

# Beyond Training: Optimizing Reinforcement Learning Based Job Shop Scheduling Through Adaptive Action Sampling

Constantin Waubert de Puiseau, Christian Dörpelkus, Jannik Peters, Hasan Tercan, Tobias Meisen

Institute for Technologies and Management of Digital Transformation  
Lise-Meitner-Strasse 27-31  
42119 Wuppertal  
{waubert, christian.doerpelkus, jpeters, tercan, meisen}@uni-wuppertal.de

## Abstract

Learned construction heuristics for scheduling problems have become increasingly competitive with established solvers and heuristics in recent years. In particular, significant improvements have been observed in solution approaches using deep reinforcement learning (DRL). While much attention has been paid to the design of network architectures and training algorithms to achieve state-of-the-art results, little research has investigated the optimal use of trained DRL agents during inference. Our work is based on the hypothesis that, similar to search algorithms, the utilization of trained DRL agents should be dependent on the acceptable computational budget. We propose a simple yet effective parameterization, called  $\delta$ -sampling that manipulates the trained action vector to bias agent behavior towards exploration or exploitation during solution construction. By following this approach, we can achieve a more comprehensive coverage of the search space while still generating an acceptable number of solutions. In addition, we propose an algorithm for obtaining the optimal parameterization for such a given number of solutions and any given trained agent. Experiments extending existing training protocols for job shop scheduling problems with our inference method validate our hypothesis and result in the expected improvements of the generated solutions.

## Introduction

Research on self-learning algorithms for scheduling problems has significantly increased in recent years, driven primarily by advancements in algorithms and the availability of more affordable and powerful computing units in both industry and academia. In academia, the standardized job shop scheduling problem (JSSP) (Pinedo 2016) is the dominant object of study for its straightforward logical constraints and comparability thanks to public benchmarks and libraries (van Hoorn 2018; Waubert de Puiseau et al. 2023). While in some successful DRL-based solution methods agents steer existing search algorithms (Zhang et al. 2022; Ni et al. 2021) and solvers (Tassel, Gebser, and Schekotihin 2023), in a majority, agents autonomously generate solution schedules imitating a construction heuristic (Zhang et al. 2020; Park et al. 2021; van Ekeris, Meyes, and Meisen 2021; Tassel et al. 2022).

In these DRL-based construction heuristics, agents make iterative decisions to schedule the next unscheduled operation until the entire schedule is complete. For each decision, an action is sampled from the agent’s action vector that consists of logits representing the current operation prioritization. Depending on how actions are sampled from the vector, different solutions for a JSSP are obtained. For our purposes, we differentiate between the following sampling strategies:

- **Deterministic sampling:** by greedily applying *argmax* policy logits we obtain one deterministic solution per problem instance.
- **Stochastic sampling:** sampling from the given distribution, where the likelihood of each action is given by the corresponding logit.
- **Exploitative/Explorative sampling:** during exploitative sampling, the likelihoods of likely actions increase even more; during explorative sampling, the likelihoods of all actions are assimilated. We achieve both through the  $\delta$ -sampling which we present in this paper.

In their extremes, exploitative sampling converges towards deterministic sampling whereas explorative sampling converges towards sampling from a uniform distribution. Stochastic sampling falls in between. In the context of scheduling, it has only very recently been shown that stochastic sampling can produce better-than-greedy schedules even with a relatively small number of sampled solutions (Iklassov et al. 2022). Motivated by this finding, our goal is to identify the optimal sampling method for any given trained agent and the desirable number of samples to be evaluated. During inference, the objective is to sample in such a way that among all sampled solutions, the completion times  $C_{\max}$ , also known as makespans, the shortest is as short as possible. To avoid confusion with the common subscript *max*, and for brevity, we refer to this minimum value as  $C^*$  throughout this paper. According to the proposed nomenclature, we formulate our working hypothesis as follows:

1. The expected  $C^*$  is a function of the sampling strategy and number of sampled solutions, i.e., the sample size.
2. For a fixed (non-deterministic) sampling strategy,  $C^*$  remains equal or decreases with increasing sample size.

3. For a small sample size, strongly exploitative sampling strategies statistically lead to lower  $C^*$  compared to very explorative strategies. In other words, it is unlikely to find a local and better-than-greedy minimum when deviating more strongly from the greedy solution.
4. Conversely, for a large sample size, strongly explorative sampling strategies statistically lead to better  $C^*$  compared to exploitative strategies, since a wider solution space is covered.

