

# Dual-Head Physics-Informed Graph Decision Transformer for Distribution System Restoration

Hong Zhao<sup>1</sup>, Jin Wei-Kocsis<sup>1\*</sup>, Adel Heidari Akhijahani<sup>2</sup>, Karen L Butler-Purry<sup>2</sup>,

<sup>1</sup>Purdue University, West Lafayette

<sup>2</sup>Texas A&M University, College Station

{zhao1211,kocsis0}@purdue.edu, adelheidari@exchange.tamu.edu, kbutler@tamu.edu

## Abstract

Driven by recent advances in sensing and computing, deep reinforcement learning (DRL) technologies have shown great potential for addressing distribution system restoration (DSR) under uncertainty. However, their data-intensive nature and reliance on the Markov Decision Process (MDP) assumption limit their ability to handle scenarios that require long-term temporal dependencies or few-shot and zero-shot decision making. Emerging Decision Transformers (DTs), which leverage causal transformers for sequence modeling in DRL tasks, offer a promising alternative. However, their reliance on return-to-go (RTG) cloning and limited generalization capacity restricts their effectiveness in dynamic power system environments. To address these challenges, we introduce an innovative Dual-Head Physics-informed Graph Decision Transformer (DH-PGDT) that integrates physical modeling, structural reasoning, and subgoal-based guidance to enable scalable and robust DSR even in zero-shot or few-shot scenarios. DH-PGDT features a dual-head physics-informed causal transformer architecture comprising Guidance Head, which generates subgoal representations, and Action Head, which uses these subgoals to generate actions independently of RTG. It also incorporates an operational constraint-aware graph reasoning module that encodes power system topology and operational constraints to generate a confidence-weighted action vector for refining DT trajectories. This design effectively improves generalization and enables robust adaptation to unseen scenarios. While this work focuses on DSR, the underlying computing model of the proposed PGDT is broadly applicable to sequential decision making across various power system operations and other complex engineering domains.

## Introduction

Recent extreme events, such as the massive power outage in Houston caused by Hurricane Beryl, have underscored the urgent need to enhance the resilience of modern power systems. The distribution system plays a pivotal role in delivering electricity to end-users, and its ability to recover swiftly from disruptions is crucial to overall grid resilience. A key indicator of distribution system resiliency is the ability to restore service to critical loads following disruptions, known as Distribution System Restoration (DSR) (Chen, Wang, and

Ton 2017). Forming microgrids (MGs) with dynamic boundaries as a service restoration strategy is a promising solution for enabling effective DSR (Liang et al. 2022; Igder, Liang, and Mitolo 2022). By incorporating various energy sources, distributed generators (DGs), along with remotely controlled switches, a distribution system can be partitioned into multiple self-sufficient MGs. This strategy enables rapid reconfiguration and localized recovery, significantly enhancing system adaptability and restoration efficiency. In this paper, we focus on advancing DSR through the sequential formation of MGs with dynamic boundaries to improve distribution system resiliency under uncertainty.

Traditional methods for DSR, such as mathematical programming, heuristic algorithms, and expert systems, often suffer from limited scalability and robustness. In recent years, machine learning, particularly deep reinforcement learning (DRL), has emerged as a powerful approach for enabling more efficient, adaptive, and robust DSR (Wu et al. 2019; Gao et al. 2020). However, DRL methods typically rely on Markov Decision Process (MDP) or partially observable MDP assumptions and require extensive training data, which limits their applicability to large-scale DSR scenarios that demand long-horizon temporal reasoning or operate under few-shot or zero-shot conditions.

The emergence of Decision Transformers (DTs), which adapt causal transformer architectures for scalable sequential decision-making while effectively capturing long-term temporal dependencies, has introduced a compelling alternative to conventional DRL methods (Chen et al. 2021). However, their application to DSR remains largely unexplored. To the best of our knowledge, our prior work (Zhao et al. 2025) is the first to investigate the use of DTs in DSR operations. Despite their potential, their reliance on RTG cloning and limited generalization capacity can restrict their effectiveness in dynamic power system environments. Recent studies have attempted to address the generalization limitation of DTs from various perspectives. A goal-conditioned DT with multi-objective pretraining was proposed using explicit goals and time-to-goal inputs (Fu et al. 2024). Although effective in surgical robotics, its core idea on goal variability limits applicability to DSR, where the restoration objective remains consistent. A conservative DT was proposed to avoid RTG at inference by using a separate transformer to estimate rewards conservatively (Kim, Noh,

\*Corresponding author.

and Jang 2023). However, the method still assumes MDP and significantly increases computational cost due to the use of dual transformers. In (Ma et al. 2023), the generalization limitation was addressed by introducing a hierarchical reformulation with auto-tuned high-level prompts and joint optimization, and, in (Hsu et al. 2024), the limitation was addressed by replacing RTG with a temporal-difference-learned steering signal. Both methods retain MDP assumptions. Trajectory prompt-based approaches were introduced in (Xu et al. 2022; Yao et al. 2025), which condition DTs on short trajectory demonstrations for few-shot generalization. Although promising, their reliance on high-quality demonstrations limits applicability to DSR in unforeseen scenarios. Therefore, while effective in their respective application scenarios, these DT variants are not well suited for DSR, motivating the development of our proposed framework, which is explicitly designed to overcome these limitations in DSR settings characterized by uncertainty, complex constraints, and limited demonstration data.

