradient of the minimax loss with respect to y , which is given by

$$\frac{d\ell}{dy} = \frac{\partial f_{\text{cont}}(z^*, y)}{\partial y} + \frac{\partial f_{\text{cont}}(z^*, y)}{\partial z^*} \frac{dz^*}{dy} + \frac{ds^*}{dy}. \quad (9)$$

197 (As the base case power production cost and the base case power flow constraint (8b) have no
 198 dependence on y , we do not need to consider these terms during the inner maximization.)

199 To calculate the Jacobians dz^*/dy and ds^*/dy , we implicitly differentiate through the third-stage
 200 optimization problem (8c). In order to do so inexpensively, we reuse the results of computations that
 201 were executed when originally obtaining z^* and s^* . More specifically, in order to obtain z^* and s^* ,
 202 we solve the non-linear KKT conditions of optimization problem (8c) using a Newton solver (see
 203 Appendix B), which entails linearizing these equations at each iteration. At convergence, we then
 204 implicitly differentiate through the linear fixed-point equation obtained at the last iteration:

$$J \begin{pmatrix} z^* \\ s^* \end{pmatrix} = b \implies \frac{dJ}{dy} \begin{pmatrix} z^* \\ s^* \end{pmatrix} + J \begin{pmatrix} \frac{dz}{dy} \\ \frac{ds}{dy} \end{pmatrix} = \frac{db}{dy} \implies \begin{pmatrix} \frac{dz}{dy} \\ \frac{ds}{dy} \end{pmatrix} = J^{-1} \left(-\frac{dJ}{dy} \begin{pmatrix} z^* \\ s^* \end{pmatrix} + \frac{db}{dy} \right), \quad (10)$$

205 where $J \in \mathbb{R}^{d \times d}$ is the Jacobian of the non-linear KKT system and $b \in \mathbb{R}^d$ is the corresponding right
 206 hand side vector, for $d = \dim(\mathcal{Z}) + \dim(\mathcal{S})$. The overall gradient of the loss is then given by

$$\frac{d\ell}{dy} = \frac{\partial f_{\text{cont}}(z^*, y)}{\partial y} + \begin{pmatrix} \frac{\partial f_{\text{cont}}(z^*, y)}{\partial z^*} \\ 1 \end{pmatrix}^T J^{-1} \left(-\frac{dJ}{dy} \begin{pmatrix} z^* \\ s^* \end{pmatrix} + \frac{db}{dy} \right). \quad (11)$$

207 As hinted earlier, we employ the ‘‘vector-Jacobian product trick’’ in order to efficiently compute these
 208 gradients. In particular, rather than computing the terms dz^*/dy and ds^*/dy explicitly via Equation (10),
 209 we directly compute their left vector-matrix product with the relevant partial derivatives of the loss –
 210 i.e., the blue term in Equation (11), with multiplications evaluated from left to right to ensure we are
 211 always taking matrix-vector (rather than matrix-matrix) products. Inspired by (11), we also reuse
 212 the LU factor from the last Newton solve we computed when obtaining z^* and s^* in order to avoid
 213 explicitly (re-)computing the matrix inverse J^{-1} . Finally, we note that for this specific problem,
 214 while the last term $-\frac{dJ}{dy} \begin{pmatrix} z^* \\ s^* \end{pmatrix} + \frac{db}{dy}$ nominally involves tensor products, the relevant terms are vastly
 215 sparse due to the structure of the underlying physics – i.e., with no more than 20 nonzero entries per
 216 potential outage – making these products relatively cheap to compute in practice.

217 4.4 Solving the defense minimization

218 After obtaining a worst-case attack y^* , our next step is to adjust our dispatch $x \in \mathcal{X}$ in response.
 219 As described in Section 3.2 for the generic setting, one option to do this involves taking a projected
 220 gradient step in x . However, we adopt a different approach for N-k SCOPF, due to the practical
 221 requirements of this setting. In particular, a general system requirement is that the base case power
 222 flow equations g_{flow} *must* remain feasible under the dispatch $x \in \mathcal{X}$ – as the most likely scenario is
 223 that no contingency will occur – but a gradient-based approach may not necessarily guarantee this.

