Democratizing Access to Power Grid Data and Models

Vladimir Dvorkin

Department of Electrical Engineering and Computer Science

University of Michigan

UM IES Energy Seminar Series

November 13, 2025

UNIVERSITY OF MICHIGAN

Why democratizing access to power grid data

- **Open-access datasets** have fueled breakthroughs in many fields:

Computer vision
(ImageNet)

Speech recognition
(LibriSpeech)

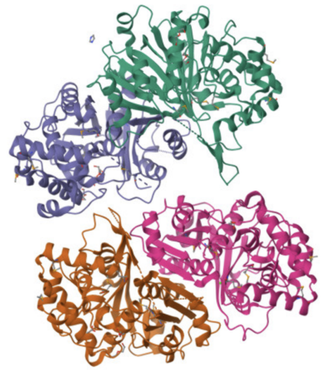

Biology
(UniProt)

- Power systems research **lags behind**:

  Real grid data are **hard to access** — due to security and regulation

- Most available datasets are **synthetic**, limiting realism and impact

- **Goal:** Enable open, realistic grid datasets

  — without compromising privacy or security of source grids

Where energy data should go?

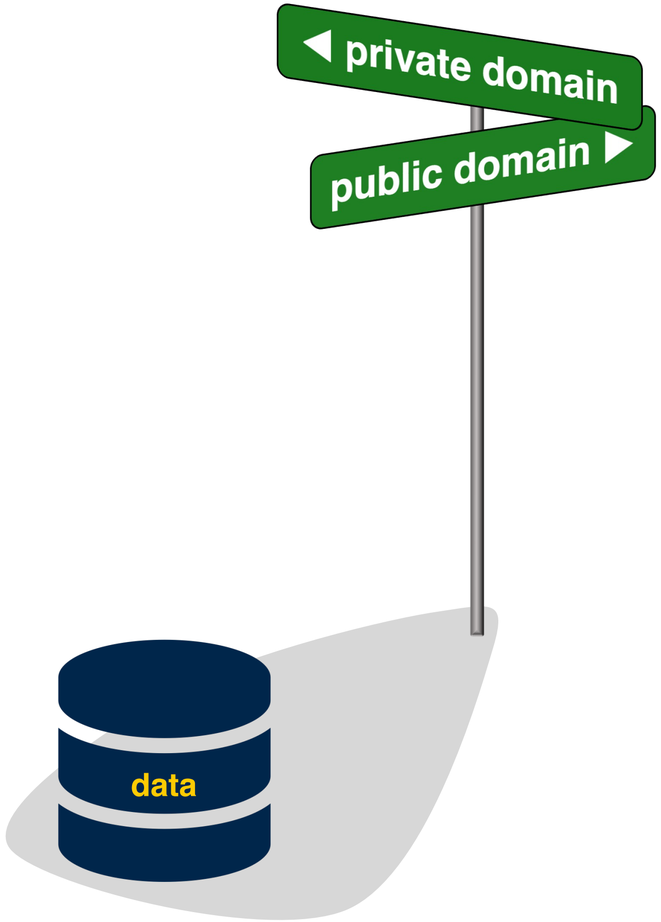

Arguments in favor of **private** data:

- Privacy and security
- Regulatory compliance
- Competitive advantage

Arguments in favor of **public** data:

- Improved decision-making
- Less barriers for entry
- Innovation, research

Is there an **non-discrete** answer to this question?

Industry is very hesitant to share data about critical infrastructure systems

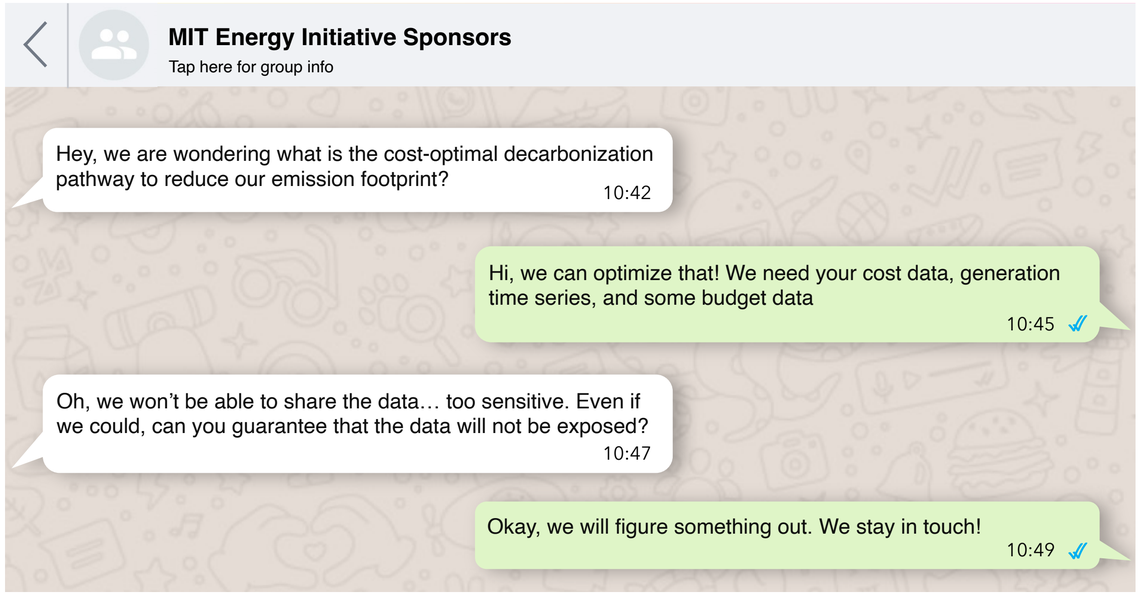

typical conversation with industry partners

- Hesitance to share operational/system data
- Data disclosure treated as binary (share/don't share) rather than algorithmic problem
- Decisions are driven more by regulators' mandates rather than by the strive to innovate

Why companies are so hesitant to share data?

ORGANIZATIONS SPEND MORE THAN EVER DEALING WITH THE COSTS AND CONSEQUENCES OF INCREASINGLY SOPHISTICATED ATTACKS

Cost of cybercrime is rising

$11.7m → **$13.0m** **+12%** **=72%**

Average cost of cybercrime in 2017

Average cost of cybercrime in 2018

Increase in the last year

Increase in the last 5 years

People-based attacks have increased the most

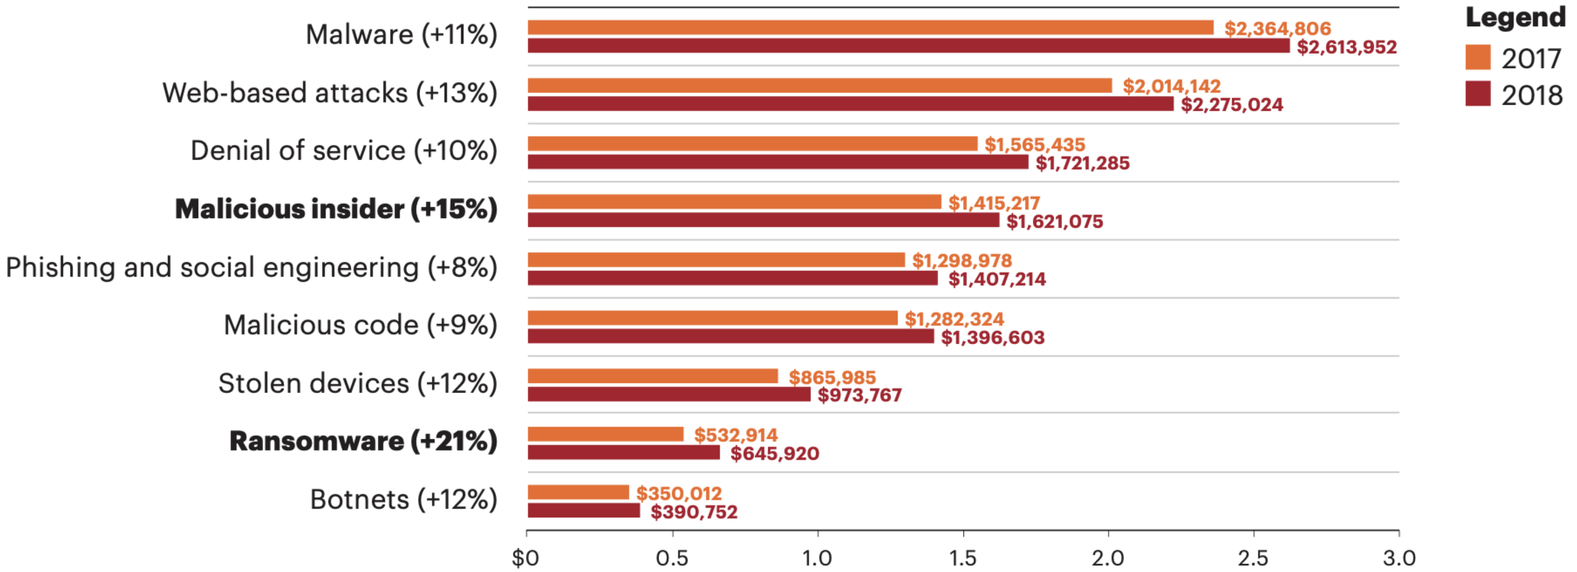

Business consequences are expensive

$4.0m

Cost of business disruption

$5.9m

Cost of information loss

36%

Proportion of spend on discovering attacks in 2018

Accenture: 9th Annual Cost of Cybercrime study

Synthetic data is a viable alternative to real data sharing

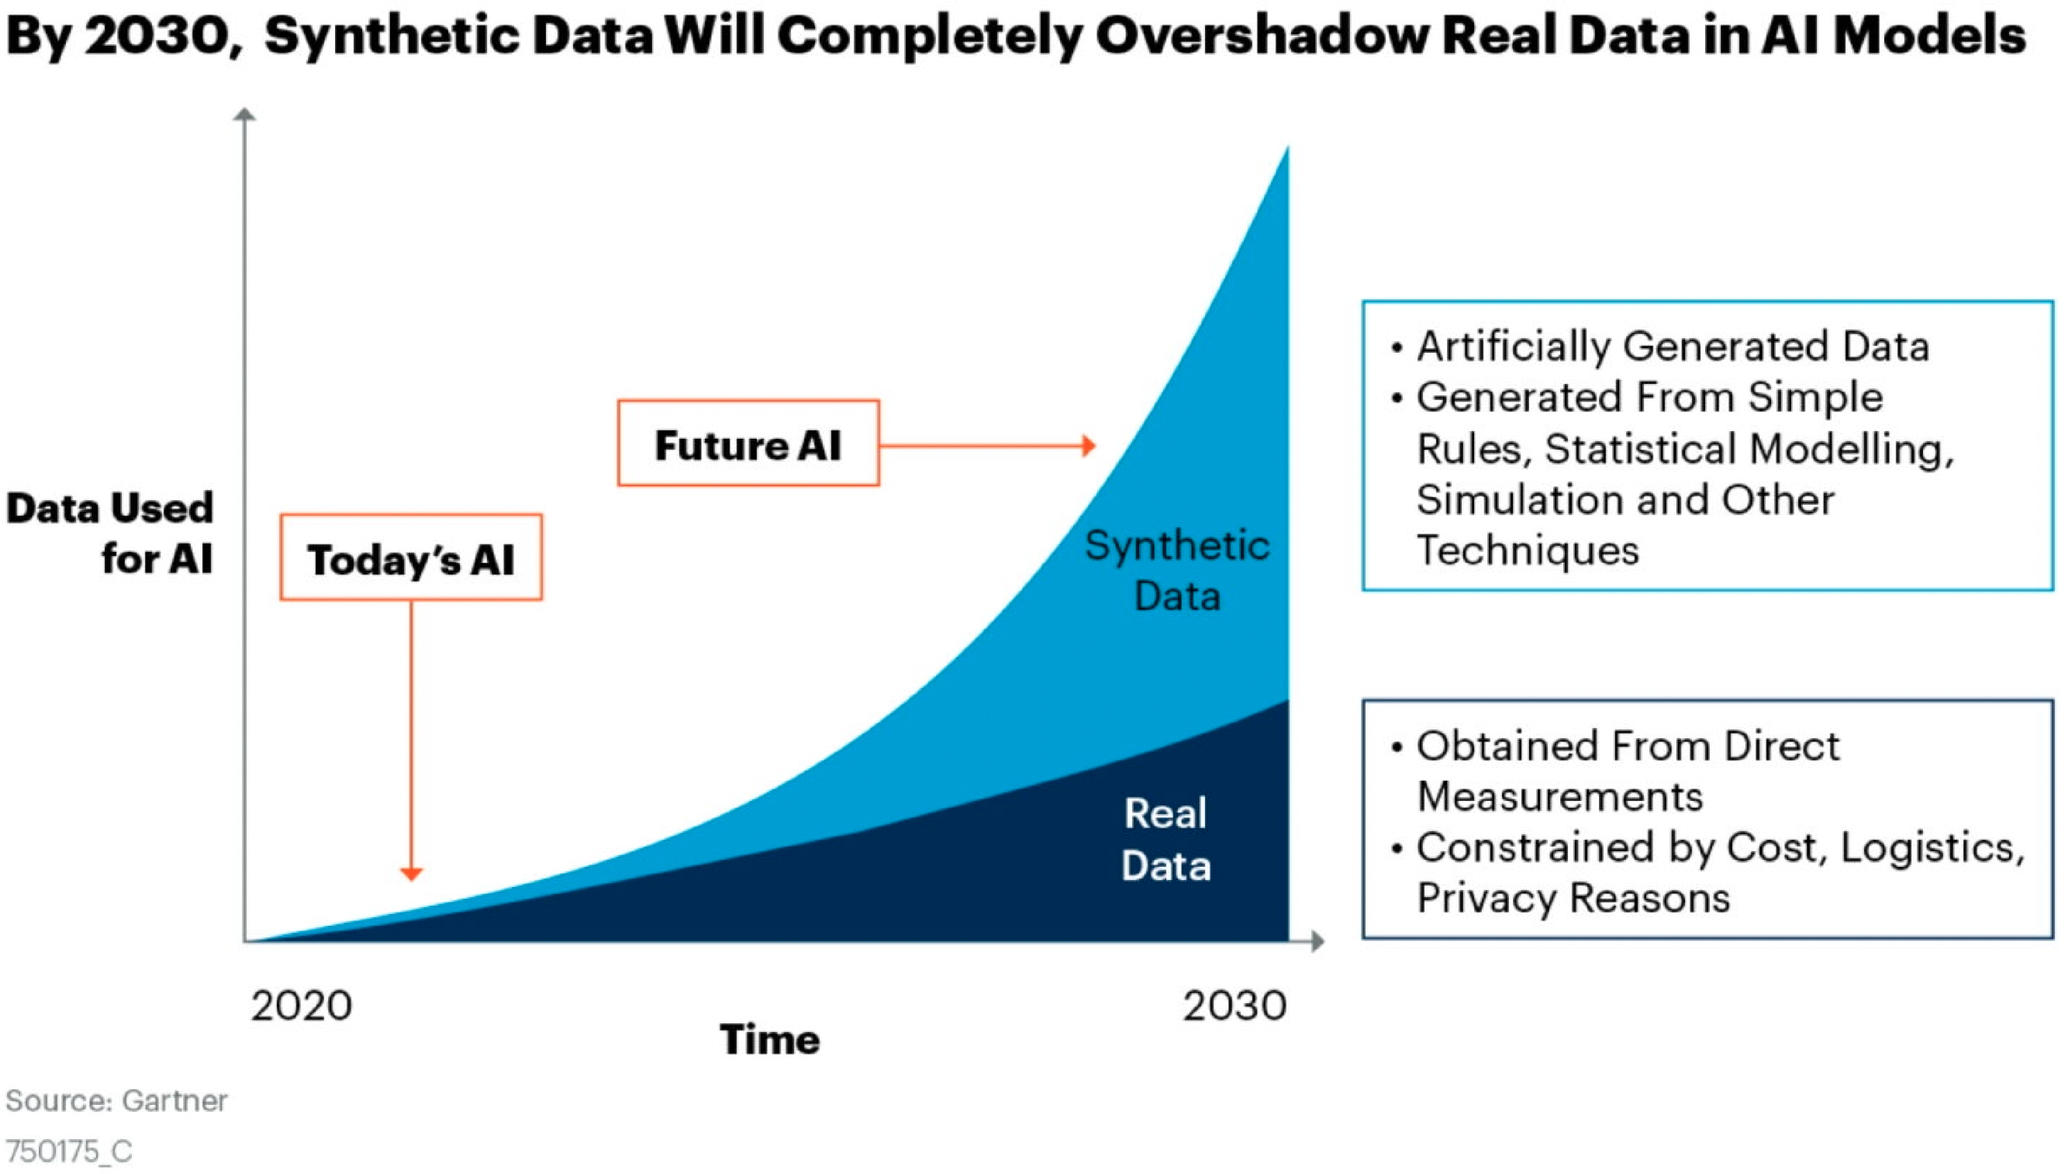

Synthetic power systems datasets

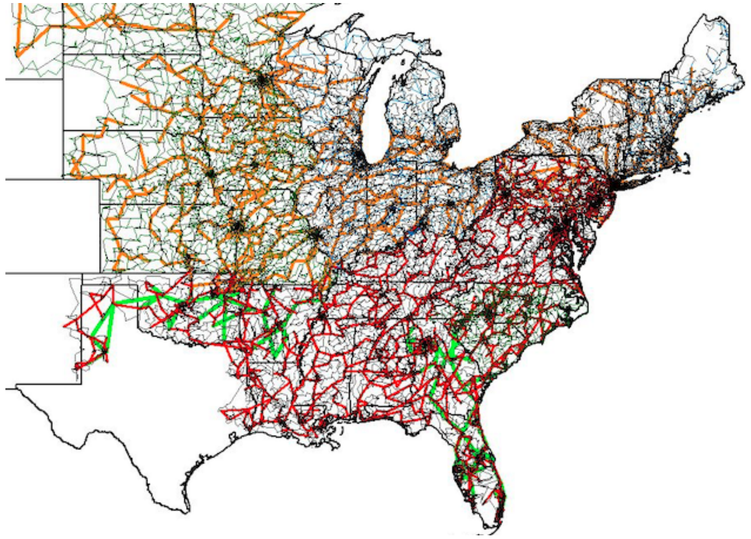

Texas A&M University Grid Datasets (from 37 to 80k+ bus networks)

