

Neural Combinatorial Optimization for Stochastic Flexible Job Shop Scheduling Problems

Igor G. Smit^{1, 2}, Yaoxin Wu^{2, 3*}, Pavel Troubil⁴, Yingqian Zhang^{2, 3}, Wim P.M. Nuijten^{1, 2}

¹Department of Mathematics and Computer Science, Eindhoven University of Technology

²Eindhoven Artificial Intelligence Systems Institute, Eindhoven University of Technology

³Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology

⁴Delmia R&D, Dassault Systèmes

i.g.smit@tue.nl, y.wu2@tue.nl, pavel.troubil@3ds.com, yqzhang@tue.nl, w.p.m.nuijten@tue.nl

Abstract

Neural combinatorial optimization (NCO) has gained significant attention due to the potential of deep learning to efficiently solve combinatorial optimization problems. NCO has been widely applied to job shop scheduling problems (JSPs) with the current focus predominantly on deterministic problems. In this paper, we propose a novel attention-based scenario processing module (SPM) to extend NCO methods for solving stochastic JSPs. Our approach explicitly incorporates stochastic information by an attention mechanism that captures the embedding of sampled scenarios (i.e., an approximation of stochasticity). Fed with the embedding, the base neural network is intervened by the attended scenarios, which accordingly learns an effective policy under stochasticity. We also propose a training paradigm that works harmoniously with either the expected makespan or Value-at-Risk objective. Results demonstrate that our approach outperforms existing learning and non-learning methods for the flexible JSP problem with stochastic processing times on a variety of instances. In addition, our approach holds significant generalizability to varied numbers of scenarios and disparate distributions.

Code — <https://github.com/ai-for-decision-making-tue/NCO-for-Stochastic-FJSP>

Introduction

Combinatorial optimization (CO) problems are prevalent in various application areas, posing significant challenges due to their complex nature and computational intractability. One critical CO problem is the flexible job shop scheduling problem (FJSP), which has wide applicability ranging from semiconductor manufacturing (Tamssaouet et al. 2022) to healthcare scheduling (Burdett and Kozan 2018) and aluminum production (Zhang, Yang, and Zhou 2016). Traditional approaches, such as constraint programming (Baptiste, Pape, and Nuijten 2001; Col and Teppan 2022), heuristics (Sels, Gheysen, and Vanhoucke 2012), or meta-heuristics (Rooyani and Defersha 2019) have made great progress in solving these problems. However, these methods generally assume deterministic problems. In practice, oppo-

sitely, there is an abundance of uncertainty, leading to vulnerable plans produced by these optimization methods.

In recent years, increasing focus has been directed to stochastic optimization methods. However, these methods remain largely underrepresented, especially for more complex optimization problems such as the FJSP. Existing methods solving the stochastic FJSP, such as a simulation-optimization framework (Ghaedy-Heidary et al. 2024) or a meta-heuristic approach including Monte-Carlo sampling (Flores-Gómez, Borodin, and Dauzère-Pérès 2023) add significant computational complexity to already expensive approaches, while also being tailored to specific use-cases and objectives, limiting their potential for adoption.

A recent branch of optimization technology is constituted by neural combinatorial optimization (NCO) methods, leveraging deep (reinforcement) learning (DRL) to solve CO problems (Bengio, Lodi, and Prouvost 2021). They learn high-quality optimization policies that fit a set of problem instances, eliminating the need for handcrafted expert rules, while also being efficient. Current NCO methods mainly target deterministic optimization, overlooking their practical use in stochastic situations. It is highly beneficial to extend NCO to the stochastic optimization domain.

Several existing works leverage learning-based methods for stochastic routing or scheduling. However, these are generally focused on online or dynamic scenarios. Conversely, in many practical cases, such as manufacturing plants, a one-time schedule is needed in advance. Moreover, existing methods only focus on implicitly learning stochastic dynamics through the Markov decision process (MDP). In stochastic optimization, on the contrary, the assumed distributions can be explicitly considered, allowing for better plans.

We address these issues by a novel attention-based *scenario processing module* (SPM) that extends deterministic NCO architectures to solve stochastic problems. The SPM, which is expressive yet easily adaptable to different base neural networks for different problems, effectively captures the representation of stochastic scenarios to favorably intervene the solution policy. We embed SPM in a novel training paradigm that can address various stochastic objectives and offers high-quality yet fast solution policies, using a sampling strategy commonly found in stochastic optimization methods. We combine SPM with an existing network archi-

*Corresponding author.

ecture to form SPM-DAN and solve the stochastic FJSP. Results show our method outperforms existing methods consistently. Our main contributions are summarized as follows:

1. We propose an attention-based SPM to solve stochastic optimization problems, which is expressive and modular, with a good transferability to different base neural networks.
2. We put forward an effective DRL training paradigm, which is used to optimize different stochastic objectives.
3. We apply SPM to solve the stochastic FJSP, which outperforms existing learning and non-learning methods.

Related Work

NCO for FJSP Graph neural networks (GNNs) for solving a variety of scheduling problems have been developed in recent years (Zhang et al. 2020; Park et al. 2021; Lei et al. 2022; Kwon et al. 2021). We refer to a recent survey (Smit et al. 2025) for a complete outline. Song et al. (2022) first proposed a competitive end-to-end DRL algorithm to construct FJSP solutions. They used a heterogeneous graph and designed a heterogeneous GNN using different GAT (Veličković et al. 2018) layers to encode machine and operation nodes. In the follow-up work, Wang et al. (2023) proposed a recent state-of-the-art dual attention network (DAN) that comprises both self- and cross-attention, which achieved superior performance over previous DRL approaches.

NCO for Dynamic FJSP Multiple papers target the dynamic FJSP (e.g. Zhang et al. 2023; Lei et al. 2023), which involve different sources of uncertainty such as random arrivals, machine availability, or processing times. These papers mainly focus on developing training strategies and MDP formulations to fit the dynamic use cases. The network architectures are not specifically tailored to address uncertainties, which are implicitly handled by approximating the expected reward in MDP trajectories. Similarly, NCO methods for other problems also mainly target deterministic or dynamic cases, in which the realized values become known during solution construction (Schmitt-Ulms et al. 2022; Kwon et al. 2020; Joe and Lau 2020).

NCO for Stochastic FJSP Similar to us, Infantes et al. (2024) consider a stochastic job shop scheduling problem (JSP) for which a full plan must be created before the realizations become known. They include three features in the state space to describe the assumed triangular distribution, and base the reward function on a single sampled realization for each instance to optimize the expected makespan. While promising, their method is limited to triangular processing time distributions and the expected makespan objective.

Preliminaries and Problem Formulation

Flexible Job Shop Scheduling Problem The FJSP consists of a pair $(\mathcal{J}, \mathcal{M})$ where \mathcal{J} and \mathcal{M} are jobs and machines, respectively. A job $J_i \in \mathcal{J}$ consists of a set $\mathcal{O}_i = \{O_{i1}, \dots, O_{in_i}\}$ of n_i operations to be performed in order. The total set of operations is $\mathcal{O} = \bigcup_i \mathcal{O}_i$. Each operation

O_{ij} must be processed by a single machine, selected from the set of compatible machines $\mathcal{M}_{ij} \subseteq \mathcal{M}$. The processing time to perform operation O_{ij} on machine $M_k \in \mathcal{M}_{ij}$ is given by $p_{ij}^k > 0$ and each machine can only process one job at a time. A solution to the FJSP is a *schedule*, which assigns a compatible machine to each operation O_{ij} and determines the order of operations on each machine. The goal is to minimize the makespan $c_{max} = \max_{O_{ij} \in \mathcal{O}} c_{ij}$, which is the maximum completion time c_{ij} of all operations.