224 As a result, we instead note that we can rewrite the N-k SCOPF minimax problem (8) as a single
 225 minimization problem:

$$\begin{aligned} & \underset{x \in \mathcal{X}, z \in \mathcal{Z}(x, y^*), s \in \mathbb{R}^{n_{\text{bus}}}}{\text{minimize}} && f_{\text{base}}(x) + f_{\text{cont}}(z, y^*) + \frac{1}{2} \|s\|_2^2 \\ & \text{subject to} && g_{\text{flow,base}}(x, w_{\text{base}}) = 0, \quad w_{\text{base}} \in \mathcal{W}_{\text{base}} \\ & && g_{\text{flow,cont}}(z, w_{\text{cont}}, x) + s = 0, \quad w_{\text{cont}} \in \mathcal{W}_{\text{cont}}(x, y^*). \end{aligned} \quad (12)$$


226 To determine our next iterate of x , our strategy is then to *partially* solve this optimization problem by
 227 running one step of a non-linear Gauss-Seidel method, and then keep the value of x obtained from
 228 that step. This allows us to incrementally update x in a direction that is more robust to the worst-case
 229 attack y^* , while still maintaining the feasibility of the base case power flow equations.

230 Importantly, we are able to run this procedure efficiently, as we can reuse the results of existing
 231 computations that were executed when obtaining the optimal attack y^* . In particular, as we describe
 232 in more detail in Appendix C, we can split the KKT conditions of the problem (12) into two groups:

$$\begin{pmatrix} d\mathcal{L}/dx \\ d\mathcal{L}/d\lambda_{\text{base}} \end{pmatrix} \equiv F_{\text{base}}(x, w_{\text{base}}, \lambda_{\text{base}}) + \begin{pmatrix} (\partial g_{\text{flow,cont}}(z, w_{\text{cont}}, x)/\partial x)^T \lambda_{\text{cont}} \\ 0 \end{pmatrix} = 0 \quad (13a)$$

$$\begin{pmatrix} d\mathcal{L}/dz \\ d\mathcal{L}/ds \\ d\mathcal{L}/d\lambda_{\text{cont}} \end{pmatrix} \equiv F_{\text{cont}}(z, w_{\text{cont}}, \lambda_{\text{cont}}) + \begin{pmatrix} 0 \\ 0 \\ g_{\text{flow,cont}}(z, w_{\text{cont}}, x) \end{pmatrix} = 0, \quad (13b)$$

233 where \mathcal{L} denotes the Lagrangian of problem (12), and λ_{base} and λ_{cont} are the dual variables on the
 234 base case and contingency power flow constraints, respectively. We note that the two terms F_{base} and
 235 F_{cont} are independent in terms of their inputs, but are weakly coupled via the additional terms (which
 236 represent a sparse set of ramping constraints and voltage setpoints that tie together the base and
 237 contingency cases). This is an ideal setup for decoupling through non-linear Gauss-Seidel solution
 238 methods. In addition, the contingency-related KKT conditions (13b) are actually identical to the
 239 KKT conditions of the problem (8c) that we solved during the last iteration of the inner maximization
 240 problem; as a result, we can reuse the result of this previous computation when executing our
 241 Gauss-Seidel step. Together, this allows us to inexpensively identify an update direction for x that
 242 nonetheless remains feasible with respect to the base case power flow constraints.



(a) Visualization of the worst-case contingency found, with the degree of the associated partial outage on each line indicated in blue. (Plot generated via PowerWorld.)

(b) Training curve for finding a worst-case contingency. The process converges within 5 iterations, and increases the loss by 3%.

Figure 1: Illustrative example of finding a worst-case N-2 contingency on a 14-node test system.

243 5 Experiments

244 We demonstrate the efficacy of our approach on the settings of N-1, N-2 and N-3 SCOPF. In particular,
 245 noting that quality of adversarial attacks is likely to have a large effect on the success of our overall
 246 procedure, we first visualize the attacks found by our inner maximization process (Sections 3.1 4.4)
 247 on a small power system test case. We then demonstrate the performance of our overall approach
 248 on a realistic 4622-node power system with approximately 6 thousand potential N-1 contingencies,
 249 almost 19 million N-2 contingencies, and over 38 billion N-3 contingencies. We show that CAN ∂ Y
 250 is able to efficiently find solutions that are competitive with other leading approaches in the N-1 case,
 251 while reducing violations in the N-2 and N-3 scenarios compared to a base case optimal power flow.

