

55th CIRP Conference on Manufacturing Systems

# Towards Standardising Reinforcement Learning Approaches for Production Scheduling Problems

Alexandru Rinciog<sup>a</sup>, Anne Meyer<sup>a</sup>

<sup>a</sup>TU Dortmund University, Faculty for Mechanical Engineering, Leonhard-Euler-Straße 5, 44227 Dortmund, Germany

## Abstract

Recent years have seen a rise in interest in terms of using machine learning, particularly reinforcement learning (RL), for production scheduling problems of varying degrees of complexity. The general approach is to break down the scheduling problem into a Markov Decision Process (MDP), whereupon a simulation implementing the MDP is used to train an RL agent. Since existing studies rely on (sometimes) complex simulations for which the code is unavailable, the experiments presented are hard, or, in the case of stochastic environments, impossible to reproduce accurately. Furthermore, there is a vast array of RL designs to choose from. To make RL methods widely applicable in production scheduling and work out their strength for the industry, the standardisation of model descriptions - both production setup and RL design - and validation scheme are a prerequisite. Our contribution is threefold: First, we standardize the description of production setups used in RL studies based on established nomenclature. Secondly, we classify RL design choices from existing publications. Lastly, we propose recommendations for a validation scheme focusing on reproducibility and sufficient benchmarking.

© 2022 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the International Programme committee of the 55th CIRP Conference on Manufacturing Systems

**Keywords:** Production Scheduling; Reinforcement Learning; Standardization; Validation

## 1. Introduction

The field of production planning and scheduling is often characterized by a high degree of dynamism, given that production reality changes constantly. As production data and computing resources become more readily available and production itself becomes more versatile, the incorporation of machine learning techniques for production scheduling is increasingly an option.

Production scheduling is NP complete for most real world cases, and, as such, difficult to solve optimally for large production instances.[1] Furthermore, unexpected events such as new job arrivals, operation duration deviations or resource availability issues, may invalidate schedules. In such cases fixed schedules need to be recomputed.

Using production simulations that include unforeseen events alongside the scheduling constraints, RL agents can be trained to optimize a target goal by taking actions in response to changing production states. The RL solutions' adaptability and speed during deployment makes them attractive for high dimensional and/or stochastic problems [2].

However, most studies suffer from reproducibility issues owing to a lack of clarity pertaining to RL modeling and/or production problem description, code and simulation input unavailability, and non-reproducible stochasticity. Furthermore, RL solutions should be benchmarked against more established approaches (e.g. local search [3]) and typical simple heuristics on a sufficient number of scheduling problem instances to avoid cherry picking [4] suspicion. A standard for discussing scheduling setups, RL solutions and RL results could alleviate the problems in the field.

This work seeks to standardise the description and validation of RL approaches for production scheduling problems. After confirming the standardisation gap in Section 2, we move towards our goal in three steps: First, we look at production setups used with RL experiments and propose a clear nomenclature scheme by extending a well established Operations Research (OR) literature standard in Section 3. Secondly, in Section 4, we investigate RL design choices available for production, isolating different possible Markov Decision Process (MDP) formulations. Thirdly, we lay down how to ensure reproducibility and avoid cherry-picking in stochastic settings in section 5. Section 6 concludes our analysis and presents some avenues for future work.

As a basis for the categories introduced in this paper, we reviewed 37 publications containing experimental results on the topic of RL production and scheduling. The papers considered were published between 1998 and 2020. This material was compiled using the Web of Knowledge portal with “shop-scheduling” and “reinforcement learning” as search topics. The need for our work is reinforced by the significant number of more recent publications on the topic [5].

To keep within the frame of this work, we omit some of the details of the introduced categories. We do not look at all possible information available for state, action and reward formulations during RL design. Similarly, we do not consider RL algorithms in detail. Instead, we briefly describe the main categories and give examples of where these were used for production scheduling.

### 2. Related Standardisation Efforts

A general standard for experiments in the field of OR is available in [6]. Here eight criteria for benchmarking are given, namely clearly stated goals, well-specified problems, suitable algorithms, adequate performance measures, thoughtful analysis, effective and efficient (experiment) designs, comprehensible presentations, and guaranteed reproducibility.

For basic scheduling problems, e.g. flow/job/open shop scheduling problems, an implicit problem definition standard is set through the benchmarking instances available in Beasley’s OR library [7]. New benchmarking instances are sometimes published alongside corresponding experiments, e.g. [8, 9].

First steps in the direction of RL project standardisation were taken by OpenAI Gym, whereby a general RL application programming interface (API) is defined. Through [10], OpenAI Gym simulations for varied combinatorial optimization problems from the field of OR, e.g. the bin packing or the traveling salesman problem, are made available.

These standardisation efforts fail to address the particularities of the RL production scheduling field. Firstly, the RL literature defines many scheduling setups using inconsistent nomenclature. Secondly, RL offers many design choices for production scheduling problems. Finally, in stochastic settings, validation is difficult because of the high instance variance.

### 3. Production Scheduling Variants

An established format for defining a production scheduling problem in OR literature is the  $\alpha|\beta|\gamma$  notation described by Pinedo in [1]. The three parameters describe the machine setup  $\alpha$ , a possibly empty set of additional constraints  $\beta$ , and the optimization target  $\gamma$ . In what follows, we introduce the values these parameters can take and the relationship between them. We define new Pinedo-style values as needed to cover the setups present in RL scheduling literature relating them back to the values in [1]. Figures 1, 2 and 3 give an overview of the categories missing from [1] but described in RL literature and their relationship with the well-defined parameter values.