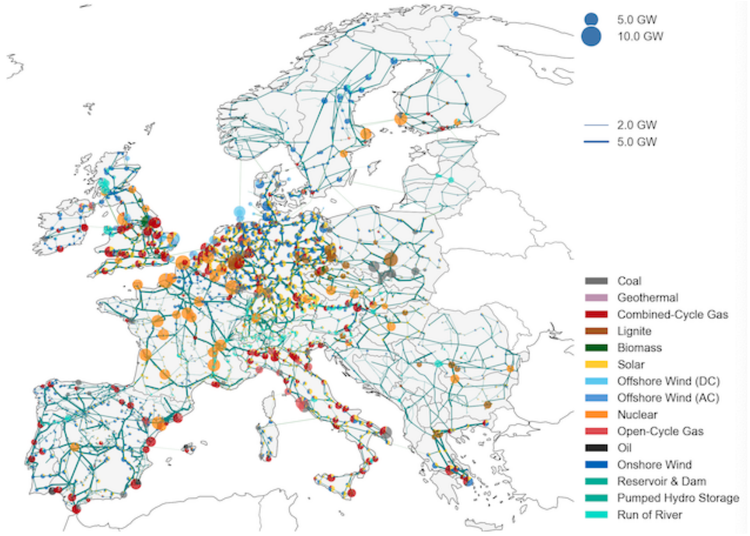

PyPSA-Eur: synthetic dataset of Europe covering the full ENTSO-E area

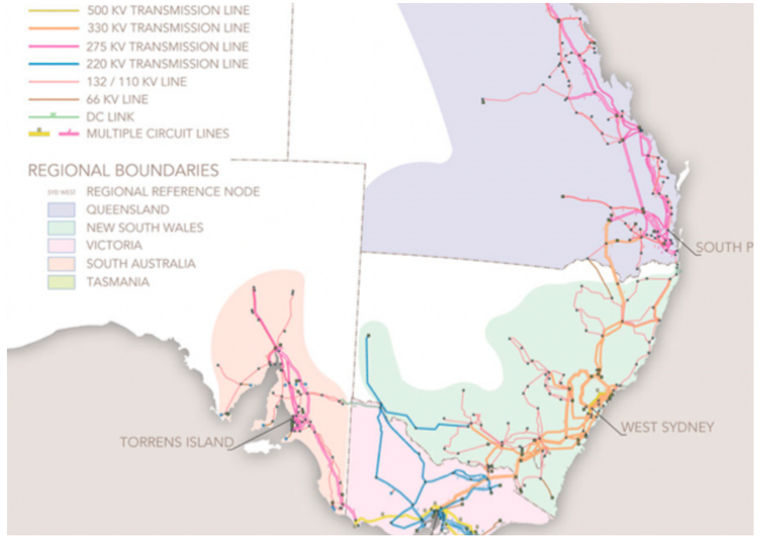

Synthetic Data of the National Electricity Market (Australia)

Why these datasets may not satisfy our needs?

- "[...] data bears **no relation** to the actual grid [...] except that generation and load profiles are similar, based on public data"
- "This test case represents a synthetic (**fictitious**) transmission"
- "This case is synthetic and **does not** model the actual grid"

This talk

We introduce algorithms to synthesize credible synthetic datasets from real power systems, while controlling privacy and security risks

- **Algorithmic formalization of trust** in energy data sharing
  - Moving from subjective decisions to quantitative guarantees
- **Private or public data?** Our algorithms provide a *non-binary* solution to this dilemma
  - Spectrum of privacy-utility tradeoffs, not just share/don't share
- **Comprehensive data scope:**
  - Optimization datasets (OPF, unit commitment, economic dispatch)
  - Dynamical models (e.g., grid frequency dynamics)
  - Operational records (power flows, market-clearing outcomes, etc.)

Core message: Our algorithms enable controlled data disclosures while resolving privacy, security, and logistics concerns

Agenda

1. Intro
2. Formalization of energy data privacy
3. Synthesizing optimization data with privacy and cyber security guarantees
4. Synthesizing power system dynamics models with privacy guarantees
5. Synthesizing power flow datasets with constrained diffusion models
6. Outro

Agenda

1. Intro
2. **Formalization of energy data privacy**
3. Synthesizing optimization data with privacy and cyber security guarantees
4. Synthesizing power system dynamics models with privacy guarantees
5. Synthesizing power flow datasets with constrained diffusion models
6. Outro

Power systems as a stroll in the fog

- Power systems are critical infrastructures with most of data being classified
- We have only a limited observably, e.g., ISO data disclosure portals
- Grid stakeholders hence act on a limited set of system data

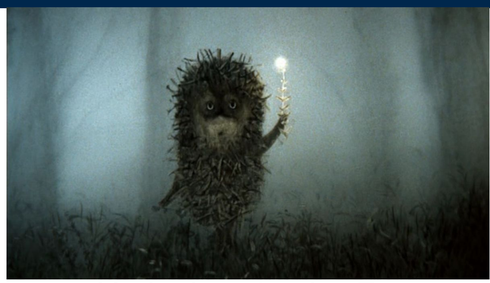

Hedgehog in the Fog
Yuri Norstein (1975)

Example: DC optimal power flow (OPF) in the (small) IEEE 118-bus system

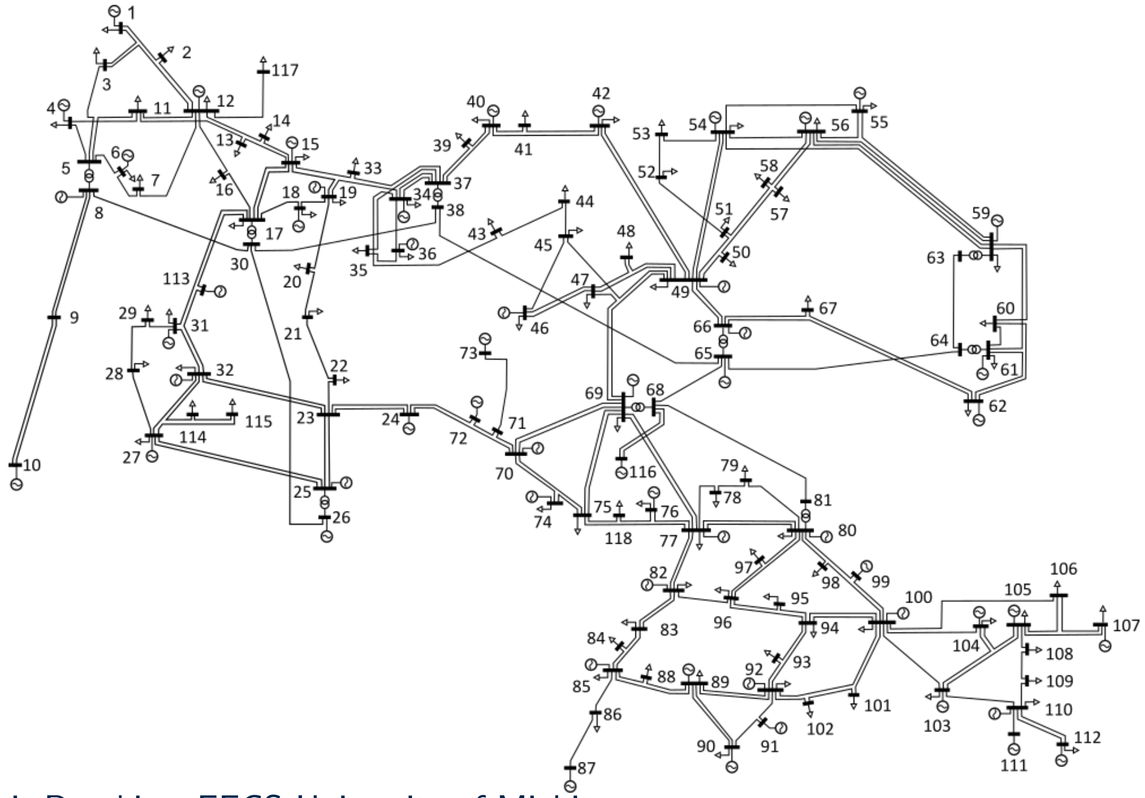

- 1079 rows of element-specific data
- Each generator owns only 2 rows
- The rest of system data is not *explicitly* disclosed to power producers

How to disclose grid data in a controllable manner?

Differential privacy (DP) enables controlled disclosure of grid data

Differential Privacy for Power Grid Obfuscation

Ferdinando Fioretto, Terrence W. K. Mak, *Member, IEEE*, and Pascal Van Hentenryck, *Member, IEEE*

Abstract—The availability of high-fidelity energy networks bring significant value to academic and commercial research. However, such releases also raise fundamental concerns related to privacy and security as they can reveal sensitive commercial information and expose system vulnerabilities. This paper investigates how to release the data for power networks where *the parameters of transmission lines and transformers are obfuscated*. It does so by using the framework of Differential Privacy (DP), that provides strong privacy guarantees and has attracted significant attention in recent years. Unfortunately, simple DP mechanisms often result in AC-infeasible networks. To address these concerns, this paper presents a novel differentially private mechanism that guarantees AC-feasibility and largely preserves the fidelity of the obfuscation significantly reduces the potential damage of an attack carried by exploiting the released dataset. We also show that the obfuscation significantly reduces the potential damage of an attack carried by exploiting the released dataset.

***Index Terms*—Data privacy, power system security, data processing.**

I. Introduction

THE AVAILABILITY of test cases representing high-fidelity power system networks is fundamental to foster research in optimal power flow, unit commitment, and transmission planning, to name only a few challenging problems. This need was recognized by ARPA-E when it initiated the *Grid Data Program* in 2015. However, the release of such rich datasets is challenging due to legal issues related to privacy and national security. For instance, the electrical load of an industrial customer indirectly exposes sensitive information on its production levels and strategic investments, and the value parameters of lines and generators may reveal how transmission operators operate their networks. Furthermore, this data could be exploited by an attacker to inflict targeted damages on the network infrastructure.

This paper explores whether differential privacy (DP) [1] is an algorithmic property that measures and bounds the privacy risks associated with answering sensitive queries or releasing a privacy-preserving dataset. It introduces carefully calibrated noise to the data to prevent the disclosure of sensitive information. An algorithm satisfying DP offers privacy protection regardless of the external knowledge of an attacker. In particular, the definition of DP adopted in this paper ensures that an attacker obtaining access to a differentially private output, cannot detect (with high probability) how close is the output to its original one.

However, DP faces significant challenges when the resulting privacy-preserving value to its original one. privacy-preserving datasets are used as inputs to complex optimization algorithms, e.g., *Optimal Power Flow* (OPF) problems. Indeed, the privacy-preserving dataset may have lost the fidelity and realism of the original data and may even not admit feasible solutions for the optimization problems of interest [2].

This paper studies how to address such a challenge when the goal is to preserve the privacy of line parameters and transformers. It presents a DP mechanism to release data for power networks that is realistic and limits the power of an attacker. More precisely, the contribution of this paper is fourfold. *(1)* It proposes the *Power Line Obfuscation* (PLO) mechanism to obfuscate the line parameters in a power system network. *(2)* It shows that PLO has strong theoretical properties: It achieves ϵ-differential privacy, it ensures that the released data can produce feasible solutions for OPF problems, and its objective value is a constant factor away from optimality. *(3)* It extends the PLO mechanism to handle time-series network data. *(4)* It demonstrates, experimentally, that the PLO mechanism improves the accuracy of existing approaches. When tested on the largest collection of OPF test cases available, it results in solutions with similar costs and optimality gaps to those obtained on the original problems, while also protecting well against an attacker that has access to the released data and uses it to damage the real power network. Interestingly, on the test cases adopted, the damage inflicted on the real power

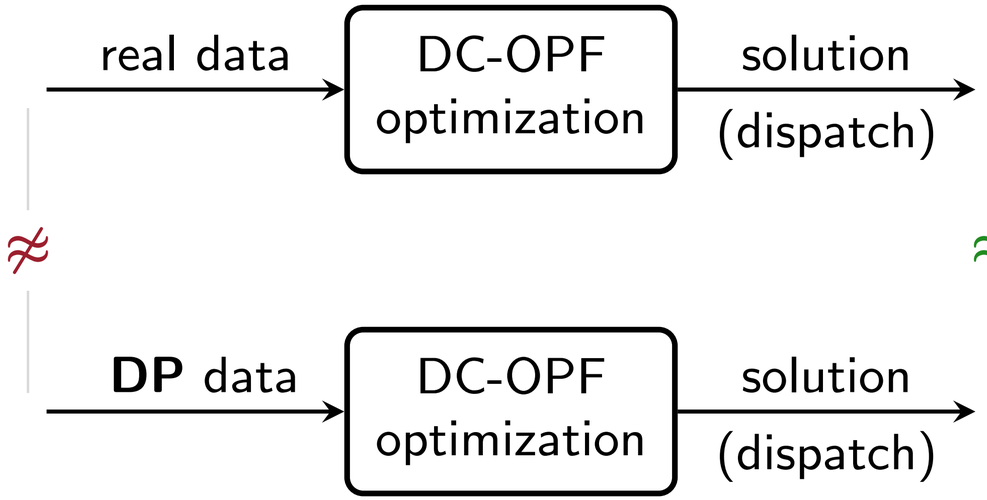

IEEE CSS IEEE CONTROL SYSTEMS LETTERS, VOL. 7, 2023

Differentially Private Algorithms for Synthetic Power System Datasets

Vladimir Dvorkin Jr., *Member, IEEE*, and Audun Botterud, *Member, IEEE*

Abstract—While power systems research relies on the availability of real-world network datasets, data owners (e.g., system operators) are hesitant to share data due to privacy risks. To control these risks, we develop privacy-preserving algorithms for synthetic generation of datasets for optimization and machine learning. Taking a real-world dataset as input, the algorithms output its noisy, synthetic version, which preserves the accuracy of the real data on a specific downstream model or even a large population of those. We control the privacy loss using Laplace and Exponential mechanisms of differential privacy and preserve data accuracy using a post-processing convex (or mixed-integer) optimization. We apply the algorithms to generate synthetic power network parameters and wind power data.

***Index Terms*—Differential privacy, machine learning, power systems optimization, synthetic datasets.**

I. Introduction

POWER system datasets are instrumental for enhancing solutions to many problems, including optimal power flow (OPF) and wind power forecasting. Releasing real data, however, is challenging due to security and privacy concerns. Indeed, detailed network datasets inform cyberattacks on SCADA systems and can be used by strategic market players to maximize profits at the expense of deteriorating social welfare. These concerns motivate producing synthetic datasets – a sanitized version of private datasets that approximately preserve accuracy of data for power system applications.

Differential privacy (DP) is an algorithmic notion of privacy preservation that enables trade-offs between data privacy and accuracy in optimization [1] and machine learning [2]. It also found applications in the context of privacy-preserving OPF computations, e.g., in distributed algorithms [3] and centralized solvers for power grids [4], [5], as well as in machine learning problems specific to power systems [6]. Models in [3], [4], [5], [6], however, only control data leakages in computations and do not provide synthetic datasets per se.

Producing synthetic datasets in a DP way is achieved by corrupting data with privacy-preserving noise [7], [8]. However, applications of the standard noise-additive DP mechanisms in power systems, such as the Laplace mechanism, may no longer admit a meaningful result. Indeed, adding noise to data may fundamentally alter important statistics and trends in machine learning datasets [9]. In the OPF context, [10] and [11] showed that the Laplacian perturbation of network data almost surely violates OPF feasibility. As a remedy, they proposed an optimization-based post-processing which restores the accuracy of synthetic OPF datasets without altering the privacy guarantee. The proposed restoration, however, renders the Laplace mechanism feasible only for a particular OPF model. Repeated applications of the Laplace mechanism instances of variable renewable production) may not be possible, as noise must be scaled drastically as per composition of DP [12].

In this letter, we introduce private synthetic dataset generation algorithms for power systems, which ensure the accuracy of synthetic datasets for downstream models. They enjoy known DP mechanisms and convex (or mixed-integer) post-processing optimization of data. Specifically, we develop:

1) Wind Power Obfuscation (WPO) algorithm, which privately generates wind power measurements, while guaranteeing DP of the real data and ensuring accuracy in terms of the outcomes of a regression analysis.
2) Transmission Capacity Obfuscation (TCO) algorithm, which generates synthetic line parameters, while ensuring their feasibility and cost-consistency on a population of OPF models. Here, we use both Laplace and Exponential mechanisms of DP to substantially reduce the noise compared to using the Laplace mechanism alone.