Stochastic Flexible Job Shop Scheduling Problem In the stochastic FJSP, we assume that the processing times are random variables P_{ij}^k . The realized processing times $p_{ij}^k \sim \mathcal{P}_{ij}^k$ follow the probability distributions \mathcal{P}_{ij}^k and only become known after the schedule is created. As a result, the realized c_{ij} and c_{max} are also realizations of the random variables C_{ij} and C_{max} . The goal is to create a schedule that optimizes an objective function $f(C_{max})$, which can be the expected value $\mathbb{E}(C_{max})$, the Value-at-Risk $VaR_\alpha(C_{max}) = \min\{c_{max} : \mathbb{P}(C_{max} \leq c_{max}) \geq \alpha\}$, or other functions of C_{max} . VaR_α indicates the value for which, with probability α , our makespan is at most this high. As such, it considers the practically relevant robustness to the uncertainty of schedules. We set and keep $\alpha = 95\%$ in this paper.

Attention Mechanism The attention function aims to map a group of n queries $Q \in \mathbb{R}^{n \times d_q}$ to their advanced representations, using n_v keys $K \in \mathbb{R}^{n_v \times d_q}$ and n_v values $V \in \mathbb{R}^{n_v \times d_v}$, such that:

$$\text{Att}(Q, K, V) = \text{softmax}(QK^T / \sqrt{d_q})V$$

where QK^T defines the similarity between Q and K , based on which weights are computed by the softmax function and then used to obtain a weighted sum of the values in V .

Vaswani et al. (2017) extended the attention mechanism by introducing the multi-head attention. By setting h heads, Q , K , and V are first projected to h matrices, respectively. Then, the attention function is applied to every three matrices in each head. The advanced representations from all heads are concatenated and linearly transformed:

$$\text{MHA}(Q, K, V) = \text{concat}(\text{head}_1, \dots, \text{head}_h)W^O$$

where $\text{head}_i = \text{Att}(QW_i^Q, KW_i^K, VW_i^V)$. $W_i^Q \in \mathbb{R}^{d_q \times d'_q}$, $W_i^K \in \mathbb{R}^{d_q \times d'_q}$, and $W_i^V \in \mathbb{R}^{d_v \times d'_v}$ are learnable matrices. We follow a typical choice for the dimensions by setting $d_q = d_v = d$ and $d'_q = d'_v = d/h$. The multi-head attention layer is often embedded in a block, alongside the normalization, skip-connection, and feed-forward layer. Given input matrices $X \in \mathbb{R}^{n \times d}$ and $Y \in \mathbb{R}^{n_v \times d}$ a multi-head attention block is defined as:

$$\text{MHAB}(X, Y) = \text{LayerNorm}(Z + \text{FF}(Z))$$

where $Z = \text{LayerNorm}(X + \text{MHA}(X, Y, Y))$. FF is a fully-connected feed-forward layer and LayerNorm is the layer normalization operation (Ba, Kiros, and Hinton 2016).

Methodology

We take the state-of-the-art FJSP method, i.e., the DAN from (Wang et al. 2023), for an example, and describe how our method SPM-DAN is developed from DAN. Our method is easily applicable to other NCO methods.

Markov Decision Process

The scheduling process can be considered as a sequential decision-making process of iteratively assigning operations to available machines. At every decision moment t , an operation-machine pair (O_{ij}, M_k) is selected such that O_{ij} can be assigned to M_k . In the MDP formulation, the agent at each step receives state s_t , representing the environment, and takes action $a_t = (O_{ij}, M_k) \in \mathcal{A}(t)$ from the set of eligible actions $\mathcal{A}(t)$, which are the possible assignments of the first unscheduled operation in each job to a compatible machine. The environment feeds back reward r_t and new state s_{t+1} . The schedule is completed after $|\mathcal{O}|$ actions.

State Space The relevant operations $\mathcal{O}_u(t) \subseteq \mathcal{O}$ for the state s_t are all operations except those that already have a successor scheduled on the same machine and, thus, do not directly influence the schedule anymore. The relevant machines $\mathcal{M}_u(t) \subseteq \mathcal{M}$ are all machines on which any of the remaining operations can be scheduled. For the deterministic method proposed in (Wang et al. 2023), the state $s_t = \{\mathcal{H}_O, \mathcal{H}_M, \mathcal{H}_{OM}\}$ consists of the operation features $\mathcal{H}_O = \{h_{O_{ij}} \in \mathbb{R}^{10} | O_{ij} \in \mathcal{O}_u(t)\}$, machine features $\mathcal{H}_M = \{h_{M_k} \in \mathbb{R}^8 | M_k \in \mathcal{M}_u(t)\}$, and operation-machine pair features $\mathcal{H}_{OM} = \{h_{(O_{ij}, M_k)} \in \mathbb{R}^8 | (O_{ij}, M_k) \in \mathcal{A}(t)\}$. These features are dynamic and interdependent. We refer to the original paper for the detailed feature descriptions. We introduce an expansion of this state space, based on sample approximation. Concretely, we sample n_{scn} independent scenarios, each representing a set of realized processing time values $p_{ij}^k \sim \mathcal{P}_{ij}^k$. For each scenario, we keep track of all the features during the scheduling process. In addition, we maintain the deterministic problem instance. In doing so, the expanded state captures both the deterministic and stochastic problem characteristics. Formally, our state at step t is $s_t = \{s_t^{det}, \mathcal{S}_t^{stoch}\}$, where $s_t^{det} = \{\mathcal{H}_O^{det}, \mathcal{H}_M^{det}, \mathcal{H}_{OM}^{det}\}$ represents the deterministic instance state, and $\mathcal{S}_t^{stoch} = \{s_t^l | 1 \leq l \leq n_{scn}\}$ is the set of states representing the scenarios $s_t^l = \{\mathcal{H}_O^l, \mathcal{H}_M^l, \mathcal{H}_{OM}^l\}$. Note that the sets $\mathcal{O}_u(t)$, $\mathcal{M}_u(t)$, and $\mathcal{A}(t)$ are the same in the deterministic and sampled scenarios.

An alternative to capture stochastic information is representing the distribution by statistical features (Infantes et al. 2024). However, these features are hard to define for complex distributions and may lead to high dimensionality. Moreover, the varied characteristics in the state formulation depend on both the scheduling logic and the processing times, which cannot be accurately described by the statistical features, so valuable information is inevitably lost. Our features of the deterministic instance and scenarios can favorably reflect the influence of stochasticity, and are applicable to different problems, feature definitions, and distributions. Using efficient multiprocessing or parallel GPU computations, the runtime overhead of processing the states of scenarios is negligible, thereby maintaining fast solution construction.

Action Space The action space $\mathcal{A}(t)$ consists of all compatible operation-machine pairs of the first unscheduled operations per job. The same action is applied to the determin-

istic instance and all the sampled scenarios.

State Transition Upon taking an action, the sets $\mathcal{O}_u(t)$, $\mathcal{M}_u(t)$, and $\mathcal{A}(t)$ are updated to reflect the new relevant operations, machines, and potential actions. Additionally, the features for s_t^{det} and all s_t^l are updated independently, resulting in a new state s_{t+1} .

