

A Learning-to-Infer Method for Real-Time Power Grid Multi-Line Outage Identification

Yue Zhao[✉], Member, IEEE, Jianshu Chen, Member, IEEE,
and H. Vincent Poor[✉], Fellow, IEEE

Abstract—Identifying a potentially large number of simultaneous line outages in power transmission networks in real time is a computationally hard problem. This is because the number of hypotheses grows exponentially with the network size. A new “Learning-to-Infer” method is developed for efficient inference of every line status in the network. Optimizing the line outage detector is transformed to and solved as a discriminative learning problem based on Monte Carlo samples generated with power flow simulations. A major advantage of the developed Learning-to-Infer method is that the labeled data used for training can be generated in an arbitrarily large amount rapidly and at very little cost. As a result, the power of offline training is fully exploited to learn very complex classifiers for effective real-time multi-line outage identification. The proposed methods are evaluated in the IEEE 30, 118, and 300 bus systems. Excellent performance in identifying multi-line outages in real time is achieved with a reasonably small amount of data.

Index Terms—Line outage detection, power system monitoring, machine learning, variational inference, Monte Carlo method.

I. INTRODUCTION

LACK of situational awareness in abnormal system conditions is a major cause of blackouts in power networks [3]. Network component failures such as transmission line outages, if not rapidly identified and contained, can quickly escalate to cascading failures. In particular, when line failures happen, the power network topology changes instantly, newly stressed areas can unexpectedly emerge, and subsequent failures may be triggered that lead to increasingly complex network topology changes. While the power system is usually protected against the so called “ $N - 1$ ” failure scenarios (i.e., only one component fails), as failures accumulate,

Manuscript received June 8, 2018; revised September 17, 2018; accepted November 6, 2018. Date of publication June 27, 2019; date of current version December 23, 2019. This work was supported in part by an SBU-BNL SEED Grant and in part by the National Science Foundation under Grant DMS-1736417 and Grant ECCS-1824710. Some preliminary results from this work were presented in part at the IEEE Workshop on Statistical Signal Processing, Palma de Mallorca, Spain, 2016 [1] and in part at the IEEE Global Conference on Signal and Information Processing (GlobalSIP), Arlington, VA, USA, 2016 [2]. Paper no. TSG-00862-2018. (Corresponding author: Yue Zhao.)

Y. Zhao is with the Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY 11794 USA (e-mail: yue.zhao.2@stonybrook.edu).

J. Chen is with the Natural Language Processing Group, Tencent AI Lab, Bellevue, WA 98004 USA (e-mail: jianshuchen@tencent.com).

H. V. Poor is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: poor@princeton.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSG.2019.2925405

effective automatic protection is no longer guaranteed. Thus, when cascading failures start developing, effective protective actions/interventions critically depend on correct and timely knowledge of the network status. Indeed, without accurate knowledge of the line outages, protective control methods have been observed to further aggravate the failure scenarios [4]. Thus, real-time line outage identification is essential to all network control decisions for mitigating failures. In particular, since the first few line outages may have already escaped the operators’ attention, the ability to identify in real time the network topology with an arbitrary number of line outages becomes critical to prevent system collapse.

Real-time line outage identification is however a very challenging problem, especially when unknown line outages in the network quickly accumulate as in scenarios that cause large-scale blackouts [3]. *The number of outage hypotheses grows exponentially with the number of possible line outages, making real-time multi-line outage identification fundamentally hard.* Other limitations in practice such as behaviors of human operators under time pressure, missing and contradictory information, and privacy concerns over data sharing can make this problem even harder. Assuming a small number of line failures, exhaustive search methods have been developed in [5], [6], [7] and [8] based on hypothesis testing, and in [9] based on logistic regression. To overcome the prohibitive computational complexity of exhaustive search methods, [10] has developed sparsity exploiting outage identification methods with overcomplete observations to identify sparse multi-line outages. Without assuming sparsity of line outages, a graphical model based approach has been developed for identifying arbitrary multi-line outages [11]. Sequential line outage detection has also been proposed [12].

On a related note, *non-real-time* power grid topology identification has also been extensively studied: the underlying topology stays the same, while many data are collected over a relatively long period of time before the topology can be identified [13], [14], [15]. A variety of types of data have been exploited for addressing this problem, e.g., data on power injections [16], voltage correlation [17], and energy prices [18]. For power distribution systems in particular, various graph learning approaches have also been developed [19], [20].

In this paper, we focus on *real-time* identification of a potentially large number of simultaneous line outages based on a set of measurements collected at any one point of time in the power system. We start with a probabilistic

model of the variables in a power system (line statuses, power injections, voltages, power flows, currents etc.) and in its monitoring system (sensor measurements of physical quantities). We then formulate the multi-line outage identification problem in a Bayesian inference framework, where we aim to compute the posterior probabilities of the post-outage topologies given any measurements at any one point of time.

To overcome the fundamental computational complexity due to the *exponentially large number of possible post-outage topologies*, we develop a learning based framework inspired by (but different from) variational inference, in which we aim to approximate the desired posterior probabilities using models that allow *computationally easy marginal inference of line statuses*. Importantly, we develop “end-to-end” predictor models for multi-line outage identification, and allow arbitrary model structures and complexities. In order to find effective end-to-end predictor models, we transform optimizing a predictor model into a discriminative learning problem leveraging a Monte Carlo approach: a) based on *full-blown power flow equations*, data samples of network topology, network states, and sensor measurements in the network can be efficiently generated according to a *generative model* of these quantities, and b) with these simulated data, *discriminative models* are learned *offline*, which then offer *real-time* prediction of the line outages based on newly observed measurements from the real network. We thus term the proposed method “*Learning-to-Infer*”. It is important to note that this *Learning-to-Infer* method is not limited by any potential lack of real-world data, as the offline training procedure can be conducted entirely based on simulated data.

A major strength of the proposed *Learning-to-Infer* method is that the *labeled data set* for training the predictor model can be generated in an arbitrarily large amount, at very little cost. As such, we can fully exploit the benefit of offline model training in order to get accurate online multi-line outage identification performance. The proposed approach is also not restricted to specific models and learning methods, but can exploit any powerful models such as deep neural networks [21]. As a result, predictor models of very high complexities can be adopted, yet without worrying about overfitting since more labeled training data can always be generated if overfitting is observed.

The developed *Learning-to-Infer* method is evaluated on the IEEE 30, 118, and 300 bus systems [22] for identifying an *arbitrary* number of line outages. It is demonstrated that, even with relatively simple predictor models and a reasonably small amount of data, the performance is surprisingly good for this very challenging task.

The remainder of the paper is organized as follows. Section II introduces the system model, and formulates real-time multi-line outage identification as a Bayesian inference problem. Section III develops the *Learning-to-Infer* method. Section IV discusses the architectures of neural networks employed in this study. Section V presents the results from our numerical experiments. Section VI concludes the paper.