In this paper, we propose a novel Dual-Head Physics-informed Graph Decision Transformer (DH-PGDT) that enables robust and scalable DSR even in unforeseen operation scenarios with complex physical and operational constraints. DH-PGDT integrates an operational constraint-aware graph reasoning module with a dual-head physics-informed causal transformer architecture to incorporate physical and operational information while reducing reliance on RTG. The main contributions of our proposed work are twofold: (1) developing a novel DH-PGDT framework that enables physical and operational constraint-aware and robust DSR decisions under uncertain and unforeseen conditions; and (2) conducting comprehensive evaluations demonstrating the effectiveness of DH-PGDT in terms of restoration efficiency, constraint satisfaction, and zero-shot/few-shot adaptability. Additionally, although our current work focuses on DSR, the underlying computing framework is broadly applicable to sequential decision-making tasks in other complex engineering domains.

The next section illustrates the problem settings for our work. After this, we describe our proposed DH-PGDT framework for the DSR decision making. The following section shows the case studies and performance evaluations of the proposed DH-PGDT framework. Conclusions are presented in the last section.

## Problem Settings

### DSR Problem Modeling

In the early stage of this research, the proposed DSR method is modeled as a sequential decision-making process involving sequences of control actions on switches to restore loads to their normal operational states. These sequences of switching actions, referred to as energization branches, are each associated with a single energy source (e.g., a single active distributed generator (DG)). Specifically, each energization branch begins at the switch connected to the associated energy source and aims to maximize load restoration by forming multiple microgrids, while ensuring compliance with all operational and physical constraints. Additionally,

to reduce the space of control actions in the modeling, we adopt the concept of node cell as introduced in Chen et al. 2019. A node cell is defined as a set of nodes that are interconnected directly by non-switchable lines. Consequently, all the lines and loads within a node cell will be energized simultaneously. Furthermore, in our current work, the physical constraints include: 1) the node voltage limits, 2) power flow constraints, 3) DGs' generation capacities, and 4) DGs' prime mover ability. The operational constraints include 1) prevention of loop formation and 2) ensuring no node cell is visited more than once.

### DSR Problem Formulation within a DT Framework

We further formulate the sequential decision-making process for energization branches in the DSR problem model. Table 1 shows the definitions of the parameters and variables that are used in this paper.

$s_t, \mathcal{S}$	State at time $t$ , the state space
$a_t, \mathcal{A}$	Action at time $t$ , the action space
$R$	Reward
$\mathcal{P}$	Transition probability
$r_t$	Reward at time $t$
$T$	Time Horizon
$L$	Set of all the loads
$N$	Set of all buses
$C,  C $	Set of all node cells, number of node cells
$E$	Set of all DGs
$x_{i,t}^L, x_{i,t}^C$	Energization status of Load $i$ at time $t$ , energization status of node cell $i$ at time $t$
$P_{l,t}^L, P_{l,t}^C, P_{l,t}$	Active power of Load $l$ at time $t$ when restored, accumulative active power of all loads in node cell $l$ at time $t$ when restored, nominal active power of Load $l$ at time $t$ .
$P_{dg,t}^E$	Active output power of DG $dg$ at time $t$ .
$P_{dg}^{\text{ramp}}$	Maximum absolute change of active power settings of DG $dg$
$H_{i,t}^L$	Squared voltage magnitude of Load $i$ at time $t$
$H_i^{\min}, H_i^{\max}$	Minimum and maximum squared nodal voltage of Load $i$
$V_{l,t}^{L,p}$	Voltage penalty function of Load $l$ at time $t$
$S_{dg,t}^{E,p}$	Ramp rate penalty function of DG $dg$ at time $t$
$\theta$	Learnable parameters for the proposed DH-PGDT
$q$	Number of subgoals for the DH-PGDT

Table 1: Definition of variables and parameters

In the context of DT, the key components in the DSR

problem are formulated as follows:

- State  $\mathbf{s}_t \in \mathcal{S}$  is defined as the available observation vector of the overall distribution system at time  $t$ . The state vector  $\mathbf{s}_t = [\mathbf{s}_{1,t} \parallel \mathbf{s}_{2,t}]$  is defined as the concatenation of two vectors,  $\mathbf{s}_{1,t}$  and  $\mathbf{s}_{2,t}$ :
  - $\mathbf{s}_{1,t}$  is a binary vector representing energization status of individual node cells. If node cell  $j$  is chosen to be energized at step  $t$ , then the  $j$ -th element will always be 1 at any time step larger than  $t$ .
  - $\mathbf{s}_{2,t}$  is another binary vector representing whether the individual node cells are activated at the current time step  $t$ . Number of 1's in  $\mathbf{s}_{2,t}$  will always equal the number of energization branches.
- Action  $\mathbf{a}_t \in \mathcal{A}$  is defined as a vector of confidence probabilities over the operable switches, where each element represents the model's confidence in activating a specific switch. During DSR, switches are sampled from  $\mathbf{a}_t$  to determine which to activate, aiming to maximize load restoration while satisfying operational and physical constraints.
- Time horizon  $T$  defines the total number of decision-making steps.
- Physics-informed Return-to-go (RTG)  $\hat{R}_t$  is defined as  $\hat{R}_t = \sum_{j=t}^T r_j$ , which is the sum of rewards  $r_t$  from a given time to the end of the time horizon. The rewards  $r_t$  evaluate the effectiveness of action  $\mathbf{a}_t$  taken in state  $\mathbf{s}_t$  in achieving the objective of our DSR problem, which is to maximize the load restoration in the time horizon  $T$  while adhering to physical constraints. Therefore, the reward will be formulated as:

$$r_t = R_t^A(\mathbf{s}_t, \mathbf{a}_t) + w_{p_1} \times R_t^V(\mathbf{s}_t, \mathbf{a}_t) + w_{p_2} \times R_t^E(\mathbf{s}_t, \mathbf{a}_t) \quad (1)$$

where:

