Solving the Paint Shop Problem with Flexible Management of Multi-Lane Buffers Using Reinforcement Learning and Action Masking

Mirko Stappert[a,*], Bernhard Lutz[a], Janis Brammer[b], Dirk Neumann[a]

[a]*University of Freiburg, Rempartstr. 16, 79098 Freiburg, Germany*
[b]*CARIAD SE, Berliner Ring 2, 38440 Wolfsburg, Germany*

Abstract

In the paint shop problem, an unordered incoming sequence of cars assigned to different colors has to be reshuffled with the objective of minimizing the number of color changes. To reshuffle the incoming sequence, manufacturers can employ a first-in-first-out multi-lane buffer system allowing store and retrieve operations. So far, prior studies primarily focused on simple decision heuristics like greedy or simplified problem variants that do not allow full flexibility when performing store and retrieve operations. In this study, we propose a reinforcement learning approach to minimize color changes for the flexible problem variant, where store and retrieve operations can be performed in an arbitrary order. After proving that greedy retrieval is optimal, we incorporate this finding into the model using action masking. Our evaluation, based on 170 problem instances with 2-8 buffer lanes and 5-15 colors, shows that our approach reduces color changes compared to existing methods by considerable margins depending on the problem size. Furthermore, we demonstrate the robustness of our approach towards different buffer sizes and imbalanced color distributions.

Keywords: Scheduling, Paint Shop Problem, Reinforcement Learning, Metaheuristics, Integer Programming

1. Introduction

The modern automotive industry relies upon a sequential manufacturing process based on assembly lines (Boysen et al., 2022). One step in this manufacturing process is to color the incoming sequence of cars arriving from the body shop in random order. The sequence of cars is generally optimized for different targets, like minimizing work overload or makespan (e.g., Brammer et al., 2022a; Mosadegh et al., 2020; Neufeld et al., 2023), so that it does not specifically account for how cars are assigned to different colors. Given that changing the color of the paint nozzle generates additional costs regarding time and material waste (Bysko et al., 2020), the paint shop problem focuses on reshuffling the incoming sequence of cars so that color changes are minimized. For this purpose, manufacturers can rely on different buffer systems like automated storage and retrieval

---

*Corresponding author.

Email addresses: `mirko.stappert@is.uni-freiburg.de` (Mirko Stappert), `bernhard.lutz@is.uni-freiburg.de` (Bernhard Lutz), `janis.adrian.brammer@cariad.technology` (Janis Brammer), `dirk.neumann@is.uni-freiburg.de` (Dirk Neumann)

system (AS/RS) buffers (e.g., Inman & Schmeling, 2003), pull-off tables (e.g., Lahmar et al., 2003), and multi-lane buffer systems (e.g., Epping & Hochstättler, 2003). We focus on a multi-lane buffer system, as illustrated in Figure 1, due to their widespread usage in industry and low implementation costs (Spieckermann et al., 2004; Taube & Minner, 2018). The buffer system can be used to store the current car from the incoming sequence to one of the buffer lanes. Each lane operates independently as a first-in-first-out (FIFO) queue.[1] In addition, the system can retrieve the car from the rightmost position of a buffer lane, which is then appended to the outgoing sequence. The goal of the paint shop problem is to find an optimal sequence of store and retrieve operations that minimizes color changes.
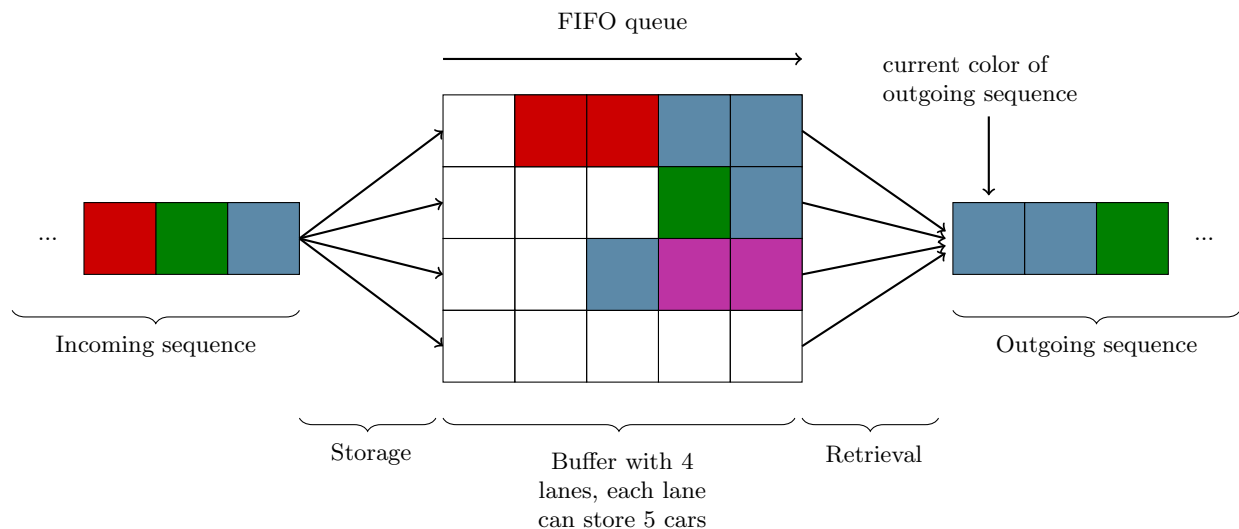