#### 3.1. Machine Setup – $\alpha$

Production scheduling is the problem of assigning start times  $s_{ji} > 0$  to a number of operations  $o_{ji}$  with processing times  $p_{ji}$  grouped into  $n$  jobs, onto one of  $m$  production resources to optimize some target value, for example the makespan. The makespan  $C_{\max}$  is defined as the maximum over the set of all job completion times  $C_j$ , with  $j$  being the job index. Operations are executed one at a time in the absence of preemption. Release dates  $r_j$  defining when job processing can commence and due dates (or deadlines)  $d_j$  complete the picture. Depending on the number and order of operations within jobs and the number and speed of resources available for different operation types, machine setups can be categorized into one of the ten hierarchically arranged groups depicted in Figure 1.

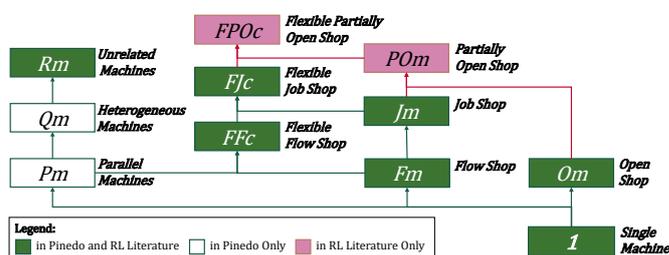


Figure 1. Hierarchy of machine setups ( $\alpha$ ).

Single machine problems (1) are defined in terms of  $n$  jobs with one operation each, and a single processing resource. By replicating the processing resource we obtain the parallel machines ( $Pm$ ) setup.  $Pms$  with different processing speed scalars  $v_i$  define the heterogeneous machines ( $Qm$ ) setup. Operation processing times are herein scaled by the machine speed to obtain the operation completion times. If the machine speed depends not only on the machine index  $i$  but also on the job index  $j$ , the unrelated machines ( $Rm$ ) setup is present, e.g. [11].

Setups with  $m$  resources and  $n$  jobs containing  $m$  operations, define shop-scheduling problems. Operations can only be processed by resources of corresponding types, and each resource is visited exactly once to complete a job. If operations within a job need to be processed in a fixed sequence, which is identical for all jobs, then the production setup is that of a flow shop ( $Fm$ ), e.g. [12]. Allowing distinct sequences for different jobs yields a job-shop environment ( $Jm$ ), e.g. [13]. No operation precedence constraints are indicative of an open-shop ( $Om$ ), e.g. [14].

Any shop setup can be extended by replicating resources and grouping them into work-centers. Operations are now fixed to a work-center, rather than a machine. If the setup prior to replication was flow or job shop, post replication the setup is that of a flexible flow shop ( $FFc$ ), e.g. [15], or flexible job shop ( $FJc$ ), e.g. [16], respectively.

The class of partially open job shop scheduling problems,  $POm$ , has an intermediary position between  $Jm$  and  $Om$ . Here job operations are partially ordered, with precedence constraints described by a directed acyclic graph. Analogously to the other flexible variants, in flexible  $POm$  (i.e.  $FPOc$ ), operations of a particular type can be executed on one of the machines of the

corresponding work-center. Examples of RL approaches over such setups can be found in [17, 18, 19].

A routing choice for jobs is required additionally to operation sequencing decisions in parallel machine setups ( $Pm$ ,  $Qm$ ,  $Rm$ ) and flexible shops. The same applies for setups where the precedence constraints allow for more than one operation to be executed immediately after a particular one has finished (i.e. in  $Om$ ,  $POc$ ). We refer to this as routing flexibility.

### 3.2. Additional Constraints – $\beta$

The  $\beta$  values modify the scheduling problem by adding constraints or changing existing ones. Figure 2 contains an overview of the constraints to be discussed. We refer to [1] for some less frequent constraints (no wait-time  $nwt$ , preemption  $prmp$ , permutation  $prmu$ , and job precedence  $prec$ ) that we do not extend in what follows.

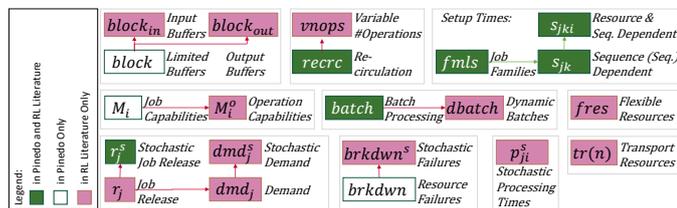


Figure 2. Overview of additional constraints ( $\beta$ ).

The *block* parameter is defined by Pinedo to model limited buffer space between production stages in  $Fm$ . Machines can only process an operation if there is enough space in the buffer following it. We define the more generic parameters  $block_{in}$ , e.g. [20], and  $block_{out}$ , e.g. [12], for input and output buffers in machine setups with routing flexibility. A full input buffer, blocks routing to the corresponding machine. Conversely, a full output buffer blocks processing on the afferent resource.

The re-circulation parameter  $recrc$ , e.g. [21], loosens the constraint whereby jobs visit each machine type exactly once. Instead some jobs may require revisiting work-centers. This results in cycles in the operation precedence graph. We introduce the heterogeneous number of operations parameter,  $vnops$ , e.g. [22], to model situations where jobs can revisit some machines while not visiting others at all, as is the case in [17, 18, 23].

Processing an operation often requires additional setup time. The parameters  $fmls$ , e.g. [15],  $S_{jk}$ , e.g. [24], and  $S_{jki}$ , e.g. [19], define different setup time penalties in an increasing order of generality. With  $fmls$  setup times are required when switching between job families.  $S_{jk}$  penalizes switches between any jobs. For  $S_{jki}$  setups, switch are additionally machine dependent.

Until now, any particular resource could only process one type of operations. This is not the case in setups of [24, 25, 26, 27]. Pinedo defines machine eligibility restrictions  $M_i$  for  $Pm$  environments only. Machine  $i$  can only process operations from jobs contained in the set  $M_i := J' \subset J$ , where  $J$  is the set of all jobs. We additionally generalize machine capability constraints for jobs to operation machine capabilities: To this end, we introduce the  $M_i^o$  parameter containing a set of operation

types where  $M_i^o := O' \subset type(O)$  with  $O$  being the set of all operations and  $type$  the mapping defining operation types.