- $R_t^A(\mathbf{s}_t, \mathbf{a}_t)$  is the reward related to the total active power restoration at  $t$ , which is defined as:

$$R_t^A(\mathbf{s}_t, \mathbf{a}_t) = \left( \sum_{l \in L} x_{l,t}^L P_{l,t}^L \right) \times \Delta t \quad (2)$$

By incorporating the concept of node cell, Eq. (2) can be rewritten as:

$$R_t^A(\mathbf{s}_t, \mathbf{a}_t) = \left( \sum_{l \in C} x_{l,t}^C P_{l,t}^C \right) \times \Delta t \quad (3)$$

- $R_t^V(\mathbf{s}_t, \mathbf{a}_t)$  is a penalty term to penalize actions that violate the voltage constraints. It is defined as:

$$R_t^V(\mathbf{s}_t, \mathbf{a}_t) = - \left( \sum_{l \in L} x_{l,t}^L V_{l,t}^{L,p} \right) \times \Delta t \quad (4)$$

where

$$V_{l,t}^{L,p} = \max(0, H_{l,t}^L - H_l^{\max}) + \max(0, H_l^{\min} - H_{l,t}^L) \quad (5)$$

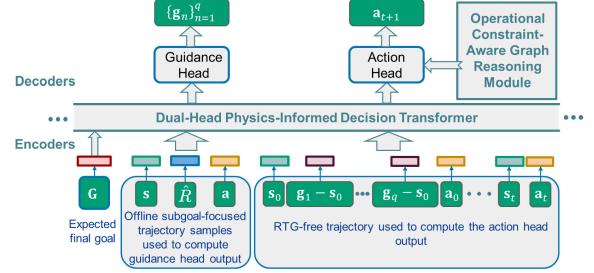


Figure 1: Overview of the architecture of our proposed DH-PGDT framework.

Where  $V_{l,t}^{L,p}$  represents the penalty terms associated with individual loads to ensure their node voltage magnitudes do not violate the constraints. These constraints are formulated using squared nodal voltage range  $[H_i^{\min}, H_i^{\max}]$ .

- $R_t^E(\mathbf{s}_t, \mathbf{a}_t)$  is a penalty term to penalize actions that violate the DG ramp rate constraints. It is defined as:

$$R_t^E(\mathbf{s}_t, \mathbf{a}_t) = - \left( \sum_{dg \in E} S_{dg,t}^{E,p} \right) \times \Delta t \quad (6)$$

where

$$S_{dg,t}^{E,p} = \max \left( 0, (P_{dg,t}^E - P_{dg,t-1}^E) - P_{dg}^{\text{ramp}} \right) + \max \left( 0, -P_{dg}^{\text{ramp}} - (P_{dg,t}^E - P_{dg,t-1}^E) \right) \quad (7)$$

Where  $S_{dg,t}^{E,p}$  represents the penalty terms associated with individual DGs to ensure their absolute change in the active power output do not exceed given maximum values. Alternatively, this is to say that the change in active power output  $P_{dg,t}^E - P_{dg,t-1}^E$  for DG  $dg$  at time  $t$  lies in the range  $[-P_{dg}^{\text{ramp}}, P_{dg}^{\text{ramp}}]$ .

The weight terms  $w_{p_1}$  and  $w_{p_2}$  are set to ensure that the penalties are comparable to the active restored power term.

- Operational constraint-aware mask function  $f(\mathbf{s}_t, \cdot)$  is formulated to automate the enforcement of operational constraints. Its output is a mask vector  $\mathbf{M}$  that is used to filter out actions that violate these constraints during decision making.

## Proposed DH-PGDT Framework for DSR Decision Makings

In this section, we will introduce our proposed DH-PGDT. As illustrated in Fig. 1, the DH-PGDT mainly consists of a dual-head physics-informed causal transformer architecture and an operational constraint-aware graph reasoning module.

### Dual-Head Physics-Informed Causal Transformer Architecture

As shown in Fig. 1, our dual-head physics-informed causal transformer architecture mainly comprises: 1) an encoder

that consists of linear layers followed by an embedding layer; 2) a GPT-based dual-head physics-informed causal transformer with a causal self-attention mask; and 3) a decoder. The encoder processes offline subgoal-focused or RTG-free trajectory samples into token embeddings, which are input to the GPT-based dual-head physics-informed causal transformer to compute the outputs of the guidance and action heads. The action head output is then passed to the decoder, which maps the predicted trajectory embeddings back to the original action space, producing the final control actions for the operational switches in DSR tasks.

The GPT-based dual-head physics-informed causal transformer lies at the core of our proposed architecture. It integrates Guidance Head and Action Head to achieve both effectiveness and robustness in decision-making for DSR. In the following, we present the detailed design of the dual-head transformer.

**Guidance Head Design:** The *Guidance Head* is designed to predict a sequence of subgoal representations  $\{\mathbf{g}_n\}_{n=1}^q$  in the state space  $\mathcal{S}$ , conditioned on the expected final goal  $\mathbf{G}$  and a subgoal-focused trajectory  $\tau^g$ , as illustrated in Fig. 1. Both the final goal  $\mathbf{G}$  and the predicted subgoals  $\mathbf{g}_n$  lie in the same state space  $\mathcal{S}$ . In our current early-stage implementation, the final goal  $\mathbf{G}$  is defined as the fully energized target state:

