

GATES: Cost-aware Dynamic Workflow Scheduling via Graph Attention Networks and Evolution Strategy

Ya Shen*, Gang Chen, Hui Ma and Mengjie Zhang

Centre for Data Science and Artificial Intelligence & School of Engineering and Computer Science,
Victoria University of Wellington, New Zealand

{ya.shen, aaron.chen, hui.ma, mengjie.zhang}@ecs.vuw.ac.nz

Abstract

Cost-aware Dynamic Workflow Scheduling (CADWS) is a key challenge in cloud computing, focusing on devising an effective scheduling policy to efficiently schedule dynamically arriving workflow tasks, represented as Directed Acyclic Graphs (DAG), to suitable virtual machines (VMs). Deep reinforcement learning (DRL) has been widely employed for automated scheduling policy design. However, the performance of DRL is heavily influenced by the design of the problem-tailored policy network and is highly sensitive to hyperparameters and the design of reward feedback. Considering the above-mentioned issues, this study proposes a novel DRL method combining Graph Attention Networks-based policy network and Evolution Strategy, referred to as **GATES**. The contributions of GATES are summarized as follows: (1) GATES can capture the impact of current task scheduling on subsequent tasks by learning the topological relationships between tasks in a DAG. (2) GATES can assess the importance of each VM to the ready task, enabling it to adapt to dynamically changing VM resources. (3) Utilizing Evolution Strategy’s robustness, exploratory nature, and tolerance for delayed rewards, GATES achieves stable policy learning in CADWS. Extensive experimental results demonstrate the superiority of the proposed GATES in CADWS, outperforming several state-of-the-art algorithms. The source code is available at: <https://github.com/YaShen998/GATES>.

1 Introduction

Real-world problems are often highly complex, requiring applications tailored to solve them to execute efficiently and deliver prompt results to meet user demands [Pallathadka *et al.*, 2022]. For example, if a weather forecasting application fails to run in real time and provide timely, accurate predictions, it may interrupt the regular management of vital industrial operations [Price *et al.*, 2025]. Cloud computing enables the

deployment of such complex applications structured as workflows consisting of multiple tasks. It facilitates the scheduling of elastic computational resources to ensure efficient and timely task execution [Jadeja and Modi, 2012]. This process is known as Workflow Scheduling (WS).

A workflow can be represented as a *Directed Acyclic Graph* (DAG), with vertices representing tasks and edges representing task dependencies [Shen *et al.*, 2024]. Workflows with different patterns have different DAG structures. Additionally, each workflow includes a *Service Level Agreement* (SLA) that defines the maximum allowed completion time for the workflow [Qazi *et al.*, 2024]. This study focuses on the Cost-aware Dynamic Workflow Scheduling (CADWS) problem [Shen *et al.*, 2024], which aims to develop an efficient scheduling policy that minimizes VM rental costs and SLA violation penalties, while handling the complexity caused by the dynamic arrival of workflows with diverse DAG structures [Huang *et al.*, 2022].

Compared to traditional scheduling problems such as job shop scheduling and vehicle routing, CADWS exhibits complex dynamic characteristics, including unpredictable workflow arrivals, varying workflow patterns, and rapidly fluctuating computational resources (e.g., the number of available VMs). Moreover, unlike job shop and vehicle routing problems, workflows in CADWS have highly diverse graph topologies across different patterns, while the others maintain similar structures regardless of scale. This study focuses on using deep reinforcement learning (DRL) techniques to automatically learn scheduling policies for solving CADWS, with an emphasis on designing effective policy networks tailored to the problem characteristics of CADWS.

The key challenges of this study are: (1) how to design an effective policy network to properly control the influence of every scheduling decision on future task execution in a dynamic environment with continuously arriving workflows of highly diverse patterns; (2) how can the designed policy network support a time-varying collection of VMs; and (3) how to reliably train this policy network to ensure robust performance across varying system settings. To address the aforementioned challenges, this paper proposes a novel DRL method that integrates Graph Attention Networks (GAT) and Evolution Strategy (ES), referred to as **GATES**.

Specifically, the main contributions of GATES are as follows: (1) GATES employs GAT to efficiently and scalably

*Corresponding author

process the local topological information of each ready task. Leveraging GAT’s attention mechanism, GATES accurately learns the global attention information between the ready task and other tasks, enabling a holistic view of current scheduling impacts on future tasks by aggregating all task embeddings. (2) Dynamic heterogeneous graphs with analogous structures but varying scales are constructed to represent the relationships between the ready task and all available VMs at each time step. Regardless of dynamic changes in VM numbers, GAT can compute attention coefficients from these graphs to identify the most suitable VM at each time step. (3) In CADWS, the reward function typically includes VM rental fees and SLA violation penalties. However, since SLA penalties can only be accurately determined after workflow execution, it’s difficult to design a reward function that quantifies SLA violations at each time step. To overcome this challenge, ES is utilized to train the policy network. This approach eliminates the need for complex reward design and hyperparameter tuning, while also preventing performance issues caused by poorly designed rewards.

2 Related Work

Currently, most methods for solving scheduling problems focus on designing effective scheduling policies [Branke *et al.*, 2015; Khalil *et al.*, 2017; Dong *et al.*, 2021; Hasanzade-Zonuzi *et al.*, 2021]. Some manually designed policies have been developed to solve CADWS [Pham and Fahringer, 2020; Wu *et al.*, 2017; Faragardi *et al.*, 2019], as these methods are easy to understand and implement. However, these methods heavily rely on expert experience and substantial domain knowledge, limiting their adaptability to dynamic environments [Geiger *et al.*, 2006], as they are developed based on historical workflows and may fail to perform effectively in rapidly changing scenarios or with unknown workflows [Shen *et al.*, 2024].

Several DRL studies [Kayhan and Yildiz, 2023; Cunha *et al.*, 2020; Ou *et al.*, 2023; Iklassov *et al.*, 2023] have shown that using neural network-based policy networks to automatically learn scheduling policies is a feasible and promising approach to solving scheduling problems. For instance, Ou *et al.* [2023] utilized DRL to solve the satellite range scheduling problem, while Song *et al.* [2022] and Zhang *et al.* [2020] employed DRL to learn effective policies for addressing the job-shop scheduling problem. However, the problems addressed by these works are either static or do not have a graph topological structure, such as the DAG of CADWS. Moreover, although Shen *et al.* [2024] and Huang *et al.* [2022] applied DRL to dynamic WS, their policy networks ignored the DAG structures of workflows, failing to leverage task dependencies for effective scheduling. Thus, the architecture of policy networks designed in these approaches did not take dynamic information or graph topological information into account to extract essential knowledge for learning scheduling policies, making them unsuitable for CADWS.