- DP principle: obfuscate real data (add noise) but preserve its value
- In the DC-OPF setting: obfuscate grid data but preserve the OPF solution
- Formal *privacy guarantee*: released DP data does NOT disclose the real data
- Many applications to synthesizing high-quality transmission, load and generation data

Formalizing differential privacy (DP)

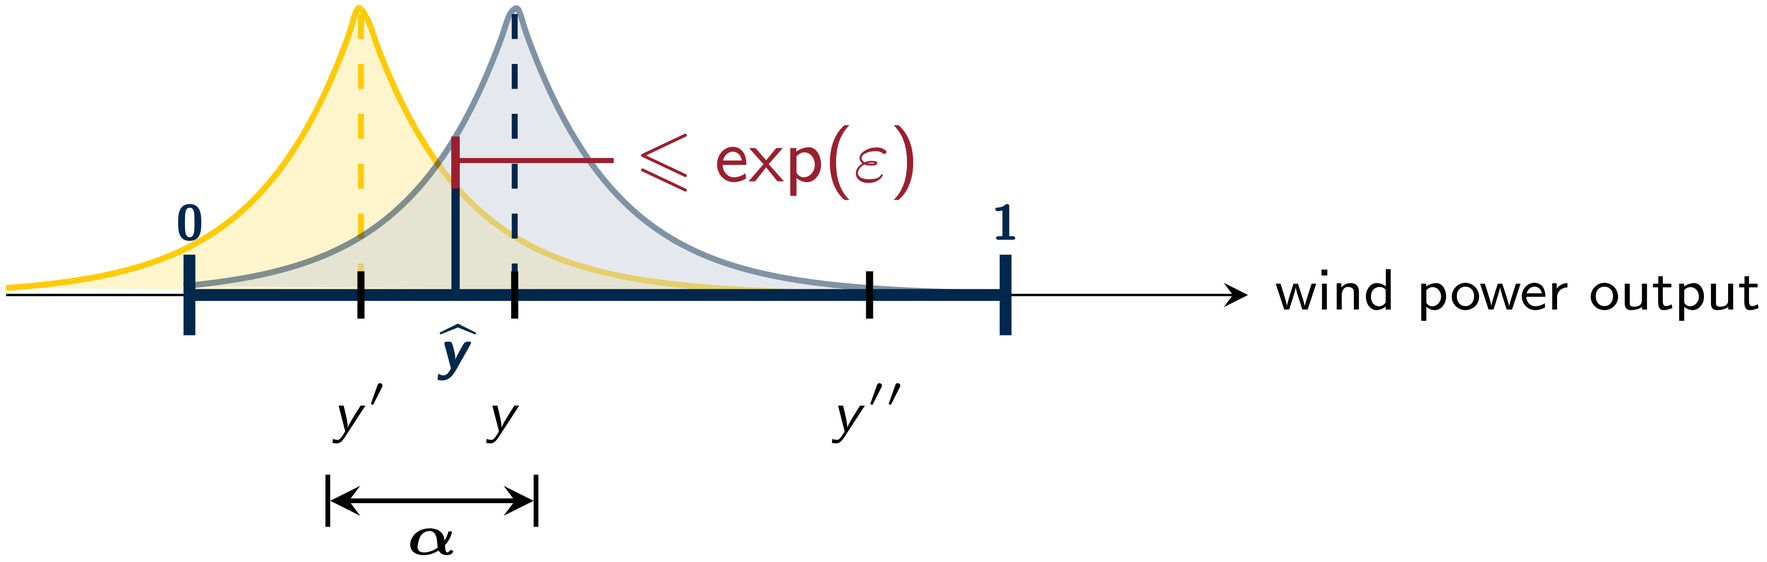

- Wind power records $y, y', y'', \ldots \in [0, 1]$
- For given $\alpha > 0$, records y and y' are $\alpha-$adjacent if $\|y - y'\| \leqslant \alpha$

- Let Lap(α/ε) be a zero-mean random Laplacian noise
- For some parameter $\varepsilon > 0$, the release is $\varepsilon-$DP if

$$\frac{\Pr[\, y + \text{Lap}(\alpha/\varepsilon) \in \widehat{y}\,]}{\Pr[\, y' + \text{Lap}(\alpha/\varepsilon) \in \widehat{y}\,]} \leqslant \exp(\varepsilon)$$

for any $\alpha-$adjacent pair (y, y') and any outcome $\widehat{y}$

Example: Synthesizing transmission line capacities using DP and optimization

- IEEE 73-RTS benchmark
- **Step 1**: add random noise to trans capacity

$$\varphi_1 = \overline{f} + \text{calibrated noise}$$

- **Step 2**: post-process φ_1 to ensure OPF feasibility and cost-consistency w.r.t. real trans capacity

$$\begin{aligned} \varphi_2 \in \underset{\varphi}{\text{argmin}} \quad & \|c(\overline{f}) - c(\varphi)\| + \|\varphi - \varphi_1\| \\ \text{s.t.} \quad & c(\varphi) = \min_{p} \quad c(p) \quad \textit{opf cost} \\ \text{s.t.} \quad & p \in \mathcal{P}(\varphi) \quad \textit{opf feas} \end{aligned}$$

- **Step 3-N**: repeat Step 2 until OPF feasibility and cost-consistency are restored across many scenarios

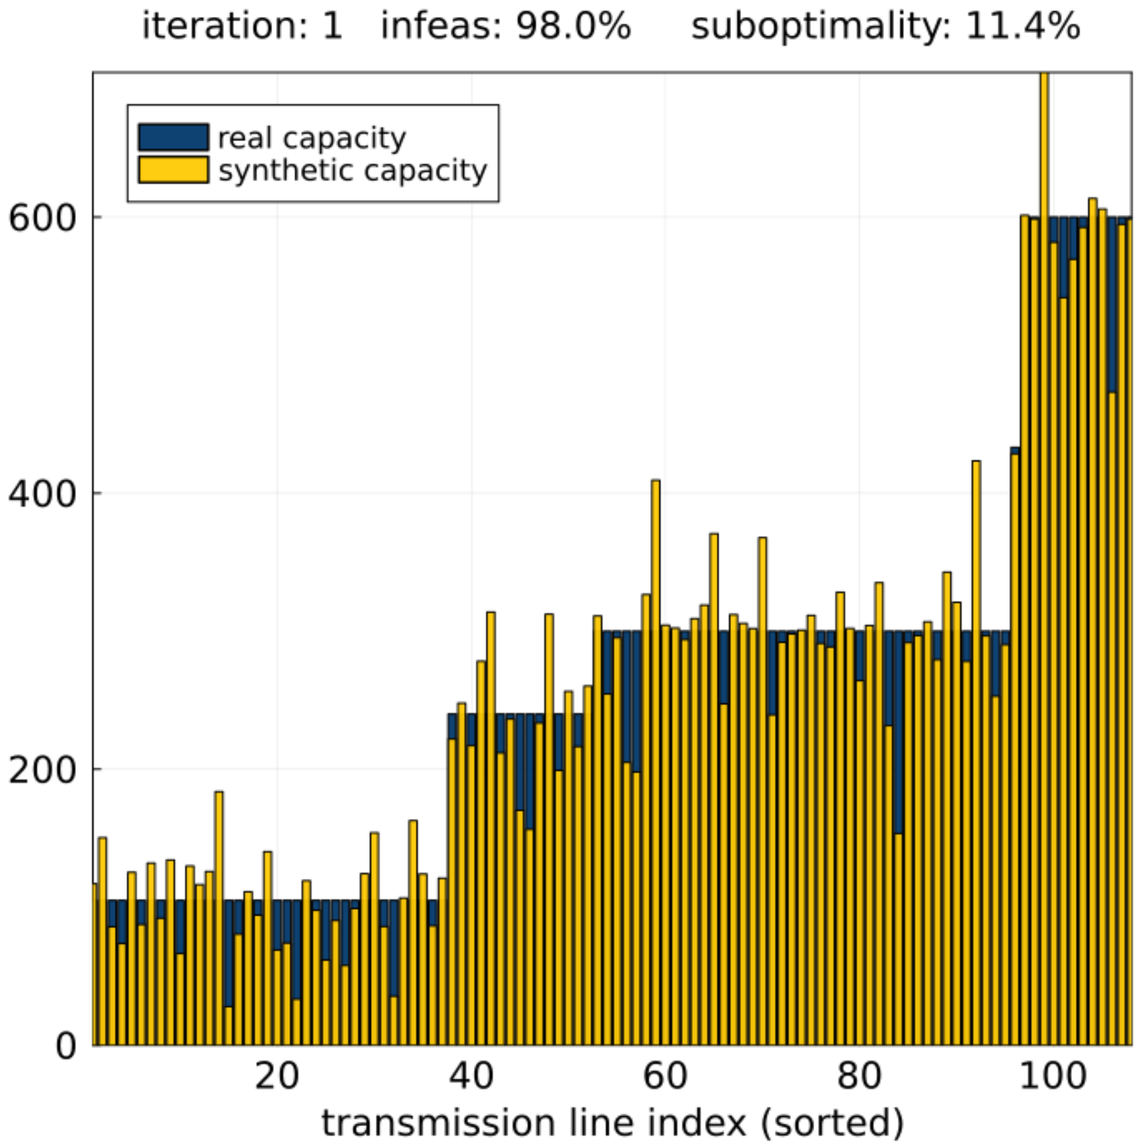

Dvorkin, V., Botterud A. Differentially private algorithms for synthetic power system datasets, IEEE Control Systems Letters, 2023.

Agenda

1. Intro
2. Formalization of energy data privacy
3. **Synthesizing optimization data with privacy and cyber security guarantees**
4. Synthesizing power system dynamics models with privacy guarantees
5. Synthesizing power flow datasets with constrained diffusion models
6. Outro

Challenge: If DP synthetic datasets are so good, do they pose any security risks?

- Grid datasets are used for calibrating cyberattacks on power grids
- Hypothesis: high-quality synthetic data $\rightarrow$ well-calibrated attacks
- Classes of cyberattacks: false data injection, line outage masking, physical attacks, ...

Contribution: We identify cyberattack risks in releasing DP grid data and propose new algorithms to guarantee both privacy and cyber resilience to source grids

Shengyang Wu

Load redistribution attack

- Given load $\mathbf{d}$, the attack corrupts the load $\mathbf{d} + \boldsymbol{\delta}$ with bus injection $\boldsymbol{\delta}$ from admissible set Δ

$$\Delta \triangleq \left\{ \boldsymbol{\delta} \;\middle|\; \begin{array}{ll} \underline{\boldsymbol{\delta}} \leqslant \boldsymbol{\delta} \leqslant \overline{\boldsymbol{\delta}} & \textit{injection limits for each bus} \\ \mathbf{1}^\top \boldsymbol{\delta} = 0 & \textit{total load remains unchanged} \end{array} \right\}$$

- Amounts solving a bi-level optimization problem:

$$\begin{aligned} C_{\text{att}}^{\text{BO}}(\mathbf{d}) = \max_{\boldsymbol{\delta} \in \Delta} \quad & C_{\text{opf}}(\mathbf{d} + \boldsymbol{\delta}) && \textit{maximize the cost} \\ \text{s.t.} \quad & C_{\text{opf}}(\mathbf{d} + \boldsymbol{\delta}) = \min_{\mathbf{x}} \quad \mathbf{c}^\top \mathbf{x} && \textit{feedback from OPF} \\ & \qquad\qquad \text{s.t.} \quad \mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top (\mathbf{d} + \boldsymbol{\delta}) + e_k \leqslant \mathbf{0} \end{aligned}$$

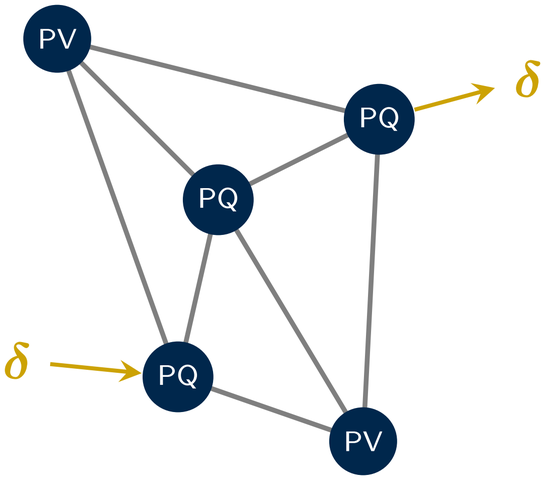

- The problem seeks a *stealthy* attack vector $\boldsymbol{\delta}$ that maximizes the OPF cost

Can DP grid data be used to successfully execute the load redistribution attack?

Numerical evidence

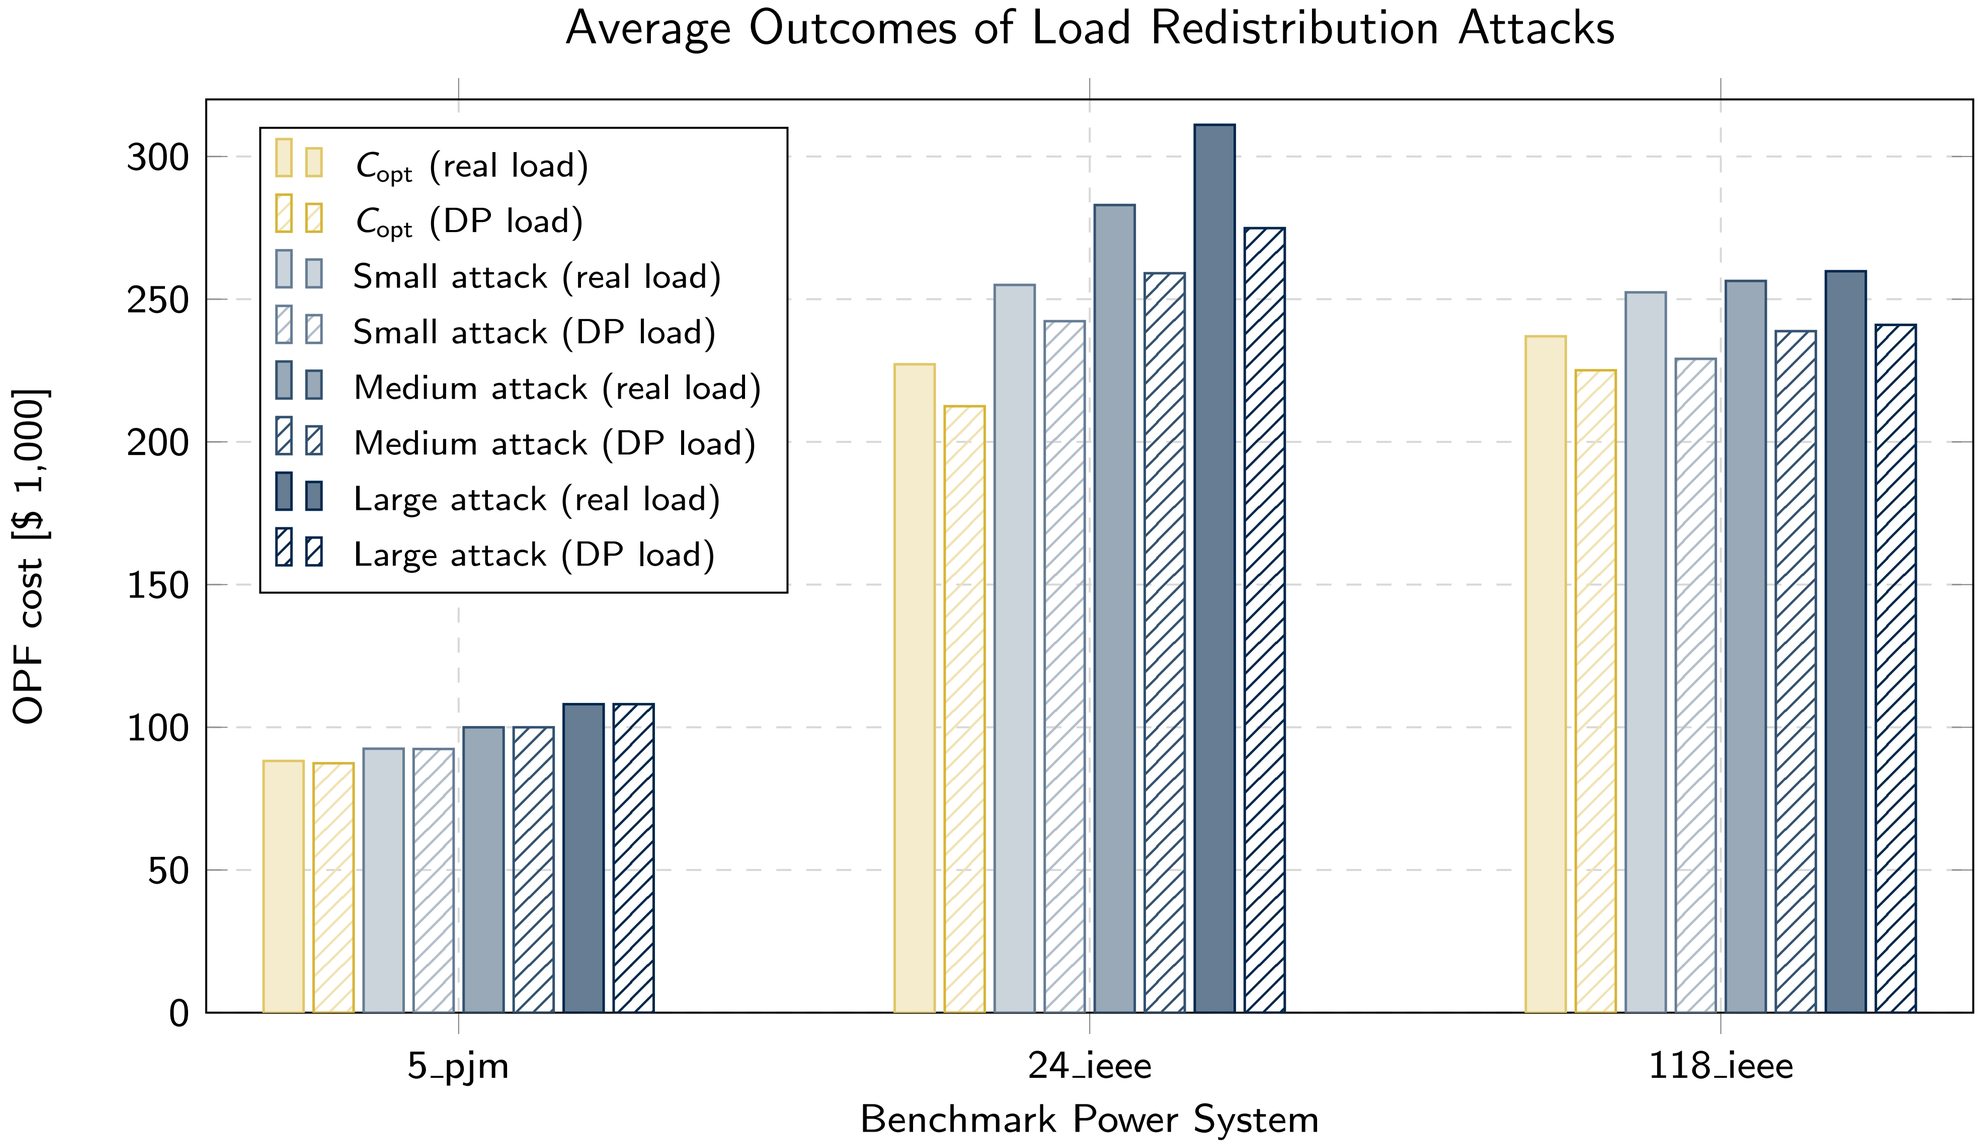

Post-processing optimization to minimize attack damage

- **Step 1**: add random noise to real loads: $\mathbf{d}_1 = \mathbf{d} +$ calibrated noise
- **Step 2**: post-process $\mathbf{d}_1$ by solving a tri-level optimization:
  - Level-1: Optimize synthetic load $\mathbf{d}$ to balance attack damage and cost-consistency
  - Level-2: Feedback from both OPF and attack optimization
  - Level-3: Embedded OPF for attack calibration
- **Result**: DP load vector $\mathbf{d}$ balancing attack damage and cost-consistency