Figure 1 summarizes our hypothesis. Moreover, if the hypothesis is valid, there should exist a sampling strategy that minimizes the expected  $C^*$  for a given agent and sample size.

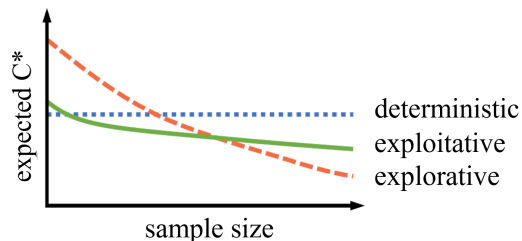


Figure 1: Expected minimal makespans  $C^*$  over sampling size for different sampling strategies

Building upon our hypothesis, we introduce a new sampling method and validate its effectiveness on agents trained with two different state-of-the-art training protocols, testing them on JSSP instances of different sizes and with different sample sizes. Our study leads to the following valuable contributions:

1. A new simple but effective tree-search-like solution generation algorithm for JSSP based on action sampling
2. An algorithm for the identification of the most effective parameterizations of action sampling methods given a trained agent and a number of samples
3. A validation of our algorithm on various JSSP instances of different sizes as well as on benchmark instances of Taillard (1993).

The remainder of this paper is structured as follows: After discussing related work, we describe our methods and experiments. We then present and discuss the experimental results, before drawing conclusions and giving an outlook on future work.

## Related Work

**DRL-based construction heuristics.** Early successful attempts to solve combinatorial optimization problems with DRL in 2016 (Bello et al. 2017) were followed by a growing wave of research tackling the JSSP (Panzer, Bender, and Gronau 2022). Recent research focuses on effective problem size agnostic neural network architectures, including graph embeddings (Zhang et al. 2020; Park et al. 2021) and recurrent networks (Iklassov et al. 2022). Another trend is the use of curriculum learning (Iklassov et al. 2022; Waubert de

Puiseau, Tercan, and Meisen 2023). Interestingly, only Iklassov et al. (2022) sample multiple solutions stochastically, although the possibility to obtain better solutions through sampling in combinatorial problems was already established by Kool, van Hoof, and Welling (2019). However, to the best of our knowledge, our study is the first to focus on improving sampling methods for generating JSSP solutions using trained agents.

**Policy-guided tree search.** Stochastic sampling from a learned policy is a common concept to guide tree search algorithms (Orseau et al. 2018). The most common search algorithm for this integration is Monte-Carlo tree search (MCTS). Famously, this combination was leveraged in AlphaGo (Silver et al. 2016) and its successors. Therein, the expansion of nodes is determined by a linear combination of learned Q-values or action predictions and the visit count of states. Lately, the approach has been successfully transferred to scheduling as well (Oren et al. 2021; Kumar and Dimitrakopoulos 2021). However, MCTS comes at a high computational expense that is accepted in many cases because it provides good expected values even in stochastic environments that are infeasible to approximate analytically. For the deterministic JSSP, however, the computational budget may be used more efficiently by sampling multiple whole trajectories depth-wise from start to finish, as done in this study.

**Solution space coverage.** The question of how closely to adhere to a learned policy is captured in the exploration-exploitation tradeoff. However, this tradeoff is usually exclusively considered during training. Therefore, methods like epsilon-greedy sampling (Sutton and Barto 2018) or added loss terms on the action confidence through entropy or curiosity terms (Schulman et al. 2017; Deepak Pathak et al. 2017) do not transfer to inference. A noteworthy technique to fine-tune exploration during inference in MCTS is to use an exponent across the action vector (Wang, Hao, and Cao 2021; Shen et al. 2018). The primary rationale is a de- or increased effect of action priors compared to visit counts, but a secondary effect is a (dis-)alignment of values in the action vector itself. Our method is inspired by this idea and leverages the secondary effect.