Reward We generalize the reward defined in (Zhang et al. 2020) and (Wang et al. 2023). In this deterministic reward function, a lower bound $\underline{C}(O_{ij}, s_t)$ of the completion time is estimated for each operation O_{ij} at time step t . This bound equals the scheduled completion time if the operation has been scheduled. For unscheduled operations, the lower bound is approximated using the recursion $\underline{C}(O_{ij}, s_t) = \underline{C}(O_{i(j-1)}, s_t) + \min_{k \in \mathcal{M}_{ij}} p_{ij}^k$. Then the reward at step t is $r_t = \max_{O_{ij} \in \mathcal{O}} \underline{C}(O_{ij}, s_t) - \max_{O_{ij} \in \mathcal{O}} \underline{C}(O_{ij}, s_{t+1})$.

For the stochastic reward, we consider n_{rew} scenarios that are independently sampled and do not overlap with the scenarios of the state, to facilitate the generalization. These scenarios require a small amount of computation for tracking their makespan lower bounds since they do not necessitate computing all features. Concretely, let \hat{s}_t^l represent the state of reward scenario l , then $\mathcal{C}_t = \{\max_{O_{ij} \in \mathcal{O}} \underline{C}(O_{ij}, \hat{s}_t^l) | 1 \leq l \leq n_{rew}\}$ are the makespan lower bounds of all scenarios at time t . The reward is computed as $r_t = f(\mathcal{C}_t) - f(\mathcal{C}_{t+1})$, where f is the objective function, which in our case is the VaR_α or mean makespan. However, other functions such as conditional VaR_α or median are also possible. If the discounting factor is $\gamma = 1$, the sum of all rewards naturally sums to $\sum_{t=0}^{|\mathcal{O}|-1} r_t = f(\mathcal{C}_0) - f(\mathcal{C}_{|\mathcal{O}|-1})$, where $\mathcal{C}_{|\mathcal{O}|-1} = \{\mathcal{C}_{max}^l | 1 \leq l \leq n_{rew}\}$ are the makespans of all scenarios. Hence, maximizing the cumulative reward corresponds to minimizing the objective.

Network Architecture

Our neural network consists of the scenario processing modules (SPMs) that take the scenarios \mathcal{S}_t^{stoch} and compute embeddings from them to represent stochasticity. These are combined with the deterministic state s_t^{det} and inputted to the base neural network. In doing so, we do not need to restructure the base neural network drastically, making our method easily applicable to different GNN architectures. An alternative would be to pass all scenarios through the GNN and do ‘post-processing’ before the actor and critic layers. However, this is significantly more computationally complex. Figure 1 outlines the proposed SPM and its application to the state-of-the-art DAN for FJSP (Wang et al. 2023).

Scenario Processing Module To efficiently and effectively extract embeddings from the scenario states to reflect the stochasticity, we design the SPM by a multi-head attention mechanism, which is inspired by (Lee et al. 2019). Consider a set \mathcal{H} of n_{scn} scenario states, which are linearly transformed into d -dimensional embeddings $H \in \mathbb{R}^{n_{scn} \times d}$. We expect to extract an embedding that represents the distribution of scenario states. To this end, we capture the interaction between scenarios using the multi-head attention. A straightforward way would be to apply the full self-attention MHAB(H, H). However, self-attention complex-

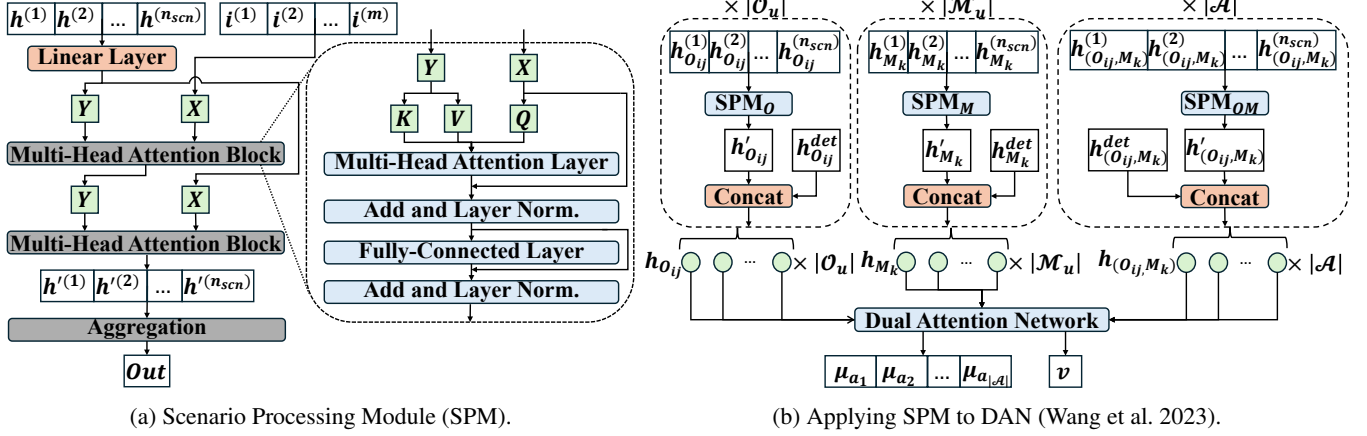


Figure 1: Overview of SPM-DAN.

ity scales quadratically with n_{scn} , making it too expensive for a large value n_{scn} . Hence, we propose a trainable set of m d -dimensional inducing point vectors $I \in \mathbb{R}^{m \times d}$ to compute the cross-attention over $H \in \mathbb{R}^{n_{scn} \times d}$, thereby maintaining a linear complexity with n_{scn} and resulting in the embeddings $J \in \mathbb{R}^{m \times d}$. We then apply the original set H to perform cross-attention with J , to get $H' \in \mathbb{R}^{n_{scn} \times d}$, which approximates $\text{MHAB}(H, H)$. To extract a single embedding of the stochasticity, we must apply permutation invariant aggregation, for which we use mean aggregation. Formally, our scenario processing module (SPM) can be expressed by:

$$\text{SPM}(\mathcal{H}) = \text{Avg}(\text{MHAB}(H, \text{MHAB}(I, H)))$$

SPM-DAN The DAN in (Wang et al. 2023) comprises L operation message attention blocks (OMBs) and machine message attention blocks (MMBs), followed by the actor and critic network. Each OMB_l computes embeddings $h_{O_{ij}}^{(l)}$ of operations $O_{ij} \in \mathcal{O}_u$ by an attention mechanism:

$$\{h_{O_{ij}}^{(l)} | O_{ij} \in \mathcal{O}_u\} = \text{OMB}_l(\{h_{O_{ij}}^{(l-1)} | O_{ij} \in \mathcal{O}_u\})$$

Similarly, MMB_l computes embeddings $h_{M_k}^{(l)}$ for all machines $M_k \in \mathcal{M}_u$ through an attention mechanism by involving both the operation and machine features, such that:

$$\{h_{M_k}^{(l)} | M_k \in \mathcal{M}_u\} = \text{MMB}_l(\{h_{M_k}^{(l-1)} | M_k \in \mathcal{M}_u\}, \{h_{O_{ij}}^{(l-1)} | O_{ij} \in \mathcal{O}_u\})$$

The actor network in the DAN is a multilayer perceptron (MLP) that calculates a score $\mu_{(O_{ij}, M_k)}$ for each action $(O_{ij}, M_k) \in \mathcal{A}$:

$$\mu_{(O_{ij}, M_k)} = \text{MLP}_\theta([h_{O_{ij}}^{(L)} || h_{M_k}^{(L)} || h_G^{(L)} || h_{(O_{ij}, M_k)}^{(L)}])$$

where $h_G^{(L)} = [\frac{1}{|\mathcal{O}_u|} \sum_{O_{ij} \in \mathcal{O}_u} h_{O_{ij}}^{(L)} || \frac{1}{|\mathcal{M}_u|} \sum_{M_k \in \mathcal{M}_u} h_{M_k}^{(L)}]$ is a graph embedding. Then, the DAN uses the softmax function over the scores to transform them into a probability distribution π_θ , from which the action can be sampled. The critic network is an MLP for outputting a scalar value:

$$v = \text{MLP}_\phi(h_G^{(L)})$$

which estimates the value of actions taken by the actor. We refer to (Wang et al. 2023) for details of each block.