The *batch* parameter, e.g. [15], relaxes the constraint whereby resources can only process one operation at a time by allowing fixed size and processing time batches on specified resources. However, both batch sizes and processing time can be dynamic as in the setup of [23]. We add the *dbatch* parameter to reflect this. In the setup of [28] resource capacity could vary dynamically, with resources being added to work-centers during the scheduling process. To cover this case, we define the flexible resources *fres* parameter. The dynamic batch, and the flexible resource constraints are fringe cases in the scheduling literature, but potentially very impactful for real world applications.

Parameters of core importance are those encoding stochasticity. Release dates ( $r_j$ ), could be deterministic, depending on the underlying planning process, but are most often stochastic ( $r_j^s$ ), e.g. [21]. Similarly, breakdowns are deterministic (*brkdwn*), if resources are taken offline for planned maintenance, and stochastic in case of unexpected failure (*brkdwn<sup>s</sup>*), e.g. [19]. The presence of the demand parameter ( $dmd_j$ ) implies that release dates are a system inherent choice. In a  $dmd_j$  system, finished jobs are consumed from a sink buffer as per incoming demand. A stochastic demand is indicated by  $dmd_j^s$ . This situation is considered in [12, 15, 29]. Operation processing times can also be seen as stochastic ( $p_{ji}^s$ ), e.g. [30], since these are often just estimates.

Transport times and resources are often not explicitly modeled. In such a case, two implicit assumptions are made: (a) distances and hence transport times are either negligible, constant or part of the processing times  $p_{ji}$ , and (b) transport resources are always available to move jobs to their resource. As soon as (a) does not hold, transport times should be modeled separately. If the production setting to be scheduled contains enough transport resources, (b) can still hold. We mark this setup by means of  $tr(\infty)$ , e.g. [28]. If conversely, transport resources are scarce, we have to model their current position explicitly  $tr(n)$  as in [21, 31], since this transport availability directly impacts the schedule.

### 3.3. Objectives – $\gamma$

As can be seen in Figure 3, there are many possible optimization targets for production scheduling, for different use cases. We categorize optimization targets as job-centric (throughput, makespan, flow-time, lateness, tardiness, unit cost) or resource-centric (machine utilization, transport resource utilization, job idle time, machine failures, buffer length, buffered times, setup times, inventory levels). These metrics target different aspects of the production system and can be combined into a joint optimization target either by constructing a score function as a function of multiple targets (scalarization), e.g. [15, 23, 29, 32], or by following a pareto front, e.g. [33].

**Job-Centric Goals:** The most often encountered metric (18 out of 37 publications) is the makespan ( $C_{max}$ ), i.e. the maximum completion time over all jobs. While  $C_{max}$  makes sense for static contexts with a fixed number of jobs, for dynamic environments a throughput measure needs to be used. In literature average job throughput  $Tpt_{ave}^j := 1/t \sum 1_{\{C_j \leq t\}}$  at time  $t$  is used in [21, 28].

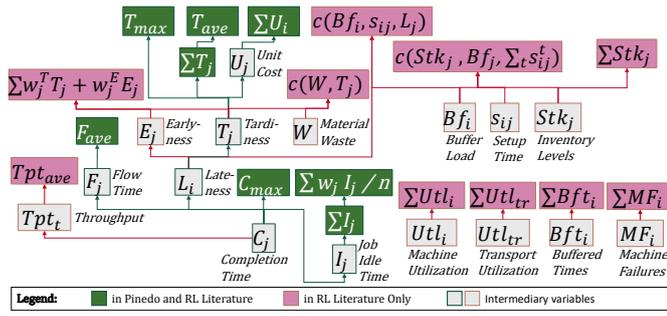


Figure 3. Optimization targets in RL scheduling literature ( $\gamma$ ).

Additionally, operation throughput  $Tpt_{ave}^o := 1/t \sum_{j,i} 1_{\{s_{ji}+p_{ji} \leq t\}}$  could be considered. Flow-time (also lead time)  $F_j$  measures the time between job release and job completion,  $F_j := C_j - r_j$ . The average flow time  $F_{ave} := 1/n \sum F_j$  is an indicator of a system’s reactivity/flexibility. Yet another measure based on the job completion time is the job idle time  $I_j$  defined as the difference between flow time and the summed job processing time:  $F_j - \sum_i p_{ji}$ .

Perhaps the most relevant metrics for manufacturing companies are timeliness related [34]. Timeliness metrics are computed using the lateness intermediary variable, i.e. the difference between completion time and due date,  $L_j := C_j - d_j$ . Tardiness extends lateness by ignoring early jobs  $T_j := \max\{0, L_j\}$ . Tardiness based optimization goals encountered are constructed by aggregating the tardiness variables, e.g.  $T_{ave}, T_{max}, \sum T_j$ . Alternatively, the number of tardy jobs can simply be counted using the unit cost variables  $U_j := 1_{\{C_j > d_j\}}$ , e.g. [35]. Finishing jobs too early can be detrimental [36]. As such, earliness  $E_j := |\min\{L_j, 0\}|$  can be another minimization target as it is considered in [32].

**Resource-Centric Goals:** Resource utilization is defined as  $Utl_i := 1/t \sum_j p_{ji} \cdot 1_{\{s_{ji}+p_{ji} \leq t\}}$ , i.e. the time the machine was working over the up-time  $t$ . Similarly one can define the transport utilization  $Utl_{tr}$  as the time spent carrying a load over the total elapsed time. The buffer composition can also be used as an optimization target. In general one would like to keep buffer lengths to a minimum. In [15] the number of buffered operations  $Bf_i$  at resource  $i$  was used as an optimization goal within a more complex cost function also involving lateness and the total number of tool switches.