For combinatorial optimization problems, depth-wise search with DRL is mostly done by stochastic policy sampling (Iklassov et al. 2022; Kool, van Hoof, and Welling 2019; Orseau et al. 2018). One exception is beam-search, which ensures a fixed tree width  $k$  during search (Vinyals, Fortunato, and Jaitly 2015). Known extensions make use of problem symmetries (Kwon et al. 2020) and entropy in the action vector for the solution space expansion (Tassel et al. 2022). Iklassov et al. (2022) compared stochastic sampling and beam-search variations, finding that stochastic sampling outperforms the other methods. However, they only reported results on a single sample size. We expand this study with a new sampling method which fine-tunes sampling for each combination of trained agent and desired sample size.

## Method and Experiments

### $\delta$ -Sampling

Our proposed sampling strategy offers the possibility to adjust the balance between exploitative and explorative sampling. To this end, we manipulate the action vector  $a$  predicted by a trained agent to obtain  $a_{sample}$  by exponentiating the predicted logits with the discrimination exponent  $\delta$  and then re-scale the vector as in equation 1.

$$a_{sample} = \frac{1}{\sum_i a_i^\delta} \cdot a^\delta \quad (1)$$

The effect is that for  $\delta > 1$ , relative differences between logits are emphasized favoring exploitation, whereas for  $\delta < 1$  differences are decreased favoring exploration. For  $\delta = 1$ , logits remain unchanged, defaulting to stochastic sampling. We call this procedure  $\delta$ -sampling. It easily scales to any problem size and allows to precisely adjust the sampling behavior.

Given a base model and sample size, we are looking for  $\delta^*$  that minimizes  $C^*$ . We expect the functions of  $C^*(\delta^*)$  to be smooth and have a single minimum based on our hypothesis. Hence, our algorithm starts with a grid search that iteratively decreases the distance between the two most promising candidate values for  $\delta$  to neighboring candidates by 50%. An example is illustrated in Figure 2. In the first iteration ( $i=0$ ), we sample with each candidate  $\delta$  from a list of candidates and the desired sample size on validation JSSP instances  $i_\delta$ . The initial list of candidates is a hyperparameter that has to be set. Sampling returns the average of  $C^*$  across all  $i_\delta$  per candidate. Then, additional candidates are added between the two most promising and their adjacent candidates, rounding to two decimal places. If the most promising candidates lie at minimum or maximum  $\delta$  values, the initial interval is extended by half of the largest distance between adjacent candidates. The function repeats until a specified number of iterations or minimal distance between adjacent candidates is reached.




Figure 2: Three iterations of searching the optimal  $\delta$  value

### Experimental Setup

**JSSP formalization** This work deals with the common JSSP formalization in which  $J$  jobs, consisting of  $O$  operations are processed on  $M$  machines. Each job visits each machine exactly once in a predefined order, such that  $O = M$ . Additionally, each operation has a specified processing time and may not be interrupted once started (no preemption). Machines can only process one job at once (no overlap) and

setup-times are ignored. The notation of problem sizes follows the structure  $J \times M$ , e.g., a  $20 \times 15$  instance consists of 20 jobs á 15 operations on 15 machines.

**Base Models.** We perform experiments on two base model designs that are adapted from Zhang et al. (2020) and Iklassov et al. (2022), two competitive and often-referenced representatives of DRL-based construction heuristics. We refer to the base models as *L2D* (Zhang et al. 2020) and *L2G* (Iklassov et al. 2022) in reference to their publication titles. From *L2D*, we can extract the openly accessible pre-trained graph neural network weights for each respective problem size  $6 \times 6$ ,  $15 \times 15$  and  $20 \times 20$ , except for  $100 \times 20$ , where we resort to the  $30 \times 20$  model as done in the original publication. The second base model, *L2G*, features a neural network architecture that uses a set2set recurrent layer (Vinyals, Bengio, and Kudlur 2016) to be problem size agnostic. Since no trained models are publicly available, we retrain our base models according to the reported protocol, one for each problem size, with the following curricula: for the  $15 \times 15$  and  $6 \times 6$  JSSP models, we choose adaptive curriculum learning and train on  $6 \times 6$  and  $15 \times 15$  JSSP instances. For the  $20 \times 20$  and  $100 \times 20$  base model, we use the RASCL method (Iklassov et al. 2022) and train on  $6 \times 6$ ,  $15 \times 15$ , and  $20 \times 20$  JSSP instances.