II. PROBLEM FORMULATION

A. Power Flow Models

We consider a power system with N buses, and its *baseline topology* (i.e., the network topology when there is no line outage) with L lines. We denote the incidence matrix of the baseline topology by $M \in \{-1, 0, 1\}^{N \times L}$ [23]. We use a binary variable s_l to denote the status of a line l , with $s_l = 1$ for a connected line l , and 0 otherwise. The actual topology of the network can then be represented by $s = [s_1, \dots, s_L]^T$. Generalizing this notation with a bit of abuse, we also employ $s_{mn} \in \{1, 0\}$ to denote whether two buses m and n are connected by a line or not. (For simplicity, we consider any two buses can be connected by at most one line.) Given a network topology s , the system’s bus admittance matrix Y can be determined accordingly with the physical parameters of the system [24]: $Y_{mn} = s_{mn}(G_{mn} + jB_{mn})$, where G_{mn} and B_{mn} denote conductance and susceptance, respectively. Note that, when two buses m and n are not connected, $Y_{mn} = s_{mn} = 0$.

We denote the real and reactive power injections at all the buses by $\mathbf{P}, \mathbf{Q} \in \mathbb{R}^N$, and the voltage magnitudes and phase angles by $\mathbf{V}, \boldsymbol{\theta} \in \mathbb{R}^N$. Given the bus admittance matrix Y , the nodal power injections and the nodal voltages satisfy the following AC power flow equations [24]: $\forall m = 1, \dots, N$,

$$\begin{aligned} P_m &= V_m \sum_{n=1}^N V_n s_{mn} (G_{mn} \cos(\theta_m - \theta_n) + B_{mn} \sin(\theta_m - \theta_n)), \\ Q_m &= V_m \sum_{n=1}^N V_n s_{mn} (G_{mn} \sin(\theta_m - \theta_n) - B_{mn} \cos(\theta_m - \theta_n)), \end{aligned} \quad (1)$$

where a subscript m denotes the m^{th} component of a vector. In particular, given the network topology s and a set of controlled input values $\{\mathbf{P}, \mathbf{Q}^{\text{in}}, \mathbf{V}^{\text{in}}\}$, (where \mathbf{Q}^{in} and \mathbf{V}^{in} consist of some subsets of \mathbf{Q} and \mathbf{V} , respectively,) the remaining values of $\{\mathbf{Q}, \mathbf{V}, \boldsymbol{\theta}\}$ can be determined by solving (1). Typically, apart from a slack bus, most buses are “*PQ buses*” at which the real and reactive power injections are controlled inputs, and the remaining buses are “*PV buses*” at which the real power injection and voltage magnitude are controlled inputs [24]. We refer the readers to [24] for further details of solving AC power flow equations.

A useful approximation of the AC power flow model is the DC power flow model: under a topology s , the nodal real power injections and voltage phase angles approximately satisfy the following equation [24]:

$$\mathbf{P} = M \boldsymbol{\Gamma} M^T \boldsymbol{\theta}, \quad (2)$$

where $S = \text{diag}(s_1, \dots, s_L)$, $\boldsymbol{\Gamma} = \text{diag}(\frac{1}{x_1}, \dots, \frac{1}{x_L})$, and x_l is the reactance of line l . We note that, in the DC power flow model, reactive power is not considered, and all voltage magnitudes are approximated by a constant.

B. Observation Models

To monitor the power system, we consider real-time measurements taken by sensors measuring nodal voltage magnitudes and phase angles, current magnitudes and phase angles

on lines, real and reactive power flows on lines, nodal real and reactive power injections, etc. In general, the observation model can be written as the following:

$$\mathbf{y} = \mathbf{h}(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}) + \mathbf{v}, \quad (3)$$

where a) $\mathbf{y} \in \mathbb{R}^K$ collects all the noisy measurements, b) $\mathbf{h}(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}) = [h_1(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}), \dots, h_K(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in})]^T$ denotes the *noiseless* values of the measured quantities, and the forms of $\{h_k(\cdot)\}$ depend on the specific locations and types of the sensors, and c) \mathbf{v} is a vector of measurement noises.

Remark 1: A noiseless measurement function $h_k(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in})$ can be an *implicit function* without a closed form expression. For example, given $\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}$ and \mathbf{V}^{in} , while the nodal voltage magnitude and phase angle at a particular PQ bus can be solved from (1), such a solution can only be obtained using numerical methods, and a closed form expression is not available. For a discussion of the existence and uniqueness of the solution to the power flow equations (1), we refer the reader to [25].

The observation models can be significantly simplified under the approximate DC power flow model (2). For example, measurements of $\boldsymbol{\theta}$ provided by phasor measurement units (PMUs) located at a subset of the buses \mathcal{M} can be modeled as

$$\mathbf{y} = \boldsymbol{\theta}_{\mathcal{M}} + \mathbf{v}, \quad (4)$$

where $\boldsymbol{\theta}_{\mathcal{M}}$ is formed by entries of $\boldsymbol{\theta}$ from buses in \mathcal{M} . From the DC power flow model (2), we have

$$\boldsymbol{\theta} = (MS\Gamma M^T)^+ \mathbf{P}, \quad (5)$$

where $(\cdot)^+$ denotes the pseudoinverse.¹ We note that, while the noiseless voltage phase angle measurements admit a closed form (5) and are linear in the power injections \mathbf{P} , they are *not* linear in the line statuses \mathbf{s} ($= \text{diag}(S)$).

C. Multi-Line Outage Identification as Bayesian Inference

We are interested in identifying the post-outage network topology \mathbf{s} *in real time* based on instant measurements \mathbf{y} collected in the power system. We formulate this multi-line outage identification problem as a *Bayesian inference* problem. First, we model $\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}$ and \mathbf{y} with a joint probability distribution,

$$\begin{aligned} p(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}, \mathbf{y}) \\ = p(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}) \cdot p(\mathbf{y}|\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}). \end{aligned} \quad (6)$$

It is important to note that, given $\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}$, and \mathbf{V}^{in} , the *noiseless* measurements \mathbf{h} (see (3)) can be *exactly computed* by solving the AC power flow equations (1). Adding noises to \mathbf{h} then leads to $p(\mathbf{y}|\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in})$.

Remark 2 (Generative Model): (6) represents a *generative model* [26] with which a) the topology and the controlled inputs of power injections and voltage magnitudes are generated according to a prior distribution $p(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in})$, and

¹For a connected network, the solution of $\boldsymbol{\theta}$ given \mathbf{P} is made unique by setting the phase angle at a reference bus to be zero.

b) all the quantities \mathbf{h} measured in the system can then be computed by solving the power flow equations (1), based on which the actual noisy measurements \mathbf{y} follow the conditional probability distribution $p(\mathbf{y}|\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in})$. We note that, as in many Bayesian inference problems, an accurate prior distribution $p(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in})$ may be difficult to obtain in practice. Nonetheless, a sharp concentration of the posterior distribution on the true post-outage network topology allows effective inference of multi-line outages even in the absence of accurate knowledge of the prior.