#### 4. RL Modeling for Production Scheduling

To solve any problem with RL first an MDP must be modeled for the particular domain. The MDP defines the agent-environment interaction. The agent, whose goal is to maximize the received reward, senses the current environment state and takes an action, whereby the environment is moved to a new state. This loop continues until an end-state is reached. The agent receives feedback by means of a reward signal, which it tries to maximize. While for many RL settings the MDP is fairly obvious, for production scheduling environments, there are many design choices to be made. In what follows we categorize the RL design options available for production scheduling problems.

#### 4.1. MDP Breakdown

We start by defining five broad categories of MDPs (i.e. MDP breakdowns) present in RL scheduling literature: operation sequencing (1), routing before sequencing (2), interlaced routing and sequencing (3) transport-centric routing (4), and re-scheduling (5). In addition, we propose the holistic routing and sequencing breakdown (6) to better tackle  $tr(n)$  settings.

**Operation sequencing:** The idea behind iterative operation sequencing is to use the moments when resources become free as the discrete time steps when decisions can be taken. An operation  $o_{ji}$  from the machine queue is then assigned to it, marking its start time  $s_{ji}$  as the current time  $t$ . This process is repeated so long as there are still jobs to complete. In setups with no routing flexibility, the iterative operation sequencing MDP breakdown is sufficient for defining all possible schedules, e.g. [12, 14, 37, 38, 39].

Iterative operation sequencing can also be used to implicitly define routes for setups with routing flexibility, if the assumption of negligible transport times and endless capacity buffers is present. After an operation finishes, the next eligible resources for the job can be seen as sharing a virtual buffer wherein the job lies. As soon as the job is assigned by the agent to one of these resources, the job is no longer available for processing at the other resources. Such an approach can be found in [2, 11, 15, 23, 27, 28].

**Routing Before Sequencing:** Within this breakdown, whenever a new job arrives, the agent iteratively assigns processing resources to it, until it has a fixed completion path through production. When the job route has been set, the scheduling process described above can continue. Such is the approach of [25, 26].

**Interlaced Routing and Sequencing:** The combined routing and sequencing problem can also be solved in an interlaced fashion. The agent decides on demand upon the next processing resource to transport a job to, when an operation is finished. This decision is then immediately followed by a sequencing decision on the freed resource. This breakdown is employed in [24].

**Transport-Centric Routing:** In  $tr(n)$  settings, there are two behaviorally distinct resource types present: processing and transport resources. Within this MDP breakdown, processing resources sequence operations following a priority rule, e.g. first-in-first-out (FIFO). Decisions are required when transport resources are idle. The agent alternates between the selection of source processing resources and destination processing resources. The oldest job in the source output buffer is chosen and transported to the destination input buffer. Examples are [21, 31, 40].

**Holistic Routing and Sequencing:** The caveat of the breakdown above is that not all schedules are possible. To complete the picture we propose that decisions be requested of the agent either when transport or processing resources are idle. In the case of transport resources, the agent behavior is the same as in the transport-centric breakdown. When dealing with processing resources, the agent should behave as in the operation sequencing breakdown.

**Re-Scheduling:** Finally, whenever a stochastic event occurs, the agents can be tasked with deploying an external solver to

compute a new plan. Such is the case in [41], where an agent selects the parameters of a variable neighborhood search (VNS) [42] on new job arrivals. VNS is used to create a new schedule which is followed until a new job arrival.

#### 4.2. Agent Types

RL agents are trained according to the generalized policy iteration principle. Herein two stages are distinguished, namely policy evaluation and policy improvement. During policy evaluation, agents select actions according to a policy function, and observe their returned reward. During policy improvement, the policy is adjusted based on the observations made. These stages are repeated until convergence (hopefully). RL algorithms can be classified along three discrete axes: value, policy or actor-critic methods (1), model-free or model-based methods (2), and on- or off-policy methods (3) (cmp. [43]).

**Value, Policy and Actor-Critic Methods:** Value methods try to estimate future reward by means of a value function, which is used to estimate the “goodness” of either states or actions from given states. During the policy evaluation stage, actions are selected by using the value function to ascertain the quality of the states reachable from the current one. The observed rewards are used during the subsequent stage to improve the value function. Examples of such methods are State Action Reward State Action (SARSA) [44], Q-Learning (QL) [45], Deep Q-learning (DQN) [46], and Double DQN (DDQN) [47]. QL is particularly popular with the production scheduling community with implementations using tables to represent value functions used in [15, 22, 24, 35, 41], or neural networks as function approximators in [26, 29, 32, 38]. DQN was used in [19] and DDQN in [11, 16, 27],

Alternatively, agent policies  $\pi_\theta$  can be used directly to select actions during the evaluation stage. Based on the observed rewards, the reward expectation under policy  $\pi_\theta$  is estimated, its gradient with respect to  $\theta$  is computed and the policy is updated using stochastic gradient ascent. Examples include REINFORCE [48] used for production scheduling in [13], Trust Region Policy Optimization (TRPO) introduced in [49] and used for instance in [21] and Proximal Policy Optimization (PPO) [50] used in [51], albeit for a deterministic scheduling problem.

Policy and value approaches can be combined into actor-critic algorithms. Instead of using environment interaction to approximate the expected reward directly, the (state) value function approximator (critic), is used to inform the policy approximator (actor) of the quality of its action. AlphaZero (AZ) [52] used in [23] and Deep Deterministic Policy Gradient (DDPG) [53] used in [30] fall in this category.

**On- vs Off-Policy:** On policy methods (e.g. SARSA) use the same policy during evaluation stage that was adjusted during the improvement stage. This leads to more stable learning at the expense of exploration, which can lead to local optima. Conversely, in off-policy methods (e.g. QL, DQN, DDQN), the policy used during the evaluation stage can differ from the one used in the improvement stage, which leads to more exploration at the expense of convergence speed.