Recently, some studies [Sun and Yang, 2023; Zhang *et al.*, 2024] proposed graph-based policy networks to extract state features for effective policy learning by leveraging the graph topological properties inherent in problems, such as travel-

ing salesman, maximal independent set, and job shop problems. Sun and Yang [2023] proposed a graph-based diffusion model to exploit the graph topological information in the traveling salesman problem and the maximal independent set problem. Similarly, Zhang *et al.* [2024] utilized Graph Neural Networks (GNNs) to extract state features for effective policy learning by leveraging the graph topological properties of the job shop scheduling problems. While the above studies considered the graph topological information, they assume static structure and fixed scheduling resources. In contrast, CADWS accounts not only for the dynamic workflow arrivals and patterns but also for the variability in scheduling resources over time [Shen *et al.*, 2024]. Specifically, the pool of active VMs can change across different system states and time steps.

Unlike existing DRL approaches, GATES is specifically designed in this study to address the dynamic nature of scheduling problems. It handles the unpredictable arrival and varying patterns of workflows, as well as fluctuations in available scheduling resources, achieving stable policy learning in complex and dynamic environments.

3 Problem Definition

3.1 Preliminary of CADWS

The CADWS problem involves a set of workflows \mathcal{W} arriving dynamically over time and a dynamically changing collection of heterogeneous VMs \mathcal{M} for task execution, meaning that $|\mathcal{M}|$ is not fixed. Each workflow $W_i \in \mathcal{W}$ is modeled as a DAG $(\mathcal{O}_{W_i}, \mathcal{C}_{W_i})$, where the node set $\mathcal{O}_{W_i} = \{O_{i1}, O_{i2}, \dots, O_{in}\}$ represents tasks, and the edge set $\mathcal{C}_{W_i} = \{(O_{ij}, O_{ik}) \mid O_{ij}, O_{ik} \in \mathcal{O}_{W_i}\}$ denotes precedence constraints (see Appendix A). An edge (O_{ij}, O_{ik}) implies that task O_{ij} must complete before task O_{ik} can begin execution. O_{ij} becomes a ready task O^* waiting for current scheduling when there are no predecessor tasks for it [Zhu *et al.*, 2016].

The complete set of tasks across all workflows is denoted as $\mathcal{O} = \mathcal{O}_{W_1} \cup \mathcal{O}_{W_2} \cup \dots \cup \mathcal{O}_{W_{|\mathcal{W}|}}$. Each task O_{ij} is associated with a workload $WL_{ij} \in \mathbb{R}^+$ and can be executed on any VM $M_k \in \mathcal{M}$. The execution time of O_{ij} on VM M_k is determined as $ET(O_{ij})^k = \frac{WL_{ij}}{PS_k}$, where PS_k denotes the processing speed of machine M_k . Hence, $CT(O_{ij}) = ST(O_{ij}) + ET(O_{ij})^k$ gives the completion time of O_{ij} on M_k , where $ST(O_{ij})$ is the start time of O_{ij} . O_{ij} is the start task of W_i , if $ST(O_{ij}) = AT(W_i)$, where $AT(W_i)$ is the arrival time of workflow W_i . The finish time of workflow W_i is defined as $FT(W_i) = \max_j \{CT(O_{ij})\}$. The total process time of workflow W_i is computed as $PT(W_i) = FT(W_i) - AT(W_i)$.

3.2 Optimization Objective of CADWS

The VM rental fees cover the total cost of all leased VMs used to execute all workflows dynamically arrived during a given time period T . This period, starting at time t_s and ending at t_e , depends on the scheduling policy π adopted. For each VM $v \in \mathcal{M}$, its usage period during T is denoted as: $UP(v, \pi, T) = [t_s(v, \pi, T), t_e(v, \pi, T)]$, where $t_s(v, \pi, T)$ represents the start time of using v , and $t_e(v, \pi, T)$ indicates the time when v is no longer in use, satisfying

$t_e(v, \pi, T) \leq t_e$. The total rental cost for all VMs under the scheduling policy π over T is computed as: $VM_{Fee}(\pi, T) = \sum_{v \in Set(\pi, T)} \left(Price(v) \times \left\lceil \frac{t_e(v, \pi, T) - t_s(v, \pi, T)}{3600} \right\rceil \right)$, where $Set(\pi, T)$ is the set of VM instances leased for task execution during T . The ceiling function ensures that the rental time during an hour will be charged as an hour, following the common practice in the cloud computing industry [Shen *et al.*, 2024; Huang *et al.*, 2022].

The SLA penalty for a workflow w_i under a scheduling policy π is defined as: $SLAPenalty(W_i, \pi) = \beta \times \max\{0, [WCT(W_i, \pi) - DL(W_i)]\}$, where $WCT(W_i, \pi)$ refers to the $PT(W_i)$ by following policy π , and $DL(W_i)$ denotes the SLA-specified deadline for W_i . The penalty factor β quantifies the severity of SLA violations. Following many existing studies [Huang *et al.*, 2022; Shen *et al.*, 2024], the deadline $DL(W_i)$ is calculated as: $DL(W_i) = AT(W_i) + \gamma \times MinMakespan(W_i)$, where $AT(W_i)$ is the workflow's arrival time, $MinMakespan(W_i)$ is the theoretical minimum time required to process W_i using the fastest VMs without delays, and γ is a relaxation coefficient. A higher γ value allows for a more lenient deadline.

The objective of CADWS is to find an optimal scheduling policy π that minimizes the total cost, including both the VM rental fees and SLA penalties, as formulated below:

$$\arg \min_{\pi} TotalCost(\pi) = \arg \min_{\pi} \{ VM_{Fee}(\pi, T) + \sum_{W_i \in \mathcal{W}} SLAPenalty(W_i, \pi) \} \quad (1)$$

4 Methodology

This section elaborates on the proposed GATES for solving CADWS in detail. CADWS is a sequential decision-making process that can be formulated as a Markov Decision Process (MDP) [Huang *et al.*, 2022; Shen *et al.*, 2024; Yang *et al.*, 2025], where scheduling decisions are made iteratively. At each step, GATES selects a compatible VM to execute a specific ready workflow task, continuing this process until all workflow tasks are scheduled. The workflow of GATES for solving CADWS is illustrated in Figure 1.