The DAN requires three distinct types of input features, $h_{O_{ij}} = h_{O_{ij}}^{(0)}$, $h_{M_k} = h_{M_k}^{(0)}$, and $h_{(O_{ij}, M_k)}$. We compute these features based on our state definition. Specifically, we apply our SPM to process the scenario states \mathcal{S}^{stoch} and concatenate the resulting embeddings with their deterministic counterparts, such that:

$$h_{O_{ij}} = [h_{O_{ij}}^{det} || \text{SPM}_O(\{h_{O_{ij}}^l | 1 \leq l \leq n_{scn}\})]$$

$$h_{M_k} = [h_{M_k}^{det} || \text{SPM}_M(\{h_{M_k}^l | 1 \leq l \leq n_{scn}\})]$$

$$h_{(O_{ij}, M_k)} = [h_{(O_{ij}, M_k)}^{det} || \text{SPM}_{OM}(\{h_{(O_{ij}, M_k)}^l | 1 \leq l \leq n_{scn}\})]$$

In summary, SPM-DAN uses the SPM to compute embeddings representing the stochasticity of the problem instance. Combined with the deterministic features and passed to DAN, they enable the learning of effective solution policies for solving the stochastic FJSP.

Training Procedure

Following (Song et al. 2022) and (Wang et al. 2023), we train our policy by an actor-critic-based proximal policy optimization (PPO) algorithm (Schulman et al. 2017), shown in Algorithm 1. We extend the default PPO by keeping track of multiple independent scenarios per problem instance and applying the same actions. We compute the reward and state using these distinct scenarios and combine them into a single transition tuple reflecting all scenarios per problem instance per step. Hence, we maintain many scenarios, accommodating our state space to reflect the stochastic problem, without increasing the PPO batch sizes. The model is trained for N_{ep} episodes. We sample batches of n_B problem instances, with n_{scn} scenarios to compute the state and n_{rew} scenarios for the reward after N_s episodes. We validate on a fixed set of n_{vali} instances every N_{eval} episodes. As a common practice, we sample actions from the distribution π_θ for training and infer greedy solutions for validation.

Algorithm 1: Training Procedure

Input: Neural network with randomly initialized parameters, fixed set of n_{eval} evaluation instances

- 1: Sample batch of n_B instances
- 2: **for** $n_{ep} = 1, 2, \dots, N_{ep}$ **do**
- 3: **for** $b = 1, 2, \dots, n_B$ **do** ▷ *In Parallel*
- 4: **for** $t = 0, 1, \dots, |\mathcal{O}| - 1$ **do**
- 5: Compute π_θ and sample $a_{t,b} \sim \pi_\theta$
- 6: Perform $a_{t,b}$ and receive $s_{t+1,b}^{det}$
- 7: **for** $l = 1, 2, \dots, n_{scn}$ **do** ▷ *In Parallel*
- 8: Perform $a_{t,b}$ and receive $s_{t+1,b}^l$
- 9: **for** $l = 1, 2, \dots, n_{rew}$ **do** ▷ *In Parallel*
- 10: Perform $a_{t,b}$ and receive $\hat{s}_{t+1,b}^l$
- 11: Compute $r_{t,b}$
- 12: Collect transition $(s_{t,b}, a_{t,b}, r_{t,b}, s_{t+1,b})$
- 13: Compute generalized advantage estimates
- 14: Compute PPO loss and update network parameters
- 15: **if** $n_{ep} \bmod N_{eval} = 0$ **then**
- 16: Validate the policy
- 17: **if** $n_{ep} \bmod N_B = 0$ **then**
- 18: Sample new batch of n_B training instances

Experiments

In this section, we numerically evaluate our method on various problem instances and compare the performance with several baseline methods.

Baselines We use the first-in-first-out (FIFO), most-operations-remaining (MOR), shortest-processing-time (SPT), and most-work-remaining (MWKR) priority dispatching rules (PDRs). In addition, we evaluate the deterministic OR-tools CP-SAT solver (Perron, Didier, and Gay 2023) with the median processing times, using the implementation by (Reijnen et al. 2023). We also extend the CP-SAT formulation to optimize over multiple sampled scenarios simultaneously, which we name CP-stoch. For these methods, we produce a plan using the deterministic information (or the used scenarios for CP-stoch) and evaluate on n_{rew} independent scenarios. We set a 1-hour time limit for the CP-SAT solver. Based on preliminary tests, we set the number of scenarios for CP-stoch to 25 for synthetic 10×5 and 20×5 instances, and 10 for the other. Lastly, we evaluate the deterministic DAN policies on stochastic instances and use DAN with our reward mechanism but without SPM (DAN-stoch).

Datasets As is common, we use synthetic datasets for training and evaluation. We use SD_1 from (Wang et al. 2023). However, SD_1 has processing times in the range $\{1, \dots, 20\}$. Since it is common to round to integers, the stochasticity is limited, especially for small numbers (e.g., 1 with 50% standard deviation will mostly be 1). We exclude SD_2 from (Wang et al. 2023) as it has an unrealistic assumption of unrelated processing times $p_{ij}^k \in \{1, \dots, 99\}$ for an operation O_{ij} , drastically limiting the effects of stochasticity (e.g., $p_{ij}^1 = 10$ will mostly be favored over $p_{ij}^2 = 90$ regardless of the variance). Instead, we propose a more realistic SD_3 for which we mimic the instance structure of SD_2 ,

but sample processing times using $\bar{p}_{ij} \sim U(1, 99)$ for each O_{ij} , after which we sample $p_{ij}^k \sim U(0.85\bar{p}_{ij}, 1.15\bar{p}_{ij})$. We use instance sizes 10×5 , 20×5 , 15×10 , 20×10 , 30×10 , and 40×10 , with $n \times m$ indicating n jobs and m machines.

We also evaluate our method using the mk (Brandimarte 1993), rdata, edata, and vdata instances (Hurink, Jurisch, and Thole 1994). Note, that the mk instances suffer from both limitations of SD_1 and SD_2 .

Creating Stochastic Instances We assume the processing times of the deterministic instances to be the median value. For each pair (O_{ij}, M_k) , we sample $CV_{ij}^k \sim U(0.1, 0.5)$ and set the standard deviation $\sigma_{ij}^k = CV_{ij}^k \times \text{median}_{ij}^k$, as different machine-operation pairs have various uncertainty levels. There is no agreement on the best probability distributions in the literature. Unless specified otherwise, we assume log-normal distributions like (Caldeira and Gnanavelbabu 2021).