**Model Based vs. Model Free:** RL algorithms can be furthermore split into model-free and model based approaches. Model-based approaches use an environment model to plan a few steps into the future before deciding on an action. The involved environment model is either estimated by the agent itself, e.g. Imagination Augmented Agent (I2A) [54], or simply given to it, e.g. AZ.

**Single- vs Multi-Agents:** We can allow multiple agents to act within the same environment. These agents can be cooperative, i.e. striving to jointly maximize the expected reward or competitive, with each agent targeting a maximization of his reward only. For production scheduling multi-agent systems are often deployed (17 of the 37 studies). Depending on the MDP breakdown, agents can be associated with different setup components. Agents posted with each processing resource for sequencing decisions can be found in [2, 24, 26, 30, 32]. Agents deployed per job for sequencing are used in [29] and for routing in [26]. Routing agents are deployed per job family by [24]) or per transport resource in [31].

#### 4.3. State, Action and Reward Modeling

**States:** The information based on which the state transition occurs, i.e. the environment state, depends on the production setup considered and the MDP breakdown chosen. Take, for instance, a standard job-shop scheduling problem  $Jm|C_{max}$  with a sequencing breakdown. Here we encode the complete state information as four  $n \times m$  dimensional matrices, the system time  $t$  and index  $i$  of the processing resource requesting a scheduling decision:  $(T, P, L, A, t, i)$ . The type matrix  $T$  keeps track of the precedence within jobs (rows) and the type of resource the job operations need. The corresponding values in the  $P$  matrix represent the remaining processing time for each operation. The location of all operations is encoded using a matrix  $L := (l_{ji})$ , with  $l_{ji} = m_k$  if operation  $o_{ji}$  is at the machine  $m_k$ . To distinguish between operations being actively processed and those waiting in the machine buffer, the boolean matrix  $A$  is used.

If we present all the environment state information to the agent together with a machine requesting a new operation, we are in a fully observable environment situation. This means that the agent could, in theory, perfectly approximate the state transition and reward functions. However, a large state space and limited computational resources can make learning very time consuming. Additionally, it could be desirable that the agent learns a good policy with only local information, to enable the production system to easily scale up (e.g. by adding a new resource), without the need of a global monitoring system.

As such, the agent is mostly only presented with a subset of the environment state information available, which corresponds to a partially observable environment. We distinguish between raw state information, e.g. any subset of the  $\{T, P, L, A, t, i\}$  in the situation above, or information in condensed form, i.e. environment state features, or both. Employed features fall into one of three categories, namely job properties, resource attributes and optimization targets.

Job property features condense job information. Examples of such are remaining job operations, remaining job processing

time [55] or number of jobs in the system [41]. Resource related features target an aggregation of machine properties, such as remaining processing time in buffers [11], resource workload [32, 55], ratio of remaining processing times for machines and buffered processing time [22], number of product types in each buffer or machine health [15]. Estimates of the target function, or of other targetable indicators, are also often included in the construction of the agent state. Such examples are estimated total tardiness [27, 35], average machine utilization [27, 28], average transport resource utilization or the aforementioned average buffer length [28].

**Actions:** In terms of actions, there are two main approaches to production scheduling. The agent either takes an action directly, or selects an appropriate priority rule for scheduling or routing.

In the operation sequencing MDP, direct actions are operation indices, e.g. [55, 56], or job indices, e.g. [26], depending on whether the job operations are totally ordered (latter) or not (former). For the routing decision, the raw actions are processing resource indices, e.g. [21].

The basic intuition for the priority rule approach is that different rules will be more effective in different production scenarios. Given a state, the agent will output a rule from a fixed set. The rule is then applied to prioritize the operations in the buffer (iterative scheduling) or the viable downstream processing resources (iterative routing). The highest priority operation/machine gets selected. Mostly, the encountered priority rules are very simple. The most frequent examples of operation prioritization rules are shortest processing time (SPT), longest processing time (LPT), earliest due date (EDD), first in first out (FIFO) or last in first out (LIFO). The rule names are self explanatory. Machine prioritization rules work analogously (e.g. [24]). One could, for instance, prioritize the downstream machine with the shortest queue (SQ) in terms of processing time, the one with the least elements (LQE) or the shortest setup times (SST).

A completely different approach to action modeling was taken in [41]. The RL agent gets deployed on occurrence of new job arrivals. At this point a VNS is started to recompute a schedule for the operations not yet started. The RL agent action outputs the VNS parameters for the given state.

**Reward:** The last modeling decision to be made pertains to the reward signal produced by the environment. For production scheduling, the reward is most often proportional to the optimization target or a value that highly correlates with it (e.g. makespan and average utilization). That being said, there is no universally accepted scheme for constructing a reward function. In [26, 39, 55] the number of queued operations at the current machine is used to build the reward signal (the shorter the queue the better), though the optimization target was makespan.

Reward functions also differ based on when the agent receives a non-zero reward, e.g., at every decision point, at the end of the game, or at any point in-between. Other design choices include whether the reward is discrete or continuous [57], strictly positive or both positive and negative and many more [43]. Finding an appropriate reward is of singular importance, given that it strongly determines learning convergence, and can be challenging task for scheduling [58].

## 5. Validation

In this section investigate what is needed in order to ensure reproducibility and provide sufficient validation of RL production scheduling experiments. Table 1 summarizes the relevant stochastic RL literature in terms of production setup clarity and RL design clarity, simulation input availability, reproducible stochasticity, train-test split, state of the art coverage through baseline algorithms and, finally, cherry picking potential.

Table 1. Reproducibility and baselining fulfillment for stochastic RL production scheduling; Filled, and empty circles for fulfilled and unfulfilled criteria.