As shown in Figure 1, the current state of the scheduling system is captured by a Dynamic Heterogeneous Graph (DHG) representation (see Figure 2). Subsequently, a heterogeneous Graph Attention Network (HGAT) employing a two-stage embedding mechanism is applied to process this graph. This step extracts feature embeddings of both tasks and VMs (see Figure 3). These embeddings are then utilized by a decision-making network to compute the action probability distribution, from which a VM is sampled or greedily selected (as detailed in Figure 4). The next subsection introduces the MDP formulation of the CADWS process.

4.1 MDP Formulation of CADWS

System State: Each system state $s_t \in \mathcal{S}$ represents a snapshot of the current state of the entire CADWS problem at decision step t , including the following elements: (1) the current ready task O^* at t ; (2) details of all active workflows, encompassing pattern information on their constituent tasks and task

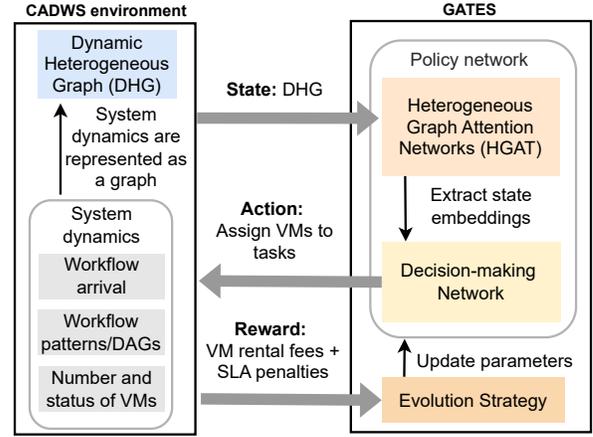


Figure 1: The overall framework of CADWS.

dependencies (modeled as DAG), and workflow arrival times; (3) the numbers and processing status of all available VMs at t . By designing an effective policy network that accurately captures the three aforementioned factors, the system state of CADWS at any time step can be precisely represented, addressing the first two challenges mentioned in Section 1. To achieve this, we define the system state of CADWS at any time step as a DHG (as shown in Figure 2) and propose the HGAT employing GAT with a two-stage embedding mechanism to extract useful state embeddings for ready tasks (details in Figure 3), enabling the learning of effective scheduling policies and making informed decisions.

Actions: At the decision step t , each action $a_t \in \mathcal{A}$ assigns the ready task O^* to the waiting queue of an eligible VM M ($M \in \mathcal{M}$). The size of the action space is $|\mathcal{A}| = |\mathcal{M}|$.

Transition: Upon taking an action a_t , the system transitions from state s_t to the next state s_{t+1} , where the focused ready task at decision step t is assigned to a specific VM for execution (Refer to Figure 1).

Rewards: The reward r_t provides a scalar feedback signal based on the action a_t taken at system state s_t . To achieve the objective of minimizing the total cost as formulated in Eq. (1), the reward is defined as $r_t = -(VM_{Fee} + SLAPenalty)$.

Policy: The policy $\pi(a_t | s_t)$ of GATES defines a probability distribution over the action set \mathcal{A} for each state s_t , modeled as a parametric neural network.

Learning Objective: The goal of the scheduling agent is to learn an optimal policy that maximizes the total rewards after the completion of the CADWS problem.

4.2 Dynamic Heterogeneous Graph (DHG)

In CADWS, we leverage the topological information of the workflows represented as DAGs to learn embeddings for ready tasks. These embeddings not only capture the features of the current ready task but also learn the global features by considering the relationships with other tasks. This DAG-based global feature learning reveals the influence of predecessor tasks on the ready task and the potential impact of the ready task on subsequent tasks.

Additionally, we incorporate dynamically changing VMs into the DAG to construct a DHG, as shown in Figure 2. Since

the number and types of VMs vary dynamically at different decision steps, the VM nodes and edges in the DHG are inherently dynamic. The DHG provides a global perspective to describe the influence of all tasks and currently available VMs on the scheduling of the ready task. Moreover, the DHG is constructed specifically for each ready task: whenever a task becomes ready and awaits scheduling, a DHG is generated to capture the essential system’s dynamics for executing the ready task.

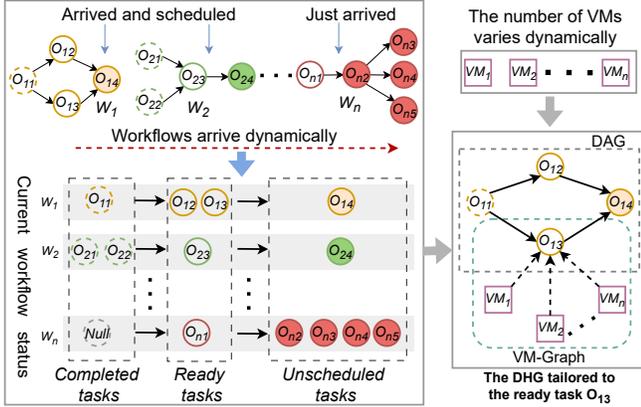


Figure 2: The proposed Dynamic Heterogeneous Graph (DHG).

4.3 Proposed Policy Network

To support scheduling policy learning in GATES (the Agent), it is crucial to extract state embedding from the constructed DHG. GNNs are well-suited for learning embeddings from graph data. Among them, GAT outperforms Graph Convolutional Networks (GCN) and Graph Isomorphism Networks (GIN) in learning node embeddings from heterogeneous graphs. This advantage arises from GAT’s use of attention mechanisms to assign dynamic weights to each neighboring node, enabling it to adapt to heterogeneous node features and diverse neighborhood relationships [Veličković *et al.*, 2017]. In contrast, GCN and GIN rely on fixed weights or mean aggregation, limiting their ability to model feature differences in heterogeneous graphs [Thekumparampil *et al.*, 2018]. Moreover, GAT does not require strict alignment of node features and naturally handles various relationships in heterogeneous graphs through its multi-head attention mechanism, offering greater flexibility and scalability [Wang *et al.*, 2019].