Figure 1: Illustration of the paint shop problem with a buffer consisting of four lanes.

So far, prior studies primarily focused on different problem variants that focus either on the retrieval phase only (e.g., Epping & Hochstättler, 2003; Spieckermann et al., 2004) or they require that storage and retrieval are performed in alternating phases, so that the buffer is first fully filled and subsequently fully emptied (e.g., Ding & Sun, 2004; Wu et al., 2021). However, it can easily be shown that solving the problem by alternating between filling the buffer first and subsequently removing all stored cars can result in arbitrarily worse solution quality compared to not restricting the order of storage and retrieval. We therefore consider the flexible problem variant, which allows store and retrieve operations to be performed at any time point whenever possible. The few studies that applied RL to the paint shop problem with multi-lane buffers (Huang et al., 2024; Leng et al., 2020, 2023) also restricted the flexibility of store and retrieve operations. Specifically, they required that store and retrieve operations alternate at each timestep so that the buffer always remains filled according a desired fill rate (e.g., 60%). To the best of our knowledge, our study is the first to formalize the fully flexible problem variant and solve it using reinforcement learning (RL).

---

[1]In the literature, such buffer systems are also referred to as *selectivity banks*, *parallel line buffer*, or *mix banks*, see Boysen et al. (2012).

We propose a deep RL approach as it is well suited for scheduling problems due to the discrete action space and time component that allows us to model the sequential decision-making process. At each time step, the RL agent can decide between storing the next incoming car in a particular buffer lane or retrieving the rightmost car from a buffer lane. The state representation encodes the buffer content, the next colors from the incoming sequence, and the current color at the reshuffled outgoing sequence. The reward function encourages retrieve actions that do not lead to a color change, while color changes are penalized with negative rewards. We increase the efficiency of the learning progress by excluding invalid actions (i.e., storing in full buffer lanes or retrieving from empty buffer lanes) through action masking. In addition, we enforce greedy retrieve operations through action masking after showing that greedy retrieval is optimal. Our evaluation based on 170 problem instances with 2–8 buffer lanes and 5–15 colors shows that the proposed RL approach decreases color changes by considerable margins on the vast majority of considered problem instances. We only observe slightly inferior performance to existing methods on small problems with 2x2 buffers and 15 colors, and on large problems instances with 8x8 buffers and 10 or 15 colors, which may be explained by the greater complexity of the state space. We finally demonstrate the robustness of our approach towards rectangular buffer sizes, initially filled buffer, and problem instances with out-of-distribution characteristics.