Configurations For a fair comparison, we keep the same training configurations $N_{ep} = 1000$, $N_B = 20$, $N_s = 20$, $N_{eval} = 10$ as (Wang et al. 2023). We also keep the same hyperparameters for the OMB and MMB blocks, the critic network, and the PPO algorithm, for which we refer to their paper. We increase the hidden layer size of the actor network to 128 to facilitate the expanded size of the input features $h_{(O_{ij}, M_k)}$. Based on preliminary tests, we use 4 attention heads and dimensions $d = 32$ for the SPMs. We set the number of inducing points $m = 16$. For training, validation, and testing we set $n_{rew} = 1000$ and use $n_{scn} = 100$. For the synthetic data, we test 100 problem instances per instance type. We use an NVIDIA A100 GPU and 18-core Intel Xeon Platinum 8360Y CPU for DRL training and a 32-core AMD Rome 7H12 CPU for the CP-SAT solver. We train policies on SD_1 and SD_3 for the smallest four instance sizes, which takes between 1 and 14 hours. For inference, we use greedy solution construction and a sampling strategy with 100 samples per instance, conforming to the existing FJSP papers. Unless specified otherwise, we assume the $Var_{95\%}$ of the makespan as the objective. We evaluate this metric using $n_{rew} = 1000$ scenarios for each instance. We also report the inference time and the gap to the CP-SAT solutions, calculated as $gap = \frac{f(C_{max}) - f(C_{max})_{CP}}{f(C_{max})_{CP}}$. The reported values are the average for each set of instances.

Synthetic Datasets Results

Table 1 presents the performance of our methods using test instances generated from the same distributions as the training instances. These results show that our method consistently outperforms all PDRs. In addition, using sampling inference, our model outperforms the deterministic CP-SAT solver on all instance sets. The CP-stoch formulation is only competitive on the 10×5 instances. In larger instances, the additional computational complexity outweighs the increased model accuracy. DAN and DAN-stoch are competitive with our model only on the smallest SD_1 instances. For other instance sets, especially larger ones, our method performs considerably better. In our implementation, SPM-DAN adds runtime of roughly 20% for greedy inference and 3 times for sampling (3.14 vs. 2.63 and 10.50 vs. 3.80

Size	PDRs				Greedy			Sample			CP-SAT	CP-stoch		
	FIFO	MOR	MWKR	SPT	DAN	DAN-stoch	SPM-DAN	DAN	DAN-stoch	SPM-DAN				
SD ₁	10×5	Obj	160.66	157.40	155.01	171.16	149.03	150.99	149.58	140.90	141.64	140.44	140.60	139.07
		Gap	14.27%	11.95%	10.25%	21.74%	6.00%	7.39%	6.39%	0.21%	0.74%	-0.11%	0.00%	-1.09%
		Time (s)	0.06	0.06	0.06	0.06	0.69	0.69	0.84	4.53	4.59	5.33	-	-
	20×5	Obj	270.60	271.81	270.93	293.45	270.14	261.82	255.47	259.75	265.59	250.81	251.26	260.14
		Gap	7.70%	8.18%	7.83%	16.79%	7.51%	4.20%	1.68%	3.38%	5.70%	-0.18%	0.00%	3.53%
		Time (s)	0.12	0.12	0.12	0.12	1.35	1.39	1.66	9.10	9.15	12.62	-	-
	15×10	Obj	254.20	241.74	239.70	266.19	226.48	228.12	220.20	216.17	218.61	211.01	216.36	224.62
		Gap	17.49%	11.73%	10.79%	23.03%	4.68%	5.44%	1.77%	-0.09%	1.04%	-2.47%	0.00%	3.82%
		Time (s)	0.19	0.19	0.19	0.18	2.10	2.07	2.45	14.41	14.42	22.50	-	-
	20×10	Obj	308.87	299.23	296.27	337.68	279.44	274.36	262.53	270.15	266.34	256.77	276.42	291.56
		Gap	11.74%	8.25%	7.18%	22.16%	1.09%	-0.75%	-5.02%	-2.27%	-3.65%	-7.11%	0.00%	5.48%
		Time (s)	0.26	0.27	0.26	0.26	2.77	2.83	3.36	19.84	19.88	34.47	-	-
SD ₃	10×5	Obj	760.14	755.55	744.38	827.24	722.77	736.05	710.63	688.74	693.49	677.24	688.10	676.19
		Gap	10.47%	9.80%	8.18%	20.22%	5.04%	6.97%	3.27%	0.09%	0.78%	-1.58%	0.00%	-1.73%
		Time (s)	0.06	0.06	0.06	0.06	0.65	0.65	0.79	0.74	0.73	1.11	-	-
	20×5	Obj	1310.67	1329.94	1320.25	1427.00	1307.48	1290.02	1279.70	1272.08	1257.08	1243.75	1253.67	1271.71
		Gap	4.55%	6.08%	5.31%	13.83%	4.29%	2.90%	2.08%	1.47%	0.27%	-0.79%	0.00%	1.44%
		Time (s)	0.12	0.12	0.12	0.12	1.28	1.37	1.58	1.46	1.48	2.99	-	-
	15×10	Obj	1215.14	1183.13	1156.25	1311.01	1111.55	1129.99	1090.92	1066.60	1080.32	1039.73	1074.88	1105.90
		Gap	13.05%	10.07%	7.57%	21.97%	3.41%	5.13%	1.49%	-0.77%	0.51%	-3.27%	0.00%	2.89%
		Time (s)	0.18	0.19	0.18	0.18	1.96	1.96	2.36	2.69	2.74	6.51	-	-
	20×10	Obj	1449.91	1435.71	1421.03	1640.16	1376.18	1379.80	1289.77	1336.97	1346.62	1261.80	1351.58	1406.68
		Gap	7.28%	6.22%	5.14%	21.35%	1.82%	2.09%	-4.57%	-1.08%	-0.37%	-6.64%	0.00%	4.08%
		Time (s)	0.25	0.26	0.26	0.25	2.63	2.60	3.14	3.80	3.78	10.50	-	-

Table 1: Results on synthetic instances of the same sizes as the instances used for training.

Size	PDRs				Greedy			Sample			CP-SAT		
	FIFO	MOR	MWKR	SPT	DAN	DAN-stoch	SPM-DAN	DAN	DAN-stoch	SPM-DAN			
SD ₁	30×10	Obj	412.05	410.85	410.41	458.12	386.09	371.58	355.85	376.92	365.36	352.16	379.06
		Gap	8.70%	8.39%	8.27%	20.86%	1.85%	-1.97%	-6.12%	-0.56%	-3.61%	-7.10%	0.00%
		Time (s)	0.40	0.41	0.41	0.40	4.16	4.15	4.93	46.80	33.40	56.67	-
	40×10	Obj	520.59	527.47	527.98	577.27	497.46	474.88	455.55	488.34	469.54	453.01	487.47
		Gap	6.79%	8.21%	8.31%	18.42%	2.05%	-2.58%	-6.55%	0.18%	-3.68%	-7.07%	0.00%
		Time (s)	0.58	0.58	0.58	0.57	5.58	5.44	6.70	42.20	41.88	89.22	-
SD ₃	30×10	Obj	1957.33	1997.42	1992.04	2251.80	1917.87	1915.07	1776.47	1882.99	1884.82	1775.37	1885.18
		Gap	3.83%	5.95%	5.67%	19.45%	1.73%	1.59%	-5.77%	-0.12%	-0.02%	-5.82%	0.00%
		Time (s)	0.41	0.41	0.41	0.40	3.95	3.95	4.76	6.17	6.40	22.26	-
	40×10	Obj	2496.87	2559.63	2572.53	2858.96	2474.38	2457.34	2294.33	2438.07	2432.05	2305.52	2442.27
		Gap	2.24%	4.81%	5.33%	17.06%	1.31%	0.62%	-6.06%	-0.17%	-0.42%	-5.60%	0.00%
		Time (s)	0.57	0.57	0.57	0.56	5.20	5.26	6.29	8.64	8.63	37.64	-

Table 2: Results on large synthetic instances using the policies trained on size 20×10.

seconds for SD₃ 20×10 per instance), but is still relatively fast and leads to clear performance improvements. For example, greedy inference using SPM-DAN considerably outperforms sampling inference using the baseline models for 20×10 instances, showcasing a better performance-runtime trade-off despite an initial runtime increase. The DAN-stoch models do not consistently outperform DAN, showcasing the value of synergizing our training procedure and SPM.