$$
\begin{aligned}
&\underset{\mathbf{d}}{\text{minimize}} && \underbrace{\|C_{\text{att}}^{\text{BO}}(\mathbf{d}) - \tilde{C}_{\text{opf}}\|}_{\text{power of attack}} + \beta \underbrace{\|C_{\text{opf}}(\mathbf{d}) - \tilde{C}_{\text{opf}}\|}_{\text{cost consistency}} + \gamma \underbrace{\|\mathbf{d} - \mathbf{d}_1\|}_{\text{regularization}} && \text{level 1}\\
&\text{subject to} && C_{\text{opf}}(\mathbf{d}) = \\
&&& \quad \underset{x}{\text{minimize}} \quad C_{\text{opf}}(x) && \text{level 2}\\
&&& \quad \text{subject to} \quad x \in \text{opf-eq}(\mathbf{d}) \\
&&& C_{\text{att}}^{\text{BO}}(\mathbf{d}) = \\
&&& \quad \underset{\delta \in \Delta}{\text{maximize}} \quad C_{\text{opf}}(\delta) \\
&&& \quad \text{subject to} \quad C_{\text{opf}}(\delta) = \\
&&&\qquad \underset{x}{\text{minimize}} \quad C_{\text{opf}}(x) && \text{level 3}\\
&&&\qquad \text{subject to} \quad x \in \text{opf-eq}(\mathbf{d} + \delta)
\end{aligned}
$$

Figure: Tri-level structure of the cyber-resilient post-processing of DP load data.

Tractable approximation of the tri-level problem

- Optimizing synthetic loads over bi-level optimization is computationally challenging
- We drawn on the connection between bi-level and robust (single-level) optimization

Bi-level attack optimization

$$C_{\text{att}}^{\text{BO}}(\mathbf{d}) = \max_{\boldsymbol{\delta} \in \Delta} C_{\text{opf}}(\mathbf{d} + \boldsymbol{\delta})$$
$$\mathbf{x} \in \operatorname*{argmin}_{\mathbf{x}} \ \mathbf{c}^\top \mathbf{x}$$
$$\text{s.t. } \mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top (\mathbf{d} + \boldsymbol{\delta}) + e_k \leqslant \mathbf{0} \quad \forall k$$

(uniform attack)

Robust optimization (RO) approximation

$$C_{\text{att}}^{\text{RO}}(\mathbf{d}) = \min_{\mathbf{x}} \ \mathbf{c}^\top \mathbf{x}$$
$$\text{s.t.} \quad \max_{\boldsymbol{\delta}_k \in \Delta} \left[\mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top (\mathbf{d} + \boldsymbol{\delta}_k) + e_k \right] \leqslant \mathbf{0} \quad \forall k$$

(per constraint attack)

Proposition: For any feasible load $\mathbf{d}$, relation $C_{\text{att}}^{\text{RO}}(\mathbf{d}) \geqslant C_{\text{att}}^{\text{BO}}(\mathbf{d})$ holds.

Cyber Resilient Obfuscation (CRO) Algorithm

- **Step 1**: Obfuscate real load data

  DP load $\mathbf{d}_1 = \mathbf{d} +$ calibrated noise

  DP estimation of OPF costs: $\tilde{C}_{\text{opf}} = C_{\text{opf}}(\mathbf{d}) +$ calibrated noise

- **Step 2**: Post-process $\mathbf{d}_1$ to balance cost-consistency and cyber resilience P

$$\underset{\mathbf{d}}{\text{minimize}} \quad \underbrace{\|C_{\text{att}}^{\text{RO}}(\mathbf{d}) - \tilde{C}_{\text{opf}}\|}_{\text{power of attack}} + \beta \underbrace{\|C_{\text{opf}}(\mathbf{d}) - \tilde{C}_{\text{opf}}\|}_{\text{cost consistency}} + \underbrace{\gamma\|\mathbf{d} - \mathbf{d}_1\|}_{\text{regularization}}$$

embedded attack optimization

embedded OPF optimization

- Replacing the two embedded optimization problems with their Karush–Kuhn–Tucker conditions leads to a single-level mixed-integer problem.

CRO-Exp: selecting only important constraints for post-processing

$$C^{\text{RO}}_{\text{att},\tau}(\mathbf{d}) = \min_{\mathbf{x}} \ \mathbf{c}^\top \mathbf{x}$$

$$\text{s.t.} \quad \max_{\delta_k \in \Delta} \left[\mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top (\mathbf{d} + \delta_k) + e_k \right] \leqslant \mathbf{0} \quad \forall k \in \mathcal{K}' \qquad \textit{RO-reformulated cons}$$

$$\left[\mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top \ \mathbf{d} \qquad + e_k \right] \leqslant \mathbf{0} \quad \forall k \in \mathcal{K} \qquad \textit{original problem cons}$$

- We select only most important τ constraints for RO reformulation
- Most important constraints in $\mathcal{K}'$ are function of original loads
- Exponential mechanism of DP to obfuscate loads when selecting $\mathcal{K}'$

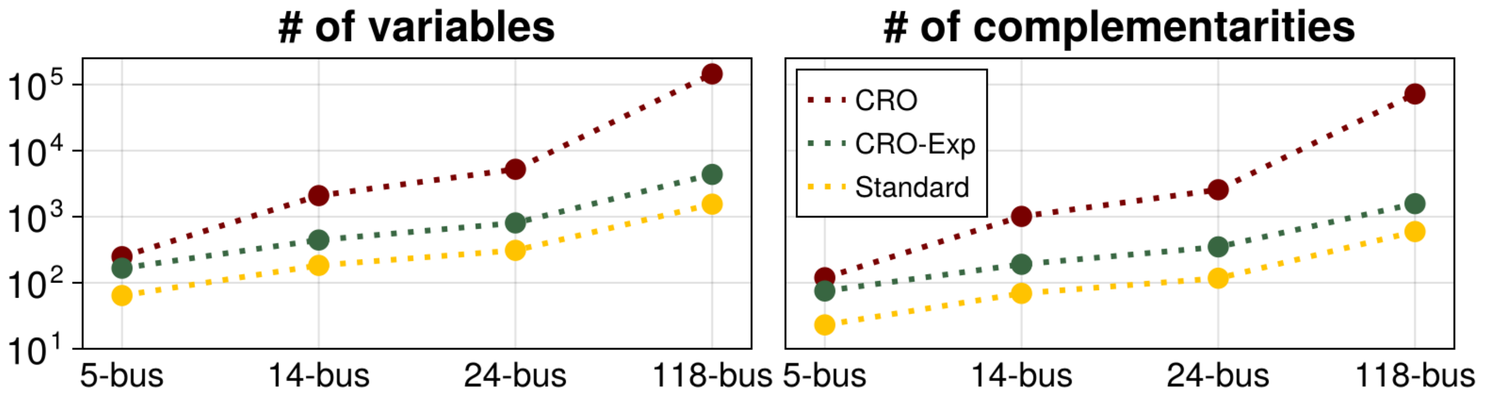

Selecting only important constraints substantially reduces the computational burden

Algorithm 1: Privacy-preserving CRO-Exp

Input: real load $\mathbf{d}$, DP parameters $(\alpha, \varepsilon_1, \varepsilon_2, \varepsilon_3)$, attack data $(\beta, \gamma, \Delta, \tau)$, $\mathcal{K} = \{\emptyset\}$

1 DP obfuscation of load and OPF costs:

$$\tilde{\mathbf{d}}^0 = \mathbf{d} + \text{Lap}\left(\frac{\alpha}{\varepsilon_1}\right)^n$$

$$\tilde{C}_{\text{opf}} = C_{\text{opf}}(\mathbf{d}) + \text{Lap}\left(\frac{\alpha\overline{c}}{\varepsilon_2}\right)$$

2 DP estimation of the set $\mathcal{K}$ of the worst-case constraints

for $t = 1, \ldots, \tau$ **do**
  for $k = 1, \ldots, K$ **do**
    $C_k = C^{\text{RO}}_{\text{att},t}(\mathbf{d}) + \text{Lap}\left(\frac{\alpha\overline{c}}{\varepsilon_3}\right)$
  end
  $k_t \leftarrow \text{argmax}_k \; C_k$
  $\mathcal{K} \leftarrow \mathcal{K} \cup \{k_t\}$
end

3 Post-processing optimization of the synthetic load vector

$$\tilde{\mathbf{d}} \in \underset{\tilde{\mathbf{d}}}{\text{argmin}} \; \|C_{\text{opf}}(\tilde{\mathbf{d}}) - \tilde{C}_{\text{opf}}\|_1 + \beta\|C^{\text{RO}}_{\text{att},\tau}(\tilde{\mathbf{d}}) - \tilde{C}_{\text{opf}}\|_1 + \gamma\|\tilde{\mathbf{d}} - \tilde{\mathbf{d}}^0\|_1$$

Output: Synthetic load vector $\tilde{\mathbf{d}}$

CRO application to PJM 5-Bus system

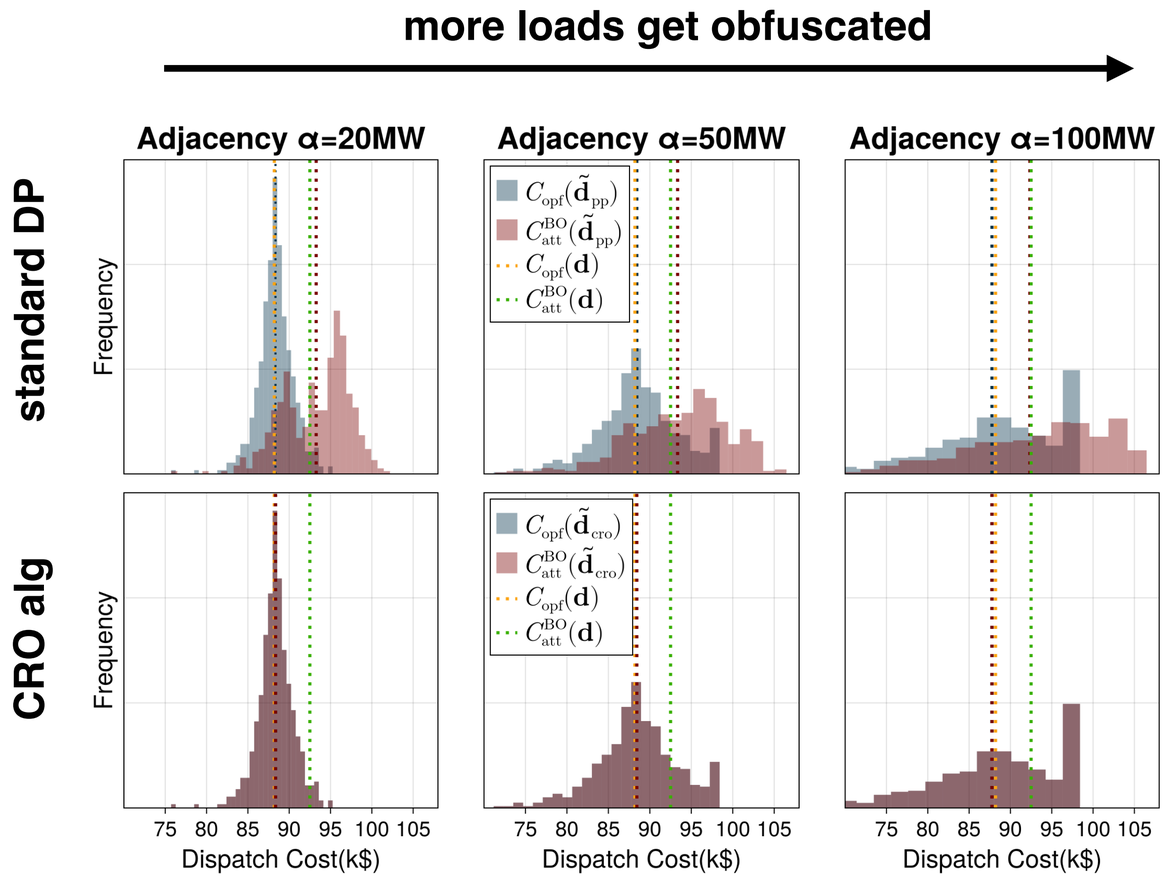

OPF cost distributions in normal and post-attack operation

- Blue distributions - normal operation
- Red distributions - post-attack operation
- Top row - standard DP post-processing
- Bottom row - proposed CRO post-processing

CRO sends important signal to attackers: **the attacks do not lead to extra OPF cost to the system**.
(normal and post-attack distributions overlap)

CRO-Exp application to larger systems

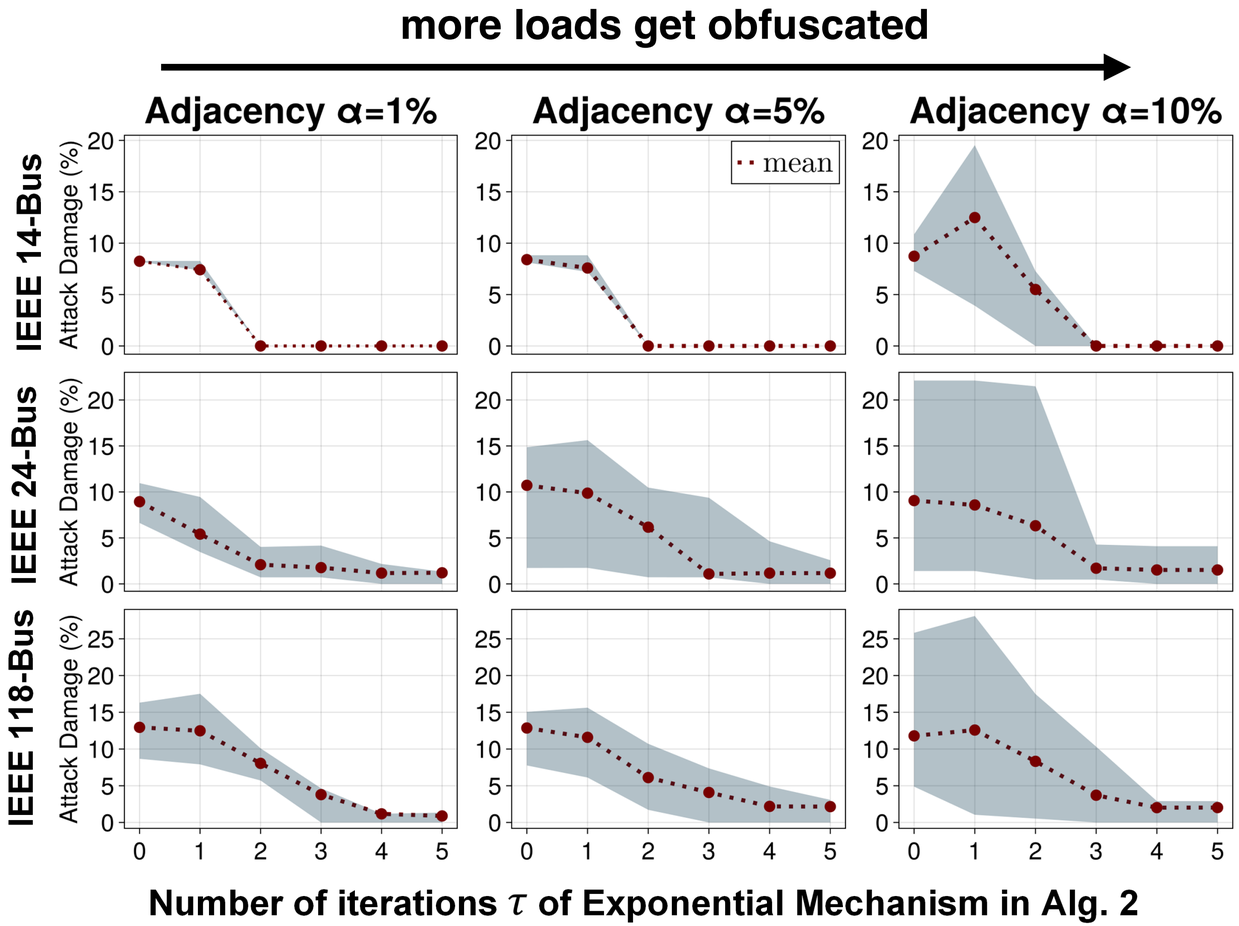

Post-attack damage as a function of constraints selected for post-processing optimization

- The more constraints under attack $\rightarrow$ more resilent the source grid to load redistribution attacks
- 5 constraints on average to minimize the damage

After 5 iterations, the CRO-Exp algorithm identified the most important constraints for post-processing synthetic loads and **ensuring grid resilience**.