Our study makes two major contributions to the literature. First, we formalize the paint shop problem with flexible storage and retrieval operations as an integer linear program (ILP) by generalizing the paint shop problem formulations from Ding & Sun (2004) and Spieckermann et al. (2004). To the best of our knowledge, no study has thus far provided a complete ILP formalization of this flexible problem variant. The problem formalization may fuel future research on the paint shop problem without restrictions on when store and retrieve operations can be performed. From a theoretical view, we show formally that the solutions of the less flexible problem variant "storage-then-retrieve" can be arbitrarily worse regarding the number of color changes than solutions of the flexible problem variant. In addition, we prove that the frequently used greedy retrieval heuristic aiming to retrieve cars of the same color as the last car in the outgoing sequence, if applicable, is indeed optimal in minimizing color changes. Second, our study contributes to the growing literature that applies machine learning to solve problems from operations research (Bengio et al., 2021). So far, prior works (Huang et al., 2024; Leng et al., 2020, 2023) only used RL for retrieve operations in a problem variant that alternates between storage and retrieval at each timestep, while store operations are still performed by heuristics. By contrast, we employ RL as a complete solution policy, including storage and retrieval. Prior studies also proposed several heuristics like greedy storage and greedy retrieval (Spieckermann et al., 2004; Sun et al., 2015; Bysko et al., 2020; Leng et al., 2020), as well as metaheuristics like genetic algorithms (Ko et al., 2016) and ant colony optimization (Lin et al., 2011). A crucial advantage of RL over metaheuristics and mathematical solvers is the fact that a trained RL policy can instantly generate solutions of decent quality to similar but unseen problem instances. By contrast, metaheuristics and solvers cannot rely on hours of pre-training. Instead, they are built to solve each problem instance individually, independently

from other problem instances within the limited cutoff time. Therefore, employing RL presents a fast and adaptive solution that can transfer prior learning experience to novel instances, which increases its applicability in real-world situations and allows manufacturers to account for unforeseen disturbances in the production process like labor shortage or failed deliveries of supplies.

The remainder of this work is structured as follows. Section 2 provides an overview of related work on the paint shop problem. Section 3 formalizes the paint shop problem with fully flexible store and retrieve operations as an integer programming problem. Section 4 presents the reinforcement learning approach, describing action space, state representation, transition and reward function, action masking, as well as the policy learning algorithm. Section 5 and Section 6 detail the evaluation procedure and results. Section 7 concludes.

2. Related work

The paint shop problem has been studied in different variants. Table 1 provides a chronological overview of prior studies about the paint shop problem with multi-lane buffers, along with the employed solution method, the problem variant that restricts the execution of storage and retrieval, and when the solution approach changes between storage and retrieval.[2] We identify three distinct variants of performing store and retrieve operations. First, the "retrieval-only" variant considers an initially filled buffer that has to be emptied by retrieve operations. Second, the "storage-then-retrieval" variant alternates between storage and retrieval phases. The solution approaches change from storage to retrieval as soon as the buffer is full. During the retrieval phase, the buffer is fully emptied. The incoming sequence is thus split into multiple parts according to the buffer size and processed by storing one part into the buffer and retrieving it afterward. Some studies (e.g., Lin et al., 2011; Wu et al., 2021) also restrict the length of the incoming sequence to be less than the buffer size. As a consequence, the solution consists of exactly one storage and one retrieval phase. Third, the problem variant "flexible storage and retrieval" allows store and retrieve operations to be performed in an arbitrary order.

Early studies focused on the retrieval-only problem variant and proposed exact algorithms. Epping & Hochstättler (2003) defined a dynamic programming algorithm inspired by multiple sequence alignment. However, this approach has an exponential time complexity in the number of buffer lanes. Shortly after, Spieckermann et al. (2004) formulated the retrieval-only problem as an ILP and developed a branch-and-bound approach. The main advantage of this approach is that it can be stopped after a given cutoff time, which increases practicability in the sense that the method finds reasonably good solutions in a feasible time.

The problem variant "storage-then-retrieval" was first formalized as an ILP by Ding & Sun

---

[2]There are also several other problems that alter car sequences (e.g., Taube & Minner, 2018), which are are out of the scope of this study. These include the M-to-1 conveyor systems (Ko et al., 2016; Hong et al., 2018), the paint shop problem for words (Epping et al., 2004; Meunier & Neveu, 2012; Winter et al., 2019), the paint shop problem with virtual resequencing (Sun & Han, 2017). A survey of various other resequencing problems is provided by Boysen et al. (2012).