$$\mathbf{G} = \mathbf{s}_T = [\mathbf{s}_{1,T} \parallel \mathbf{s}_{2,T}],$$

where  $\mathbf{s}_{1,T}$  is a vector of all ones (indicating that all node cells are energized), and  $\mathbf{s}_{2,T}$  is masked. To construct subgoals, in our current work, we predefine a fixed number  $q$  of intermediate subgoals. Each subgoal  $\mathbf{g}_n$  corresponds to the state  $\mathbf{s}_{m_n}$  when  $\lceil nN/(q+1) \rceil$  node cells are energized, where  $N$  is the total number of node cells to be restored. Thus, the subgoals represent meaningful intermediate milestones toward the final restoration goal  $\mathbf{G}$ . If a trajectory terminates before reaching certain subgoal thresholds (i.e., fewer than  $\lceil sN/(q+1) \rceil$  nodes energized for some  $s \leq q$ ), we assign those remaining subgoals to the final timestep  $T$  for simplicity. The trajectory input to the Guidance Head,  $\tau^g$ , is a *subgoal-focused trajectory* sparsely sampled from the dataset of offline trajectories. It only includes the initial state  $\mathbf{s}_0$ , action  $\mathbf{a}_0$ , and physics-informed RTG  $\hat{R}_0$  and the subset of transitions associated with the subgoal steps. The RTG values  $\hat{R}$  is described in the previous section. Formally, the trajectory is defined as:

$$\tau^g = \{\mathbf{G}, \mathbf{s}_0, \mathbf{a}_0, \hat{R}_0, \tau_s^g\}, \quad (8)$$

where the subgoal-related sequence  $\tau_s^g$  is defined as:

$$\tau_s^g = \{\mathbf{s}_{t_{m_1}}, \mathbf{a}_{t_{m_1}}, \hat{R}_{t_{m_1}}, \dots, \mathbf{s}_{t_{m_q}}, \mathbf{a}_{t_{m_q}}, \hat{R}_{t_{m_q}}\}, \quad (9)$$

and  $\hat{R}$  is the physics-informed RTG defined in the previous section.

**Action Head Design:** Guided by the subgoal representations generated from the Guidance Head, the Action Head is designed to predict control actions at each step independently of individual-action-level RTG values. This decoupling enhances robustness by avoiding RTG overfitting and enables more flexible trajectory stitching.

Suppose we have a state-action-RTG trajectory of length  $K$ , denoted as:

$$\tau_{-K:(K-1)} = \{\mathbf{s}_0, \mathbf{a}_0, \hat{R}_0, \dots, \mathbf{s}_{K-1}, \mathbf{a}_{K-1}, \hat{R}_{K-1}\}.$$

To generate inputs for the Action Head, we augment this trajectory by appending the final goal  $\mathbf{G}$  and the subgoal offsets  $\mathbf{g}_i - \mathbf{s}_0$  to the beginning of the sequence. The full input sequence used for training is:

$$\begin{aligned} \tau'_{-K:(K-1);q} = & \{\mathbf{G}, \mathbf{s}_0, \mathbf{g}_1 - \mathbf{s}_0, \mathbf{g}_2 - \mathbf{s}_0, \dots, \mathbf{g}_q - \mathbf{s}_0, \mathbf{a}_0, \\ & \mathbf{s}_1, \mathbf{a}_1, \dots, \mathbf{s}_{K-1}, \mathbf{a}_{K-1}\}, \end{aligned} \quad (10)$$

where  $q$  is the predefined number of subgoals.

For generality and ease of notation, we also define a time-shifted sequence:

$$\begin{aligned} \tau'_{-K:t;q} = & \{\mathbf{G}, \mathbf{s}_{t-K+1}, \mathbf{g}_1 - \mathbf{s}_{t-K+1}, \mathbf{g}_2 - \mathbf{s}_{t-K+1}, \dots, \\ & \mathbf{g}_q - \mathbf{s}_{t-K+1}, \mathbf{a}_{t-K+1}, \\ & \mathbf{s}_{t-K+2}, \mathbf{a}_{t-K+2}, \dots, \mathbf{s}_t, \mathbf{a}_t\}, \end{aligned} \quad (11)$$

which represents the same sequence structure aligned with a later time index  $t$ . This representation  $\tau'_{-K:t;q}$  will be used in the subsequent sections describing the model training and testing pipelines.

## Operational Constraint-Aware Graph Reasoning Module

To enforce operational constraints during action selection, we introduce an operational constraint-aware graph reasoning module integrated into the action head. This module automates operational constraint enforcement by leveraging real-time topological and operational information from the distribution system. Specifically, we formulate an adjacency matrix  $\mathbf{A}_t$  to represent the system topology at time step  $t$ , capturing the connectivity between energized node cells and candidate nodes that may be energized. The previously defined state vector  $\mathbf{s}_t$ , which encodes dynamic attributes for each node cell, such as energization status and whether it has been activated at the current step, is treated as a node feature matrix in the graph module. Using  $\mathbf{A}_t$  and  $\mathbf{s}_t$ , we define an operational constraint-aware mask function  $f(\mathbf{s}_t, \mathbf{A}_t)$  that outputs a mask vector  $\mathbf{M}_t$ . This mask is applied to the raw logit vector  $\tilde{a}_{t+1}$  generated by the action head, effectively filtering out infeasible actions, such as those that would create network loops or revisit previously energized nodes. The masked logits are then passed through a softmax layer to produce the final confidence probability vector, which constitutes the action head's output  $a_{t+1}$ .

## Model Training and Inference Pipeline

**Model Training:** The overall model training procedure is described as follows:

- 1) We prepare a dataset  $\mathcal{D}$  consisting of offline trajectories for DSR operations on the targeted distribution system. These trajectories can be collected from domain experts or generated via simple offline random walks that apply sequences of control actions to the operable switches.