Table 2 shows the generalization of our policies to larger instances, using the 20×10 policies on 30×10 and 40×10 instances. We do not report CP-stoch, as memory limits are reached before finding sensible solutions for these instance sizes. In these larger instances, our method consistently and

considerably outperforms all baselines, even achieving a 6% improvement over CP-SAT and DAN using greedy sampling. DAN-stoch also beats DAN in these instances but remains far from SPM-DAN performance. These results show that SPM-DAN can scale to new instance sizes without the need for retraining.

Public Datasets Results

We assess the cross-distribution performance of our method on the public benchmark datasets, shown in Table 3. We report the policy trained on 15×10 SD₃ instances, but policies trained on 10×5 SD₃, 15×10 SD₁, and 10×5 SD₁ instances show similar results. We find that our method maintains per-

Dataset		PDRs				Greedy			Sample			CP-SAT	CP-stoch
		FIFO	MOR	MWKR	SPT	DAN	DAN-stoch	SPM-DAN	DAN	DAN-stoch	SPM-DAN		
mk	Obj	257.45	253.37	255.39	295.30	232.22	238.81	232.51	223.61	229.21	224.01	230.20	346.91
	Gap	11.84%	10.07%	10.94%	28.28%	0.88%	3.74%	1.00%	-2.86%	-0.43%	-2.69%	0.00%	50.70%
	Time (s)	0.18	0.18	0.18	0.18	1.98	1.94	2.37	14.89	14.96	21.71	-	-
rdata	Obj	1408.60	1397.46	1391.50	1568.94	1370.80	1366.73	1345.01	1318.70	1311.90	1294.62	1316.22	1549.72
	Gap	7.02%	6.17%	5.72%	19.20%	4.15%	3.84%	2.19%	0.19%	-0.33%	-1.64%	0.00%	17.74%
	Time (s)	0.19	0.19	0.19	0.19	2.12	2.07	2.46	16.06	16.08	23.98	-	-
edata	Obj	1586.08	1570.56	1560.18	1695.89	1534.73	1544.10	1520.25	1467.36	1483.50	1457.05	1423.01	1553.96
	Gap	11.46%	10.37%	9.64%	19.18%	7.85%	8.51%	6.83%	3.12%	4.25%	2.39%	0.00%	9.20%
	Time (s)	0.19	0.19	0.19	0.19	2.10	2.11	2.49	15.96	16.09	24.11	-	-
vdata	Obj	1295.16	1293.13	1286.46	1444.67	1281.67	1271.83	1227.05	1235.49	1225.49	1189.42	1300.97	1336.54
	Gap	-0.45%	-0.60%	-1.12%	11.05%	-1.48%	-2.24%	-5.68%	-5.03%	-5.80%	-8.57%	0.00%	2.73%
	Time (s)	0.20	0.20	0.19	0.19	2.06	2.05	2.52	16.19	16.11	24.27	-	-

Table 3: Results on benchmark instances using policies trained on SD_3 instances of size 15×10 .

Size		PDRs				Greedy			Sample			CP-SAT	CP-stoch
		FIFO	MOR	MWKR	SPT	DAN	DAN-stoch	SPM-DAN	DAN	DAN-stoch	SPM-DAN		
10×5	Obj	646.99	643.20	634.07	709.00	616.62	621.76	619.44	595.87	596.10	591.32	586.89	591.72
	Gap	10.24%	9.59%	8.04%	20.81%	5.07%	5.94%	5.55%	1.53%	1.57%	0.75%	0.00%	0.82%
	Time (s)	0.06	0.06	0.06	0.06	0.65	0.65	0.81	0.74	0.79	1.55	-	-
20×5	Obj	1172.19	1187.68	1181.30	1280.42	1169.24	1149.28	1142.68	1146.42	1145.32	1130.61	1120.84	1156.04
	Gap	4.58%	5.96%	5.39%	14.24%	4.32%	2.54%	1.95%	2.28%	2.18%	0.87%	0.00%	3.14%
	Time (s)	0.12	0.12	0.12	0.12	1.28	1.33	1.61	1.46	1.64	5.05	-	-
15×10	Obj	1074.56	1044.78	1021.75	1168.17	983.03	986.66	968.87	951.44	952.00	931.59	949.46	991.76
	Gap	13.18%	10.04%	7.61%	23.04%	3.54%	3.92%	2.04%	0.21%	0.27%	-1.88%	0.00%	4.46%
	Time (s)	0.18	0.19	0.18	0.18	1.96	2.00	2.43	2.69	2.96	11.19	-	-
20×10	Obj	1302.11	1288.43	1277.09	1485.58	1237.78	1231.46	1181.44	1209.97	1203.39	1171.56	1213.78	1286.98
	Gap	7.28%	6.15%	5.22%	22.39%	1.98%	1.46%	-2.66%	-0.31%	-0.86%	-3.48%	0.00%	6.03%
	Time (s)	0.25	0.26	0.25	0.26	2.63	2.68	3.20	3.80	4.25	18.96	-	-

Table 4: Results on synthetic SD_3 instances using the expected makespan objective.

formance across the benchmark instances. On the mk data, we see no performance improvement of our approach over the default DAN, caused by the previously mentioned instance limitations. On the other datasets, we maintain a clear performance improvement over the baseline neural methods and outperform CP-SAT up to 8% on vdata instances.

Flexibility to Different Objectives

We train and evaluate policies for optimizing the expected makespan to show that our method works for different objective functions. Table 4 shows the results of these policies for SD_3 instances. We see a similar pattern as for the $VaR_{95\%}$ where there is a limited effect for the smallest instances but our policies outperform the other methods toward larger problem sizes. The overall improvement is slightly less compared to the $VaR_{95\%}$ objective. This is expected as the expected value and deterministic values tend to be more related to each other than the $VaR_{95\%}$ and deterministic values as the $VaR_{95\%}$ is more dependent on the distributional range of processing times. However, the improvements of our method are still considerable and for the larger instances we improve over CP-SAT, as well as the other baselines.