252 All experiments are run on a single core on Macbook Pro with a 2.6 GHz Core i7 CPU. We implement
 253 our approach in Python, using a custom optimal power flow solver called [redacted] to compute the
 254 optimization (8c), and CVXPY [36] to compute the convex projections for projected gradient descent.
 255 We evaluate the strength of all dispatch solutions using PowerWorld, a commercial power flow tool.

256 5.1 Illustrative adversarial attack

257 Finding worst-case contingencies is often beneficial to power systems engineers, who try to identify
 258 fragile areas of their grid for future development. Traditionally, engineers use linear approximations
 259 of the grid physics [37-39] to identify single outages that pose a significant risk to system stability.
 260 However, given that the underlying physics are fundamentally non-linear, such linear approximations
 261 quickly become inaccurate when trying to identify the risks associated with multiple simultaneous
 262 outages. Our implicit differentiation approach, on the other hand, employs accurate gradient informa-
 263 tion from the physics of the network to quickly identify contingencies that maximally increase our
 264 loss function, or an alternative loss function of choice that also captures system infeasibilities.

265 For ease of visualization, we demonstrate this attack-identification approach on the IEEE 14-node test
 266 system. In particular, we identify an adversarial N-2 contingency on this system in just 5 iterations
 267 (approximately one minute), increasing the value of the loss function by 3% over the base case
 268 scenario, as shown in Figure 1b. This worst-case contingency represents a combination of multiple
 269 partial outages on different lines, and (perhaps surprisingly) does not include any generator outages,
 270 as shown in Figure 1a. This is likely to present a stronger attack than those obtained via the “standard”
 271 linear approximation approach, serving as a potential benefit to power system planners who are trying
 272 to reinforce their grid, as well as to “adversarially robust training” procedures like ours.

273 5.2 Validating N-1 security

274 Today, most grid operators in the United States require that their dispatch be N-1 secure, i.e., secure
 275 against any single outage, prompting the development of associated methods. In particular, the recent

	gollnl	GO-SNIP	GMI-GO	BAT	gravityx	CAN ∂ Y*
GO Challenge 1 Rank	1	2	3	4	5	-
Score for 5K network	546,302	553,152	553,328	545,783	550,020	552,032

Table 1: Comparison of the performance of our method against top-performing submissions to the ARPA-E GO Competition, which addresses N-1 SCOPF. Results are shown for the 4622-node Challenge 1 test case (“5K network”). While the score comparisons shown are inexact due to subtleties of the evaluation metric (see Appendix D), at a high level, we see that CAN ∂ Y performs competitively with all top-scoring methods.




Figure 2: Loss vs. iterations in adversarial training. Each attack stage is at most 10 iterations and each defense stage is one iteration. The loss value after the defense step is in green.

Contingency type	N-1	N-2	N-3
Scenarios tested	6,133	359,712	428,730
OPF violations	36	10,572	4,086
CAN ∂ Y violations*	59	3,580	1,122

Table 2: Number of feasibility violations incurred by the N-3 SCOPF version of our method, and a baseline optimal power flow (OPF) method on randomly selected N-1, N-2, N-3 contingency scenarios for a 4622-node test case. While CAN ∂ Y incurs slightly more N-1 scenario violations than the baseline, we see that it reduces the number of N-2 and N-3 violations by a factor of 3-4 \times over the baseline.

276 Grid Optimization (GO) Competition [40], hosted by ARPA-E, focused on finding algorithms to
 277 solve N-1 SCOPF. Each participating team used a variety of methods to produce a dispatch that was
 278 evaluated on the basis of power cost and feasibility in both the base and contingency cases. In order to
 279 validate that our method works well in the N-1 setting, we solve a particular case from the competition
 280 – namely, the 4622-node test case with a sub-selection of 3071 N-1 potential contingencies, provided
 281 as part of the Challenge 1 stage – by constructing our relaxed contingency set \mathcal{Y} with $k = 1$.