Our objective is to infer the topology of the power grid \mathbf{s} given the observed measurements \mathbf{y} . Thus, under a Bayesian inference framework, we are interested in computing the posterior *conditional* probabilities: $\forall \mathbf{s}$,

$$p(\mathbf{s}|\mathbf{y}) = \frac{\int p(\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}) p(\mathbf{y}|\mathbf{s}, \mathbf{P}, \mathbf{Q}^{in}, \mathbf{V}^{in}) d\mathbf{P} d\mathbf{Q}^{in} d\mathbf{V}^{in}}{p(\mathbf{y})}. \quad (7)$$

Given the observations \mathbf{y} , a maximum a-posteriori probability (MAP) detector would pick $\text{argmax}_{\mathbf{s}} p(\mathbf{s}|\mathbf{y})$ as the topology/multi-line outage identification decision, which minimizes the identification error probability [27]. However, as the number of hypotheses of \mathbf{s} grows exponentially with the number of unknown line statuses, performing such a hypothesis testing based on an exhaustive search becomes computationally intractable. *In general, as there are up to 2^L possibilities for \mathbf{s} , computing, or even listing the probabilities $p(\mathbf{s}|\mathbf{y})$, $\forall \mathbf{s}$ has an exponential complexity.*

Posterior Marginal Probabilities: As an initial step towards addressing the fundamental challenge of computational complexity, instead of computing $p(\mathbf{s}|\mathbf{y})$, we focus on computing the *posterior marginal conditional probabilities* $p(s_l|\mathbf{y})$, $l = 1, \dots, L$. We note that the posterior marginals are characterized by just L numbers, $\mathbb{P}(s_l = 1|\mathbf{y})$, $l = 1, \dots, L$, as opposed to $2^L - 1$ numbers required for characterizing $p(\mathbf{s}|\mathbf{y})$. Accordingly, the hypothesis testing problem on \mathbf{s} is decoupled into L separate *binary* hypothesis testing problems: for each line l , the MAP detector identifies $\text{argmax}_{s_l \in \{0, 1\}} p(s_l|\mathbf{y}, \mathbf{P})$. As a result, instead of minimizing the identification error probability of the vector \mathbf{s} , the binary MAP detectors minimize the identification error probability of each line status s_l .

Although listing the posterior marginals $p(s_l|\mathbf{y})$ is tractable, computing them, however, still remains intractable. In particular, even with $p(\mathbf{s}|\mathbf{y})$ given, summing out all s_k , $k \neq l$, to obtain $p(s_l|\mathbf{y})$ still requires exponential computational complexity [28]. As a result, even a *binary* MAP decision about s_l cannot be made in a computationally tractable way. This challenge will be addressed by a novel method we will develop in the next section.

III. A LEARNING-TO-INFER METHOD

A. A Variational Inference-Inspired Framework

In this section, we develop a variational inference-inspired method for approximate inference of the posterior marginal conditional probabilities $p(s_l|\mathbf{y})$, $l = 1, \dots, L$. The general idea is to find a conditional distribution $q(\mathbf{s}|\mathbf{y})$ that

- a) approximates the original $p(\mathbf{s}|\mathbf{y})$ very closely, and

b) offers fast and accurate multi-line outage identification results based on easily computable $q(s_l|y)$, $\forall l$.

In particular, we consider that $q(s|y)$ is modeled by some parametric form (e.g., a neural network), and is hence chosen from some family of parametrized conditional probability distributions $\{q_{\beta}(s|y)\}$, where β is a vector of model parameters. It is worth noting that $q(s|y)$ is a *function of both s and y*, and the parameters β associate both s and y with the probability value $q_{\beta}(s|y)$, for all possible s and y .

To achieve the two goals above, we aim to choose a family of probability distributions $\{q_{\beta}(s|y)\}$ to satisfy the following:

- The parametric form of $\{q_{\beta}(s|y)\}$ has sufficient expressive power to represent very complicated functions, so that our approximation to the true $p(s_l|y)$ can be made sufficiently precise.
- It is easy to compute the marginal $q_{\beta}(s_l|y)$, so that we can use it to infer s_l with low computational complexity in real time based on the observed y .

From a family of parametrized distributions $\{q_{\beta}(s|y)\}$, we would like to choose a $q_{\beta}(s|y)$ that approximates $p(s|y)$ as closely as possible. For this, we employ the *Kullback-Leibler (KL) divergence* as a metric of closeness between two probability distributions,

$$D(p\|q_{\beta}) \triangleq \sum_s p(s|y) \log \frac{p(s|y)}{q_{\beta}(s|y)}. \quad (8)$$

Note that, for any particular realization of observations y , a KL divergence $D(p\|q_{\beta})$ can be computed. Thus, $D(p\|q_{\beta})$ can be viewed as a function of y . Since we would like the parametrized conditional $q_{\beta}(s|y)$ to closely approximate $p(s|y)$ for all y , we would like to minimize the *expected KL divergence* as follows:

$$\begin{aligned} & \min_{\beta} \mathbb{E}_y[D(p\|q_{\beta})] \\ & \Leftrightarrow \min_{\beta} \sum_y p(y) \sum_s p(s|y) \log \frac{p(s|y)}{q_{\beta}(s|y)} \\ & \Leftrightarrow \min_{\beta} \sum_{s,y} p(s,y) \log \frac{p(s|y)}{q_{\beta}(s|y)} \\ & \Leftrightarrow \max_{\beta} \mathbb{E}_{s,y}[\log q_{\beta}(s|y)], \end{aligned} \quad (9)$$

where the expectation is taken with respect to the *true* distribution $p(s,y)$.

B. From Generative Model to Discriminative Learning

Evaluating $\mathbb{E}_{s,y}[\log q_{\beta}(s|y)]$ is, however, very difficult, primarily because it again requires the summation of an exponentially large number of terms. To address this, the key step forward is that we can *approximate the expectation by the empirical mean* of $\log q_{\beta}(s|y)$ over a large number of *Monte Carlo samples*, generated according to (ideally) the *true* joint probability $p(s, P, Q^{in}, V^{in}, y)$ (see (6)). We denote the relevant Monte Carlo samples by $\{s^i, y^i; i = 1, \dots, I\}$. Accordingly, (9) is approximated by the following:

$$\max_{\beta} \frac{1}{I} \sum_{i=1}^I \log q_{\beta}(s^i|y^i). \quad (10)$$

TABLE I
THE LEARNING-TO-INFER METHOD

Offline computation:

1. Generate a *labeled* data set $\{s^i, y^i\}$ using Monte Carlo simulations with the complete power flow and sensor models.
2. Select a parametrized predictor model $\{q_{\beta}(s|y)\}$.
3. Train the model parameters β using the generated data set.

Online inference (in real time):

1. Collect instant measurements y from the system.
2. Compute the approximate posterior marginals $q_{\beta^*}(s_l|y)$, $l = 1, \dots, L$, and infer the line statuses $\{s_l\}$.

With a data set $\{s^i, y^i\}$ generated using Monte Carlo simulations, (10) can then be solved as a deterministic optimization problem. The optimal solution of the model parameters β^* approaches that for the original problem (9) as $I \rightarrow \infty$.

In fact, the problem (10) can be viewed as an *empirical risk minimization* problem in machine learning [29], as it trains a *discriminative model* $q_{\beta}(s|y)$ with a data set $\{s^i, y^i\}$ generated from a *generative model* $p(s, P, Q^{in}, V^{in}, y)$ (see Remark 2). As a result of this offline learning/training process (10), an approximate posterior function $q_{\beta^*}(s|y)$ is obtained. Furthermore, it can be shown that (10) is equivalent to finding the maximum likelihood estimate of β on the data set $\{s^i, y^i\}$.

C. Offline Learning for Online Inference

It is important to note that,

- a) the training process to obtain the function $q_{\beta^*}(s|y)$ is conducted completely *offline*;
- b) the *use* of the trained function $q_{\beta^*}(s|y)$ is, however, *in real time*, i.e., *online*.