**JSSP Instances.** We test our sampling method on four problem sizes.  $15 \times 15$ ,  $20 \times 20$ , and  $100 \times 20$  are common benchmark sizes in the operations research domain.  $6 \times 6$  is also common in the literature since it is useful for the visual interpretations of results in Gantt charts. To avoid overfitting on test data, we create separate JSSP instances for training and for the search of suitable sampling parameterizations. These Taillard-like instances are created with separate random seeds, assigning random perturbations of machine sequences to each job and sampling processing times uniformly from the interval  $[1, 99]$ .

## Results

### Hypothesis Validation

To empirically test our hypotheses illustrated in Figure 1, we create 10,000 solutions ( $C_{10k}^*$ ) on ten  $6 \times 6$  JSSP instances with our *L2G* base model and choose four sampling parameterizations:

1. random sampling, which draws actions from a uniform distribution;
2.  $\delta = 1$  as baseline of how sampling is usually performed;
3.  $\delta = 0.05$  as extreme example of explorative sampling;
4.  $\delta = 10$  as extreme example of exploitative sampling.

Figure 3 displays the averaged minimal makespan achieved over different sample sizes (logarithmic scale). To make up for statistical outliers in small sample sizes, we average the minimal makespans achieved over multiple sampling runs. Using Python array slicing notation, this can be formulated as in Equation 2,

$$\frac{1}{10,000/s} \sum_{i=0}^{10,000/s} \min(C_{10k}^*[i : s(i+1)]), \quad (2)$$




Figure 3: Minimal makespans over sampling size for different sampling strategies

base	problem size	sample size		
		32	128	512
L2G	6x6	0.30	0.10	0.11
	15x15	0.70	0.86	0.43
	20x20	1.25	1.46	0.80
	100x20	0.58	0.43	n.a.
L2D	6x6	2.06	1.37	0.93
	15x15	12.25	6.02	3.32
	20x20	10.35	10.38	10.05
	100x20	14.0	10.88	n.a.

Table 1: Best found delta values for different base models, problem sizes and sample sizes

where  $s$  is the sample size.

In Figure 3, we divide the plot into four zones at interesting line intersections. In zone A, which in this example only includes a sample size of one, strongly exploitative sampling yields the best solutions. In zones B and C, stochastic sampling with  $\delta = 1$  yields better makespans than strongly exploitative or explorative sampling. For sampling sizes  $< 50$ , strongly explorative sampling with  $\delta = 0.05$  performs better than stochastic sampling. The switch from zone B to C indicates, where explorative sampling outperforms exploitative sampling. The results in Figure 3 thereby verify our hypothesis.

### Optimal $\delta$ -Values

In the following we present the results of finding and using the optimal parameterizations of our action sampling approach, adhering to the sample size of 128 used in (Iklavov et al. 2022) for comparison and adding 32 and 512 as significantly smaller and larger exemplary sample sizes. We omitted experiments on 100x20 JSSP with 512 samples, since we expect little additional insights compared to the considerable required computational times for solving large instances. Table 1 shows the results obtained from our parameter search algorithm per base model, problem size, and sample size.

The results show two general trends. Firstly, in accordance with the above hypothesis, larger  $\delta$ -values are favorable for smaller sample sizes. Secondly, more exploitative sampling seems favorable with increasing problem size.

However, we also observe exceptions to these general trends. For *L2G* and problem sizes 15x15 and 20x20, there is a small increase of found  $\delta$ -values from sample size 32 to 128, before they decrease again for 512. Also for *L2G*, the trend of increasing  $\delta$ -values with increasing problem sizes discontinues for problem size 100x20. Overall, the found  $\delta$ -values of *L2G* are smaller in comparison to *L2D*. These discontinuing trends, exceptions, and differences between base models lead us to believe that analytical approaches to finding  $\delta$ -values are intractable and they underscore the necessity to determine the sampling parameterization empirically and individually for any trained model and sample size.

For a deeper analysis, we examine the function  $C^*(\delta)$  which was tested on the evaluation instances during  $\delta$ -value search. Noteworthy representative examples of this function are presented in Figure 4. Across all examples, the functions satisfy the assumed smoothness and the hypothesis of a single minimum is valid, although some noise exists. Our iterative algorithm to find these minima converges towards reasonable minima, which are marked by the green triangles. Note that the small local minima, e.g. in Figure 4a) at  $\delta \approx 0.8$ , could render gradient-based methods ineffective. However, as Figure 4a) also exemplifies,  $\delta$  can have a minimal overall effect, and in cases such as Figure 4b), the function plateaus across large  $\delta$  ranges. In light of these observations it is likely that the rare exceptions to decreasing  $\delta$ -values with increasing sample size in *L2G* in Table 1, which were discussed above, are relics of statistical noise.