Heterogeneous Graph Attention Networks (HGAT)

We propose an HGAT to extract the state embedding from the DHG, as shown in Figure 3. The embedding learning process of HGAT consists of two stages: (1) Use GAT to learn feature embedding for each task from the DAG (Stage 1 of Figure 3), in which the **Ready-task embedding 1** refers to the feature embedding of the ready task. Meanwhile, the feature embeddings of all tasks are aggregated using mean pooling to obtain the **All tasks embedding**; (2) Another GAT is employed to learn the **Ready-task embedding 2** from the VM-graph in the DHG, as well as the **VMs embeddings**, enabling GATES to

assess the relative importance of each VM to the ready task. Finally, these learned embeddings are concatenated to form the state embedding in GATES that comprehensively captures critical dynamics in CADWS, including the dynamic arrival of workflows, the changing DAG structures, and the dynamic variation in the number and status of VMs.

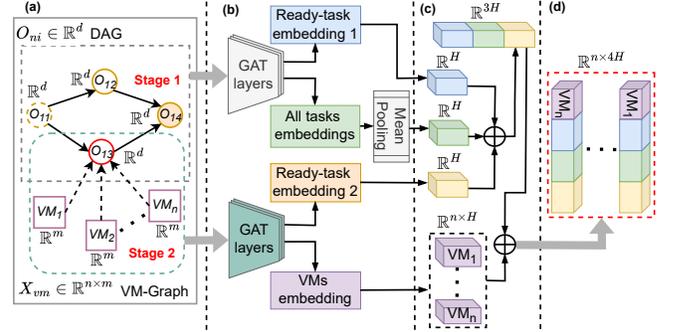


Figure 3: The proposed HGAT: (a) a DHG tailored to the ready task, (b) the embedding learning process, (c) feature concatenation, (d) output the learned state embedding.

The Decision-making Networks

We design the decision-making networks, as shown in Figure 4, to calculate the VM selection probabilities (i.e., the action distribution). Although the DHG of Figure 2 can capture the significance of dynamically changing VMs, an individual VM cannot reference the status information of others to determine its scheduling decisions. To address this, we employ a self-attention mechanism to learn global information among VMs, enabling each VM’s feature embedding to account for the scheduling status of other VMs, thereby identifying the optimal VM to handle ready tasks. The computation process of the decision-making network consists of two stages: (1) the feature matrix of all VMs is first fed into a Transformer [Vaswani, 2017], where a self-attention mechanism captures the global information among VMs; (2) each VM’s embedding is then concatenated with the state embedding to jointly determine the corresponding action score. The softmax function then computes the probability of selecting each VM based on these action scores. The detailed embedding learning process of Figure 3 and Figure 4 is described further in Appendix B.

4.4 Raw Features of CADWS

The scheduling policy learned by GATES for CADWS relies on the state embedding derived from both the ready task (denoted as O^*) and all VMs at the current time step. The state embedding is constructed using three categories of raw features: **Workflow-related**, **Task-related**, and **VM-related** raw features. For any task ready for execution, these features are extracted from the current state of workflows, tasks, and VMs to construct the input to GATES, as summarized below.

Workflow-related raw features:

- *Forthcoming successors*: the count of tasks dependent on O^* .

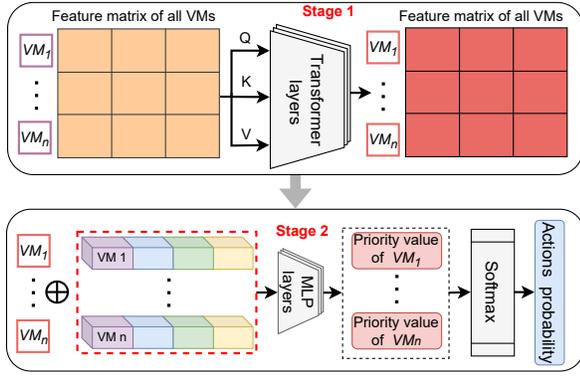


Figure 4: The decision-making networks.

- **Completion ratio:** the ratio of completed tasks (C_{tasks}) to total tasks (T_{tasks}) in \mathcal{W}_i , calculated as C_{tasks}/T_{tasks} .
- **Estimated arrival rate:** an estimation of future workflow arrival rates based on the current execution status of \mathcal{W}_i .

Task-related raw features:

- **Predecessor:** the number of predecessor tasks of each task in its corresponding DAG.
- **Successor:** the number of successor tasks of each task in its corresponding DAG.
- **Processing time:** the average duration required to execute a task on all types of VMs.
- **Task size:** the workload of each task in DAG.
- **Task deadline:** the task deadline is calculated using the method in [Wu *et al.*, 2017]. The deadline is assigned to each task, subject to the task size. A large task will be assigned a more relaxed task deadline.
- **Task status:** the task status is represented by an integer: 0 for unscheduled, 1 for scheduled, and 2 for ready to be scheduled.

VM-related raw features:

- **Task deadline feasibility:** an indicator that shows whether a VM can meet the deadline of O^* , assigned based on workflow size.
- **Incurred cost:** the sum of the VM rental fee and SLA violation penalties for using the VM to execute O^* .
- **Remaining rental time:** the remaining rental time on a VM after executing O^* .
- **Fittest VM:** an indicator that shows whether a VM has the lowest incurred cost among all candidates while meeting the task deadline.

4.5 Training Process via Evolution Strategy

The performance of gradient-based DRLs is often influenced by hyperparameter choices and the design of reward functions, which can result in unstable or suboptimal learning outcomes [Salimans *et al.*, 2017]. Evolution Strategy (ES) has been proposed as a robust alternative to mitigate these issues

and achieve stable performance [Khadka and Tumer, 2018; Salimans *et al.*, 2017]. In this study, we employ ES to train the policy network denoted as π_{GATES} . The pseudo-code of the training algorithm is presented in Algorithm 1. The training process is described below:

1. **Policy Sampling:** Let $\hat{\theta}$ represent the current parameter of the policy π_{GATES} . In each generation of ES, a population of N individuals is sampled. Specifically, for each individual θ_i , where $i = 1, 2, \dots, N$, θ_i is drawn from an isotropic multivariate Gaussian distribution with mean $\hat{\theta}$ and covariance $\sigma^2 I$, i.e., $\theta_i \sim \mathcal{N}(\hat{\theta}, \sigma^2 I)$. This can be expressed equivalently as $\theta_i = \hat{\theta} + \sigma \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, I)$. Here, θ_i denotes the parameters of the corresponding policy π_i .