2) Once  $\mathcal{D}$  contains a sufficient number of offline trajectories, we preprocess each trajectory through two main steps. First, for each trajectory  $\tau \in \mathcal{D}$ , we extract a subgoal-focused trajectory  $\tau_s^g$  as described in the Guidance Head Design. Second, using the expected final goal  $\mathbf{G}$ , we construct a modified trajectory  $\tau_m$  as:

$$\tau_m = (U_1, U_2, \dots, U_T), \quad (12)$$

where each element is defined as

$$U_t = (\mathbf{G}, \mathbf{s}_t, \mathbf{a}_t, \hat{R}_t, \tau_s^g), \quad t = 1, 2, \dots, T.$$

The resulting modified trajectories are combined into a new training dataset  $\mathcal{D}' = \{\tau_m\}$ .

3) Minibatches  $B$  of length  $K$  are sampled from  $\mathcal{D}'$  and fed into the DH-PGDT framework to compute control actions for DSR. The model parameters are updated by minimizing the joint loss:

$$\mathcal{L} = \mathcal{L}_g + \mathcal{L}_a,$$

where  $\mathcal{L}_g$  is defined as the mean squared error between the predicted subgoal representations and the ground-truth subgoals, and  $\mathcal{L}_a$  is defined as the cross-entropy loss between the predicted and true control actions.

4) Repeat Step 3 for  $M$  training episodes.

The above procedure is summarized in Algorithm 1.

---

#### Algorithm 1: DH-PGDT Model Training

**Initialize:** Offline dataset  $\mathcal{D}$ , predefined goal  $\mathbf{G}$ , DH-PGDT model with parameters  $\theta$ , maximum number of episodes  $M$ , and minibatch size  $b$ .

- 1: Construct the processed dataset  $\mathcal{D}'$  from  $\mathcal{D}$  and  $\mathbf{G}$  using Eq. (12).
- 2: **for** episode  $m = 1$  to  $M$  **do**
- 3:   Sample a random minibatch  $B$  of  $b$  sequences of length  $K$  from  $\mathcal{D}'$ .
- 4:   Compute action and guidance predictions using DH-PGDT on batch  $B$ .
- 5:   Calculate the loss  $\mathcal{L} = \mathcal{L}_g + \mathcal{L}_a$ .
- 6:   Update model parameters  $\theta$ .
- 7: **end for**

---

**Model Inference** The inference process of our DH-PGDT model consists of two stages. First, given the initial state of the distribution system  $\mathbf{s}_0$ , the predefined final goal  $\mathbf{G}$ , and the initial physics-informed RTG  $\hat{R}_0$ , the Guidance Head of DH-PGDT generates a sequence of subgoal representations  $\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_q$ . For brevity, we denote this computation as  $\text{GH\_PGDT}(\cdot)$ , referring to the Guidance Head module of the DH-PGDT model. Next, these subgoals are fed into the same DH-PGDT model along with the system states, actions, and the final goal  $\mathbf{G}$  to generate the next control action. We denote this computation as  $\text{AH\_PGDT}(\cdot)$ , representing the Action Head module of DH-PGDT. The overall inference procedure is outlined in Algorithm 2.

---

#### Algorithm 2: Inference of the Trained DH-PGDT Model

**Initialize:** DH-PGDT model with trained parameters  $\theta$ , predefined goal  $\mathbf{G}$ , initial RTG  $\hat{R}_0$ , and initial state  $\mathbf{s}_0$ .

- 1: Set  $\mathbf{s}_1 \leftarrow \mathbf{s}_0, \tau_1 \leftarrow (\mathbf{s}_1, \hat{R}_0)$
- 2: Obtain  $\mathbf{g}_1, \dots, \mathbf{g}_q \leftarrow \text{GH\_PGDT}(\tau_1)$
- 3: Set  $\tau_2 \leftarrow \{\mathbf{G}, \mathbf{s}_1, \mathbf{g}_1 - \mathbf{s}_1, \dots, \mathbf{g}_q - \mathbf{s}_1\}$
- 4: **for** time step  $t = 1$  to  $T - 1$  **do**
- 5:   Compute action:  $\mathbf{a}_t \leftarrow \text{AH\_PGDT}(\tau_2)$
- 6:   Apply action  $\mathbf{a}_t$  and observe next state  $\mathbf{s}_{t+1}$
- 7:   Append  $(\mathbf{a}_t, \mathbf{s}_{t+1})$  to  $\tau_2$
- 8:   Keep the last  $K$  time steps of  $\tau_2$  (i.e.,  $\tau_2 \leftarrow \tau'_{-K:t+1;q}$ ).
- 9: **end for**

---

## Case Studies

In this section, we evaluate the performance of our proposed DH-PGDT framework using two large-scale distribution system testbeds simulated in the Open Distribution Simulator Software (OpenDSS): the modified IEEE 123-node test feeder (IEEEStd 2014), and a modified Iowa 240-node test feeder (Bu et al. 2019), respectively. To evaluate the performance of our DH-PGDT framework, we compare it with a DT framework, PIDT, developed in our previous work (Zhao et al. 2025), and two benchmark DRLs, the Proximal Policy Optimization (PPO) algorithm (Schulman et al. 2017), and the Advantage Actor-Critic (A2C) algorithm (Mnih et al. 2016), in the DSR tasks.

### Modified IEEE 123-Node Test Feeder

The system topology is shown in Fig. 2. As shown in Fig. 2, in the modified 123-node test feeder, there are five energy sources in the system, two of which are substations located in Buses 150 and 350, and the three other sources are DGs located in Buses 95, 250, and 450. To evaluate the effectiveness of our DH-PGDT framework under complex operational scenarios, we introduce an additional operation constraint related to the energization capability of the DG at Bus 250. Specifically, this DG is limited to fully energizing only Node Cells 2 and 3. If it attempts to energize any other node cells beyond these, the DG automatically shuts down.