### Performance Improvements

**Test Instances.** The effectiveness of the obtained parameterization is tested on 100 test instances. The average  $C^*$  values are presented in Table 2. We report the results for  $\delta$ -sampling (*ours*), using stochastic sampling as in Iklavov et al. (2022) (*stochastic*) as a baseline, both for three sample sizes, and the deterministic solution (*deterministic*). Relative improvements between *ours* and stochastic sampling ( $(\text{stochastic} - \text{ours})/\text{stochastic}$ ) are depicted in the last column.

Both sampling methods find significantly better solutions than the deterministic solution, confirming previous results (Iklavov et al. 2022; Kool, van Hoof, and Welling 2019; Orseau et al. 2018). Our parameterized method outperforms stochastic sampling in 17 out of 22 cases with up to 3.2% improvement. Notably, improvements are achieved across all tested sample sizes. The most notable exception is  $\delta$  sampling for the *L2D* base model on the 6x6 JSSP, which we partially attribute to statistical uncertainty given the small sample size of 32. An exception that persists across sample sizes is *L2G* on the 20x20 JSSP, where  $\delta$ -sampling performs slightly worse than the stochastic baseline. Overall, we observe larger performance differences where  $\delta$  deviates more from one, i.e. stochastic sampling. This implies that our parameter search is effective in finding stronger manipulations where they are useful. On the other hand, constellations where  $\delta$ -values are close to one and reside within plateaus of the function  $C^*(\delta)$ , as shown in Figure 4 c) and d), only minor performance differences can be obtained, as observed for *L2G* on the 20x20 JSSP.

sample size	base	problem size	ours	stochastic	deterministic	improvement		
32	L2G	6x6	496.0	505.5	517.0	1.9%		
		15x15	1310.4	1311.4	1364.5	0.1%		
		20x20	1782.9	1782.3	1866.2	0.0%		
		100x20	5759.6	5764.1	5867.3	0.1%		
	L2D	6x6	519.3	516.1	577.1	-0.6%		
		15x15	1406.9	1434.6	1534.7	1.9%		
		20x20	1891.4	1957.4	2012.5	3.4%		
		100x20	5865.8	5928.4	5996.7	1.1%		
	128	L2G	6x6	491.7	503.6	517.0	2.4%	
			15x15	1300.4	1300.9	1364.5	0.0%	
20x20			1769.3	1767.6	1866.2	-0.1%		
100x20			5733.6	5745.1	5867.3	0.2%		
L2D		6x6	504.8	506.2	577.1	0.3%		
		15x15	1385.7	1410.1	1534.7	1.7%		
		20x20	1865.8	1927.6	2012.5	3.2%		
		100x20	5869.8	5899.7	5996.7	0.5%		
512		L2G	6x6	489.9	501.6	517.0	2.3%	
			15x15	1281.6	1291.7	1364.5	0.8%	
	20x20		1752.0	1750.6	1866.2	-0.1%		
	L2D	6x6	498.0	498.7	577.1	0.1%		
		15x15	1369.1	1386.9	1534.7	1.3%		
		20x20	1842.2	1879.6	2012.5	2.0%		