282 We find that our score is comparable against the top approaches submitted to the GO Competition,
 283 as shown in Table I. We note that these score comparisons are not exact, as our power flow solver
 284 uses a more realistic model for coupling between the base and contingency cases than was posed
 285 in Challenge 1, which affects the way that the evaluation metric is computed (see Appendix D).
 286 Nonetheless, at a high level, these results demonstrate that our method performs competitively with
 287 respect to the top-performing methods for solving N-1 SCOPF.

288 5.3 Improving N-3 SCOPF

289 We now describe the performance of our method on our main setting of interest: N-3 SCOPF. While
 290 other competitive methods exist for solving N-1 SCOPF, previous work has struggled to approach
 291 settings allowing for larger numbers of simultaneous outages (e.g., N-k SCOPF for $k = 2$ or 3) due
 292 to the associated combinatorial explosion in problem size. Our method, however, scales gracefully
 293 with respect to the number of allowable simultaneous outages, as we need only tweak the value of k
 294 used within our attack set $\mathcal{Y} := \{y : y \in [0, 1]^{n_o}, \|y\|_1 \leq k\}$. More specifically, each iteration of
 295 the attack maximization and each defense step calculation take approximately the same amount of
 296 time *regardless of the value of k* , given that the costs of the gradient computations, projections, and
 297 (optimal) power flow solves are independent of k . (The total number of iterations it takes for our
 298 method to converge may vary between settings, though we do not notice a substantive difference in
 299 this respect between the N-1, N-2, and N-3 versions of our approach during our experiments.)

300 We use our method to attempt to solve N-3 SCOPF (i.e., set $k = 3$) on a 4622-node test case over
301 all 6,133 potential outages (i.e., over 38 billion N-3 contingencies); the associated training curve is
302 shown in Figure 2. In total, our approach takes only 21 minutes to converge.

303 We evaluate the strength of our obtained dispatch in maintaining feasibility against a combination
304 of N-1, N-2, and N-3 contingency scenarios, which are all technically contained within the threat
305 model represented by our choice of \mathcal{Y} . We note that while full security against all these contingencies
306 is likely impossible with a single dispatch – e.g., we can very often construct an N-3 contingency
307 that isolates, or *islands*, some non-self-sustaining part of the electrical grid – we aim to demonstrate
308 that our method can improve upon existing methods in terms of providing robustness against a wide
309 variety of scenarios. As there remain a lack of available N-2 or N-3 SCOPF methods against which
310 we can readily compare, we compare our performance against that of a base case optimal power flow
311 (OPF) solver. Due to the intractability of evaluating these dispatches on *all* possible contingency
312 scenarios, we randomly sub-select the set of N-1, N-2, and N-3 scenarios on which we evaluate, and
313 run these evaluations in PowerWorld over the course of several days.

314 The results of our evaluation are shown in Table 2. Overall, we see that CAN ∂ Y significantly reduces
315 the number of total contingency violations as compared to the OPF solution by a factor of 3-4 \times .
316 Analyzing the N-3 contingencies in more detail, we find that of the 4086 specific violations incurred
317 by OPF, 931 of those were also incurred by CAN ∂ Y (while the remaining 191 violations incurred by
318 our method were disjoint). Overall, these results indicate that our method is much more effective
319 than OPF at guarding against N-2 and N-3 contingencies, though the actual distribution of specific
320 contingencies that are guarded against may differ between the two methods.

321 6 Conclusion

322 In this paper, we have described our approach, CAN ∂ Y, for N-k security-constrained optimal power
323 flow. Specifically, we formulate N-k SCOPF as a minimax problem over power system dispatches and
324 potential outages by forming a continuous outer approximation to the contingency space. This enables
325 us to compute “worst-case” contingencies via projected gradient descent (using tricks from the implicit
326 layers literature) and then employ these contingencies to update our proposed dispatch. Notably,
327 our formulation scales gracefully in the number of contingencies – requiring only a minor tweak to
328 the projection set during the attack step – even as the underlying problem scales combinatorially.
329 In particular, our approach takes only 21 minutes to converge on a standard laptop for a realistic
330 4622-node test case with over 38 billion potential N-3 contingencies, and reduces feasibility violations
331 by a factor of 3-4 \times compared to a baseline optimal power flow method. Overall, we believe this
332 demonstrates the promise of our approach in enabling scalable N-k security-constrained optimization.