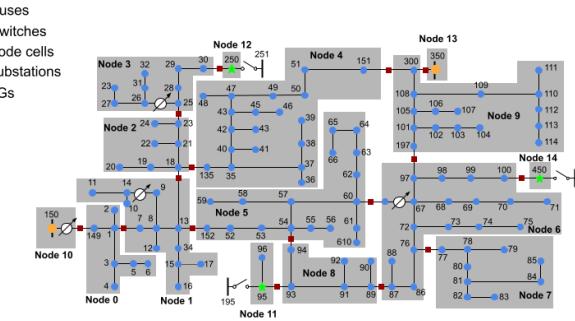


Figure 2: The physical topology of the modified IEEE 123-node test feeder.

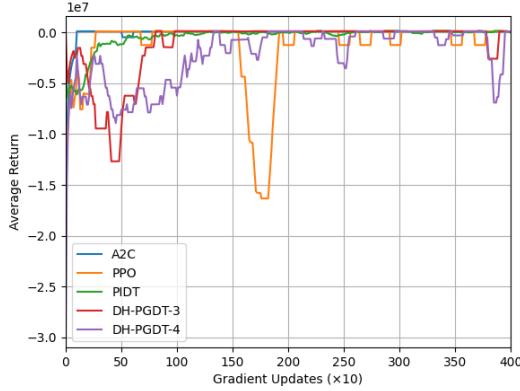


Figure 3: Learning curve for the average return of the first 4000 gradient updates using our proposed DH-PGDT (DH-PGDT- $q$  denotes our DH-PGDT frameworks with hyperparameter  $q$ , representing the number of subgoals), PIDT, PPO, and A2C methods for the DSR operations in the modified 123-node test feeder. The curve is updated per ten gradient updates, and is smoothed using a moving average per 10 data points.

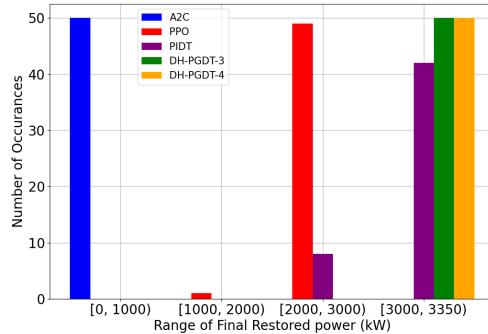


Figure 4: Distribution of power restoration levels in the 50 independent trials using the four methods, respectively, in the modified 123-node test feeder.

The actual total power demand of the system is about 3350 kW if all the loads are energized properly, which is considered as the objective of the DSR operation. The learning curves of the first 6000 gradient updates using the DH-PGDT, PIDT, A2C, and PPO methods are shown in Fig. 3. As shown in Fig. 3, for the modified IEEE 123-node test feeder with a complex and large-scale topology, our method has convergence speed comparable to the PIDT and two DRL methods. Fig. 3 also shows that the convergence speed of our method is influenced by the hyperparameter  $q$ , the predefined number of subgoals. As  $q$  increases, the convergence becomes slower. Further evaluation results of these three methods across 50 independent trials in the inference stage are presented in Table 2. As shown in Table 2, our proposed DH-PGDT achieves optimal solutions in all 50 in-

dependent trials, outperforming PIDT, PPO and A2C.

Method	Average Return	Std. Return	APR (kW)	SDPR (kW)	# Opt. Sols.
A2C	8696.511	0.000	486.619	0.000	0
PPO	-163818.290	1878998.794	2097.300	83.751	0
PIDT	158850.797	9911.682	3228.657	206.415	32
DH-PGDT-3	164978.536	0.000	3349.696	0.000	50
DH-PGDT-4	164978.536	0.000	3349.696	0.000	50

Table 2: Further performance comparison between our proposed DH-PGDT, PIDT, A2C, and PPO methods for the DSR operation in the modified IEEE 123-node test feeder. Metrics include average return, standard deviation of returns (Std. Return), average power restoration (APR), standard deviation of power restorations (SDPR), and number of expected optimal solutions (# Opt. Sols.).

Furthermore, Fig. 4 provides deeper insights into the simulation results shown in Table 2 by illustrating the distribution of power restoration levels across 50 independent trials using the four methods. Our method consistently achieves optimal power restoration in all 50 trials. In contrast, the A2C method predominantly generates results in the range 1000 kW to 2000 kW. The output of the PPO method mainly falls within the range of 2000 kW to 3000 kW, occasionally dropping to the range of 1000 kW to 2000 kW. The output of the PIDT method mainly falls within the range of 3000 kW to 3350 kW, occasionally dropping to the range of 2000 kW to 3000 kW, which is better than both PPO and A2C, but worse than our currently proposed DH-PGDT method. These findings highlight the superior performance of our method compared to the PIDT method in our previous work and two emerging DRL methods.

## Modified Iowa 240-Node Test Feeder

The system topology is shown in Fig. 5. As shown in Fig. 5, the modified Iowa 240-node test feeder includes five energy sources in the system: one substation located at the source bus and four DGs located at Buses 1009, 3030, 2027, and 3082. To evaluate the zero-shot generalization capability of our DH-PGDT framework, we configure the system with dynamic loads whose active power varies over the time horizon. The model is trained solely on a static-load condition but tested directly on unseen dynamic load conditions without any fine-tuning.

Unlike the 123-node test case, the dynamic nature of the 240-node test case makes it impossible to determine a single total power value without fixing the time horizon. Therefore, we set the time horizon to  $T = 20$ , which results in a final total restored power of approximately 1565 kW.