2. **Fitness Evaluation:** The fitness of each sampled parameter θ_i is evaluated as $F(\theta_i)$, which corresponds to the total cost achieved by the policy π_i on a given CADWS problem instance. The fitness function is defined as $F(\theta_i) = F(\hat{\theta} + \sigma \epsilon_i) = -\text{TotalCost}(\pi_i)$, where $\text{TotalCost}(\pi_i)$ represents the overall cost of executing π_i as shown in Eq. (1).

3. **Policy Update:** To optimize the policy parameter $\hat{\theta}$, ES maximizes the expected fitness $\mathbb{E}_{\theta_i \sim \mathcal{N}(\hat{\theta}, \sigma^2 I)}[F(\theta_i)]$, effectively minimizing the total cost. The gradient is estimated as:

$$\begin{aligned} \nabla_{\hat{\theta}} \mathbb{E}_{\theta_i \sim \mathcal{N}(\hat{\theta}, \sigma^2 I)}[F(\theta_i)] &= \nabla_{\hat{\theta}} \mathbb{E}_{\epsilon_i \sim \mathcal{N}(0, I)}[F(\hat{\theta} + \sigma \epsilon_i)] \quad (2) \\ &\approx \frac{1}{N\sigma} \sum_{i=1}^N \{F(\hat{\theta} + \sigma \epsilon_i) \epsilon_i\} \end{aligned}$$

Algorithm 1 The training process of Evolution Strategy

Input: Population size: N , max number of generation: Gen , initial parameters of π_{GATES} : $\hat{\theta}$, initial learning rate: α , and the Gaussian standard noise deviation: σ

Output: The trained π_{GATES}

- 1: **while** the current number of generation $\leq Gen$ **do**
 - 2: Randomly sample a CADWS training instance: Pro .
 - 3: **for** each individual ($i=1,2,\dots$) **in** N **do**
 - 4: Sample a $\epsilon_i \sim \mathcal{N}(0, I)$.
 - 5: The parameters of π_i represented by individual i : $\theta_i = \hat{\theta} + \sigma \epsilon_i$
 - 6: Evaluate the fitness value of $F(\theta_i)$ using Eq. (1) based on Pro
 - 7: **end for**
 - 8: Estimate the policy gradient $\nabla_{\hat{\theta}} \mathbb{E}_{\theta_i \sim \mathcal{N}(\hat{\theta}, \sigma^2 I)} F(\theta_i)$ using Eq. (2).
 - 9: Update parameters of π_{GATES} : $\hat{\theta} \leftarrow \hat{\theta} + \alpha \frac{1}{N\sigma} \sum_{i=1}^N \{F(\hat{\theta} + \sigma \epsilon_i) \epsilon_i\}$.
 - 10: **end while**
 - 11: **return** π_{GATES}
-

5 Experiments

We conduct experiments using a dedicated simulator designed for the CADWS problem [Yang *et al.*, 2022; Yang *et al.*, 2025], leveraging Amazon EC2-based virtual machine (VM) configurations and four workflow patterns/DAGs: CyberShake, Montage, Inspiral, and SIPHT [Deelman *et al.*,

2015], as shown in Appendix C. To reflect real-world resource heterogeneity, we employ six VM types with varying computational capacities and costs. Workflows are categorized into three scales, including small, medium, and large, based on the number of tasks, representing different levels of computational complexity. SLA deadlines are defined using a coefficient $\gamma \in \{1.00, 1.25, 1.50, 1.75, 2.00, 2.25\}$, guiding resource selection under varying time constraints. More details on training and testing can be found in Appendix D.

The proposed GATES method is compared against four baseline algorithms: **ProLis** [Wu *et al.*, 2017], **GRP-HEFT** [Faragardi *et al.*, 2019], **ES-RL** [Huang *et al.*, 2022], and **SPN-CWS** [Shen *et al.*, 2024], spanning recently developed heuristic and DRL-based approaches. Key hyperparameters are configured to ensure fair comparison and computation efficiency.

5.1 Experiment Configuration

VM Configuration: We have six VM types, ranging from m5.large to m5.12xlarge, with varying computational resources (e.g., vCPUs, memory) and hourly costs (e.g., \$0.096 to \$2.304). Each type of VM can be rented on demand to accommodate dynamically changing resource requirements. The detailed configuration of VMs is shown in Table 1.

VM name	vCPU	Memory (GB)	Cost (\$ per hour)
m5.large	2	8	0.096
m5.xlarge	4	16	0.192
m5.2xlarge	8	32	0.384
m5.4xlarge	16	64	0.768
m5.8xlarge	32	128	1.536
m5.12xlarge	48	192	2.304

Table 1: The configuration of VMs.

Workflow Patterns: As shown in Table 2, workflows were categorized into three scales according to task size:

Workflow set	Name of pattern/number of task				
<i>Small</i>	CyberShake/30	Montage/25	Inspiral/30	SIPHT/30	
<i>Medium</i>	CyberShake/50	Montage/50	Inspiral/50	SIPHT/60	
<i>Large</i>	CyberShake/100	Montage/100	Inspiral/100	SIPHT/100	

Table 2: The three workflow sets used in this study.

SLA Penalties: SLA deadlines were defined using γ , which influenced resource selection. Larger SLA penalties encouraged renting cheaper VMs to manage cost while adhering to deadlines. The penalty weight β was set to $0.24/hour$ [Shen *et al.*, 2024].

Baseline Algorithms: ProLis [Wu *et al.*, 2017] and GRP-HEFT [Faragardi *et al.*, 2019] are heuristic methods. ES-RL [Huang *et al.*, 2022] is a DRL-based scheduling policy, developed based on OpenAI’s code¹. SPN-CWS [Shen *et al.*, 2024] is also a DRL-based method using the self-attention mechanism incorporated within the Transformer [Vaswani, 2017].

Hyperparameters of GATES: The hyperparameters of GATES mainly involve the parameters of HGAT (as shown in Figure 3), decision-making networks (as illustrated in Figure 4), and the training method ES (see Algorithm 1).