333 We note that the success of our minimax optimization approach is likely highly reliant on the strength
334 of the adversarial attacks that we generate, just as in the adversarially robust training literature
335 [3, 4]. As such, a fruitful direction for future work may involve developing improved procedures
336 for obtaining adversarial attacks in the context of SCOPF. In addition, given the large scale of the
337 power networks we consider, it is generally impossible to evaluate proposed dispatches against the
338 full suite of potential contingencies in order to check whether they are indeed N-k secure. Given that,
339 another fruitful direction may entail developing better evaluation metrics or verification procedures to
340 inexpensively evaluate whether a proposed SCOPF dispatch is (likely) robust, perhaps again drawing
341 inspiration from the literature on verification methods for adversarially robust deep learning [41].

342 7 Broader impacts

343 To address climate and sustainability goals, many power grids are starting to integrate larger amounts
344 of time-varying renewable energy, such as solar and wind. As described in [42], this means power
345 systems optimization problems must generally be solved both more quickly and at larger scales.
346 We believe our work makes an important contribution to this area by providing a scalable and
347 efficient method to address the problem of N-k SCOPF. One potential concern, however, is that our
348 method relies on identifying “worst-case” power system attacks, and our associated approach could
349 theoretically be exploited by adversarial actors; that said, we believe it unlikely that our method in
350 particular would be used for this purpose, as it employs a continuous “partial outage” approximation
351 that is useful for the purposes of our algorithm, but does not map neatly to actual real-world attacks.

References

- 352
353 [1] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*. Princeton university
354 press, 2009.
- 355 [2] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. Theory and applications of robust
356 optimization. *SIAM review*, 53(3):464–501, 2011.
- 357 [3] Zico Kolter and Aleksander Madry. Adversarial robustness - theory and practice. [https://
358 adversarial-ml-tutorial.org/](https://adversarial-ml-tutorial.org/), 12 2018.
- 359 [4] Han Xu, Yao Ma, Hao-Chen Liu, Debayan Deb, Hui Liu, Ji-Liang Tang, and Anil K Jain. Adversarial
360 attacks and defenses in images, graphs and text: A review. *International Journal of Automation and
361 Computing*, 17(2):151–178, 2020.
- 362 [5] Eric Wong, Frank R Schmidt, Jan Hendrik Metzen, and J Zico Kolter. Scaling provable adversarial defenses.
363 In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages
364 8410–8419, 2018.
- 365 [6] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples.
366 *arXiv preprint arXiv:1801.09344*, 2018.
- 367 [7] Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial
368 polytope. In *International Conference on Machine Learning*, pages 5286–5295. PMLR, 2018.
- 369 [8] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples.
370 *International Conference on Learning Representations*, 2015.
- 371 [9] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. To-
372 wards deep learning models resistant to adversarial attacks. In *International Conference on Learning
373 Representations*, 2018.
- 374 [10] Zico Kolter, David Duvenaud, and Matthew Johnson. Tutorial: Deep implicit layers - neural odes, deep
375 equilibrium models, and beyond. <http://implicit-layers-tutorial.org/>, 12 2020.
- 376 [11] Brandon Amos and J Zico Kolter. OptNet: Differentiable Optimization as a Layer in Neural Networks. In
377 *International Conference on Machine Learning*, pages 136–145. PMLR, 2017.
- 378 [12] Josip Djolonga and Andreas Krause. Differentiable learning of submodular models. In *Advances in Neural
379 Information Processing Systems (NeurIPS)*, 2017.
- 380 [13] Sebastian Tschiatschek, Aytunc Sahin, and Andreas Krause. Differentiable submodular maximization.
381 *Preprint arXiv:1803.01785*, 2018.
- 382 [14] Priya L Donti, Inês Lima Azevedo, and J Zico Kolter. Inverse optimal power flow: Assessing the
383 vulnerability of power grid data. *Workshop on Modeling the Physical World: Perception, Learning, and
384 Control at NeurIPS*, 2018.
- 385 [15] Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J Zico Kolter. Dif-
386 ferentiable convex optimization layers. In *Advances in Neural Information Processing Systems (NeurIPS)*,
387 2019.
- 388 [16] Stephen Gould, Richard Hartley, and Dylan Campbell. Deep declarative networks: A new hope. *Preprint
389 arXiv:1909.04866*, 2019.
- 390 [17] Po-Wei Wang, Priya L Donti, Bryan Wilder, and Zico Kolter. SATNet: Bridging deep learning and logical
391 reasoning using a differentiable satisfiability solver. In *International Conference on Machine Learning
392 (ICML)*, 2019.
- 393 [18] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential
394 equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- 395 [19] Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. End-to-end
396 differentiable physics for learning and control. In *Advances in Neural Information Processing Systems
397 (NeurIPS)*, 2018.
- 398 [20] Chun Kai Ling, Fei Fang, and J Zico Kolter. What game are we playing? End-to-end learning in normal
399 and extensive form games. *Preprint arXiv:1805.02777*, 2018.