The learning curves of the first 5000 gradient updates using our DH-PGDT method and two benchmark DRL methods are shown in Fig. 6. Additional evaluation results across 50 independent inference trials are presented in Table 3. As shown in Fig. 6 and Table 3, although our method

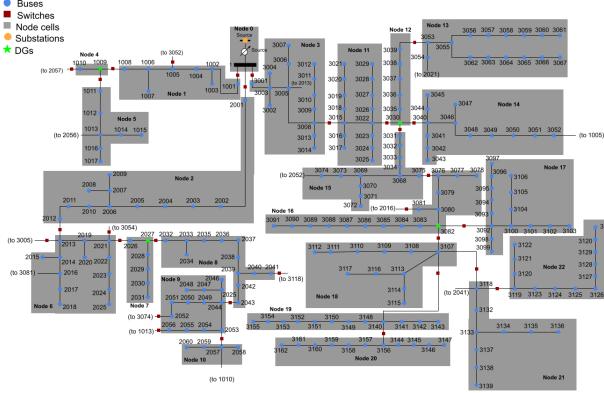


Figure 5: The physical topology of the modified Iowa 240-node test feeder.

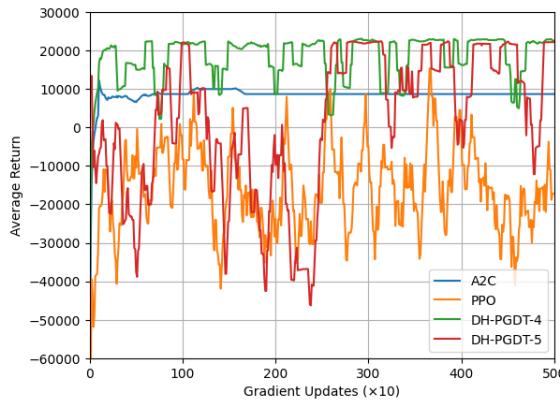


Figure 6: Learning curve for the average return of the first 5000 gradient updates using our DH-PGDT method and two benchmark DRL methods, PPO and A2C, for the DSR operations in the modified Iowa 240-node test feeder. The curve is updated per ten gradient updates and is smoothed using moving average per 10 data points.

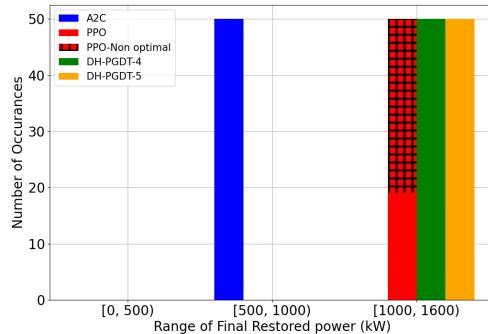


Figure 7: Distribution of power restoration levels in the 50 independent trials using the three methods, respectively, in the modified Iowa 240-bus test feeder.

Method	Average Return	Std. Return	APR (kW)	SDPR (kW)	# Opt. Sols.
A2C	47414.036	1675.582	1077.232	0.000	0
PPO	-18523.242	33103.943	1565.522	0.007	19
DH-PGDT-4	23203.563	0.000	1565.526	0.000	50
DH-PGDT-5	23330.984	0.000	1565.526	0.000	50

Table 3: Further performance comparison between our DH-PGDT method and benchmark DRL methods for the DSR operation in the modified Iowa 240-Node test feeder.

achieves a convergence rate comparable to the two DRL methods, it consistently reaches optimal solutions in all 50 trials, which outperforms both the DRL methods. Furthermore, Fig. 7 provides deeper insights into the simulation results presented in Table 3 by illustrating the distribution of power restoration levels across 50 independent trials using the three methods. It shows that our method achieves optimal power restoration in all 50 trials. In contrast, the A2C-based method predominantly produces solutions within the 500–1000 kW range. While PPO mainly generates solutions in the 1000–1600 kW range, it is worth noting that 31 of these solutions violate the ramp rate constraints, rendering them non-optimal.

Moreover, to further evaluate the zero-shot generalization of our DH-PGDT framework, we test it in two additional scenarios involving unexpected switch contingencies during inference. In the first scenario, a switch contingency causes the power source at the source bus in Fig. 5 to become isolated from the system. For clarity, we use the node-cell-based graph representation to illustrate the contingency as shown in Fig. 8 (a), where Node Cell 0 is isolated from the rest of the graph. We would like to mention that the node in the graph corresponds to the node cells in the system topology in Fig. 5. The solution generated by our method to adapt to this zero-shot scenario is shown in Fig. 8 (b), achieving a final restored power of 1565.526 kW and a final return of 22743.817. In the second scenario, a switch contingency causes the DG at Bus 3030 in Fig. 5 to become isolated from the system. The contingency is illustrated using the node-cell-based graph representation as shown in Fig. 9 (a), where Node Cell 12 is isolated from the rest of the graph. The solution generated by our method to adapt to this zero-shot scenario is shown in Fig. 9 (b), achieving a final restored power of 1565.526 kW and a final return of 22308.007.

## Conclusions

In this paper, we proposed DH-PGDT, an innovative framework for complex DSR tasks under uncertainty. By integrating a dual-head physics-informed causal transformer architecture with an operational constraint-aware graph reasoning module, the proposed DH-PGDT effectively enables scalable and reliable DSR even in zero-shot and few-shot scenarios. Extensive evaluations with different large-scale benchmark distribution systems demonstrate that DH-PGDT

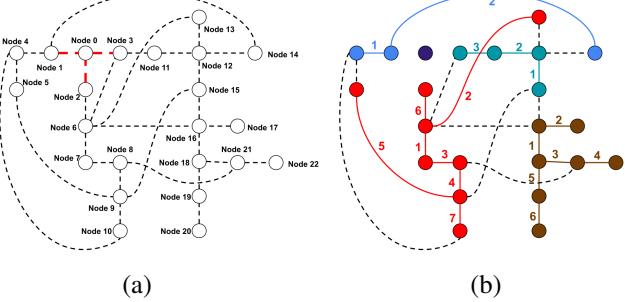


Figure 8: (a): Graph representation of Scenario 1, with red lines indicating switch contingencies; (b): Solution generated from DH-PGDT for Scenario 1, with numbers on the sides representing the step when the switch is closed.