Table 2: Average  $C^*$  values on 100 generated test instances per problem size

s.-size	base	p.-size	ours	stochastic	deterministic	improvement	optimal	
			$C^*$ (opt. gap)	$C^*$ (opt. gap)	$C^*$ (opt. gap)	ours vs. stoch.	C	
32	L2G	15x15	1345.9 (9.5%)	1345.1 (9.5%)	1417.4 (15.3%)	-0.1%		1228.9
		20x20	1858.9 (14.9%)	1858.7 (14.9%)	1935.3 (19.7%)	0.0%		1617.3
		100x20	5759.7 (7.3%)	5784.6 (7.8%)	5878.4 (9.6%)	0.4%		5365.7
	L2D	15x15	1452.5 (18.2%)	1465.9 (19.3%)	1530.5 (24.5%)	0.9%		1228.9
		20x20	1953.2 (20.8%)	2027.1 (25.3%)	2081.8 (28.7%)	3.6%		1617.3
		100x20	5928.6 (10.5%)	5983.1 (11.5%)	6089 (13.5%)	0.9%		5365.7
128	L2G	15x15	1340.3 (9.1%)	1340.9 (9.1%)	1417.4 (15.3%)	0.0%		1228.9
		20x20	1846.4 (14.2%)	1837.4 (13.6%)	1935.3 (19.7%)	-0.5%		1617.3
		100x20	5751.1 (7.2%)	5766.5 (7.5%)	5878.4 (9.6%)	0.3%		5365.7
	L2D	15x15	1416.2 (15.2%)	1442 (17.3%)	1530.5 (24.5%)	1.8%		1228.9
		20x20	1943.4 (20.2%)	2002.7 (23.8%)	2081.8 (28.7%)	3.0%		1617.3
		100x20	5888.2 (9.7%)	5930.4 (10.5%)	6089 (13.5%)	0.7%		5365.7
512	L2G	15x15	1323.1 (7.7%)	1333.7 (8.5%)	1417.4 (15.3%)	0.8%		1228.9
		20x20	1825.5 (12.9%)	1828.9 (13.1%)	1935.3 (19.7%)	0.2%		1617.3
	L2D	15x15	1409.0 (14.7%)	1427.6 (16.2%)	1530.5 (24.5%)	1.3%		1228.9
		20x20	1895.3 (17.2%)	1946.2 (20.3%)	2081.8 (28.7%)	2.6%		1617.3

Table 3: Results on Taillard Instances

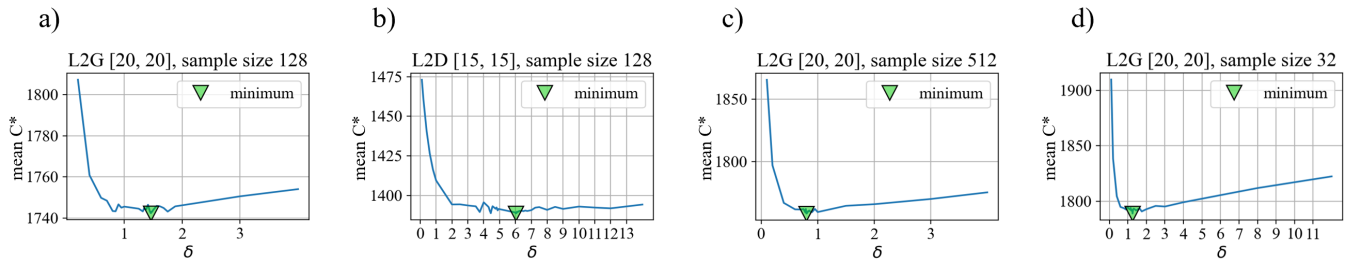


Figure 4: Examples of results in the  $\delta$  value search algorithm with found minima




Figure 5: Comparison of sampling methods over sample sizes

To understand those cases where very similar average  $C^*$  are generated, we analyze the resulting schedule differences. Within the schedules of *L2G* on the 15x15 JSSP with 128 samples, the most striking example, we counted in how many cases the same schedules, measured by the same  $C^*$ , were found per instance. Surprisingly,  $C^*$  was equal in only 4/100 test instances. This shows that even though the average values of  $C^*$  were less than 0.00% apart, different solutions were found, such that  $\delta$  still had a significant impact. In contrast, for 6x6 JSSPs of *L2G* with a sample size of 128, we obtained the same  $C^*$  in roughly one third of the instances. This is most likely an effect of the smaller total solution space compared to the other problem sizes. The conclusion is that even if only a few different solutions are found, they can still have a significant impact, as the 1.9% improvement demonstrates. Therefore, our sampling method is effective for small problem sizes with smaller solution spaces, too.

**Benchmark Instances.** We further report the results on benchmark instances of Taillard (1993), which we abbreviate as TA instances, in Table 3. In comparison, there are only ten instances per size, but the optimal solutions are known, enabling us to additionally provide the relative gap (*opt. gap*) for reference. The effects of sampling methods are fundamentally similar. As before, both sampling approaches outperform the deterministic solution. While the results on *L2D* indicate even greater improvements through our sampling method, the results on *L2G* are more ambiguous. We attribute the differences of the above results to the fact that only ten TA instances exist per problem size, introducing a larger statistical error. Note that overall we find slightly larger makespans than those reported in (Iklassov et al. 2022), as we could not reproduce the reported perfor-

mance on benchmarks with our trained model. However, our contribution is a relative improvement per model and sample size applicable to any learned construction heuristic, not the creation of a new learning strategy or network architecture. For details on the most recent absolute results and a comprehensive comparison of learning-based, priority-rule-based or heuristic methods, we refer the interested reader to (Iklassov et al. 2022), (Corsini et al. 2024) and (Falkner et al. 2022).