In particular, in real time, given *newly observed measurements* y of the system, based on $q_{\beta^*}(s|y)$, the L *approximate posterior marginals* $q_{\beta^*}(s_l|y)$, $l = 1, \dots, L$, will be computed. Based on such quickly computed $q_{\beta^*}(s_l|y)$, a detection decision of whether line l ($= 1, \dots, L$) is connected or not in the current topology will be made. For example, a MAP detector would make the following decision:

$$\forall l = 1, \dots, L, \hat{s}_l = \begin{cases} 0, & \text{if } q_{\beta^*}(s_l = 0|y) > 0.5, \\ 1, & \text{otherwise.} \end{cases} \quad (11)$$

Accordingly, we name our proposed methodology Learning-to-Infer: To perform *real time inference* of multi-line outages, we exploit *offline learning* to train a detector based on *labeled* data simulated from the complete physical model of the power system. The methodology is summarized in Table I. A system diagram is plotted in Figure 1.

Remark 3 (Training Binary Classifiers): Any detector that identifies the status of a line l (e.g., a binary MAP detector) can also be viewed as a *binary classifier* $\hat{s}_l(y) \in \{0, 1\}$: For each possible realization of y , this classifier outputs an inferred status of line l . From this perspective, solving (10) is exactly a *supervised learning* process based on a *labeled* data set, $\{s^i, y^i\}$, where $\{s^i\}$ are the output labels that correspond to the

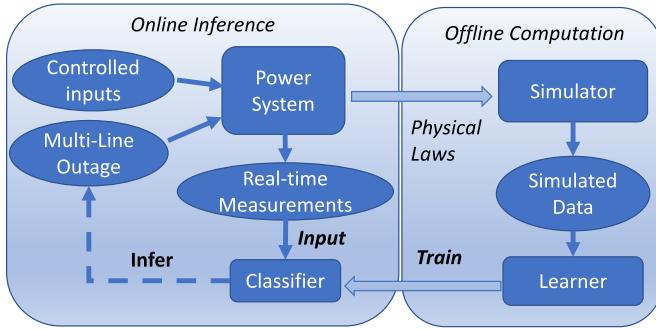


Fig. 1. Overall architecture of the Learning-to-Infer method.

input data $\{y^i\}$. As a result, the rich literature on supervised learning for training binary classifiers directly applies to our problem under this Learning-to-Infer framework.

Remark 4 (Difference From Variational Inference): It is worth noting the fundamental difference between the proposed Learning-to-Infer method and variational inference methods. Importantly, for every new inference instance given a new observation y , variational inference methods need to call an optimization procedure to solve for a new variational model. In contrast, Learning-to-Infer only trains the predictor $q_\beta(s|y)$ once in an offline fashion, and simply calls the trained $q_\beta(s|y)$ for any new inference instance given a new observation y . As such, the online computation time needed by Learning-to-Infer is very little (e.g., performing a forward pass in a neural network), whereas that needed by variational inference methods is much more significant. In essence, Learning-to-Infer exploits the underlying lower dimensional structure of $p(s|y)$ to achieve *generalizability* of the trained predictor $q_\beta(s|y)$ to all possible new observations y .

D. Advantages of the Proposed Method

One significant advantage of this Learning-to-Infer method is that we can generate *labeled data* very efficiently. Specifically, we can efficiently sample from the generative model of $p(s, P, Q^{in}, V^{in}, y)$ (see (6)) as long as we have some prior $p(s, P, Q^{in}, V^{in})$ that is easy to sample from. While historical data and expert knowledge would surely help in forming such priors, using simple uninformative priors can already suffice as will be shown later in the numerical examples. As a result, we can obtain *an arbitrarily large set of data at very little cost* to train the discriminative model. This is quite different from the typical situations encountered in machine learning problems, where obtaining a large amount of labeled data is usually expensive as it requires extensive human annotation effort.

Furthermore, once the approximate posterior distribution $q_\beta(s|y)$ is learned, it can be deployed to infer the multi-line outages *in real-time* as the computation complexity of $q_\beta(s|y)$ is very low by design. This is especially important in monitoring large-scale power grids in real time, because, although training $q_\beta(s|y)$ could take a reasonably long time, the inference speed is very fast. Therefore, the learned predictor q can be used in real time with low-cost hardware.

Limitations of Historical Data and Power of Simulated Data: In overcoming the computational complexity challenges

of real-time multi-line outage identification, it is particularly worth noting the fundamental limitation of using real historical data. Even with the explosion of data available from pervasive sensors in power systems, the data are often collected under a very limited set of system scenarios. For example, most historical data are collected under normal system topologies. Even with data collected under slowly updated systems or faulty systems, *the underlying topologies in these real world cases only represent an extremely small fraction of the entire, exponentially large set of all topologies*. Consequently, historical data are fundamentally insufficient for real-time multi-line outage identification especially under rare failure events.

Simulated data, as evidenced in the proposed Learning-to-Infer framework, offer significant potential beyond what historical data can offer. An *orders of magnitude richer set of scenarios* can be generated, and a learning procedure based on these simulated data can provide very powerful classifiers for identifying arbitrary multi-line outages that may appear in the future, but have not appeared in the past including the simulated scenarios. Last but not least, it is important to note that the simulated scenarios needed for the proposed Learning-to-Infer method would still be *a very small fraction* of the entire, exponentially large model space, as will be demonstrated later in the numerical experiments. As such, it is the good *generalizability* of the classifiers trained using the simulated data that enables effective outage inference under new failure events.

Remark 5 (Learning from the Physical Model): In the proposed Learning-to-Infer method, the training process is at heart learning from the underlying power system physical model. Instead of manually deriving outage detection rules from analyzing the physical model, the proposed method uses a training procedure to learn such rules from massive data generated *according to* the physical model. As such, the rich information embedded in the physical model is carried by the data simulated with it, and then learned by the predictor from training with these simulated data. The Learning-to-Infer method is thus a systematic “indirect” way of learning and using the information from the physical model.

Remark 6 (Side Information and Change of Settings): An interesting question on generalizing the Learning-to-Infer method is how additional information (other than the observed y) may be incorporated. For example, the system operator may receive the side information that certain lines are active for sure. Furthermore, there can also be more *systematic* changes on what information is collected, notably, change of the measurement set y due to installation of additional sensors. For incorporating additional information, one way is to introduce additional inputs to the predictor during the offline training process. For example, we can let each line have a “prior” (even though in reality it can come from a posterior knowledge source) which is fed into the predictor. The data set generation and training would then need to include varying priors of these. Furthermore, a systematic way of dealing with slowly updating priors as well as changes in the measurement sets is to employ “Transfer Learning”. Specifically, the changes in the measurement sets tend not to be so dramatic over a short period of time. Thus, the previously trained neural network can serve as a good initial point when we tune

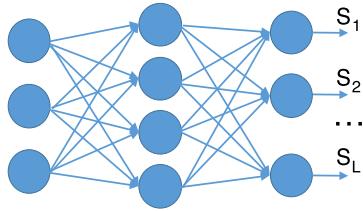


Fig. 2. A single jointly trained neural network (which could have multiple hidden layers) whose features are shared for inferring all L line statuses.

the neural network for an updated measurement set. The additional training time needed would be much shorter than if we train from scratch. These extensions are however beyond the scope of this paper, and are left for future investigations.