Source	Reproducibility				Evaluation		
	Setup Clarity	RL Clarity	Input Availability	Reproducible Stochasticity	Train-Test Split	Sufficient Baselines	Cherry Picking Potential
[2, 11, 24]	●	●	○	○	○	○	○
[15]	●	●	○	○	○	○	○
[16]	●	●	○	○	○	○	○
[19]	●	●	○	○	○	○	○
[20]	●	●	○	○	○	○	○
[21]	●	●	○	○	○	○	○
[22]	●	●	○	○	○	○	○
[26]	●	●	○	○	○	○	○
[27, 35]	●	●	○	○	○	○	○
[28, 38, 41]	●	●	○	○	○	○	○
[30]	●	●	○	○	○	○	○
[32]	●	●	○	○	○	○	○
[40]	●	●	○	○	○	○	○
[59]	●	○	○	○	○	○	○

### 5.1. Reproducibility

**Setup and RL Design Clarity:** In the absence of standardisation, it becomes difficult not to forget to include all production setups details. In the literature at hand the presented setups were described transparently with respect to most of their details save for the exact problem size (e.g. number of machines in each work-center, total number of scheduled jobs) and sometimes some constraint details. Similarly, more effort should be put into a clear presentation of all the RL design elements. Particularly the state and action designs suffer from a lack of clarity.

**Input Availability:** Ideally, in order to compare production scheduling approaches, the exact inputs, e.g. the job sequence with the operation precedence constraints and duration, should be provided. Unfortunately, this is mostly not the case. Instead the input sampling scheme is given. Running experiments on randomly sampled inputs has the caveat of requiring a large number of experiments and statistical testing to ensure comparability with other approaches in the absence of reproducibility. Such an explicit statistical analysis is absent from the RL scheduling literature. Moreover, the simulation code associated with the experiments should be provided as in [21].

**Reproducible Stochasticity:** Code availability alone does not suffice to ensure reproducible stochasticity, though it is a necessary condition. There is a simple way to ensure that a sequence of sampling actions produces the same output irre-

spective of the system the simulation is ran on: random number generator seeding (RNG). To ensure that the occurrence time of the stochastic events is independent of simulation control, one should sample *all* the stochastic events together with their occurrence time during simulation initialization. These events can then be queued and triggered at the appropriate time. A simulation implementing this scheme is described in [60].

## 5.2. Evaluation

**Train-Test Split:** As with all machine learning approaches, RL schedulers are susceptible to overfitting. In deterministic setups, this refers to a situation where a strategy is developed that is very well suited for the training instances, but cannot generalize well to unseen situations. Only reporting results on the production instances used for training is tantamount to assuming that the scheduling problem optimized is recurrent. In such recurrent, deterministic situations one may be better served by using other approaches, seen as whenever the comparison is available, RL fails to outperform the literature lower bounds, e.g. [13, 25, 56]. In the dynamic setting, the test-train separation is not mandatory, since the stream of product specification induces a test-train separation, with the beginning of the stream being used for training and the latter part for testing.

**Sufficient Baselines:** RL solutions for stochastic settings are almost exclusively baselined against simple heuristic solutions. It could be, however, that constructing schedules as if there were no stochasticity and simply ignoring plan deviations thereafter could be a competitive approach to RL or heuristics. Yet another approach would be to recompute a schedule on occurrence of unforeseen events. Constraint programming solvers such as OR-Tools [61] allow for the use of a time limit (e.g. [62]) for the solution space search. Therefore, RL solvers for large production instances could at least be benchmarked against search, which is a standard OR strategy.

**Cherry Picking Potential:** Cherry picking as a concept, is very simple: One runs  $n$  experiments, but only reports the results on  $m < n$  of those that fit the own hypothesis. Whenever the results are reported on only a small number of experiments, the experiment design is not reproducible, and there is no statistical testing involved, this lingering suspicion remains.

## 6. Conclusion

In this work we took steps towards standardizing RL experiments for production scheduling problems by introducing and extending an OR standard for the problem formulation (Section 3), describing the core elements of production scheduling RL design (Section 4), and proposing a scheme guaranteeing reproducibility and sufficient validation (Section 5).

The categories presented here were constructed to curtail three common problems of the RL scheduling literature. Firstly, possibly owing to the interdisciplinary nature of the field, the nomenclature used to describe scheduling setups is mostly inconsistent. This leads to a lack of clarity and redundant descriptive work, which in turn renders experiment reproducibility more

difficult. Secondly, because of the large design space associated with RL production scheduling, necessary MDP details are sometimes imprecisely formulated. Thirdly, there is a validation gap for RL solutions to stochastic problems owing to a lack of experiment reproducibility and insufficient baselining.

To consolidate the proposed standard, three further steps should be taken. To begin with, the more recently published RL scheduling literature should be indexed along the categories presented here. Additionally, the information used in the definition of states, actions and rewards should be exactly extracted and formalized. For instance, we should define individual features and rewards and not just their categories, as was done here. Finally, the standard should be used to implement or extend a benchmarking simulation such as the one in [60].

## Acknowledgements

This work was funded by the German Research Foundation through the Research Training Group 2193 “Adaption Intelligence of Factories in a Dynamic and Complex Environment”.