¹<https://github.com/openai/evolution-strategies-starter>

In Figure 3, both stage 1 and stage 2 employ two GAT layers. The first layer has two attention heads and the second layer has one. The output dimensions H of all GATs are 16. In Figure 4, the Transformer layers consist of 2 encoder layers, each with 2 self-attention heads, an input and output dimension of 16, and a hidden layer dimension of 128. In Stage 2 of the decision-making network, the MLP consists of 4 layers with hidden dimensions of 128 and output dimension of 1, and uses ReLU as the activation function. For ES in Algorithm 1, the hyperparameter settings are based on [Huang *et al.*, 2022] and [Shen *et al.*, 2024]. Specifically, we set the population size to 40, the maximum number of generations to 2000, the initial learning rate to 0.01, and the standard deviation of Gaussian noise to 0.05.

5.2 Results and Analysis

The results of all algorithms are presented in Table 3. The notation $\langle 1.00, S \rangle$ indicates that $\gamma = 1.00$ and the algorithms are tested on the small-scale CADWS instance. ProLis and GRP-HEFT, being deterministic heuristic methods, do not have associated standard deviations. The symbols (+), (−), or (\approx) denote that the results are significantly better, worse, or statistically equivalent, respectively, compared to the corresponding algorithm. These comparisons are based on the Wilcoxon test with a significance level of 0.05.

As shown in Table 3, GATES achieves the lowest average total cost across all tested scenarios, encompassing various SLA coefficients (γ) as well as small (S), medium (M), and large (L) workflow sizes. It also demonstrates relatively small standard deviations, indicating more stable performance across multiple independent runs. Compared with existing approaches (ProLis, GRP-HEFT, ES-RL, and SPN-CWS), GATES exhibits notable cost advantages in multiple scenarios. For instance, under the $\langle 1.50, M \rangle$ scenario, the total cost is reduced to 196.90, which is notably lower than the 276.09 achieved by SPN-CWS and also surpasses the results of other algorithms. Even under tight γ ($=1$) and with large workflow size, GATES consistently maintains the lowest cost, highlighting its adaptability across both relaxed and stringent SLA constraints. A Wilcoxon rank-sum test ($p < 0.05$) based on 30 independent runs confirms the statistical significance of these differences, indicating that GATES significantly outperforms the current state-of-the-art methods for workflows with varying patterns and scales. Considering the average cost, variance, and statistical evidence, GATES offers a distinct combined advantage in reducing overall cost while preserving performance stability, providing a more practical approach for solving CADWS.

5.3 Trade-off between VM Fees and SLA Penalties

Figure 5 compares the VM rental fees and corresponding SLA penalty costs for all algorithms under various scenarios. The y-axis is plotted on a logarithmic scale for better visualization, with each algorithm’s bar representing VM fees and the diagonally striped portion indicating SLA penalties. GRP-HEFT enforces strict SLA deadlines, leading to high VM fees but negligible penalties, while ProLis exhibits a similar pattern. ES-RL employs reinforcement learning for adaptive allocation, obtaining moderately high fees and relatively

Scenarios	ProLis	GRP-HEFT	ES-RL	SPN-CWS	GATES (our)
(1.00, S)	773.69	1685.01(-)	223.68(51.51)(+)(+)	158.28(14.45)(+)(+)(+)	143.95(18.94)(+)(+)(+)(+)
(1.00, M)	1829.25	2867.64(-)	422.26(100.92)(+)(+)	308.93(75.88)(+)(+)(+)	225.41(37.21)(+)(+)(+)(+)
(1.00, L)	3641.79	5873.20(-)	829.64(262.91)(+)(+)	524.83(163.15)(+)(+)(+)	324.93(65.65)(+)(+)(+)(+)
(1.25, S)	923.15	1967.31(-)	320.75(83.69)(+)(+)	143.12(15.73)(+)(+)(+)	124.14(7.10)(+)(+)(+)(+)
(1.25, M)	2068.78	3578.04(-)	565.53(213.00)(+)(+)	272.77(60.65)(+)(+)(+)	193.26(24.06)(+)(+)(+)(+)
(1.25, L)	4213.29	6868.22(-)	989.79(458.04)(+)(+)	496.63(155.02)(+)(+)(+)	300.19(76.29)(+)(+)(+)(+)
(1.50, S)	901.40	1963.39(-)	270.93(67.20)(+)(+)	144.13(18.97)(+)(+)(+)	121.56(7.11)(+)(+)(+)(+)
(1.50, M)	2035.52	3567.97(-)	505.44(222.45)(+)(+)	276.09(60.63)(+)(+)(+)	196.90(33.40)(+)(+)(+)(+)
(1.50, L)	4066.84	6863.62(-)	886.87(380.26)(+)(+)	511.69(161.85)(+)(+)(+)	316.23(99.12)(+)(+)(+)(+)
(1.75, S)	804.33	1959.71(-)	241.09(51.43)(+)(+)	137.36(14.10)(+)(+)(+)	118.45(4.85)(+)(+)(+)(+)
(1.75, M)	1977.10	3560.22(-)	491.97(229.48)(+)(+)	263.12(57.61)(+)(+)(+)	195.25(35.83)(+)(+)(+)(+)
(1.75, L)	3898.42	6854.40(-)	828.06(359.74)(+)(+)	502.67(155.64)(+)(+)(+)	315.84(92.48)(+)(+)(+)(+)
(2.00, S)	789.71	1957.25(-)	225.70(50.89)(+)(+)	129.75(12.72)(+)(+)(+)	113.48(5.30)(+)(+)(+)(+)
(2.00, M)	1921.42	3554.07(-)	427.16(213.71)(+)(+)	251.83(52.01)(+)(+)(+)	185.55(30.66)(+)(+)(+)(+)
(2.00, L)	4024.03	6847.49(-)	780.77(408.14)(+)(+)	497.65(158.52)(+)(+)(+)	322.01(110.05)(+)(+)(+)(+)
(2.25, S)	710.18	1954.71(-)	201.54(50.56)(+)(+)	126.18(13.02)(+)(+)(+)	110.90(5.70)(+)(+)(+)(+)
(2.25, M)	1924.27	3551.31(-)	374.96(190.21)(+)(+)	238.35(48.26)(+)(+)(+)	179.17(24.54)(+)(+)(+)(+)
(2.25, L)	3957.08	6842.88(-)	678.94(396.79)(+)(+)	475.46(139.99)(+)(+)(+)	318.19(112.21)(+)(+)(+)(+)

Table 3: Average (standard deviation) total cost of each algorithm over 30 independent runs.

low penalties. SPN-CWS focuses on reducing VM rental costs by using cheaper VMs at the price of high SLA penalties. In comparison, GATES demonstrates strong capabilities in balancing VM rental fees and SLA penalties to effectively reduce the total cost. Although GATES incurs slightly higher SLA penalties than ES-RL and SPN-CWS in some scenarios, it significantly lowers the total cost by effectively leveraging cheap VMs for task scheduling.