Resilience to Varying Probability Distributions

An important criterion for stochastic optimization methods is that they perform under different assumed probability distributions. Therefore, we create one instance set with beta distributions (B) similar to (Flores-Gómez, Borodin, and Dauzère-Pérès 2023), one set with a mixture of log-normal and beta distributions (LB), and one set with a mixture of log-normal, beta, and gamma distributions (LBG). The results are shown in Table 5. It is clear that our method maintains performance across all instance sets. We see that DAN-stoch struggles to outperform DAN, as it cannot capture the dynamics that it observes through the rewards. Oppositely, our SPM-DAN model is aware of the types of stochasticity, which helps it distinguish different stochastic states. As a result, SPM-DAN achieves a gap compared to CP-SAT of 6% and improves performance up to 5% over the other models.

Trade-Off Performance and Number of Scenarios

We assess how n_{scn} affects the performance of SPM-DAN. To this end, we evaluate models trained with 100 (SPM-DAN) and 200 (SPM-DAN-200) scenarios using 50, 100, and 200 scenarios for inference. Table 6 shows that, at the cost of higher (lower) inference times, using a higher

Distribution		PDRs				Greedy			Sample			CP-SAT	CP-stoch
		FIFO	MOR	MWKR	SPT	DAN	DAN-stoch	SPM-DAN	DAN	DAN-stoch	SPM-DAN		
B	Obj	1267.59	1251.52	1236.12	1452.37	1186.82	1187.25	1161.30	1156.79	1159.41	1133.56	1170.30	1233.08
	Gap	8.31%	6.94%	5.62%	24.10%	1.41%	1.45%	-0.77%	-1.15%	-0.93%	-3.14%	0.00%	5.36%
	Time (s)	0.25	0.26	0.26	0.25	2.70	2.66	3.40	5.84	5.85	18.88	-	-
LB	Obj	1371.43	1356.26	1343.65	1559.23	1298.64	1303.83	1229.54	1254.58	1265.35	1194.14	1275.49	1322.15
	Gap	7.52%	6.33%	5.34%	22.25%	1.81%	2.22%	-3.60%	-1.64%	-0.79%	-6.38%	0.00%	3.66%
	Time (s)	0.25	0.26	0.26	0.25	2.77	2.70	3.25	5.84	5.81	19.17	-	-
LBG	Obj	1359.61	1344.19	1329.91	1546.49	1289.39	1278.59	1228.24	1247.24	1242.22	1197.75	1263.08	1321.16
	Gap	7.64%	6.42%	5.29%	22.44%	2.08%	1.23%	-2.76%	-1.25%	-1.65%	-5.17%	0.00%	4.60%
	Time (s)	0.25	0.26	0.26	0.25	2.74	2.66	3.25	5.50	5.87	19.17	-	-

Table 5: Results on synthetic SD_3 instances of size 20×10 using different processing time distributions.

Size		Greedy						Sample						CP-SAT	CP-stoch
		SPM-DAN			SPM-DAN-200			SPM-DAN			SPM-DAN-200				
		50	100	200	50	100	200	50	100	200	50	100	200		
10×5	Obj	719.14	710.63	709.58	713.64	709.69	712.11	674.51	677.24	672.29	673.36	671.75	675.07	688.10	676.19
	Gap	4.51%	3.27%	3.12%	3.71%	3.14%	3.49%	-1.98%	-1.58%	-2.30%	-2.14%	-2.38%	-1.89%	0.00%	-1.73%
	Time (s)	0.80	0.79	0.81	0.80	0.81	0.78	1.10	1.11	2.64	1.12	1.55	1.52	-	-
20×5	Obj	1283.09	1279.70	1277.45	1312.16	1307.93	1297.82	1242.07	1243.75	1238.65	1268.96	1265.04	1266.04	1253.67	1271.71
	Gap	2.35%	2.08%	1.90%	4.67%	4.33%	3.52%	-0.93%	-0.79%	-1.20%	1.22%	0.91%	0.99%	0.00%	1.44%
	Time (s)	1.65	1.58	1.61	1.66	1.62	1.57	3.20	2.99	8.72	3.19	5.03	4.87	-	-
15×10	Obj	1084.99	1090.92	1077.59	1092.26	1088.52	1089.67	1035.24	1039.73	1035.22	1041.44	1040.63	1044.55	1074.88	1105.90
	Gap	0.94%	1.49%	0.25%	1.62%	1.27%	1.38%	-3.69%	-3.27%	-3.69%	-3.11%	-3.19%	-2.82%	0.00%	2.89%
	Time (s)	2.41	2.36	2.47	2.43	2.47	2.37	6.64	6.51	16.50	6.65	11.22	11.17	-	-
20×10	Obj	1290.38	1289.77	1285.83	1296.15	1295.71	1285.53	1261.14	1261.80	1256.07	1262.78	1259.25	1259.41	1351.58	1406.68
	Gap	-4.53%	-4.57%	-4.86%	-4.10%	-4.13%	-4.89%	-6.69%	-6.64%	-7.07%	-6.57%	-6.83%	-6.82%	0.00%	4.08%
	Time (s)	3.24	3.14	3.33	3.24	3.30	3.17	11.10	10.50	28.22	11.09	19.09	18.61	-	-
30×10	Obj	1783.31	1776.47	1776.10	1788.87	1778.37	1775.00	1775.52	1775.37	1768.42	1765.85	1759.35	1754.91	1885.18	-
	Gap	-5.40%	-5.77%	-5.79%	-5.11%	-5.67%	-5.84%	-5.82%	-5.82%	-6.19%	-6.33%	-6.67%	-6.91%	0.00%	-
	Time (s)	4.89	4.76	5.22	4.98	4.86	4.95	23.40	22.26	62.53	23.40	33.93	61.56	-	-
40×10	Obj	2297.15	2294.33	2289.39	2303.53	2292.60	2286.71	2309.34	2305.52	2299.30	2283.10	2276.22	2269.88	2442.27	-
	Gap	-5.94%	-6.06%	-6.26%	-5.68%	-6.13%	-6.37%	-5.44%	-5.60%	-5.85%	-6.52%	-6.80%	-7.06%	0.00%	-
	Time (s)	6.58	6.29	7.00	6.66	6.67	7.18	40.06	37.64	107.75	40.06	58.94	107.20	-	-

Table 6: Results on synthetic SD_3 instances of models trained with different n_{scn} using various numbers of n_{scn} for inference.

(lower) n_{scn} to generate solutions generally improves (decreases) performance. The differences in performance gap are roughly between 0.5% and 1%. However, training using more scenarios does not per se lead to better performance, especially for smaller instances. Namely, SPM-DAN with 200 evaluation scenarios outperforms SPM-DAN-200 in some instance sets. In larger instances, the performance difference between SPM-DAN and SPM-DAN-200 is limited, although SPM-DAN-200 performs slightly better. Hence, computational requirements can be decreased by training with fewer scenarios at the cost of a limited potential performance loss. Then, in inference, n_{scn} can be varied to trade off runtime and performance. The increased runtime of using a higher n_{scn} is also limited by increasing parallelism.

Conclusion

This study proposes the scenario processing module SPM, a novel attention-based module to extend NCO methods to stochastic problems. SPM captures embeddings from sampled scenarios and feeds those to the base neural network,

allowing it to learn effective policies under uncertainty. We integrate SPM into a training procedure that works well with different stochastic objectives and apply it to the stochastic FJSP, leading to SPM-DAN. We experimentally show that SPM-DAN outperforms existing learning and non-learning methods. SPM-DAN works well across a range of synthetic and benchmark instances while handling various processing time distributions. Future work may focus on developing and adapting sampling techniques for stochastic optimization to enhance the performance of constructive policies. Moreover, integrating data from historical schedules with their realizations or contextual factors influencing processing times is a worthwhile avenue for exploration.