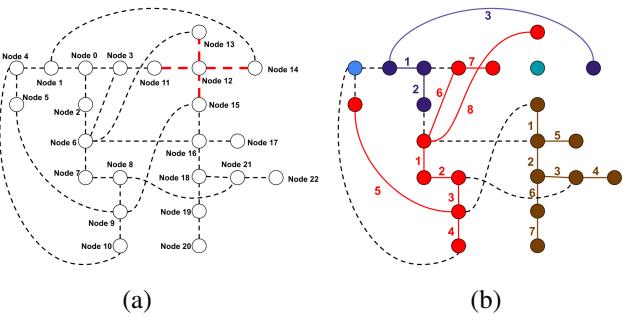


Figure 9: (a): Graph representation of Scenario 2, with red lines indicating the switch contingencies; (b): Solution generated from DH-PGDT for Scenario 2, with numbers on the sides representing the step when the switch is closed.

significantly improves restoration efficiency, constraint satisfaction, and generalization to unforeseen contingencies. While this framework is specifically designed for DSR in power systems, its underlying computing model is generalizable to a wide range of sequential decision-making tasks in engineering domains with complex physical and operational constraints. In ongoing work, we are evaluating the performance of DH-PGDT in even larger-scale DSR scenarios with diverse contingencies and extending the framework to other power system operational tasks.

## References

Bu, F.; Yuan, Y.; Wang, Z.; Dehghanpour, K.; and Kimber, A. 2019. A Time-Series Distribution Test System Based on Real Utility Data. *2019 North American Power Symposium (NAPS)*, 1–6.

Chen, B.; Ye, Z.; Chen, C.; and Wang, J. 2019. Toward a MILP Modeling Framework for Distribution System Restoration. *IEEE Transactions on Power Systems*, 34(3): 1749–1760.

Chen, C.; Wang, J.; and Ton, D. T. 2017. Modernizing Distribution System Restoration to Achieve Grid Resiliency Against Extreme Weather Events: An Integrated Solution. *Proceedings of the IEEE*, 105: 1267–1288.

Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision Transformer: Reinforcement Learning via Sequence Modeling. *arXiv preprint arXiv:2106.01345v2*.

Fu, J.; Long, Y.; Chen, K.; Wei, W.; and Dou, Q. 2024. Multi-objective Cross-task Learning via Goal-conditioned GPT-based Decision Transformers for Surgical Robot Task Automation. *arXiv preprint arXiv:2405.18757v1*.

Gao, Y.; Wang, W.; Shi, J.; and Yu, N. 2020. Batch-Constrained Reinforcement Learning for Dynamic Distribution Network Reconfiguration. *IEEE Transactions on Smart Grid*, 11(6): 5357–5369.

Hsu, H.-L.; Bozkurt, A. K.; Dong, J.; Gao, Q.; Tarokh, V.; and Pajic, M. 2024. Steering Decision Transformers via Temporal Difference Learning. *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

IEEEStd. 2014. IEEE Recommended Practice for Electric Power Distribution System Analysis. *IEEE Std 1729-2014*, 1–20.

Igder, M. A.; Liang, X.; and Mitolo, M. 2022. Service Restoration Through Microgrid Formation in Distribution Networks: A Review. *IEEE Access*, 10: 46618–46632.

Kim, S.; Noh, S.; and Jang, I. 2023. Numerical Goal-based Transformers for Practical Conditions. In *NeurIPS 2023 Workshop GCRL*.

Liang, X.; Saaklayen, M. A.; Igder, M. A.; Shawon, S. M. R. H.; Faried, S. O.; and Janbakhsh, M. 2022. Planning and Service Restoration Through Microgrid Formation and Soft Open Points for Distribution Network Modernization: A Review. *IEEE Transactions on Industry Applications*, 58(2): 1843–1857.

Ma, Y.; Xiao, C.; Liang, H.; and Hao, J. 2023. Rethinking Decision Transformer via Hierarchical Reinforcement Learning. *arXiv preprint arXiv:2311.00267v1*.

Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T. P.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *International Conference on Machine Learning*.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *ArXiv*, abs/1707.06347.

Wu, J.; Fang, B.; Fang, J.; Chen, X.; and Tse, C. K. 2019. Sequential topology recovery of complex power systems based on reinforcement learning. *Physica A: Statistical Mechanics and its Applications*, 535: 122487.

Xu, M.; Shen, Y.; Zhang, S.; Lu, Y.; Zhao, D.; Tenenbaum, J. B.; and Gan, C. 2022. Prompting Decision Transformer for Few-Shot Policy Generalization. *Proceedings of the 39th International Conference on Machine Learning*.

Yao, T.; Chen, X.; Yao, Y.; Huang, W.; and Chen, Z. 2025. Offline prompt reinforcement learning method based on feature extraction. *PeerJ Computer Science*.

Zhao, H.; Wei-Kocsis, J.; Akhijahani, A. H.; and Butler-Purry, K. L. 2025. Advancing Distribution System Restoration via an Innovative Physics-Informed Decision Transformer. In *4th Annual AAAI Workshop on AI to Accelerate Science and Engineering (AI2ASE)*.