IV. NEURAL NETWORK ARCHITECTURES FOR LEARNING CLASSIFIERS

To perform binary MAP inference of each line status, the decision boundary of the MAP detector is highly nonlinear (see Remark 3). We investigate classifiers based on neural networks to capture such complex nonlinear decision boundaries. In other words, we employ neural networks as the parametric models $q_{\beta}(s|y)$: given the input data y , the output layer of the neural network will produce the probabilities $q_{\beta}(s_l|y)$, $l = 1, \dots, L$, (based on which identification decisions are then made).

In particular, we employ a neural network architecture that allows classifiers for different lines to *share features*. Specifically, instead of training L separate neural networks each with one node in its output layer, we train *one* neural network whose output layer consists of L nodes each predicting a different line's status. An illustration of this architecture is depicted in Figure 2: a) the input layer of the neural network consists of y , b) the hidden layers of neurons compute a number of *nonlinear features* of the input y , and c) the output layer applies binary classifiers to these features to predict $s_l \in \{0, 1\}$, $l = 1, \dots, L$. Specifically, logistic functions are employed in the output layer whose outputs correspond to $q_{\beta}(s_l = 1|y)$, $l = 1, \dots, L$. As a result, the features computed by the hidden layers can all be used in classifying any line's status. The intuition of using shared features is that certain *common* features may provide good predictive power in inferring *many different* lines' statuses in a power network. For training and testing, we generate labeled data $\{s^i, y^i\}$ randomly that satisfy the power flow equations and the observation models. Each $s^i = [s_1^i, \dots, s_L^i]^T$ then consists of L labels used by the L output classifiers respectively.

With the proposed Learning-to-Infer method, since labeled data can be generated in an arbitrarily large amount using Monte Carlo simulations, whenever overfitting is observed, it can in principle always be overcome by generating more labeled data for training. Thus, as long as the computation time allows, we can use neural network models of very high complexity for approximating the binary MAP detectors, without worrying about overfitting.

V. NUMERICAL EXPERIMENTS

We evaluate the proposed Learning-to-Infer method for multi-line outage identification with three benchmark systems of increasing sizes, the IEEE 30, 118, and 300 bus systems, as the baseline topologies. As opposed to considering only a small number of simultaneous line outages as in existing works, we allow *any number* of line outages, and investigate whether the learned discriminative classifiers can successfully recover the post-outage topologies in real time.

A. Data Set Generation

In our experiments, the data sets are primarily generated with the DC power flow model (2). Here, our focus is to examine whether the proposed Learning-to-Infer method can effectively overcome the fundamental challenge of exponential computational complexity due to the potentially large number of simultaneous line outages. For this, the DC power flow model offers sufficient modeling details. We will then at the end of the section run experiments with data sets generated with the AC power flow model (1), and verify that the lessons learned from the DC power flow experiments continue to hold.

With the DC power model, the set of controlled inputs $\{P, Q^{in}, V^{in}\}$ reduce to $\{P\}$, and the generative model (6) reduces to $p(s, P, y) = p(s, P)p(y|s, P)$. To generate a data set $\{s^i, P^i, y^i, i = 1, \dots, I\}$, we assume the prior distribution $p(s, P)$ factors as $p(s)p(P)$. As such, we generate the post-outage network topologies s and the power injections P independently:

- We generate the line statuses $\{s_l\}$ using independent and identically distributed (IID) Bernoulli random variables, so that the average numbers of line outages are 7.8, 13.4 and 11.6 for the IEEE 30, 118 and 300 bus systems, respectively. These numbers of simultaneous line outages are significantly higher than those typically assumed in sparse line outage studies. We do not consider disconnected networks in this study, and exclude the line status samples if they lead to disconnected networks. As such, considering that some lines must always be connected to ensure network connectivity, after some network reduction, the equivalent networks for the IEEE 30, 118, and 300 bus systems have 38, 170, and 322 lines that can possibly be in outage, respectively.
- We would like our predictor to be able to identify multi-line outages for *arbitrary values of power injections* as opposed to fixed ones. Accordingly, we generate P using the following procedure: For each data sample, we first generate bus voltage phase angles θ as IID uniformly distributed random variables in $[0, 0.2\pi]$, and then compute P according to (2) under the baseline topologies. We note that, the spread of the phase angles in the generated data sets can cover nearly all possible power injection cases in real power transmission networks.

With each pair of generated s^i and P^i , we consider two types of measurements that constitute y : nodal voltage phase angle measurements and nodal power injection measurements. For these, a) we generate IID Gaussian voltage phase angle measurement noises with a standard deviation of 0.01 degree, the

TABLE II
DATA SET SIZE VS. THE ENTIRE SEARCH SPACE

<i>The (reduced) IEEE 30 bus system with 38 lines</i>	
Number of all possible post-outage topologies	$2^{38} = 2.75 \times 10^{11}$
Number of topologies with 8 line outages	$\binom{38}{8} = 4.89 \times 10^7$
The generated data set	3×10^5
<i>The (reduced) IEEE 118 bus system with 170 lines</i>	
Number of all possible post-outage topologies	$2^{170} = 1.50 \times 10^{51}$
Number of topologies with 13 disconnected lines	$\binom{170}{13} = 9.94 \times 10^{18}$
The generated data set	8×10^5
<i>The (reduced) IEEE 300 bus system with 322 lines</i>	
Number of all possible post-outage topologies	$2^{322} = 8.54 \times 10^{96}$
Number of topologies with 12 disconnected lines	$\binom{322}{12} = 2.11 \times 10^{21}$
The generated data set	2.2×10^6

state-of-the-art PMU accuracy [30], and b) we assume power injections are measured accurately. In the following experiments, we consider that measurements of voltage phase angles and power injections are collected at all the buses. The effect of number and locations of sensors will be discussed later in this section.

In this study, we generate 300K, 800K, and 2.2M data samples for the IEEE 30, 118, and 300 bus systems, respectively. These 300K/800K/2.2M data are further divided into 200K/600K/1.8M, 50K/100K/200K, and 50K/100K/200K samples for training, validation, and testing, respectively. We note that *over 99%* of the generated 300K 30-bus multi-line outages are distinct from each other, so are those of the generated 800K 118-bus multi-line outages and those of the 2.2M 300-bus multi-line outages. As a result, these generated data sets can very well evaluate the *generalizability* of the trained classifiers, as (almost) all data samples in the test set have post-outage topologies *unseen* in the training set.

Furthermore, we would like to compare the *size of the generated data set to the total number of possible outage hypotheses*, as highlighted in Table II. Clearly, a) it is computationally prohibitive to perform line outage inference based on exhaustive search, and b) *the generated 300K, 800K and 2.2M data sets are only a tiny fraction of the entire space of all multi-line outages*. Yet, we will show that the classifiers trained with the generated data sets exhibit excellent inference performance and generalizability.

B. Neural Network Structure and Training

We employ three-layer (i.e., one hidden layer) fully connected neural networks with the feature sharing architecture (see Figure 2). Rectified Linear Units (ReLUs) are employed as the activation functions in the hidden layer. In training the classifiers, we use stochastic gradient descent (SGD) with momentum update and Nesterov's acceleration [31]. While this

optimization algorithm works sufficiently well for our experiments, we note that other algorithms may further accelerate the training procedure [32].