Acknowledgements

The LEO (Learning and Explaining Optimization) project is co-funded by Holland High Tech | TKI HSTM via the PPS allowance scheme for public-private partnerships. This work used the Dutch national e-infrastructure with the support of the SURF Cooperative using grant no. EINF-10518.

References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. arXiv:1607.06450.
- Baptiste, P.; Pape, C. L.; and Nuijten, W. 2001. *Constraint-based scheduling: applying constraint programming to scheduling problems*, volume 39. Springer Science & Business Media. ISBN 0792374088.
- Bengio, Y.; Lodi, A.; and Prouvost, A. 2021. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 290(2): 405–421.
- Brandimarte, P. 1993. Routing and scheduling in a flexible job shop by tabu search. *Annals of Operations Research*, 41: 157–183.
- Burdett, R. L.; and Kozan, E. 2018. An integrated approach for scheduling health care activities in a hospital. *European Journal of Operational Research*, 264: 756–773.
- Caldeira, R. H.; and Gnanavelbabu, A. 2021. A simheuristic approach for the flexible job shop scheduling problem with stochastic processing times. *SIMULATION*, 97(3): 215–236.
- Col, G. D.; and Teppan, E. C. 2022. Industrial-size job shop scheduling with constraint programming. *Operations Research Perspectives*, 9: 100249.
- Flores-Gómez, M.; Borodin, V.; and Dauzère-Pérès, S. 2023. Maximizing the service level on the makespan in the stochastic flexible job-shop scheduling problem. *Computers & Operations Research*, 157: 106237.
- Ghaedy-Heidary, E.; Nejati, E.; Ghasemi, A.; and Torabi, S. A. 2024. A simulation optimization framework to solve Stochastic Flexible Job-Shop Scheduling Problems—Case: Semiconductor manufacturing. *Computers & Operations Research*, 163: 106508.
- Hurink, J.; Jurisch, B.; and Thole, M. 1994. Tabu search for the job-shop scheduling problem with multi-purpose machines. *OR Spektrum*, 15: 205–215.
- Infantes, G.; Roussel, S.; Pereira, P.; Jacquet, A.; and Benazera, E. 2024. Learning to Solve Job Shop Scheduling Under Uncertainty. In Dilkina, B., ed., *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 329–345. Cham: Springer Nature Switzerland. ISBN 978-3-031-60597-0.
- Joe, W.; and Lau, H. C. 2020. Deep Reinforcement Learning Approach to Solve Dynamic Vehicle Routing Problem with Stochastic Customers. *Proceedings of the International Conference on Automated Planning and Scheduling*, 30(1): 394–402.
- Kwon, Y.-D.; Choo, J.; Kim, B.; Yoon, I.; Gwon, Y.; and Min, S. 2020. POMO: Policy Optimization with Multiple Optima for Reinforcement Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 21188–21198. Curran Associates, Inc.
- Kwon, Y.-D.; Choo, J.; Yoon, I.; Park, M.; Park, D.; and Gwon, Y. 2021. Matrix encoding networks for neural combinatorial optimization. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 5138–5149. Curran Associates, Inc.
- Lee, J.; Lee, Y.; Kim, J.; Kosiorek, A.; Choi, S.; and Teh, Y. W. 2019. Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 3744–3753. PMLR.
- Lei, K.; Guo, P.; Wang, Y.; Zhang, J.; Meng, X.; and Qian, L. 2023. Large-scale dynamic scheduling for flexible job-shop with random arrivals of new jobs by hierarchical reinforcement learning. *IEEE Transactions on Industrial Informatics*.
- Lei, K.; Guo, P.; Zhao, W.; Wang, Y.; Qian, L.; Meng, X.; and Tang, L. 2022. A multi-action deep reinforcement learning framework for flexible Job-shop scheduling problem. *Expert Systems with Applications*, 205: 117796.
- Park, J.; Chun, J.; Kim, S. H.; Kim, Y.; and Park, J. 2021. Learning to schedule job-shop problems: representation and policy learning using graph neural network and reinforcement learning. *International Journal of Production Research*, 59(11): 3360–3377.
- Perron, L.; Didier, F.; and Gay, S. 2023. The CP-SAT-LP Solver. In Yap, R. H. C., ed., *29th International Conference on Principles and Practice of Constraint Programming (CP 2023)*, volume 280 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 3:1–3:2. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-300-3.
- Reijnen, R.; van Straaten, K.; Bukhsh, Z.; and Zhang, Y. 2023. Job Shop Scheduling Benchmark: Environments and Instances for Learning and Non-learning Methods. arXiv:2308.12794.
- Rooyani, D.; and Defersha, F. M. 2019. An Efficient Two-Stage Genetic Algorithm for Flexible Job-Shop Scheduling. *IFAC-PapersOnLine*, 52: 2519–2524.
- Schmitt-Ulms, F.; Hottung, A.; Sellmann, M.; and Tierney, K. 2022. Learning to Solve a Stochastic Orienteering Problem with Time Windows. In Simos, D. E.; Rasskazova, V. A.; Archetti, F.; Kotsireas, I. S.; and Pardalos, P. M., eds., *Learning and Intelligent Optimization*, 108–122. Cham: Springer International Publishing. ISBN 978-3-031-24866-5.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347.
- Sels, V.; Gheysen, N.; and Vanhoucke, M. 2012. A comparison of priority rules for the job shop scheduling problem under different flow time- and tardiness-related objective functions. *International Journal of Production Research*, 50: 4255–4270. Doi: 10.1080/00207543.2011.611539.
- Smit, I. G.; Zhou, J.; Reijnen, R.; Wu, Y.; Chen, J.; Zhang, C.; Bukhsh, Z.; Zhang, Y.; and Nuijten, W. 2025. Graph Neural Networks for Job Shop Scheduling Problems: A Survey. *Computers & Operations Research*, 176: 106914.

- Song, W.; Chen, X.; Li, Q.; and Cao, Z. 2022. Flexible Job-Shop Scheduling via Graph Neural Network and Deep Reinforcement Learning. *IEEE Transactions on Industrial Informatics*, 19: 1600–1610.
- Tamssaouet, K.; Dauzère-Pérès, S.; Knopp, S.; Bitar, A.; and Yugma, C. 2022. Multiobjective optimization for complex flexible job-shop scheduling problems. *European Journal of Operational Research*, 296: 87–100.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. arXiv:1710.10903.
- Wang, R.; Wang, G.; Sun, J.; Deng, F.; and Chen, J. 2023. Flexible job shop scheduling via dual attention network-based reinforcement learning. *IEEE Transactions on Neural Networks and Learning Systems*.
- Zhang, C.; Song, W.; Cao, Z.; Zhang, J.; Tan, P. S.; and Chi, X. 2020. Learning to Dispatch for Job Shop Scheduling via Deep Reinforcement Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 1621–1632. Curran Associates, Inc.
- Zhang, J.; Yang, J.; and Zhou, Y. 2016. Robust scheduling for multi-objective flexible job-shop problems with flexible workdays. *Engineering Optimization*, 48: 1973–1989. Doi: 10.1080/0305215X.2016.1145216.
- Zhang, L.; Feng, Y.; Xiao, Q.; Xu, Y.; Li, D.; Yang, D.; and Yang, Z. 2023. Deep reinforcement learning for dynamic flexible job shop scheduling problem considering variable processing times. *Journal of Manufacturing Systems*, 71: 257–273.