- 400 [21] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural*
401 *Information Processing Systems (NeurIPS)*, 2019.
- 402 [22] Priya L Donti, Brandon Amos, and J Zico Kolter. Task-based end-to-end model learning in stochastic
403 optimization. *Preprint arXiv:1703.04529*, 2017.
- 404 [23] Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused
405 learning for combinatorial optimization. In *AAAI Conference on Artificial Intelligence*, 2018.
- 406 [24] Jan Larsen, Lars Kai Hansen, Claus Svarer, and M Ohlsson. Design and regularization of neural networks:
407 the optimal use of a validation set. In *Neural Networks for Signal Processing VI. Proceedings of the 1996*
408 *IEEE Signal Processing Society Workshop*, pages 62–71. IEEE, 1996.
- 409 [25] Yoshua Bengio. Gradient-based optimization of hyperparameters. *Neural computation*, 12(8):1889–1900,
410 2000.
- 411 [26] Fabian Pedregosa. Hyperparameter optimization with approximate gradient. In *International conference*
412 *on machine learning*, pages 737–746. PMLR, 2016.
- 413 [27] Matthew MacKay, Paul Vicol, Jon Lorraine, David Duvenaud, and Roger Grosse. Self-tuning net-
414 works: Bilevel optimization of hyperparameters using structured best-response functions. *arXiv preprint*
415 *arXiv:1903.03088*, 2019.
- 416 [28] Wanxin Jin, Zhaoran Wang, Zhuoran Yang, and Shaoshuai Mou. Pontryagin differentiable programming:
417 An end-to-end learning and control framework. *arXiv preprint arXiv:1912.12970*, 2019.
- 418 [29] Mohammadhafez Bazrafshan, Kyri Baker, and Javad Mohammadi. Computationally efficient solutions for
419 large-scale security-constrained optimal power flow. *arXiv preprint arXiv:2006.00585*, 2020.
- 420 [30] Leonel de Magalhães Carvalho, Armando Martins Leite da Silva, and Vladimiro Miranda. Security-
421 constrained optimal power flow via cross-entropy method. *IEEE Transactions on Power Systems*,
422 33(6):6621–6629, 2018.
- 423 [31] Mingyu Yan, Mohammad Shahidehpour, Aleks Paaso, Liuxi Zhang, Ahmed Alabdulwahab, and Abdullah
424 Abusorrah. A convex three-stage scopf approach to power system flexibility with unified power flow
425 controllers. *IEEE Transactions on Power Systems*, 36(3):1947–1960, 2021.
- 426 [32] Qianfan Wang, Jean-Paul Watson, and Yongpei Guan. Two-stage robust optimization for n-k contingency-
427 constrained unit commitment. *IEEE Transactions on Power Systems*, 28(3):2366–2375, 2013.
- 428 [33] Alexandre Moreira, Alexandre Street, and José M Arroyo. An adjustable robust optimization approach
429 for contingency-constrained transmission expansion planning. *IEEE Transactions on Power Systems*,
430 30(4):2013–2022, 2014.
- 431 [34] Shaoyun Hong, Haozhong Cheng, and Pingliang Zeng. N-k constrained composite generation and
432 transmission expansion planning with interval load. *IEEE Access*, 5:2779–2789, 2017.
- 433 [35] Steven G Krantz and Harold R Parks. *The implicit function theorem: history, theory, and applications*.
434 Springer Science & Business Media, 2012.
- 435 [36] Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex optimiza-
436 tion. *The Journal of Machine Learning Research*, 17(1):2909–2913, 2016.
- 437 [37] P. Kaplunovich and K. Turitsyn. Fast and reliable screening of n-2 contingencies. *IEEE Transactions on*
438 *Power Systems*, 31(6):4243–4252, 2016.
- 439 [38] Liang Che, Xuan Liu, and Zuyi Li. Screening hidden n- k line contingencies in smart grids using a
440 multi-stage model. *IEEE Transactions on Smart Grid*, 10(2):1280–1289, 2019.
- 441 [39] Saqib Hasan, Amin Ghafouri, Abhishek Dubey, Gabor Karsai, and Xenofon Koutsoukos. Heuristics-based
442 approach for identifying critical n — k contingencies in power systems. In *2017 Resilience Week (RWS)*,
443 pages 191–197, 2017.
- 444 [40] ARPA-E. Grid Optimization (GO) Competition. <https://gocompetition.energy.gov/> 2019.
- 445 [41] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian
446 Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint*
447 *arXiv:1902.06705*, 2019.
- 448 [42] David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran,
449 Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. Tackling
450 climate change with machine learning. *arXiv preprint arXiv:1906.05433*, 2019.