C. Evaluation Results

1) *Performance of the Learning-to-Infer Method*: We employ 300, 1000 and 3000 neurons in the hidden layer for the IEEE 30, 118 and 300 bus systems, respectively. For all the three systems, we plot in Figure 3(a) the achieved training and validation losses for every epoch, and in Figure 3(b) the achieved testing accuracies for every epoch. It is clear that the training and validation losses stay very close to each other for all the three systems, and thus no overfitting is observed. Moreover, very high testing accuracies, 0.989, 0.990 and 0.997 are achieved for the IEEE 30, 118 and 300 bus systems, respectively.

The testing accuracies can be equivalently understood by the *average numbers of misidentified line statuses*, plotted in Figure 3(c). We observe that, *at the beginning* of the training procedures, the average numbers of misidentified line statuses are 7.8, 13.4 and 11.6 for the IEEE 30, 118 and 300 bus systems, which are exactly the *average numbers of disconnected lines* in the respective generated data sets (see Section V-A). Indeed, this coincides with the result from a naive identification decision rule of always claiming all the lines as connected (i.e., a trivial majority guess). As the training procedures progress, the average numbers of misidentified line statuses are drastically reduced to eventually 0.4, 1.7 and 1.0. In other words, for the IEEE 300 bus system for example, *facing on average 11.6 simultaneous line outages, only 1 line status would be misidentified* on average by the learned classifier. We note that such a performance is achieved with outage identification decisions made *in real time, under a millisecond*. While the training process can potentially be time consuming, it is however done completely offline.

It is worth noting that we have generated the training, validation and testing data sets with uniformly random voltage phase angles, and hence considerably variable power injections. In practice, there is often more informative prior knowledge about the power injections based on historical data and load forecasts. With such information, the model can be trained with much less variable samples of power injections, and the outage identification performance can be further improved.

2) *Model Size, Sample Complexity, and Scalability*: In the proposed Learning-to-Infer method, obtaining labeled data is not an issue since data can be generated in an arbitrarily large amount using Monte Carlo simulations. This leads to two questions that are of particular interest: to learn a good classifier, a) what size of a neural network is needed? and b) how much data needs to be generated? To answer these questions, we vary the sizes of the hidden layer of the neural networks as well as the training data size, and evaluate the learned classifiers for the three benchmark systems. We plot the testing results for the IEEE 30, 118 and 300 bus systems in Figure 4(a), 4(b) and 4(c), respectively. It is observed that the best performance is achieved with 200K/600K/1.8M data and with 300/1000/3000 neurons for the 30/118/300 bus systems, respectively. Further increasing the data size or the neural network size would see much diminished returns.

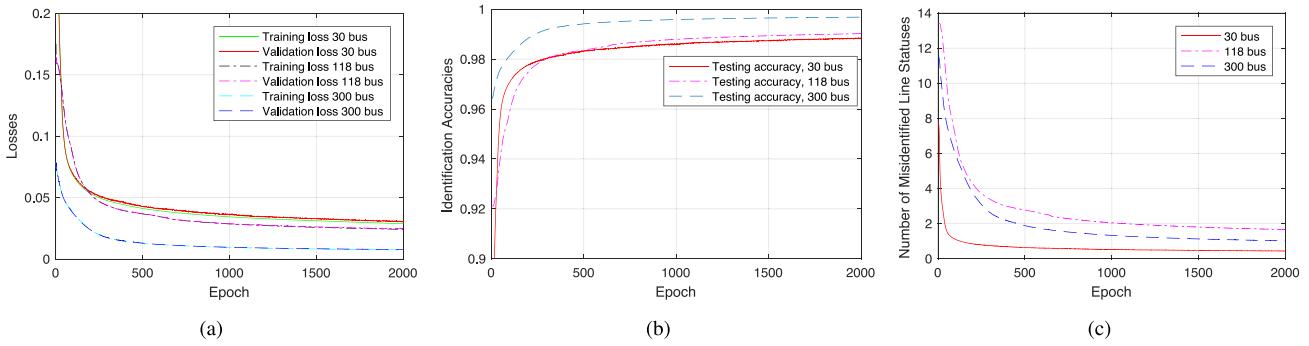


Fig. 3. Progressions of a) training and validation losses, b) testing accuracies, and c) average numbers of misidentified line statuses in IEEE 30, 118 and 300 bus systems.

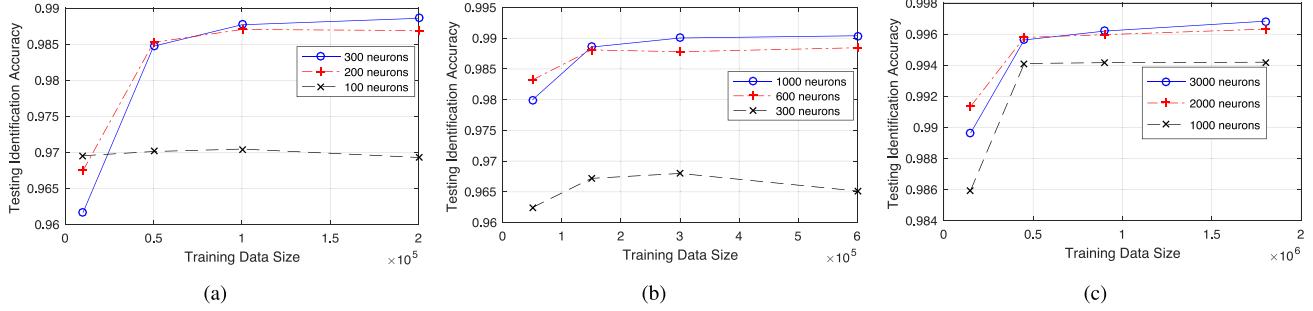


Fig. 4. Effect of model size and sample size, (a) IEEE 30 bus system, (b) IEEE 118 bus system, (c) IEEE 300 bus system.

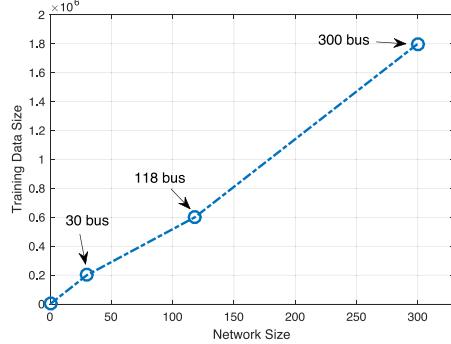


Fig. 5. Scalability of the Learning-to-Infer method, from the IEEE 30 bus system to the IEEE 300 bus system.

Based on all these experiments, we now examine the *scalability* of the proposed Learning-to-Infer method as the problem size increases. We observe that training data sizes of 200K, 600K and 1.8M and neural network models of sizes 300, 1000 and 3000 ensure very high and comparable performance with no overfitting for the IEEE 30, 118 and 300 bus systems, respectively. When these data sizes are reduced by a half, some levels of overfitting then appeared for these models in all the three systems. We plot the training data sizes compared to the problem sizes for the three systems in Figure 5. We observe that the required training data size increases approximately *linearly* with the problem size. This linear scaling behavior implies that the proposed Learning-to-Infer method can be effectively implemented for large-scale systems with reasonable computation resources.