In Figure 5, relative improvements on TA instances through our adjusted sampling method can be compared to improvements achievable by increasing sample sizes. In many cases, we find that the achieved  $C^*$  through  $\delta$ -sampling is similar or better than the  $C^*$  achieved via stochastic sampling with a four times larger sample size. In practical terms, the same results can be achieved with a quarter of the computational budget through better sampling in these cases, rendering  $\delta$ -sampling four times more efficient.

## Conclusion and Outlook

In this paper, we proposed a sampling method for the more effective usage of trained DRL-based construction heuristics for the JSSP. We showed that an optimal balance between exploration and exploitation during inference exists, which is specific to the trained model and the chosen sample size. Based on this insight, a method was developed that allows the tradeoff between exploration and exploitation to be easily and smoothly parameterized during solution generation of learned construction heuristics. The effectiveness of this method was experimentally evaluated and showed promising results across different base models and computational budgets on common JSSP sizes. We believe that our approach is neither limited to the studied base models nor the

application to the JSSP, but can be used to improve any other learned construction heuristic on both the JSSP and other related combinatorial optimization problems.

In the future, we would like to transfer the insights gained into training protocols that steer agents to find the best possible solutions given a previously set computational budget to enforce a learned exploration-exploitation tradeoff. In addition, we will combine our sampling method with Monte Carlo Tree Search (MCTS) for a more effective policy-guided search during inference.