Overall, Figure 5 illustrates distinct trade-offs between cost and SLA compliance. While GRP-HEFT and ProLis favor meeting tight deadlines at the expense of higher VM fees, and SPN-CWS accepts greater penalties to minimize costs, GATES achieves a strong balance between these objectives by keeping both VM fees and SLA penalties within a favorable range to reduce the total cost.

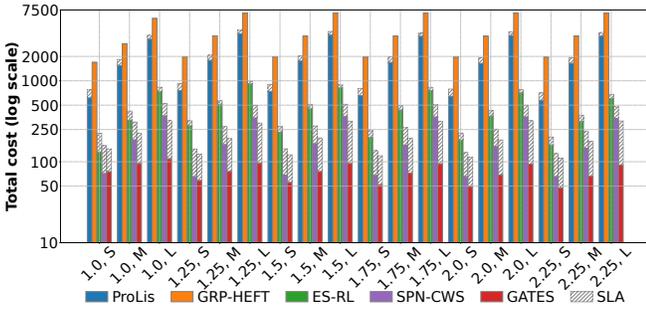


Figure 5: The trade-off between VM rental fees and SLA penalty.

5.4 Convergence

As shown in Figure 6, all three DRL-based algorithms (ES-RL, SPN-CWS, and GATES) converge during both training and testing, but their final performances vary significantly. GATES achieves superior results compared to the other algorithms, with its total cost curve significantly lower, indicating its ability to learn better scheduling policies in complex dynamic environments to reduce the total scheduling cost of CADWS. SPN-CWS ranks second, with convergence levels lower than GATES but better than ES-RL. In contrast, ES-RL exhibits lower curves with higher fluctuations, suggesting

poorer policy quality and less stability.

From the perspective of confidence intervals, both GATES and SPN-CWS demonstrate smaller fluctuations in training and testing curves, showing greater robustness to random seeds and environmental disturbances. Although some oscillations occur in the early iterations, their overall convergence variance is notably better than that of ES-RL. In comparison, ES-RL falls short in both solution quality and stability. Overall, GATES excels in solution quality, with its learned scheduling policies offering superior stability and significantly outperforming the other two DRL algorithms. This analysis further validates that the DHG constructed in GATES effectively captures the dynamics of CADWS, and the designed HGAT network extracts meaningful state embeddings to learn efficient scheduling policies.

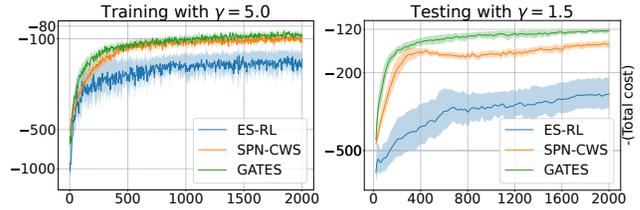


Figure 6: The convergence on the small-scenario instance. The horizontal axis denotes the training generation, while the vertical axis indicates the negative total cost.

6 Conclusion

This paper addresses the complex characteristics of CADWS in cloud environments, such as dynamic workflow arrivals and resource variations, by proposing a novel DRL method called GATES, which integrates GAT and ES. GATES learns task-level topological dependencies and VM importance to dynamically allocate tasks to optimal VMs while leveraging ES to avoid the challenges of complex reward design and hyperparameter tuning. Experimental results demonstrate that GATES can learn an efficient and stable scheduling policy, significantly outperforming several existing methods.

Ethical Statement

There are no ethical issues. The name of “GATES” is used coincidentally and has no relation to [Ning *et al.*, 2020] in terms of problem setting or technical content.

Acknowledgments

The authors would like to express their gratitude to the anonymous reviewers for their valuable feedback and constructive suggestions. The authors also acknowledge Victoria Huang, Chen Wang, Yifan Yang, and their collaborators for their prior works, on which the simulation environment used in this study is based. This research is partially supported by Grant VUW-FSRG- 410114, administered by Victoria University of Wellington.