Conclusions

- Synthetic grid data is optimized to guarantee privacy, quality and cyber resilience simultaneously
- Trade-offs under linear cost functions are "flat": resilience is achieved with little to no impact on data quality
- The tri-level post-processing optimization can be efficiently collapsed to single-level optimization under reasonable and judicious approximations (connecting bi-level and robust optimization techniques)

Synthesizing Grid Data With Cyber Resilience and Privacy Guarantees

Shengyang Wu, *Student Member, IEEE*, and Vladimir Dvorkin, *Member, IEEE*

***Abstract*—Differential privacy (DP) provides a principled approach to synthesizing data (e.g., loads) from real-world power systems while limiting the exposure of sensitive information. However, adversaries may exploit synthetic data to calibrate cyberattacks on the source grids. To control these risks, we propose new DP algorithms for synthesizing data that provide the source grids with both cyber resilience and privacy guarantees. The algorithms incorporate both normal operation and attack optimization models to balance the fidelity of synthesized data and cyber resilience. The resulting post-processing optimization is reformulated as a robust optimization problem, which is compatible with the exponential mechanism of DP to moderate its computational burden.**

***Index Terms*—Power systems, synthetic dataset, differential privacy, cyber security.**

with such releases remain largely unexplored. Possible cyber attacks include *false data injection*, which subtly alters state estimation results [13], *line outage masking*, which disconnects a transmission line and misguides a control center to seek outage elsewhere [14], and *load redistribution*, which manipulates demand measurements to increase OPF cost and constraint violation [15]. The latter is of main interest to this letter. Executing such attacks requires some grid knowledge [16], which is traditionally difficult to obtain. However, the availability of synthetic grid data may unintentionally inform adversaries and help them calibrate the attack.

Contribution: Recognizing the risks that synthetic grid parameters may inform cyber adversaries, we develop new DP algorithms that simultaneously guarantee cyber resilience and privacy for the source power grids. Our algorithms

Check paper for details on:

- DP guarantees of CRO and CRO-Exp
- Connection between Bi-level and RO
- Experiment settings, data and code

Agenda

1. Intro
2. Formalization of energy data privacy
3. Synthesizing optimization data with privacy and cyber security guarantees
4. **Synthesizing power system dynamics models with privacy guarantees**
5. Synthesizing power flow datasets with constrained diffusion models
6. Outro

The Challenge

- **Stability** of power grids is critical

  Renewables, inverter-based resources, and data centers reshape grid dynamics

- We need a **shared dynamical model** so independent parties can design optimal control

- But releasing the real model **exposes sensitive system details**

- **Goal:** Synthesize a model that behaves nearly identically — without disclosing the original model

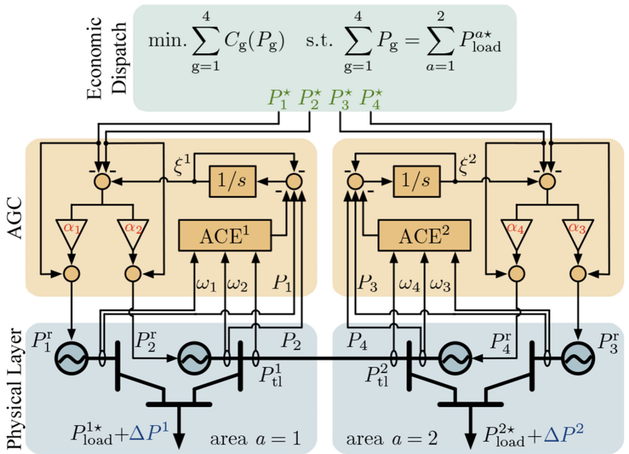

picture courtesy of Sairaj Dhople

Grid frequency dynamic model

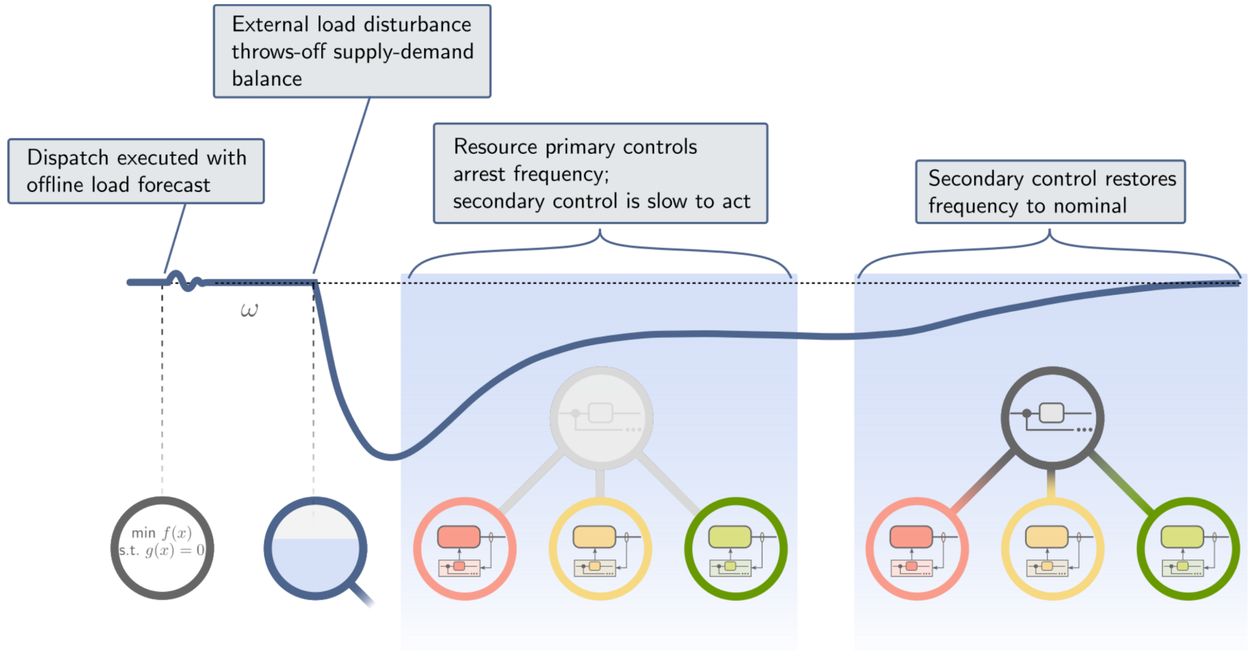

picture courtesy of Sairaj Dhople

$$\dot{\boldsymbol{\omega}} = \mathbf{M}^{-1}(\mathbf{K}\boldsymbol{\delta} + \mathbf{p} - \mathbf{d} - \mathbf{D}\boldsymbol{\omega}) \qquad \textit{swing equation}$$

$$\dot{\boldsymbol{\delta}} = \boldsymbol{\omega} \qquad \textit{phase angle dynamics}$$

$$\dot{\mathbf{p}} = -\mathbf{T}(\mathbf{p} + \mathbf{r} + \mathbf{R}\boldsymbol{\omega}) \qquad \textit{generator control}$$

network topology and parameters

How to release the parameters of system dynamics in a privacy-preserving way?

Classic solution based on Kron reduction

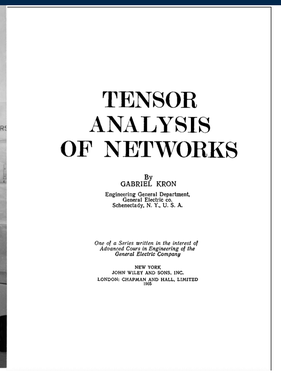

TENSOR
ANALYSIS
OF NETWORKS

By
GABRIEL KRON
Engineering General Department,
General Electric Co.
Schenectady, N. Y., U. S. A.

One of a Series written in the interest of Advanced Cours in Engineering of the General Electric Company

NEW YORK
JOHN WILEY AND SONS, INC.
LONDON: CHAPMAN AND HALL, LIMITED
1955

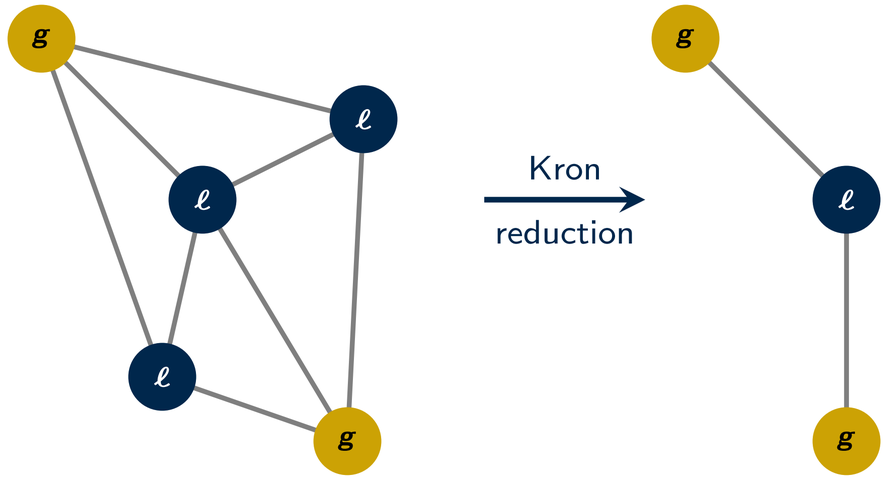

- Partitioning of the swing equation:

$$\begin{bmatrix} \mathbf{M}_g & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\omega}_g \\ \boldsymbol{\omega}_\ell \end{bmatrix} = \begin{bmatrix} \mathbf{K}_{gg} & \mathbf{K}_{g\ell} \\ \mathbf{K}_{\ell g} & \mathbf{K}_{\ell\ell} \end{bmatrix} \begin{bmatrix} \boldsymbol{\delta}_g \\ \boldsymbol{\delta}_\ell \end{bmatrix} + \begin{bmatrix} \mathbf{p}_g \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{d}_g \\ \mathbf{d}_\ell \end{bmatrix} - \begin{bmatrix} \mathbf{D}_g & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_\ell \end{bmatrix} \begin{bmatrix} \boldsymbol{\omega}_g \\ \boldsymbol{\omega}_\ell \end{bmatrix}$$

- Kron-reduced swing equation with the identical dynamics

$$\mathbf{M}_g \dot{\boldsymbol{\omega}}_g = \mathbf{K}_{\text{red}} \boldsymbol{\delta}_g - \mathbf{D}_g \boldsymbol{\omega}_g + (\mathbf{p} - \mathbf{d}_g) - \mathbf{K}_{\text{ac}} \mathbf{d}_\ell$$

- Does the release of the reduced equation preserves privacy?
  No! [SH Low (2024); Deka, Kekatos & Cavraro (2024)]

Differentially private Kron reduction

- We synthesize the reduced topology matrix $\mathbf{K}_{\text{red}}$
- **Step 1:** Perturb the topology with random noise (Laplace mechanism)

$$\tilde{\mathbf{K}}_{\text{red}} = \mathbf{K}_{\text{red}} + \text{Lap}(\alpha/\varepsilon)$$

- **Step 2:** Post-process the perturb topology by optimizing the perturbed model of frequency dynamics

$$\min_{\widehat{\mathbf{K}}_{\text{red}}} \quad \int_0^T \|\boldsymbol{\omega}_g(t|\widehat{\mathbf{K}}_{\text{red}}) - \boldsymbol{\omega}_{\text{ref}}(t)\|_2^2 \, dt + \|\tilde{\mathbf{K}}_{\text{red}} - \widehat{\mathbf{K}}\|_F$$

Frequency on synthetic params

Frequency on real params

$$\text{s.t.} \quad \mathbf{M}_g\dot{\boldsymbol{\omega}}_g = \widehat{\mathbf{K}}_{\text{red}}\boldsymbol{\delta}_g - \mathbf{D}_g\boldsymbol{\omega}_g + (\mathbf{p} - \mathbf{d}_g) - \mathbf{K}_{\text{ac}}\mathbf{d}_\ell$$

Differentially private way to solve this optimization: Adjoint method + DP gradient descent

- The optimal $\widehat{\mathbf{K}}$ preserves the privacy of the original topology in $\mathbf{K}$ but behaves similarly

Differentially private Kron reduction

- We synthesize the reduced topology matrix $\mathbf{K}_{\text{red}}$
- **Step 1:** Perturb the topology with random noise (Laplace mechanism)

$$\tilde{\mathbf{K}}_{\text{red}} = \mathbf{K}_{\text{red}} + \text{Lap}(\alpha/\varepsilon)$$

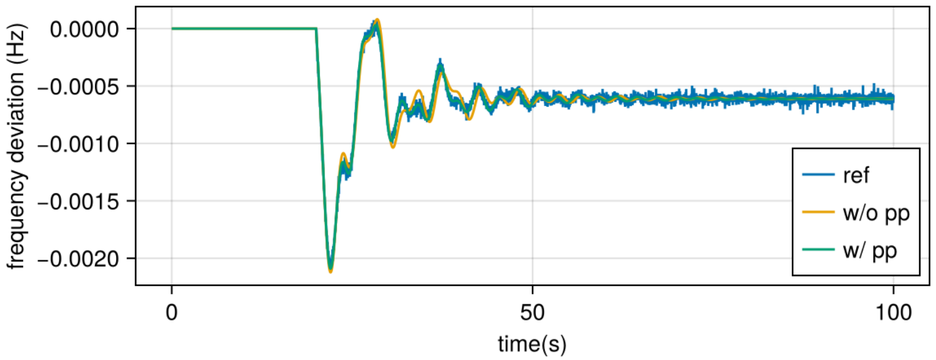

$$\tilde{\mathbf{K}}_{\text{red}} = \mathbf{K}_{\text{red}} + \text{Lap}(1/\varepsilon)$$

Less privacy protection

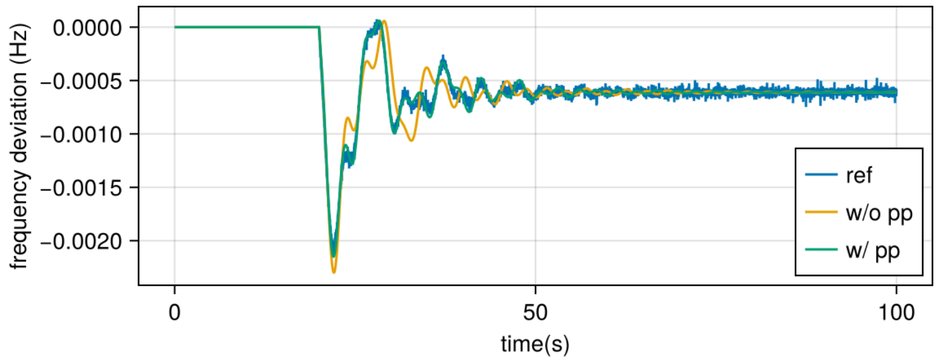

$$\tilde{\mathbf{K}}_{\text{red}} = \mathbf{K}_{\text{red}} + \text{Lap}(3/\varepsilon)$$

More privacy protection

Agenda

1. Intro
2. Formalization of energy data privacy
3. Synthesizing optimization data with privacy and cyber security guarantees
4. Synthesizing power system dynamics models with privacy guarantees
5. **Synthesizing power flow datasets with constrained diffusion models**
6. Outro

Motivation and Goal

- There is an increasing demand for high-quality power flow datasets for machine learning (ML) tasks in power systems: state estimation, optimal power flow solvers, etc.
- Challenges of data availability: privacy and security concerns, legal barriers
- Given a dataset of real power flow data points, a system operator aims to generate:
  (1) **statistically representative** and
  (2) **physically meaningful** synthetic power flow data points.

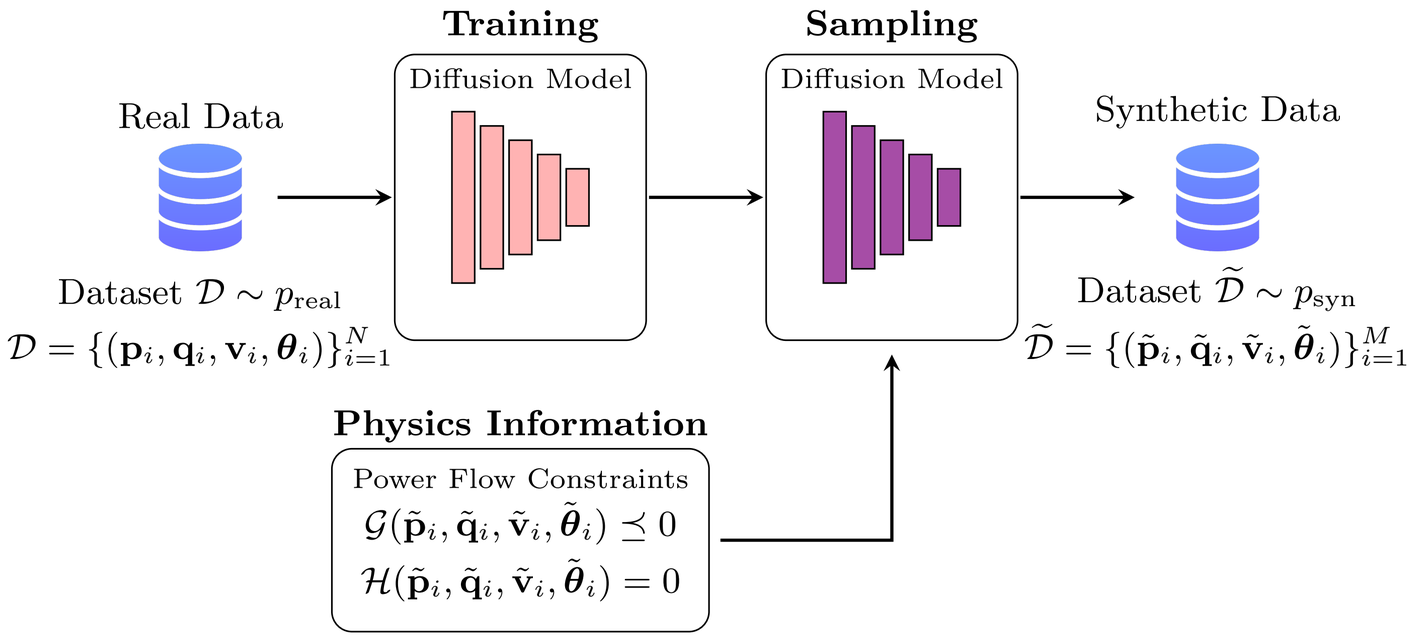

Milad Hoseinpour

Related Work

(1) Generic random sampling approaches:

Iteratively perturbing system parameters (e.g., demand level) around the nominal value and solving the corresponding OPF problem.