(2004). They also suggested greedy heuristics for storage and retrieval. The ILP formulation was later used by Lin et al. (2011), who developed a nested ant colony optimization approach. Further heuristics, like the shuffling heuristic and the arraying heuristic procedure, were proposed by Sun et al. (2015). The shuffling heuristic is a retrieval heuristic that aims to sequentially merge all buffer lanes by constructing a minimal spanning tree of an associated graph. The array heuristic procedure, by contrast, is a storage heuristic that aims to group together as many cars of the color that occurs most often in the incoming sequence as possible.

Table 1: Prior studies on the paint shop problem with multi-lane buffers.

Study	Solution method	Storage and retrieval	
		Variant	Change of phases
Epping & Hochstättler (2003)	Dynamic programming	Retrieval only / Flexible	– / Model decision
Spieckermann et al. (2004)	Branch-and-bound	Retrieval only	–
Ding & Sun (2004)	Heuristics	Store-then-retrieve	Buffer is full
Lin et al. (2011)	Ant colony optimization	Store-then-retrieve	Buffer is full
Sun et al. (2015)	Heuristics	Store-then-retrieve	Buffer is full
Sun & Han (2017)	Heuristics	Store-then-retrieve	Buffer is full
Bysko & Krystek (2019)	Game theoretic	Flexible	External signal
Bysko et al. (2020)	Heuristics	Flexible	External signal
Leng et al. (2020)	Reinforcement learning	Store-then-retrieve	Each timestep
Wu et al. (2021)	Heuristics	Store-then-retrieve	Buffer is full
Leng et al. (2023)	Reinforcement learning	Store-then-retrieve	Each timestep
Huang et al. (2024)	Reinforcement learning	Store-then-retrieve	Each timestep
This study	Reinforcement learning	Flexible	Model decision

The problem variant with flexible store and retrieve operations was first studied by Epping & Hochstättler (2003). The authors proposed an exact dynamic programming algorithm for small problem instances with up to 5 colors and 3 buffer lanes. However, the exponential time complexity in the size of the buffer makes this algorithm impractical for any real-sized problem instance. Bysko & Krystek (2019) developed an approach based on game theory, where the storage and retrieval phases are both formulated as (different) two-player games, where the set of strategies corresponds to the set of buffer lanes. The Nash equilibrium strategy is then proposed as a suitable storage or retrieval operation. In their more recent study (Bysko et al., 2020), the authors developed heuristic approaches that dynamically assign priorities to certain colors. Although the studies by Bysko and colleagues also considered the flexible problem variant, the decision of whether or not to change from retrieval to storage and vice versa is given by an external signal, e.g., from production and not by the solution approach itself.

The studies by Leng et al. (2020), Leng et al. (2023), and Huang et al. (2024) are closest to ours as they also applied RL to the paint shop problem with multi-lane buffers. However, these studies require that store and retrieve operations alternate at each timestep, while the buffer always

remains filled according to a given percentage (e.g., 60%). Furthermore, they rely on a heuristic to perform store operations, which aims to store the next car to a lane, where the leftmost position matches its color. Otherwise, the heuristic inserts the car into the lane with least cars and a random lane in case of ties. In particular, the studies by Leng et al. and Huang et al. train the RL policy only for the retrieve operations. By contrast, we consider the fully flexible problem version, where store and retrieve operations can be performed in an arbitrary order, and the RL policy is trained to perform both, store and retrieve operations.