References

- [Branke *et al.*, 2015] Jürgen Branke, Su Nguyen, Christoph W Pickardt, and Mengjie Zhang. Automated design of production scheduling heuristics: A review. *IEEE Transactions on Evolutionary Computation*, 20(1):110–124, 2015.
- [Cunha *et al.*, 2020] Bruno Cunha, Ana M Madureira, Benjamim Fonseca, and Duarte Coelho. Deep reinforcement learning as a job shop scheduling solver: A literature review. In *18th International Conference on Hybrid Intelligent Systems*, pages 350–359. Springer, 2020.
- [Deelman *et al.*, 2015] Ewa Deelman, Karan Vahi, Gideon Juve, Mats Rynge, Scott Callaghan, Philip J Maechling, Rajiv Mayani, Weiwei Chen, Rafael Ferreira Da Silva, Miron Livny, et al. Pegasus, a workflow management system for science automation. *Future Generation Computer Systems*, 46:17–35, 2015.
- [Dong *et al.*, 2021] Hang Dong, Boshi Wang, Bo Qiao, Wenqian Xing, Chuan Luo, Si Qin, Qingwei Lin, Dongmei Zhang, Gurpreet Virdi, and Thomas Moscibroda. Predictive job scheduling under uncertain constraints in cloud computing. In *IJCAI*, pages 3627–3634, 2021.
- [Faragardi *et al.*, 2019] Hamid Reza Faragardi, Mohammad Reza Saleh Sedghpour, Saber Fazliahmadi, Thomas Fahringer, and Nayereh Rasouli. Grp-heft: A budget-constrained resource provisioning scheme for workflow scheduling in iaas clouds. *IEEE Transactions on Parallel and Distributed Systems*, 31(6):1239–1254, 2019.
- [Geiger *et al.*, 2006] Christopher D Geiger, Reha Uzsoy, and Haldun Aytuğ. Rapid modeling and discovery of priority dispatching rules: An autonomous learning approach. *Journal of Scheduling*, 9:7–34, 2006.
- [HasanzadeZonuzy *et al.*, 2021] Aria HasanzadeZonuzy, Dileep Kalathil, and Srinivas Shakkottai. Model-based reinforcement learning for infinite-horizon discounted constrained markov decision processes. *IJCAI 2021*, 2021.
- [Huang *et al.*, 2022] Victoria Huang, Chen Wang, Hui Ma, Gang Chen, and Kameron Christopher. Cost-aware dynamic multi-workflow scheduling in cloud data center using evolutionary reinforcement learning. In *International Conference on Service-Oriented Computing*, pages 449–464. Springer, 2022.
- [Iklassev *et al.*, 2023] Zangir Iklassev, Dmitrii Medvedev, Ruben Solozabal Ochoa De Retana, and Martin Takac. On the study of curriculum learning for inferring dispatching policies on the job shop scheduling. In *IJCAI*, pages 5350–5358, 2023.
- [Jadeja and Modi, 2012] Yashpalsinh Jadeja and Kirit Modi. Cloud computing-concepts, architecture and challenges. In *2012 international conference on computing, electronics and electrical technologies (ICCEET)*, pages 877–880. IEEE, 2012.
- [Kayhan and Yildiz, 2023] Behice Meltem Kayhan and Gokalp Yildiz. Reinforcement learning applications to machine scheduling problems: a comprehensive literature review. *Journal of Intelligent Manufacturing*, 34(3):905–929, 2023.
- [Khadka and Tumer, 2018] Shauharda Khadka and Kagan Tumer. Evolution-guided policy gradient in reinforcement learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- [Khalil *et al.*, 2017] Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial optimization algorithms over graphs. *Advances in neural information processing systems*, 30, 2017.
- [Ning *et al.*, 2020] Xuefei Ning, Yin Zheng, Tianchen Zhao, Yu Wang, and Huazhong Yang. A generic graph-based neural architecture encoding scheme for predictor-based nas. In *European Conference on Computer Vision*, pages 189–204. Springer, 2020.
- [Ou *et al.*, 2023] Junwei Ou, Lining Xing, Feng Yao, Mengjun Li, Jimin Lv, Yongming He, Yanjie Song, Jian Wu, and Guoting Zhang. Deep reinforcement learning method for satellite range scheduling problem. *Swarm and Evolutionary Computation*, 77:101233, 2023.
- [Pallathadka *et al.*, 2022] Harikumar Pallathadka, Guna Sekhar Sajja, Khongdet Phasinam, Mahyudin Ritonga, Mohd Naved, Rajni Bansal, and Jose Quiñonez-Choquecota. An investigation of various applications and related challenges in cloud computing. *Materials Today: Proceedings*, 51:2245–2248, 2022.
- [Pham and Fahringer, 2020] Thanh-Phuong Pham and Thomas Fahringer. Evolutionary multi-objective workflow scheduling for volatile resources in the cloud. *IEEE Transactions on Cloud Computing*, 10(3):1780–1791, 2020.
- [Price *et al.*, 2025] Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Probabilistic weather forecasting with machine learning. *Nature*, 637(8044):84–90, 2025.
- [Qazi *et al.*, 2024] Faiza Qazi, Daehan Kwak, Fiaz Gul Khan, Farman Ali, and Sami Ullah Khan. Service level agreement in cloud computing: Taxonomy, prospects, and challenges. *Internet of Things*, page 101126, 2024.

- [Salimans *et al.*, 2017] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- [Shen *et al.*, 2024] Ya Shen, Gang Chen, Hui Ma, and Mengjie Zhang. Cost-aware dynamic cloud workflow scheduling using self-attention and evolutionary reinforcement learning. In *International Conference on Service-Oriented Computing*, pages 3–18. Springer, 2024.
- [Song *et al.*, 2022] Wen Song, Xinyang Chen, Qiqiang Li, and Zhiguang Cao. Flexible job-shop scheduling via graph neural network and deep reinforcement learning. *IEEE Transactions on Industrial Informatics*, 19(2):1600–1610, 2022.
- [Sun and Yang, 2023] Zhiqing Sun and Yiming Yang. Difusco: Graph-based diffusion solvers for combinatorial optimization. *Advances in Neural Information Processing Systems*, 36:3706–3731, 2023.
- [Thekumparampil *et al.*, 2018] Kiran K Thekumparampil, Chong Wang, Sewoong Oh, and Li-Jia Li. Attention-based graph neural network for semi-supervised learning. *arXiv preprint arXiv:1803.03735*, 2018.
- [Vaswani, 2017] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- [Veličković *et al.*, 2017] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [Wang *et al.*, 2019] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph attention network. In *The world wide web conference*, pages 2022–2032, 2019.
- [Wu *et al.*, 2017] Quanwang Wu, Fuyuki Ishikawa, Qingsheng Zhu, Yunni Xia, and Junhao Wen. Deadline-constrained cost optimization approaches for workflow scheduling in clouds. *IEEE Transactions on Parallel and Distributed Systems*, 28(12):3401–3412, 2017.
- [Yang *et al.*, 2022] Yifan Yang, Gang Chen, Hui Ma, and Mengjie Zhang. Dual-tree genetic programming for deadline-constrained dynamic workflow scheduling in cloud. In *Service-Oriented Computing*, pages 433–448. Springer, 2022.
- [Yang *et al.*, 2025] Yifan Yang, Gang Chen, Hui Ma, Cong Zhang, Zhiguang Cao, and Mengjie Zhang. Graph assisted offline-online deep reinforcement learning for dynamic workflow scheduling. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [Zhang *et al.*, 2020] Cong Zhang, Wen Song, Zhiguang Cao, Jie Zhang, Puay Siew Tan, and Xu Chi. Learning to dispatch for job shop scheduling via deep reinforcement learning. *Advances in neural information processing systems*, 33:1621–1632, 2020.
- [Zhang *et al.*, 2024] Cong Zhang, Zhiguang Cao, Wen Song, Yaixin Wu, and Jie Zhang. Deep reinforcement learning guided improvement heuristic for job shop scheduling. In *The Twelfth International Conference on Learning Representations*, 2024.
- [Zhu *et al.*, 2016] Xiaomin Zhu, Ji Wang, Hui Guo, Dakai Zhu, Laurence T Yang, and Ling Liu. Fault-tolerant scheduling for real-time scientific workflows with elastic resource provisioning in virtualized clouds. *IEEE Transactions on Parallel and Distributed Systems*, 27(12):3501–3517, 2016.