- Drawback 1: Datasets do NOT represent the **true underlying distribution** of real operating records.
- Drawback 2: The number of **required** samples to cover the feasible region grows **exponentially**.

(2) Historical data-driven approaches: Generative models (e.g., VAE, GAN)

The historical data-driven approaches leverage real operational records to learn the data distribution.

- Drawback 1: Poor generation quality (VAE), mode collapse (GAN)
- Drawback 2: No rigorous method to **control** their output (e.g., enforcing power flow constraints)

Ref.:

- S. Lovett et al., "OPFData: Large-scale datasets for AC optimal power flow with topological perturbations," arXiv preprint arXiv:2406.07234, 2024.
- A. Jabbar et al., "A survey on generative adversarial networks: Variants, applications, and training," ACM CSUR, vol. 54, no. 8, pp. 1–49, 2021.
- Z. Pan et al., "Data-driven EV load profiles generation using a variational auto-encoder," Energies, vol. 12, no. 5, p. 849, 2019.

Preliminary: Diffusion Models

- Forward diffusion process: iteratively add noise to data sample

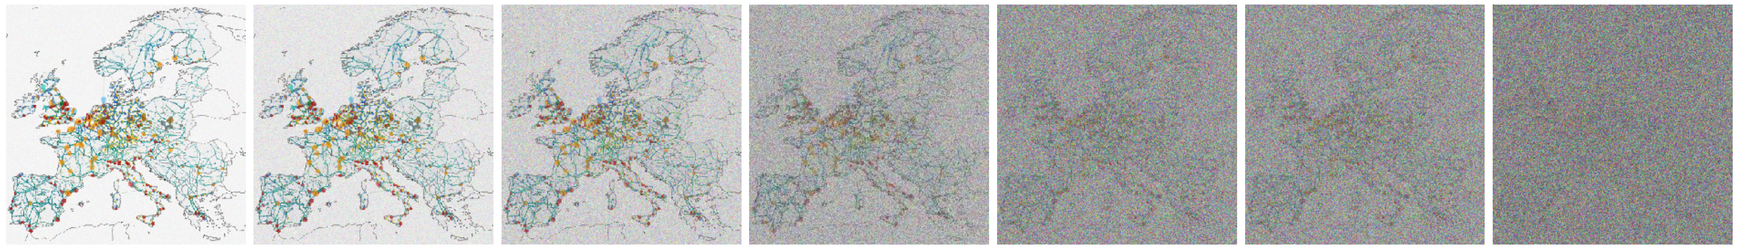

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}\left(\mathbf{x}_t; \sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I}\right)$$

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\epsilon_t, , \quad \epsilon_t \sim \mathcal{N}(0, \mathbb{I}).$$

- Reverse denoising process: denoise to restore the data sample

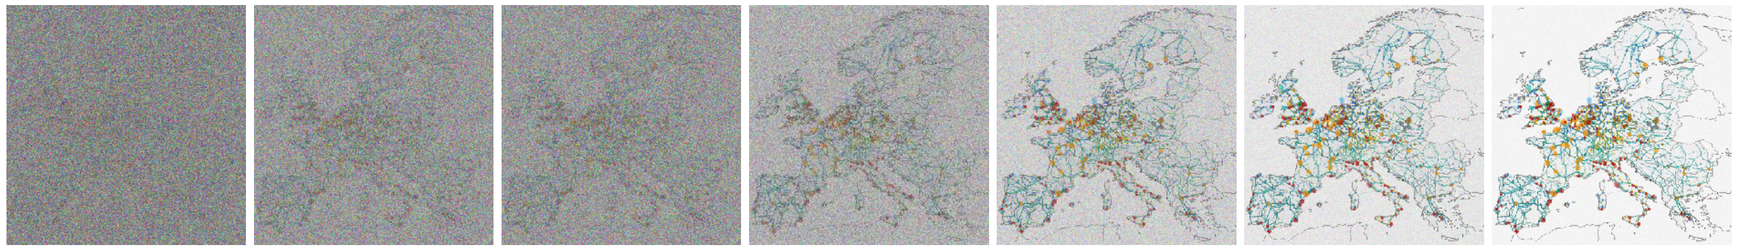

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}\left(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \sigma_t^2\mathbf{I}\right),$$

$$\mathbf{x}_{t-1} = \mu_\theta(\mathbf{x}_t, t) + \sigma_t\epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \mathbb{I}).$$

Power flow constraints

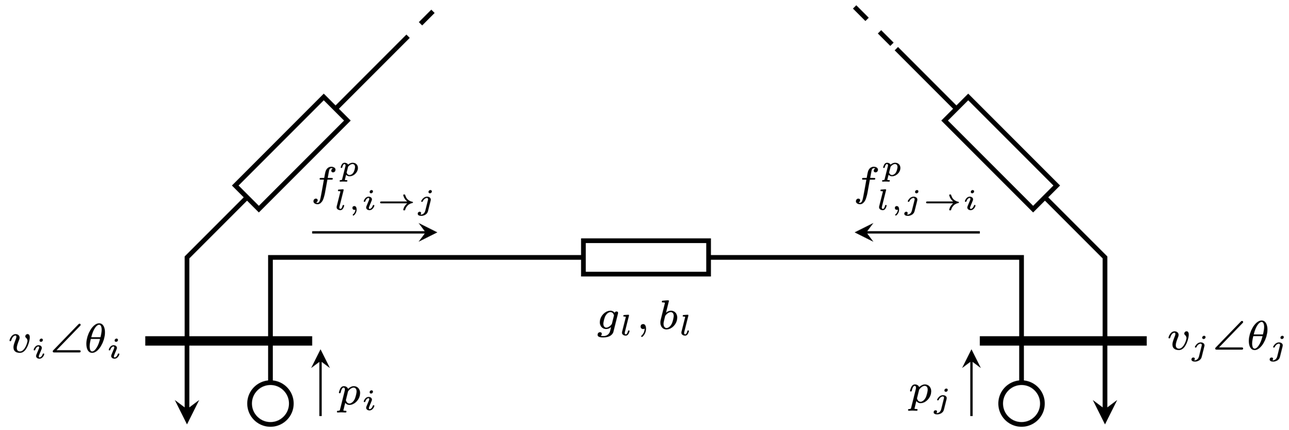

$$p_b - \sum_{l\in\mathcal{L}:i=b} f^p_{l,i\to j} - \sum_{l\in\mathcal{L}:j=b} f^p_{l,j\to i} = 0, \quad \forall b \in \mathcal{B},$$

$$q_b - \sum_{l\in\mathcal{L}:i=b} f^q_{l,i\to j} - \sum_{l\in\mathcal{L}:j=b} f^q_{l,j\to i} = 0, \quad \forall b \in \mathcal{B}.$$

- The expressions for $f^p_{l,i\to j}$ and $f^q_{l,i\to j}$ are

$$f^p_{l,i\to j} = v_i v_j \left[g_l \cos(\theta_i - \theta_j) + b_l \sin(\theta_i - \theta_j)\right], \quad \forall l \in \mathcal{L},$$

$$f^q_{l,i\to j} = v_i v_j \left[g_l \sin(\theta_i - \theta_j) - b_l \cos(\theta_i - \theta_j)\right], \quad \forall l \in \mathcal{L},$$

  where g_l and b_l are the real and imaginary parts of $Y = G + jB$.

Ref.: D. K. Molzahn et al., "A survey of relaxations and approximations of the power flow equations," Found. Trends Electr. Energy Syst., vol. 4, no. 1-2, pp. 1–221, 2019.

Diffusion guidance based on power flow constraints

- **Theoretically**, a diffusion model trained on a dataset of **feasible** power flow data points should satisfy the power flow constraints.
- In practice, a diffusion model may generate power flow data points that are **infeasible** due to learning and sampling errors.

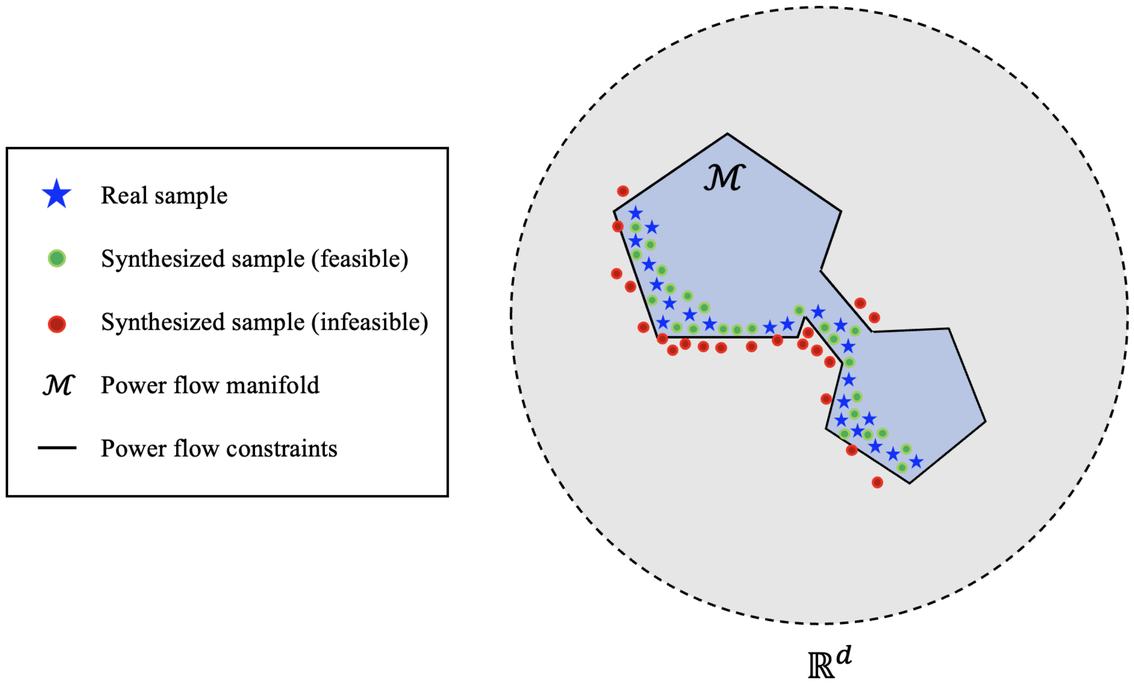

How can ensure that the generated samples satisfy power flow constraint?

Ref.:

- Feng, Berthy T., Ricardo Baptista, and Katherine L. Bouman. "Neural approximate mirror maps for constrained diffusion models." arXiv preprint arXiv:2406.12816, 2024.

- G. Daras et al., "Consistent diffusion models: Mitigating sampling drift by learning to be consistent," Adv. Neural Inf. Process. Syst., vol. 36, pp. 42 038–42 063, 2023.

Geometry of Sampling without Guidance

- Sampling steps as transitions from noisy ($\mathcal{M}_i$) to less noisy ($\mathcal{M}_{i-1}$) data manifolds:
  - (1) Denoise based on $\mathbf{x}_t$ and estimate the clean data $\hat{\mathbf{x}}_0$,
  - (2) add noise w.r.t. the corresponding noise schedule and obtain $\mathbf{x}_{t-1}$.

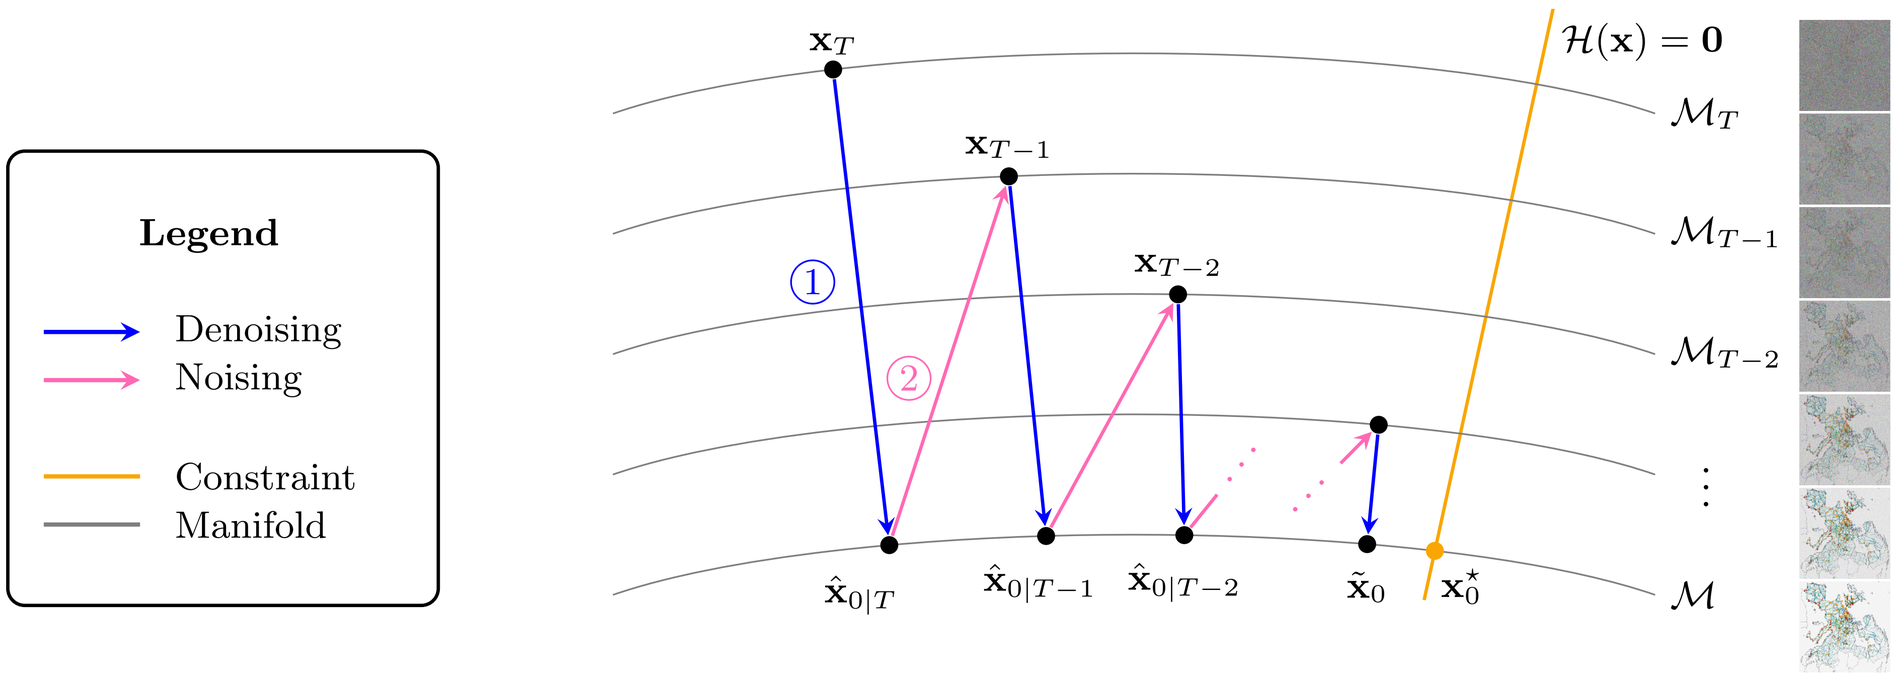

The sampling trajectory is **oblivious** to the power flow constraints.

Diffusion guidance based on power flow constraints

Geometry of sampling with guidance

- Sampling steps as transitions from noisy ($\mathcal{M}_i$) to less noisy ($\mathcal{M}_{i-1}$) data manifolds:
  - (1) Denoise based on $\mathbf{x}_t$ and estimate the clean data $\hat{\mathbf{x}}_0$,
  - (2) add the gradient guidance term,
  - (3) add noise w.r.t. the corresponding noise schedule and obtain $\mathbf{x}_{t-1}$.