3) *Effect of Number and Locations of Sensors:* We now discuss the effect of sensor placement in real-time multi-line outage identification. It is clear that the performance of line

outage identification would closely depend on where and what types of sensor measurements are collected. Given limited sensing resources, optimizing the sensor placement is a hard problem which many studies have addressed (see, e.g., [7] among others). Here, we present the results from a case study on the IEEE 30 bus system, for which voltage phase angles are collected only at 19 buses (as opposed to all the buses as in the previous experiments), as depicted in Figure 6. Interestingly, the achieved average identification accuracy only drops to 0.978 (from 0.989 when all the buses are monitored). This translates to on average only 0.83 misidentified line statuses among a total of 38 lines. A more comprehensive study of sensor placement for real-time multi-line outage identification is left for future work.

D. Experiments with the AC Power Flow Model

We close this section by verifying the performance of the proposed Learning-to-Infer method with data generated from the AC power flow model. Specifically, we consider the IEEE 118-bus system with 18 generators and 99 loads. Similar to the earlier data set generation process with the DC power flow model, we randomly generate 1M distinct connected post-outage topologies with an average number of 16.2 line outages. We then significantly and randomly vary the power generation and loads in the system with standard deviations equal to 50% of the means, and generate 1M distinct generation and load profiles.

For each data point, which includes a post-outage topology and a generation and load profile, we solve the AC power flow equations (1). To have a consistent comparison with the earlier experiments with the DC power flow model, we continue to rely on measurements of nodal voltage phase angles,

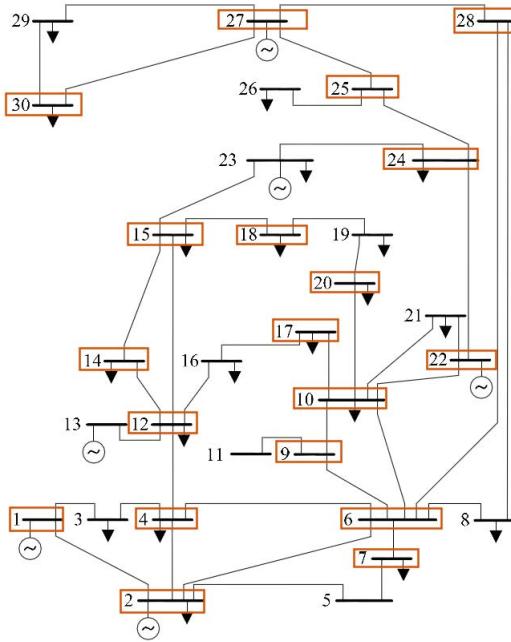


Fig. 6. The IEEE 30 bus system, and a set of locations of PMUs.

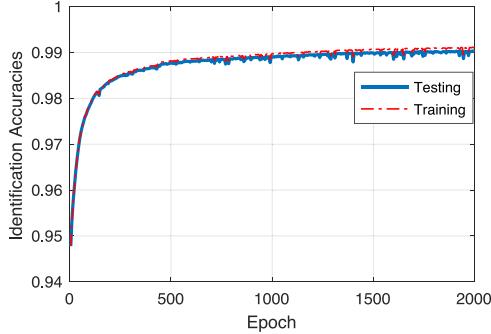


Fig. 7. Progressions of the training and testing accuracies, the IEEE 118-bus system, with the AC power flow model employed.

real power generation, and real power loads to infer the multi-line outages in real time. We will demonstrate that, with the AC power flow model, very high performance similar to that with the DC power flow model can be achieved. Undoubtedly, other types of measurements (e.g., voltage magnitudes, reactive power) may be used to further improve the performance, which is left for future investigation.

The 1M data are divided into 800K, 100K, and 100K for training, validation, and testing, respectively. Similarly to the DC power flow experiments, we employ a three-layer fully connected neural network with 1000 neurons in the hidden layer for learning to infer multi-line outages. The same training algorithm is applied. We plot the training and testing accuracies for every epoch in Figure 7. We observe that a 0.990 testing accuracy is achieved, (recall that the same accuracy, 0.990, is achieved in the earlier experiments on the 118-bus system with the DC power flow model). This translates to on average 1.74 mis-identified line statuses.

Furthermore, we looked into the types of mis-identification errors, and observed that a) the rate of *missed detection* (i.e., missing a line outage when it actually occurred among other

simultaneous line outages) is 8.4%, and b) the rate of *false alarm* (i.e., identifying a line as in outage when it is in fact connected) is a much lower 0.24%. As a result, we observe that nearly 80% of the on average 1.74 mis-identified line statuses are from *missing to detect* 8.4% of the on average 16.2 simultaneous line outages, resulting in 1.36 ($= 16.2 \times 8.4\%$) missed line outages.

E. On Computation Times for Data Generation and Training

As discussed above, a major advantage of the Learning-to-Infer method is that offline computation is exploited for achieving fast and accurate online inference. Specifically, the offline computation consists of two components: a) data generation based on the physical model, and b) predictor training based on the generated data. We discuss in the following several aspects of the offline computation times for data generation and predictor training.

The time consumed for generating the 1M data with the AC power flow on the IEEE 118 bus system (see Section V-D) is a little over an hour using MATPOWER [33]. The training time with 2000 epochs on these data is a little over two hours. Both are run on a laptop with an Intel Core i7 3.1-GHz CPU and 8 GB of RAM. Various approaches can be applied to reduce both times. On the one hand, data generation can be *trivially parallelized* and significantly accelerated as such. It is worth re-emphasizing that data generation via simulations, while still may take a non-trivial amount of time for large systems, is regardless many orders of magnitude faster than collecting and manually labeling historical data from real-world systems. On the other hand, the experiments conducted in this section have achieved very high identification accuracies around or above 99%. In practice, if the performance requirement is not as high (e.g., 97%), then a significantly smaller amount of data (see Figures 4(a) 4(b) and 4(c)) and fewer training epochs (see Figure 3(b)) would be sufficient. The sizes of the neural networks can also be reduced which will lead to faster training. Leveraging the above approaches, much lower computation times can be achieved for offline data generation and training.

VI. CONCLUSION

We have developed a new Learning-to-Infer method for real-time multi-line outage identification in power grids. The computational complexity due to the exponentially large number of outage hypotheses is overcome by efficient marginal inference with optimized predictor models. Optimization of the predictor model is transformed to and solved as a discriminative learning problem, based on Monte Carlo samples efficiently generated with full-blown power flow models. The developed Learning-to-Infer method has the major advantages that a) the training process takes place completely offline, and b) labeled data sets can be generated in an arbitrarily large amount fast and at very little cost. As a result, very complex predictor models can be employed without concern for overfitting, as more labeled training data can always be generated if overfitting is observed. With the classifiers learned offline, their actual use is in real time, and outage identification decisions are made in under a millisecond. We have evaluated the

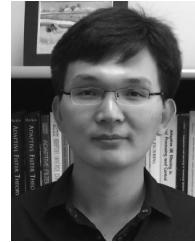
proposed method with the IEEE 30, 118 and 300 bus systems. It has been demonstrated that arbitrary multi-line outages can be identified in real time with excellent performance using classifiers trained with a reasonably small amount of generated data.