451 **Checklist**

- 452 1. For all authors...
- 453 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
454 contributions and scope? [Yes]
- 455 (b) Did you describe the limitations of your work? [Yes] We note throughout the paper that
456 our method does not provide provable N-k security because it is a heuristic method.
- 457 (c) Did you discuss any potential negative societal impacts of your work? [Yes] Section 7
458 discusses the broader impacts of our work, including potential negative aspects.
- 459 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
460 them? [Yes]
- 461 2. If you are including theoretical results...
- 462 (a) Did you state the full set of assumptions of all theoretical results? [N/A] While we
463 do not include theoretical results, we do – to the best of our knowledge – state all
464 assumptions associated with the computations we present.
- 465 (b) Did you include complete proofs of all theoretical results? [N/A]
- 466 3. If you ran experiments...
- 467 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
468 mental results (either in the supplemental material or as a URL)? [No] As parts of the
469 power flow solver, [redacted], that we use are proprietary, we only include the select
470 portion of the code that shows how we execute our minimax optimization procedure.
- 471 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
472 were chosen)? [N/A] The only hyperparameter we explicitly set is the convergence
473 tolerance of the method. All hyperparameters associated with the [redacted] power
474 flow solver were determined outside the context of this work.
- 475 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
476 ments multiple times)? [N/A] All of our experiments are deterministic and were only
477 run once each.
- 478 (d) Did you include the total amount of compute and the type of resources used (e.g.,
479 type of GPUs, internal cluster, or cloud provider)? [Yes] We include these details in
480 Section 5.
- 481 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 482 (a) If your work uses existing assets, did you cite the creators? [No] We evaluate our meth-
483 ods on the GO Competition dataset (which we cite) and the IEEE 14 node test system
484 (for which there is no canonical citation). We use commercial software, PowerWorld,
485 to evaluate the strength of our dispatches, as we describe in the text. We also use an
486 existing, custom power flow solver called [redacted], but do not cite it for now in order
487 to preserve the anonymity of this submission.
- 488 (b) Did you mention the license of the assets? [N/A]
- 489 (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
- 490
- 491 (d) Did you discuss whether and how consent was obtained from people whose data you’re
492 using/curating? [N/A]
- 493 (e) Did you discuss whether the data you are using/curating contains personally identifiable
494 information or offensive content? [N/A]
- 495 5. If you used crowdsourcing or conducted research with human subjects...
- 496 (a) Did you include the full text of instructions given to participants and screenshots, if
497 applicable? [N/A]
- 498 (b) Did you describe any potential participant risks, with links to Institutional Review
499 Board (IRB) approvals, if applicable? [N/A]
- 500 (c) Did you include the estimated hourly wage paid to participants and the total amount
501 spent on participant compensation? [N/A]