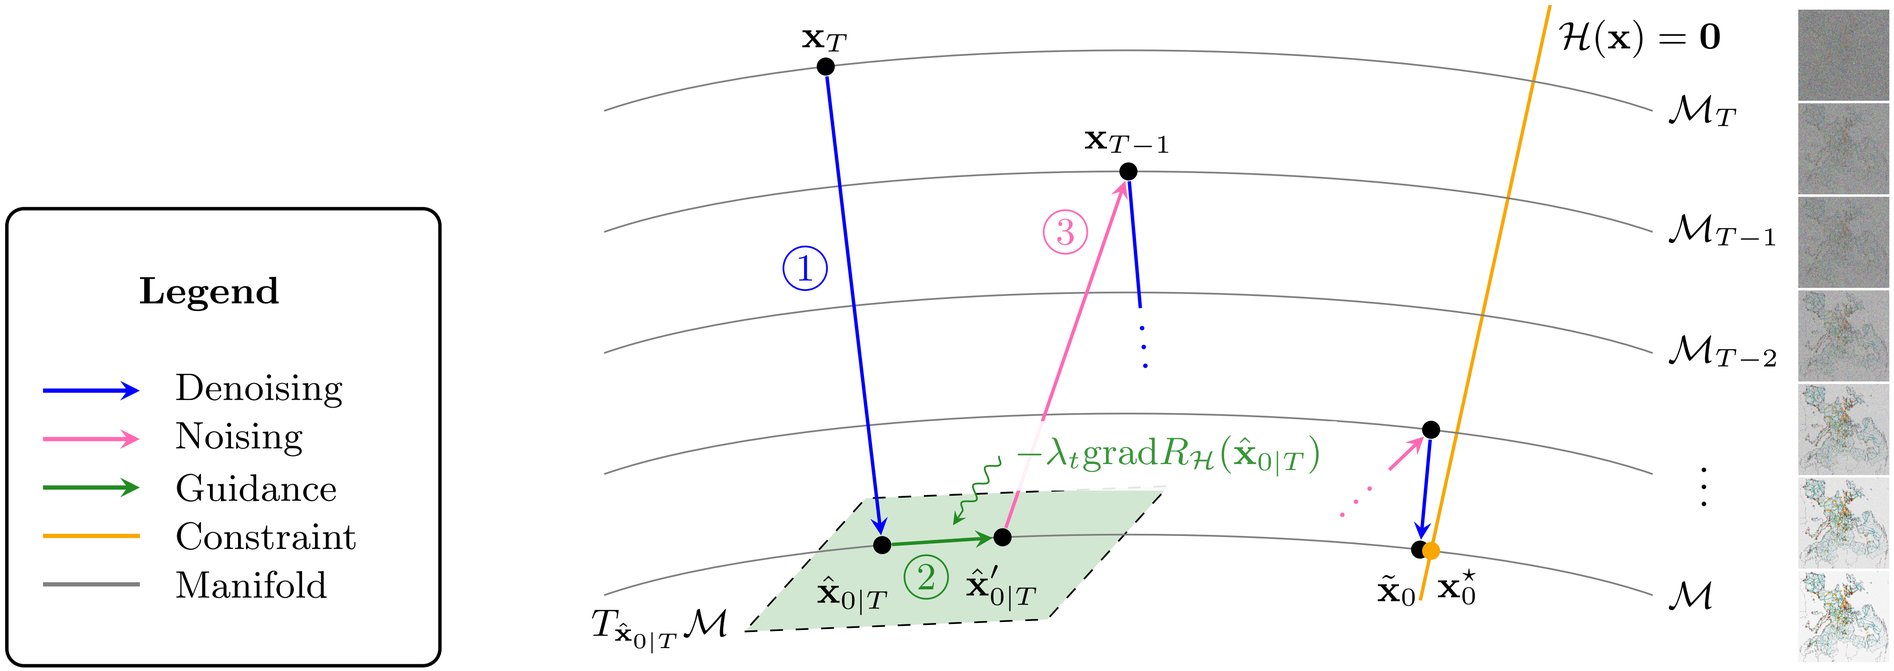

The **gradient guidance** steers the sampling trajectory toward **feasible** power flow data points.

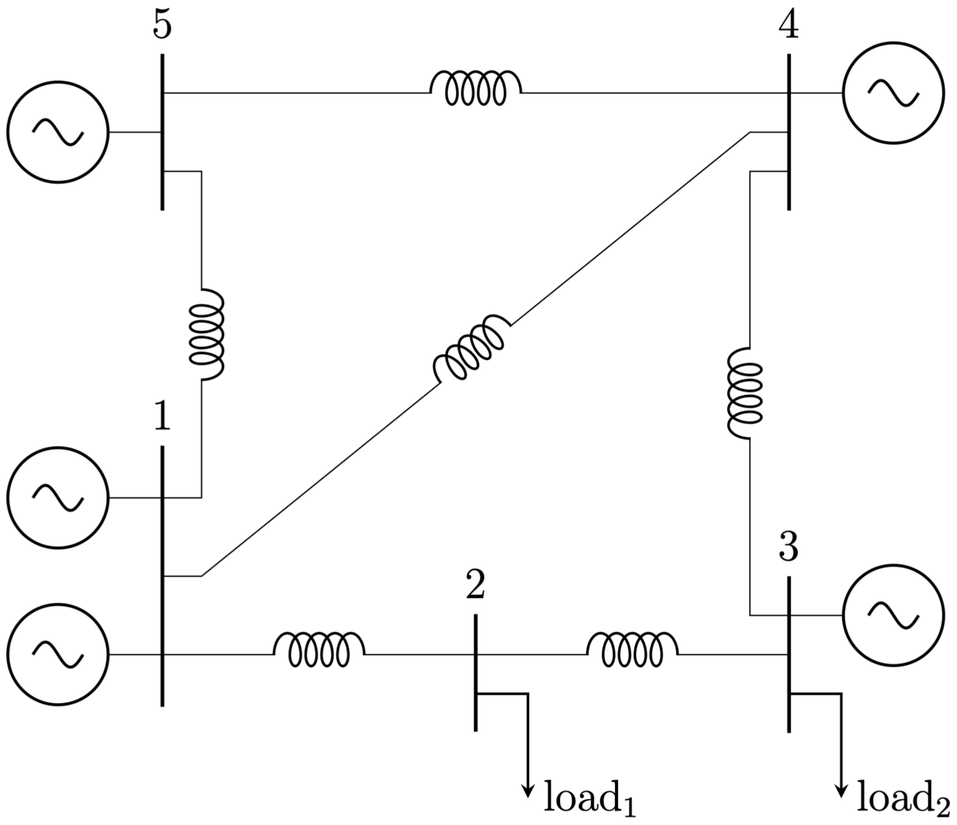

5-Bus PJM system

What usually gets published?

- Electricity prices
- Aggregated statistics (demand, gen mix)

What we can publish using diffusion?

- High-granular streams of operational data
- Active and reactive power injections
- Voltage magnitudes and phase angles

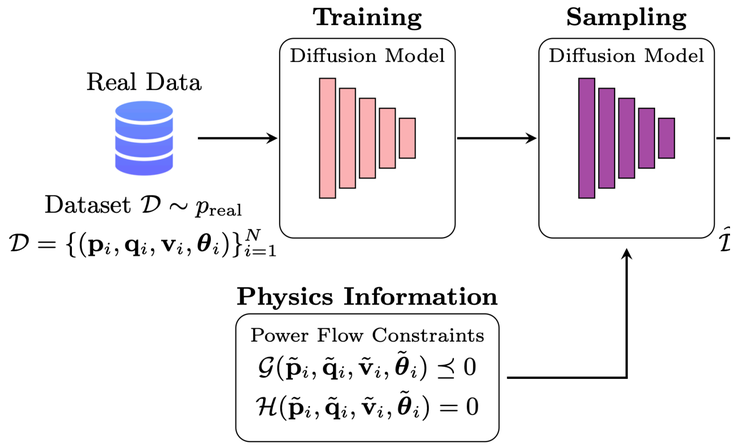

Statistical similarity: Marginal distributions

- Histograms of the ground truth versus synthetic power flow data points:

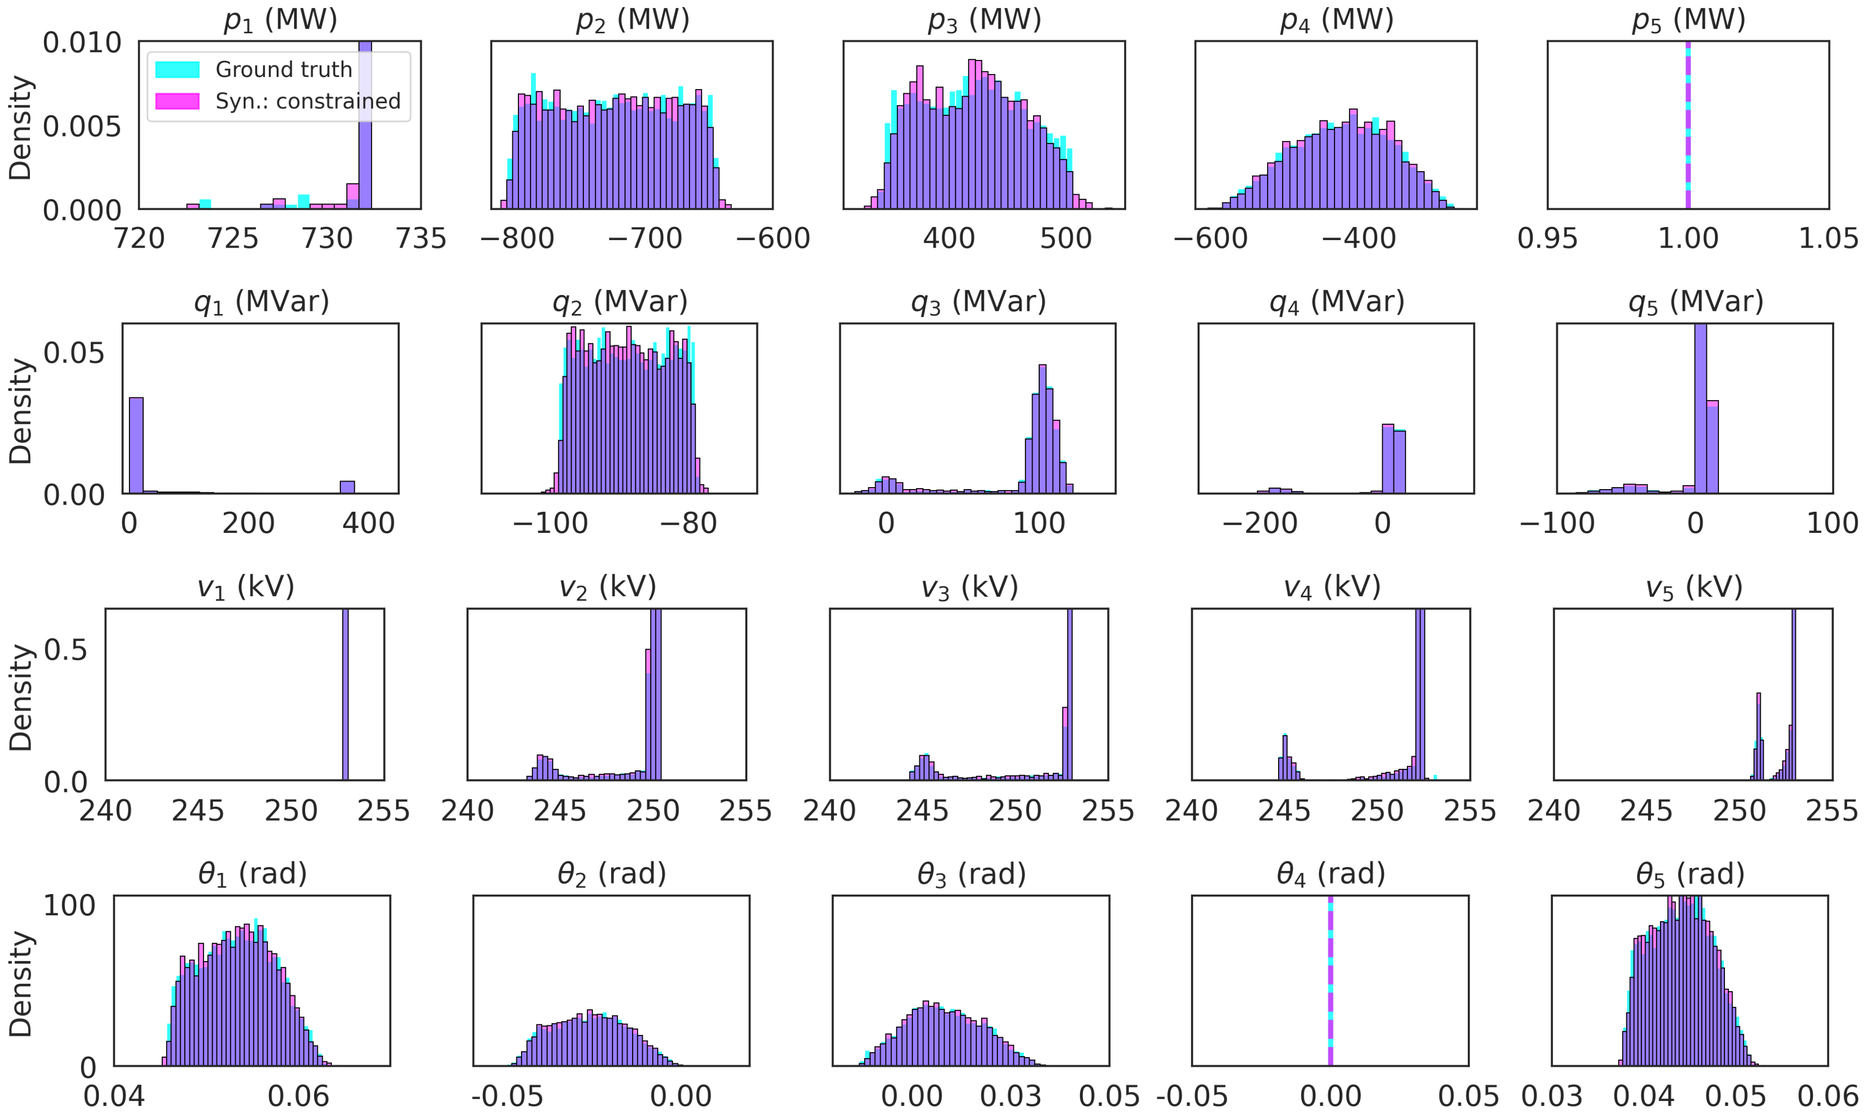

Statistical similarity: Joint distribution

- 2D scatter plots with density estimates of $\mathbf{p} - \mathbf{q}$ and $\mathbf{v} - \boldsymbol{\theta}$:

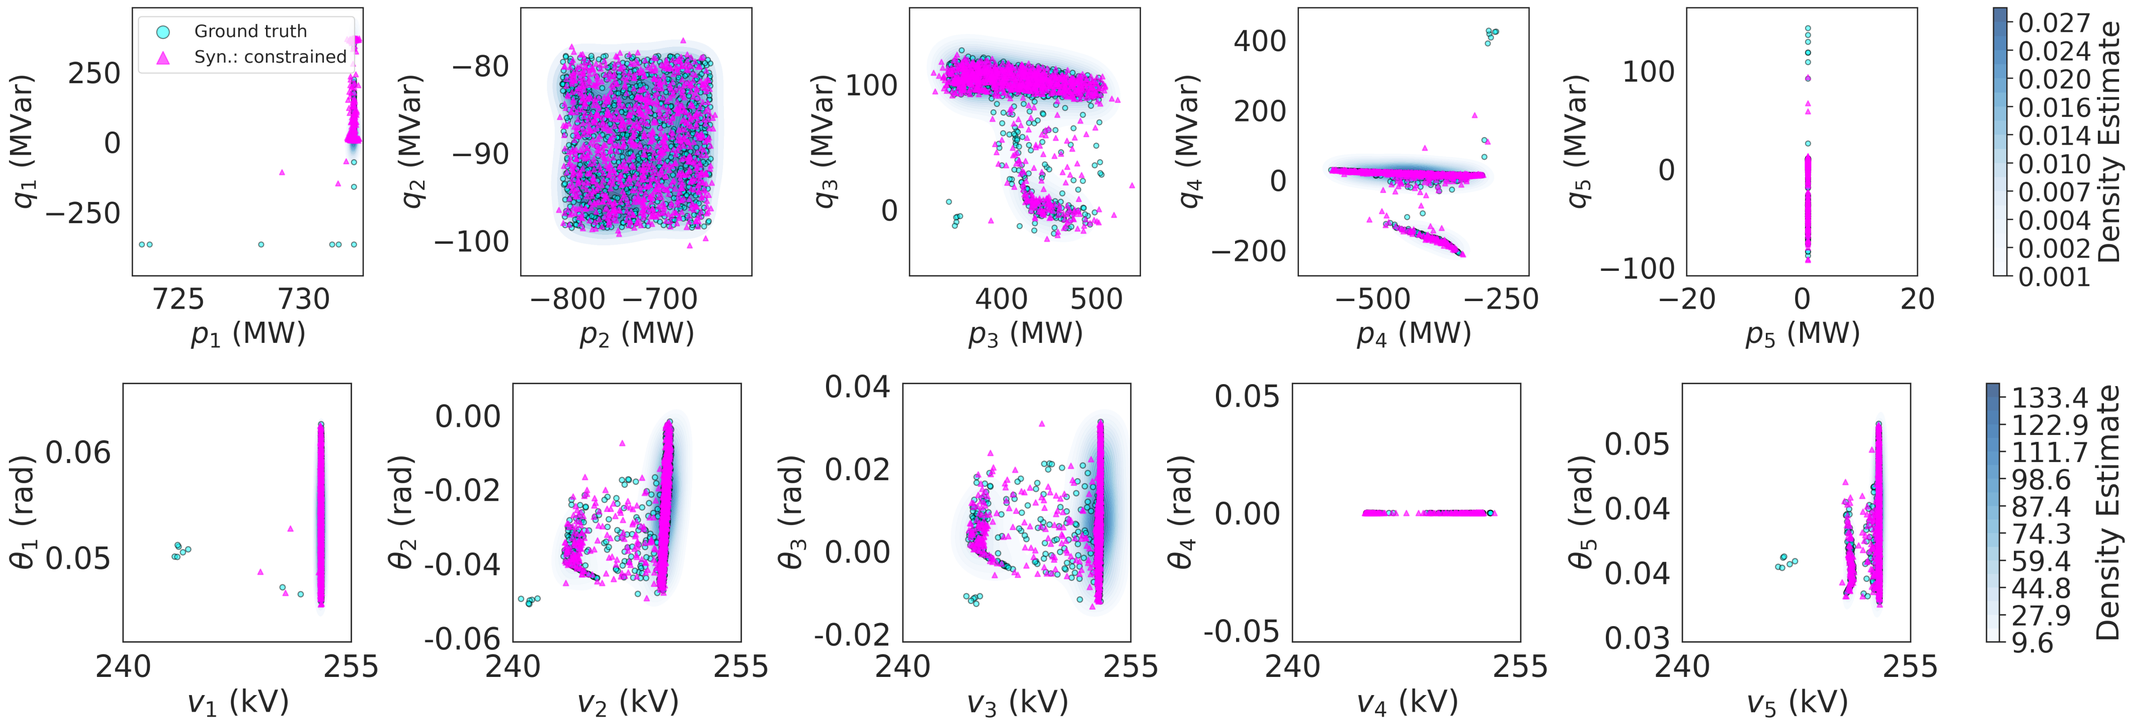

- The synthetic data points
  - closely follow the distributional **pattern** of the real data.
  - closely span the entire **domain** of the real joint probability distributions.
  - captures the **multi-modal structure** of the real data distribution.