REFERENCES

- [1] Y. Zhao, J. Chen, and H. V. Poor, "Learning to infer: A new variational inference approach for power grid topology identification," in *Proc. IEEE Workshop Stat. Signal Process.*, Jun. 2016, pp. 1–5.
- [2] Y. Zhao, J. Chen, and H. V. Poor, "Efficient neural network architecture for topology identification in smart grid," in *Proc. IEEE Glob. Conf. Signal Inf. Process. (GlobalSIP)*, Washington, DC, USA, Dec. 2016, pp. 811–815.
- [3] "Final report on the August 14, 2003 blackout in the United States and Canada," U.S. Canada Power Syst. Outage Task Force, Rep., 2004.
- [4] *Arizona-Southern California Outages on September 8, 2011: Causes and Recommendations*. Washington, DC, USA: FERC and NERC, 2012.
- [5] J. E. Tate and T. J. Overbye, "Line outage detection using phasor angle measurements," *IEEE Trans. Power Syst.*, vol. 23, no. 4, pp. 1644–1652, Nov. 2008.
- [6] J. E. Tate and T. J. Overbye, "Double line outage detection using phasor angle measurements," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, Jul. 2009, pp. 1–5.
- [7] Y. Zhao, J. Chen, A. Goldsmith, and H. V. Poor, "Identification of outages in power systems with uncertain states and optimal sensor locations," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 6, pp. 1140–1153, Dec. 2014.
- [8] Y. Zhao, A. Goldsmith, and H. V. Poor, "On PMU location selection for line outage detection in wide-area transmission networks," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, Jul. 2012, pp. 1–8.
- [9] M. Garcia, T. Catanach, S. V. Wiel, R. Bent, and E. Lawrence, "Line outage localization using phasor measurement data in transient state," *IEEE Trans. Power Syst.*, vol. 31, no. 4, pp. 3019–3027, Jul. 2016.
- [10] H. Zhu and G. B. Giannakis, "Sparse overcomplete representations for efficient identification of power line outages," *IEEE Trans. Power Syst.*, vol. 27, no. 4, pp. 2215–2224, Nov. 2012.
- [11] J. Chen, Y. Zhao, A. Goldsmith, and H. V. Poor, "Line outage detection in power transmission networks via message passing algorithms," in *Proc. 48th Asilomar Conf. Signals Syst. Comput.*, Pacific Grove, CA, USA, 2014, pp. 350–354.
- [12] J. Heydari and A. Tajer, "Quickest localization of anomalies in power grids: A stochastic graphical framework," *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 4679–4688, Sep. 2018.
- [13] X. Li, H. V. Poor, and A. Scaglione, "Blind topology identification for power systems," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Vancouver, BC, Canada, 2013, pp. 91–96.
- [14] Y. Yuan, O. Ardakianian, S. Low, and C. Tomlin, "On the inverse power flow problem," *arXiv preprint arXiv:1610.06631*, 2016.
- [15] I. Gera, Y. Yakoby, and T. Routhenberg, "Blind estimation of states and topology (BEST) in power systems," in *Proc. IEEE Glob. Conf. Signal Inf. Process. (GlobalSIP)*, Montreal, QC, Canada, 2017, pp. 1080–1084.
- [16] M. He and J. Zhang, "A dependency graph approach for fault detection and localization towards secure smart grid," *IEEE Trans. Smart Grid*, vol. 2, no. 2, pp. 342–351, Jun. 2011.
- [17] S. Bolognani, N. Bof, D. Michelotti, R. Muraro, and L. Schenato, "Identification of power distribution network topology via voltage correlation analysis," in *Proc. IEEE Conf. Decis. Control*, 2013, pp. 1659–1664.
- [18] V. Kekatos, G. B. Giannakis, and R. Baldick, "Online energy price matrix factorization for power grid topology tracking," *IEEE Trans. Smart Grid*, vol. 7, no. 3, pp. 1239–1248, May 2016.
- [19] Y. Weng, Y. Liao, and R. Rajagopal, "Distributed energy resources topology identification via graphical modeling," *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 2682–2694, Jul. 2017.
- [20] D. Deka, S. Backhaus, and M. Chertkov, "Structure learning in power distribution networks," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 3, pp. 1061–1074, Sep. 2017.
- [21] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [22] *Power Systems Test Case Archive*, Univ. Washington Elect. Eng., Seattle, WA, USA. [Online]. Available: <https://www.ee.washington.edu/research/pstca/>
- [23] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Upper Saddle River, NJ, USA: Prentice-Hall, 1993.
- [24] J. D. Glover, M. Sarma, and T. Overbye, *Power System Analysis & Design*. Boston, MA, USA: Cengage Learn., 2011.
- [25] R. Baldick, *Applied Optimization: Formulation and Algorithms for Engineering Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2006.
- [26] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [27] H. V. Poor, *An Introduction to Signal Detection and Estimation*. New York, NY, USA: Springer-Verlag, 1994.
- [28] M. Mezard and A. Montanari, *Information, Physics, and Computation*. Oxford, U.K.: Oxford Univ. Press, 2009.
- [29] V. Vapnik, *Statistical Learning Theory*. New York, NY, USA: Wiley, 1998.
- [30] A. von Meier, D. Culler, A. McEachern, and R. Arghandeh, "Micro-synchrophasors for distribution systems," in *Proc. IEEE Innovative Smart Grid Technol. (ISGT)*, Washington, DC, USA, 2013, pp. 1–5.
- [31] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*, vol. 87. New York, NY, USA: Springer, 2013.
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [33] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas, "MATPOWER: Steady-state operations, planning and analysis tools for power systems research and education," *IEEE Trans. Power Syst.*, vol. 26, no. 1, pp. 12–19, Feb. 2011.



Yue Zhao (S'06–M'11) received the B.E. degree in electronic engineering from Tsinghua University, Beijing, China, in 2006, and the M.S. and Ph.D. degrees in electrical engineering from the University of California at Los Angeles, Los Angeles, in 2007 and 2011, respectively. He is an Assistant Professor of electrical and computer engineering, Stony Brook University. His current research interests are in the areas of statistical signal processing, machine learning, mechanism design, optimization, renewable energy integration, and smart grid.



with Tencent AI Lab, Bellevue, WA, USA. His research interests include machine learning, optimization theory, and signal processing.



H. Vincent Poor (S'72–M'77–SM'82–F'87) received the Ph.D. degree in EECS from Princeton University in 1977.

From 1977 to 1990, he was on the faculty of the University of Illinois at Urbana–Champaign. Since 1990, he has been on the faculty at Princeton, where he is the Michael Henry Strater University Professor of Electrical Engineering. From 2006 to 2016, he served as Dean of Princeton's School of Engineering and Applied Science. He has also held visiting appointments at several other universities, including most recently at Berkeley and Cambridge. His research interests are in the areas of information theory and signal processing, and their applications in wireless networks, energy systems and related fields. Among his publications in these areas is the book *Information Theoretic Security and Privacy of Information Systems* (Cambridge University Press, 2017).

Recent recognition of Dr. Poor's work includes the 2017 IEEE Alexander Graham Bell Medal, the 2019 ASEE Benjamin Garver Lamme Award, a D.Sc. *honoris causa* from Syracuse University in 2017, and a D.Eng. *honoris causa* from the University of Waterloo in 2019. He is a member of the National Academy of Engineering and the National Academy of Sciences, and a Foreign Member of the Chinese Academy of Sciences, the Royal Society, and other national and international academies.