## References

- [1] M. Pinedo, Scheduling, Vol. 29, Springer, 2012.
- [2] B. Waschneck, A. Reichstaller, L. Belzner, T. Altenmüller, T. Bauernhansl, A. Knapp, A. Kyek, Optimization of global production scheduling with deep reinforcement learning, *Procedia CIRP* 72 (1) (2018) 1264–1269.
- [3] R. J. M. Vaessens, E. H. Aarts, J. K. Lenstra, Job shop scheduling by local search, *Informatics Journal on computing* 8 (3) (1996) 302–317.
- [4] J. M. Morse, “cherry picking”: Writing from thin data, *Qualitative health research* 20 (1) (2010) 3–3.
- [5] B. M. Kayhan, G. Yildiz, Reinforcement learning applications to machine scheduling problems: a comprehensive literature review, *Journal of Intelligent Manufacturing* (2021) 1–25.
- [6] T. Bartz-Beielstein, C. Doerr, J. Bossek, S. Chandrasekaran, T. Eftimov, A. Fischbach, P. Kerschke, M. Lopez-Ibanez, K. M. Malan, J. H. Moore, et al., Benchmarking in optimization: Best practice and open issues, *arXiv preprint arXiv:2007.03488* (2020).
- [7] J. E. Beasley, Or-library: distributing test problems by electronic mail, *Journal of the operational research society* 41 (11) (1990) 1069–1072.
- [8] J. Barnes, J. Chambers, Flexible job shop scheduling by tabu search, Graduate Program in Operations and Industrial Engineering, The University of Texas at Austin, Technical Report Series, ORP96-09 (1996).
- [9] S. Dauzère-Pères, W. Roux, J. B. Lasserre, Multi-resource shop scheduling with resource flexibility, *European Journal of Operational Research* 107 (2) (1998) 289–305.
- [10] C. D. Hubbs, H. D. Perez, O. Sarwar, N. V. Sahinidis, I. E. Grossmann, J. M. Wassick, Or-gym: A reinforcement learning library for operations research problem, *arXiv preprint arXiv:2008.06319* (2020).
- [11] L. Zhou, L. Zhang, B. K. Horn, Deep reinforcement learning-based dynamic scheduling in smart manufacturing, *Procedia CIRP* 93 (2020) 383–388.
- [12] S. Mahadevan, G. Theocharous, Optimizing production manufacturing using reinforcement learning., in: *FLAIRS Conference*, Vol. 372, 1998, p. 377.
- [13] T. Gabel, M. Riedmiller, Distributed policy search reinforcement learning for job-shop scheduling tasks, *International Journal of production research* 50 (1) (2012) 41–61.
- [14] G. H. Kim, C. G. Lee, Genetic reinforcement learning approach to the heterogeneous machine scheduling problem, *IEEE Transactions on Robotics and Automation* 14 (6) (1998) 879–893.
- [15] S. Qu, T. Chu, J. Wang, J. Leckie, W. Jian, A centralized reinforcement learning approach for proactive scheduling in manufacturing, in: *ETFA, IEEE*, 2015, pp. 1–8.