Constraint satisfaction of synthetic power flow records

- Histograms of violation magnitudes for the active and reactive power balance constraints

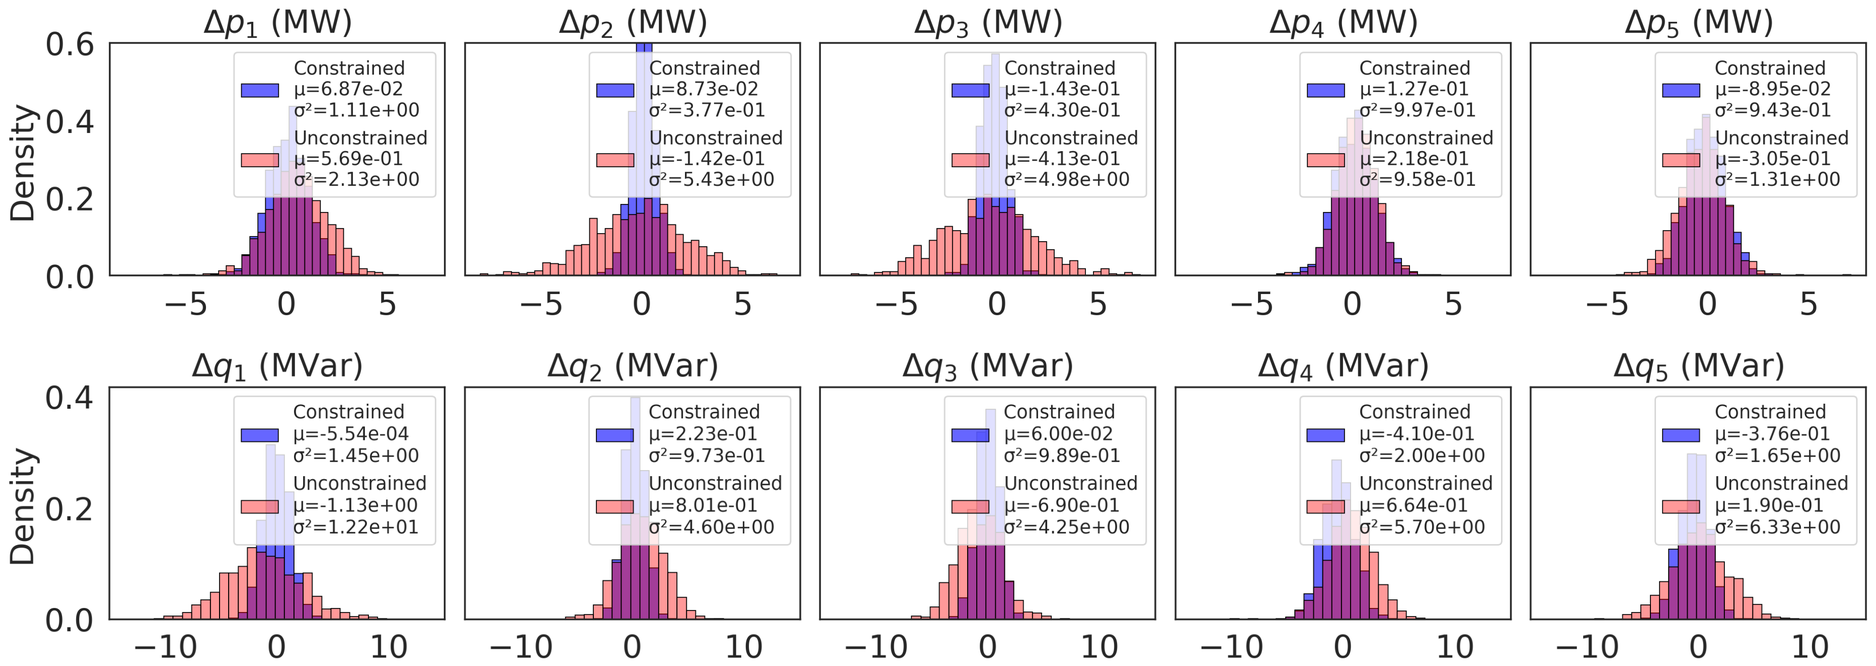

- Wasserstein distances between the ground truth $\mathcal{D}$ and synthetic $\widetilde{\mathcal{D}}$ datasets

Distance between...	5-Bus	24-Bus	118-Bus
...$\mathcal{D}$ and $\widetilde{\mathcal{D}}$ w/o guidance	0.442	0.607	0.622
...$\mathcal{D}$ and $\widetilde{\mathcal{D}}$ w/ guidance	0.382	0.585	0.597

Agenda

1. Intro
2. Formalization of energy data privacy
3. Synthesizing optimization data with privacy and cyber security guarantees
4. Synthesizing power system dynamics models with privacy guarantees
5. Synthesizing power flow datasets with constrained diffusion models
6. **Outro**

Where energy data should go?

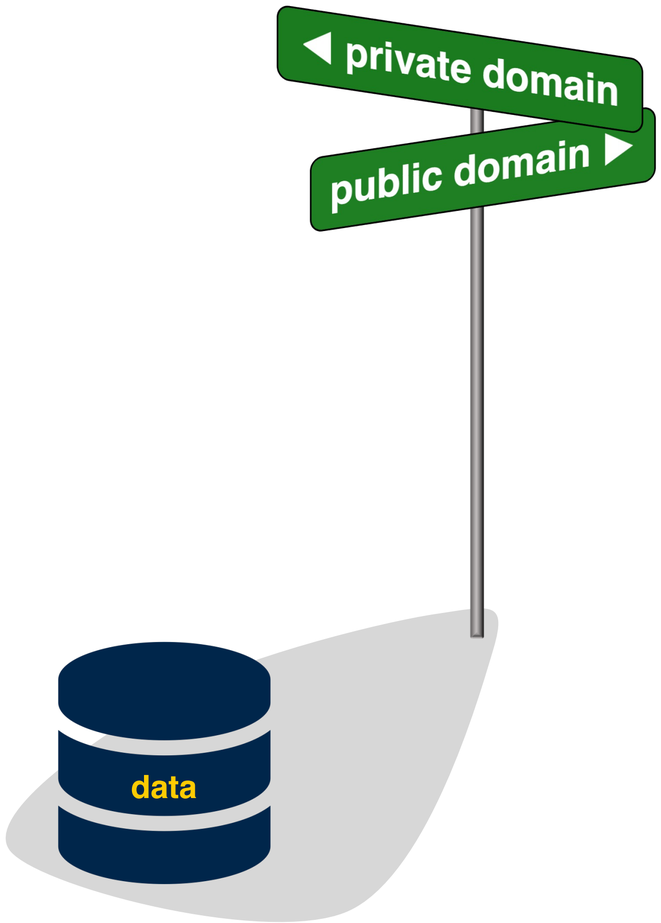

Our privacy-preserving algorithms provide a **non-discrete** answer to this question!

- Controlled disclosure of optimization models, system dynamics, and machine learning datasets
- Rigorous quantification of privacy (for some algorithms), verifiable synthetic datasets
- Performance guarantees across statistical consistency, grid resilience, and physical realism

What we have in other fields….

Computer vision
(ImageNet)

Speech recognition
(LibriSpeech)

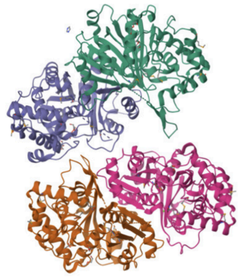

Biology
(UniProt)

What we *can* have in power systems

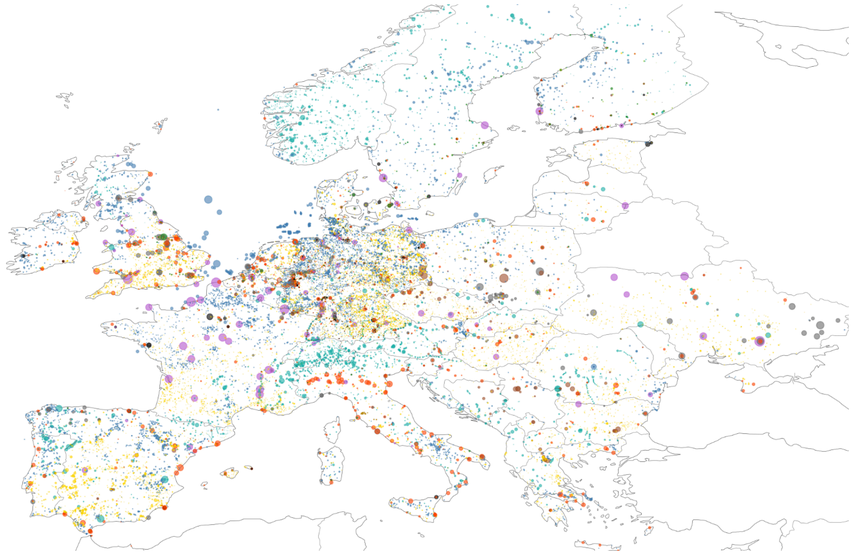

- Statistically credible and operation-feasible models of power grid dispatch
- High-fidelity models of power systems dynamics
- Arbitrary large, credible training datasets for machine learning applications in power systems

Future of synthetic power system datasets

What we used to say about synthetic datasets:

- "[...] data bears **no relation** to the actual grid [...]"
- "This test case represents [...] **fictitious** transmission"
- "This case is synthetic and **does not** model the actual grid"

What we will say about synthetic datasets:

- "This synthetic dataset is produced based on the data from a real-world power grid"
- "It is not possible to infer the real data from this synthetic dataset"
- "Computational results on this data are consistent with the real data"

From today's talk

Synthetic optimization data

438 IEEE CONTROL SYSTEMS LETTERS, VOL. 9, 2025

Synthesizing Grid Data With Cyber Resilience and Privacy Guarantees

Shengyang Wu, *Student Member, IEEE*, and Vladimir Dvorkin, *Member, IEEE*

***Abstract*—Differential privacy (DP) provides a principled approach to synthesizing data (e.g., loads) from real-world power systems while limiting the exposure of sensitive information. However, adversaries may exploit synthetic data to calibrate cyberattacks on the source grids. To control these risks, we propose new DP algorithms for synthesizing data that provide the source grids with both cyber resilience and privacy guarantees. The algorithms incorporate both normal operation and attack optimization models to balance the fidelity of synthesized data and cyber resilience. The resulting post-processing optimization is reformulated as a robust optimization problem, which is compatible with the exponential mechanism of DP to moderate its computational burden.**

***Index Terms*—Power systems, synthetic dataset, differential privacy, cyber security.**

I. Introduction

OPTIMAL power flow (OPF) analysis in power systems requires realistic grid models with accurate network, generation, and load parameters—data that is difficult to source from real-world grids due to privacy and (cyber-)security concerns. While the lack of such models has inspired the development of artificial grids [1], [2], a more principled approach leverages the theory of differential privacy (DP) [3] to release grid models directly from real-world systems.

The DP theory asserts that it is impossible—up to prescribed privacy parameters—to infer the original parameters from their DP release. Such strong privacy guarantees originate from Laplacian perturbations [4] of real grid parameters, followed by post-processing optimization of the perturbed parameters to restore their modeling fidelity to the source grid, e.g., in terms of similarity of the OPF outcomes [5], [6], [7]. The DP theory also lies at the core of modern privacy-preserving OPF solvers [8], [9], [10], the release of aggregated grid statistics [11], and related grid information [12].

However, the privacy guarantees alone may not suffice to release grid parameters, as cybersecurity risks associated with such releases remain largely unexplored. Possible cyber attacks include *false data injection*, which subtly alters state estimation results [13], *line outage masking*, which disconnects a transmission line and misguides a control center to seek outage elsewhere [14], and *load redistribution*, which manipulates demand measurements to increase OPF cost and constraint violation [15]. The latter is of main interest to this letter. Executing such attacks requires some grid knowledge [16], which is traditionally difficult to obtain. However, the availability of synthetic grid data may unintentionally inform adversaries and help them calibrate the attack.

Contribution: Recognizing the risks that synthetic grid parameters may inform cyber adversaries, we develop new DP algorithms that simultaneously guarantee cyber resilience and privacy for the source power grids. Our algorithms build on [5], [6], [7] and leverage the Laplace mechanism and post-processing optimization to tune synthetic data while anticipating cyber risks through embedded attack optimization.

The contributions of this letter are summarized as follows:

1) We formulate a Cyber Resilient Obfuscation (CRO) algorithm, an optimization-based algorithm to release electric load data with a guarantee to preserve the privacy of the original data and ensure the resilience of the source grid to load redistribution attacks. The algorithm post-processes synthetic loads to balance their fidelity with the potential damage to the grid; other grid

Received 16 March 2025; revised 2 May 2025; accepted 15 May 2025. Date of publication 27 May 2025; date of current version 11 June 2025. Recommended by Senior Editor S. Olaru. *(Corresponding author: Shengyang Wu.)*

The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: syseanwu@umich.edu; dvorkin@umich.edu).

Digital Object Identifier 10.1109/LCSYS.2025.3574146

Synthetic machine learning datasets

Constrained Diffusion Models for Synthesizing Representative Power Flow Datasets

Milad Hoseinpour, *Student Member, IEEE*, Vladimir Dvorkin, *Member, IEEE*

***Abstract*—High-quality power flow datasets are essential for training machine learning models in power systems. However, security and privacy concerns restrict access to real-world data, making statistically accurate and physically consistent synthetic datasets a viable alternative. We develop a diffusion model for generating synthetic power flow datasets from real-world power grids that both replicate the statistical properties of the real-world data and ensure AC power flow feasibility. To enforce the constraints, we incorporate gradient guidance based on the power flow constraints to steer diffusion sampling toward feasible samples. For computational efficiency, we further leverage insights from the fast decoupled power flow method and propose a variable decoupling strategy for the training and sampling of the diffusion model. These solutions lead to a physics-informed diffusion model, generating power flow datasets that outperform those from the standard diffusion in terms of feasibility and statistical similarity, as shown in experiments across IEEE benchmark systems.**

***Index Terms*—Diffusion model, generative AI in power systems, physics-informed machine learning, power flow, synthetic data.**

I. Introduction

POWER flow datasets [1]–[3] are essential for training and benchmarking machine learning (ML) models for optimal power flow (OPF) [4] and state estimation [5]. However, the real-world power flow datasets are rarely available due to privacy, security, and legal barriers [6]–[10]. Recent advances in generative AI, capable of producing synthetic data with distributions similar to the original data [11]–[20], have partially lifted these barriers, yet statistical consistency alone cannot guarantee adherence to physical grid constraints [21]. Consequently, ML models trained on constraint-agnostic synthetic datasets are likely to perform substantially worse than those trained on original data. This paper introduces a data generation framework to synthesize statistically consistent and physically meaningful power flow datasets. To achieve this, we develop a constrained diffusion model to learn the underlying distribution of power flow data and generate synthetic samples that are both statistically representative and feasible with respect to the AC power flow constraints. This constrained diffusion model can be trained internally by system operators to publicly release high-quality synthetic power flow data to support a wide range of downstream ML applications.

A. Related Work

The literature on generating synthetic datasets for power systems broadly falls into two categories: generic random sampling and historical data-driven approaches.

The former focuses on power flow data generation through iterative uniform sampling of loads followed by solving the OPF problem [22], [23]. In [24], authors use a truncated Gaussian distribution as another variation of sampling, which also accounts for correlations between power injections at different locations. However, the datasets based on generic sampling only represent a small portion of the feasibility region. To solve this, [6] uniformly samples loads from a convex set, containing the feasible region, and iteratively refines this set using infeasibility certificates. In [25], a bilevel optimization is proposed to sample operating conditions close to the boundaries of the feasible region, which is more informative that a random sampling. A basic requirement for ML-based OPF solvers is robustness to grid topology variations, e.g., network topology switching [26]. To meet this requirement, authors in [27] incorporate topological perturbations in addition to load perturbations in their synthetic data generation framework.

Although straightforward, random sampling comes with certain limitations. The resulting datasets do not represent the true underlying distribution of real-world operating conditions. That is, the synthetic data points may fail to capture correlations, patterns, or variability present in historical data. ML-based OPF solvers trained on such data may generalize poorly, leading to inaccurate predictions and erroneous uncertainty quantification [28], [29]. Moreover, the required number of

The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA. E-mail: {miladh, dvorkin}@umich.edu.

Synthetic dynamics models

Coming soon

Thank you for your attention!