## References

- Bello, I.; Pham, H.; Le V, Q.; Norouzi, M.; and Bengio, S. 2017. Neural Combinatorial Optimization with Reinforcement Learning. *ICLR Workshop*. <http://arxiv.org/pdf/1611.09940v3>.
- Corsini, A.; Porrello, A.; Calderara, S.; and Dell’Amico, M. 2024. Self-Labeling the Job Shop Scheduling Problem. arXiv:2401.11849.
- Deepak Pathak; Pulkit Agrawal; Alexei A. Efros; and Trevor Darrell. 2017. Curiosity-driven Exploration by Self-supervised Prediction. *International Conference on Machine Learning*, 34(70): 2778–2787.
- Falkner, J. K.; Thyssens, D.; Bdeir, A.; and Schmidt-Thieme, L. 2022. Learning to Control Local Search for Combinatorial Optimization. *Falkner, Jonas K., et al. Learning to Control Local Search for Combinatorial Optimization. Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 13717: 361–376.
- Iklasov, Z.; Medvedev, D.; Solozabal, R.; and Takac, M. 2022. Learning to generalize Dispatching rules on the Job Shop Scheduling. *Advances in Neural Information Processing Systems*, 33: 1621–1632.
- Kool, W.; van Hoof, H.; and Welling, M. 2019. Attention, Learn to Solve Routing Problems! *International Conference on Learning Representations*, 7.
- Kumar, A.; and Dimitrakopoulos, R. 2021. Production scheduling in industrial mining complexes with incoming new information using tree search and deep reinforcement learning. *Applied Soft Computing*, 110: 107644.
- Kwon, Y.-D.; Choo, J.; Kim, B.; Yoon, I.; Gwon, Y.; and Min, S. 2020. Pomo: Policy optimization with multiple optima for reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 21188–21198. <https://arxiv.org/pdf/2010.16011.pdf>.
- Ni, F.; Hao, J.; Lu, J.; Tong, X.; Yuan, M.; Duan, J.; Ma, Y.; and He, K. 2021. A Multi-Graph Attributed Reinforcement Learning based Optimization Algorithm for Large-scale Hybrid Flow Shop Scheduling Problem. In Zhu, F., ed., *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, ACM Digital Library, 3441–3451. New York, NY, United States: Association for Computing Machinery. ISBN 9781450383325.
- Oren, J.; Ross, C.; Lefarov, M.; Richter, F.; Taitler, A.; Feldman, Z.; Daniel, C.; and Di Castro, D. 2021. SOLO: Search Online, Learn Offline for Combinatorial Optimization Problems. *Proceedings of the Fourteenth International Symposium on Combinatorial Search (SoCS 2021)*.
- Orseau, L.; Lelis, L. H. S.; Lattimore, T.; and Weber, T. 2018. Single-Agent Policy Tree Search With Guarantees. *32nd Conference on Neural Information Processing Systems (NIPS)*, 32: 3201–3211.
- Panzer, M.; Bender, B.; and Gronau, N. 2022. Neural agent-based production planning and control: An architectural review. *Journal of Manufacturing Systems*, 65: 743–766.
- 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.
- Pinedo, M. 2016. *Scheduling: Theory, algorithms, and systems*. Springer International Publishing, fifth edition edition. ISBN 978-3-319-26578-0.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. <http://arxiv.org/pdf/1707.06347v2>.
- Shen, Y.; Chen, J.; Huang, P.-S.; Guo, Y.; and Gao, J. 2018. M-Walk: Learning to Walk over Graphs using Monte Carlo Tree Search. <https://arxiv.org/pdf/1802.04394.pdf>.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587): 484–489.
- Sutton, R. S.; and Barto, A. 2018. *Reinforcement learning: An introduction*. Adaptive computation and machine learning. Cambridge, Massachusetts, London, England: The MIT Press, second edition edition. ISBN 9780262039246.
- Taillard, E. 1993. Benchmarks for basic scheduling problems. *European Journal of Operational Research*, 64(2): 278–285.
- Tassel, P.; Gebser, M.; and Schekotihin, K. 2023. An End-to-End Reinforcement Learning Approach for Job-Shop Scheduling Problems Based on Constraint Programming. *Proceedings of the Thirty-Third International Conference on Automated Planning and Scheduling (ICAPS 2023)*, 614–622.
- Tassel, P.; Kovács, B.; Gebser, M.; Schekotihin, K.; Kohlenbrein, W.; and Schrott-Kostwein, P. 2022. Reinforcement Learning of Dispatching Strategies for Large-Scale Industrial Scheduling. *Proceedings of the International Conference on Automated Planning and Scheduling*, 32: 638–646.
- van Ekeris, T.; Meyes, R.; and Meisen, T. 2021. Discovering Heuristics And Metaheuristics For Job Shop Scheduling From Scratch Via Deep Reinforcement Learning. *Proceedings of the Conference on Production Systems and Logistics : CPSL 2021*, 709–718.
- van Hoorn, J. J. 2018. The Current state of bounds on benchmark instances of the job-shop scheduling problem. *Journal of Scheduling*, 21(1): 127–128.

- Vinyals, O.; Bengio, S.; and Kudlur, M. 2016. Order Matters: Sequence to sequence for sets. *4th International Conference on Learning Representations 2016*.
- Vinyals, O.; Fortunato, M.; and Jaitly, N. 2015. Pointer Networks. *Advances in Neural Information Processing Systems*, 29: 2692–2700.
- Wang, Q.; Hao, Y.; and Cao, J. 2021. Learning to traverse over graphs with a Monte Carlo tree search-based self-play framework. *Engineering Applications of Artificial Intelligence*, 105: 104422.
- Waubert de Puiseau, C.; Peters, J.; Dörpelkus, C.; Tercan, H.; and Meisen, T. 2023. schlably: A Python framework for deep reinforcement learning based scheduling experiments. *SoftwareX*, 22: 101383.
- Waubert de Puiseau, C.; Tercan, H.; and Meisen, T. 2023. Curriculum Learning in Job Shop Scheduling using Reinforcement Learning. *Proceedings of the Conference on Production Systems and Logistics: CPSL 2023*, 34–43.
- Zhang, C.; Song, W.; Cao, Z.; Zhang, J.; Tan, P. S.; and Xu, C. 2020. Learning to Dispatch for Job Shop Scheduling via Deep Reinforcement Learning. *Advances in Neural Information Processing Systems*, 33: 1621–1632.
- Zhang, C.; Song, W.; Cao, Z.; Zhang, J.; Tan, P. S.; and Xu, C. 2022. Learning to Search for Job Shop Scheduling via Deep Reinforcement Learning. <https://arxiv.org/pdf/2211.10936>.