- [16] I.-B. Park, J. Huh, J. Kim, J. Park, A reinforcement learning approach to robust scheduling of semiconductor manufacturing facilities, *IEEE Transactions on Automation Science and Engineering* (2019).
- [17] W. Zhang, T. G. Dietterich, High-performance job-shop scheduling with a time-delay td ( $\lambda$ ) network, in: *Advances in neural information processing systems*, 1996, pp. 1024–1030.
- [18] W. Zhang, T. G. Dietterich, A reinforcement learning approach to job-shop scheduling, in: *IJCAI*, Vol. 95, 1995, pp. 1114–1120.
- [19] C. Hofmann, C. Krahe, N. Stricker, G. Lanza, Autonomous production control for matrix production based on deep q-learning, *Procedia CIRP* 88 (2020) 25–30.
- [20] A. Kuhnle, L. Schäfer, N. Stricker, G. Lanza, Design, implementation and evaluation of reinforcement learning for an adaptive order dispatching in job shop manufacturing systems, *Procedia CIRP* 81 (2019) 234–239.
- [21] A. Kuhnle, J.-P. Kaiser, F. Theiß, N. Stricker, G. Lanza, Designing an adaptive production control system using reinforcement learning, *Journal of Intelligent Manufacturing* (2020) 1–22.
- [22] X. Chen, X. Hao, H. W. Lin, T. Murata, Rule driven multi objective dynamic scheduling by data envelopment analysis and reinforcement learning, in: *2010 IEEE International Conference on Automation and Logistics*, IEEE, 2010, pp. 396–401.
- [23] A. Rinciog, C. Mieth, P. M. Scheikl, A. Meyer, Sheet-metal production scheduling using alphago zero, in: *CPSL*, 2020.
- [24] W. Bouazza, Y. Sallez, B. Beldjilali, A distributed approach solving partially flexible job-shop scheduling problem with a q-learning effect, *IFAC-PapersOnLine* 50 (1) (2017) 15890–15895.
- [25] Y. Martínez, A. Nowé, J. Suárez, R. Bello, A reinforcement learning approach for the flexible job shop scheduling problem, in: *International Conference on Learning and Intelligent Optimization*, Springer, 2011, pp. 253–262.
- [26] Y. M. Jiménez, A generic multi-agent reinforcement learning approach for scheduling problems, PhD, Vrije Universiteit Brussel (2012) 1–128.
- [27] S. Luo, Dynamic scheduling for flexible job shop with new job insertions by deep reinforcement learning, *Applied Soft Computing* (2020) 106208.
- [28] T. E. Thomas, J. Koo, S. Chaterji, S. Bagchi, Minerva: A reinforcement learning-based technique for optimal scheduling and bottleneck detection in distributed factory operations, in: *2018 10th International Conference on Communication Systems & Networks (COMSNETS)*, IEEE, 2018, pp. 129–136.
- [29] C. D. Paternina-Arboleda, T. K. Das, A multi-agent reinforcement learning approach to obtaining dynamic control policies for stochastic lot scheduling problem, *Simulation Modelling Practice and Theory* 13 (5) (2005) 389–406.
- [30] C.-L. Liu, C.-C. Chang, C.-J. Tseng, Actor-critic deep reinforcement learning for solving job shop scheduling problems, *IEEE Access* 8 (2020) 71752–71762.
- [31] K. Arviv, H. Stern, Y. Edan, Collaborative reinforcement learning for a two-robot job transfer flow-shop scheduling problem, *International Journal of Production Research* 54 (4) (2016) 1196–1209.
- [32] Y.-F. Wang, Adaptive job shop scheduling strategy based on weighted q-learning algorithm, *Journal of Intelligent Manufacturing* 31 (2) (2020) 417–432.
- [33] B. M. Méndez-Hernández, E. D. Rodríguez-Bazan, Y. Martínez-Jimenez, P. Libin, A. Nowé, A multi-objective reinforcement learning algorithm for jssp, in: *International Conference on Artificial Neural Networks*, Springer, 2019, pp. 567–584.
- [34] G. Schuh, V. Stich, H. Wienholdt, *Logistikmanagement*, Springer, 2013.
- [35] Y.-C. Wang, J. M. Usher, Application of reinforcement learning for agent-based production scheduling, *Engineering Applications of Artificial Intelligence* 18 (1) (2005) 73–82.
- [36] K. R. Baker, Minimizing earliness and tardiness costs in stochastic scheduling, *European Journal of Operational Research* 236 (2) (2014) 445–452.
- [37] S. Riedmiller, M. Riedmiller, A neural reinforcement learning approach to learn local dispatching policies in production scheduling, in: *IJCAI*, Vol. 2, Citeseer, 1999, pp. 764–771.
- [38] M. E. Aydın, E. Öztemel, Dynamic job-shop scheduling using reinforcement learning agents, *Robotics and Autonomous Systems* 33 (2-3) (2000) 169–178.
- [39] T. Gabel, M. Riedmiller, On a successful application of multi-agent reinforcement learning to operations research benchmarks, in: *2007 IEEE International Symposium on ADP and RL*, IEEE, 2007, pp. 68–75.
- [40] N. Stricker, A. Kuhnle, R. Sturm, S. Friess, Reinforcement learning for adaptive order dispatching in the semiconductor industry, *CIRP Annals* 67 (1) (2018) 511–514.
- [41] J. Shahrabi, M. A. Adibi, M. Mahootchi, A reinforcement learning approach to parameter estimation in dynamic job shop scheduling, *Computers & Industrial Engineering* 110 (2017) 75–82.
- [42] N. Mladenović, P. Hansen, Variable neighborhood search, *Computers & operations research* 24 (11) (1997) 1097–1100.
- [43] R. S. Sutton, A. G. Barto, *Reinforcement learning: An introduction*, MIT press, 2018.
- [44] G. A. Rummery, M. Niranjan, *On-line Q-learning using connectionist systems*, Vol. 37, University of Cambridge, Department of Engineering Cambridge, UK, 1994.
- [45] C. Watkins, *Learning from delayed rewards*, Ph. D. thesis, King's College, University of Cambridge (1989).
- [46] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, M. Riedmiller, Playing atari with deep reinforcement learning, *arXiv preprint arXiv:1312.5602* (2013).
- [47] H. Van Hasselt, A. Guez, D. Silver, Deep reinforcement learning with double q-learning, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30, 2016.
- [48] R. Williams, A class of gradient-estimation algorithms for reinforcement learning in neural networks, in: *Proceedings of the International Conference on Neural Networks*, 1987, pp. II–601.
- [49] J. Schulman, S. Levine, P. Abbeel, M. Jordan, P. Moritz, Trust region policy optimization, in: *International conference on machine learning*, 2015, pp. 1889–1897.
- [50] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms, *arXiv preprint arXiv:1707.06347* (2017).
- [51] C. Zhang, W. Song, Z. Cao, J. Zhang, P. S. Tan, X. Chi, Learning to dispatch for job shop scheduling via deep reinforcement learning, *Advances in Neural Information Processing Systems* 33 (2020).
- [52] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, et al., Mastering chess and shogi by self-play with a general reinforcement learning algorithm, *arXiv preprint arXiv:1712.01815* (2017).
- [53] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, D. Wierstra, Continuous control with deep reinforcement learning, *arXiv preprint arXiv:1509.02971* (2015).
- [54] S. Racanière, T. Weber, D. P. Reichert, L. Buesing, A. Guez, D. J. Rezende, A. P. Badia, O. Vinyals, N. Heess, Y. Li, et al., Imagination-augmented agents for deep reinforcement learning., in: *NIPS*, 2017, pp. 5690–5701.
- [55] T. Gabel, Multi-agent reinforcement learning approaches for distributed job-shop scheduling problems, PhD, Universität Osnabrück (2009) 1–173.
- [56] Y. C. Fonseca-Reyna, Y. Martínez-Jiménez, A. Nowé, Q-learning algorithm performance for m-machine, n-jobs flow shop scheduling problems to minimize makespan, *Investigación Operacional* 38 (3) (2018) 281–290.
- [57] L. Maignon, G. J. Laurent, N. Le Fort-Piat, Reward function and initial values: Better choices for accelerated goal-directed reinforcement learning, in: *International Conference on Artificial Neural Networks*, Springer, 2006, pp. 840–849.
- [58] S. Lang, F. Behrendt, N. Lanzerath, T. Reggelin, M. Müller, Integration of deep reinforcement learning and discrete-event simulation for real-time scheduling of a flexible job shop production, in: *2020 Winter Simulation Conference (WSC)*, IEEE, 2020, pp. 3057–3068.
- [59] M. Zhao, X. Li, L. Gao, L. Wang, M. Xiao, An improved q-learning based rescheduling method for flexible job-shops with machine failures, in: *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, IEEE, 2019, pp. 331–337.
- [60] A. Rinciog, A. Meyer, Fabricatio-rl: a reinforcement learning simulation framework for production scheduling, in: *2021 Winter Simulation Conference (WSC)*, IEEE, 2021, pp. 1–12.
- [61] L. Perron, V. Furnon, Or-tools, <https://developers.google.com/optimization> 5411 (2019).
- [62] G. Da Col, E. C. Teppan, Industrial size job shop scheduling tackled by present day cp solvers, in: *International Conference on Principles and Practice of Constraint Programming*, Springer, 2019, pp. 144–160.