3. Paint shop problem

In this section, we first formalize the paint shop problem. Subsequently, we show that allowing flexibility in performing store and retrieve operations can lead to arbitrarily better solution quality than the less complex problem variant "store-then-retrieve".
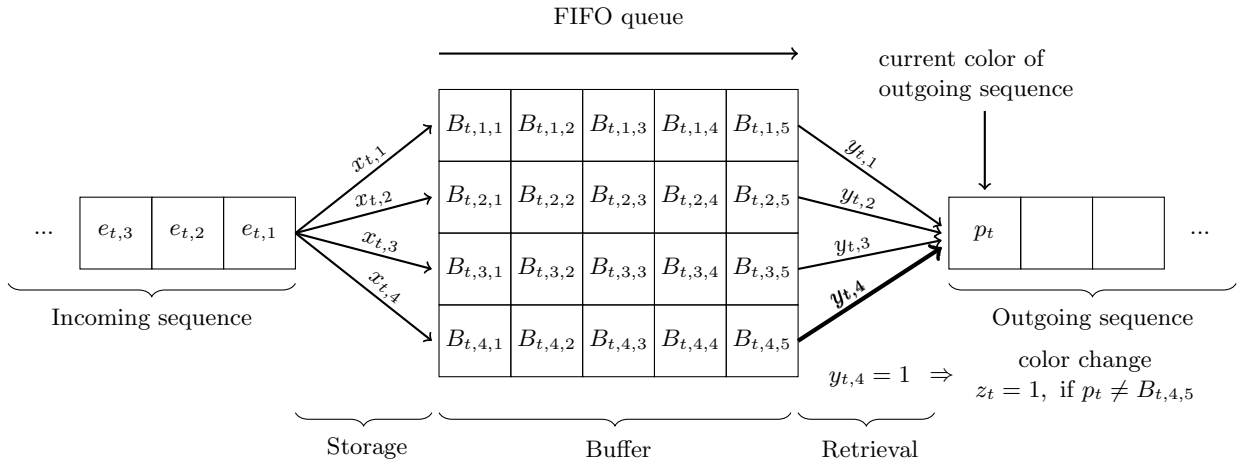

Figure 2: Paint shop problem with a 4x5 buffer ($L = 4$ lanes of width $W = 5$) and corresponding notation. The binary decision variables are given by $x_{t,i}$ (storage) and $y_{t,i}$ (retrieval).

3.1. Formalization

We formalize the paint shop problem with flexible storage and retrieval operations as an ILP. To ease comprehensibility of the notation, we illustrate an instance of the paint shop problem, along with the problem parameters and variables in Figure 2. An incoming sequence of N cars needs to be painted in their assigned color $c_1, c_2, \ldots, c_N$. We introduce a discrete time $t = 1, \ldots, 2N$ to denote the entire sequence of store and retrieve operations. At each timestep t, the system can either store the current car from the incoming sequence in one of L buffer lanes or retrieve from a buffer lane. The decision variables $x_{t,i}$ are equal to 1 if the current car is stored in lane i at timestep t, and 0 otherwise. A store operation to lane i puts the car in the rightmost empty position of the buffer lane. The color of the car in lane i at position j and timestep t is denoted by $B_{t,i,j}$. The width of each lane is denoted by W. The decision variables $y_{t,i}$ are equal to 1 if the last car of buffer lane i is retrieved at timestep t. A retrieved car is added to the outgoing sequence. The

6

color of the most recently added car is denoted by p_t. A color change occurs if the last retrieve operation selected a car at the end of buffer lane i with a different color than p_t, which is denoted by the dummy variable z_t. Therefore, the goal of the problem is to minimize the sum over all z_t for $t = 1, \ldots, 2N - 1$. An overview of all parameters and variables is provided in Table 2.

We do not impose length constraints for the outgoing sequence. In addition, we do not impose any time constraints on the speed of store and retrieve operations. In particular, we assume that painting a car requires more time than storing, moving, and retrieving a car through a buffer lane.

Table 2: Overview of parameters and variables.

Parameters	
C	Number of colors
N	Length of incoming sequence
$c_n \in \{1, \ldots, C\}$	Color of car at initial sequence position n
L	Number of buffer lanes
W	Width of buffer
Variables	
$x_{t,i}$	Binary decision variable indicating whether the next incoming car is placed in buffer lane i at time t
$y_{t,i}$	Binary decision variable indicating whether the last car in buffer lane i is retrieved at time t
$B_{t,i,j} \in \{0, \ldots, C\}$	Color of j-th car in buffer lane i at time t, 0 if empty
$e_{t,n} \in \{0, \ldots, C\}$	Color of car in position n of incoming sequence at time t, 0 if empty
$p_t \in \{0, \ldots, C\}$	Current color at time t, i.e., color that the last car was painted in
z_t	Binary variable indicating if a color change occurs at time t

$$\text{Minimize} \sum_{t=1}^{2N-1} z_t \tag{1}$$

Subject to:

$$x_{t,i}, y_{t,i}, z_t \in \{0,1\}, \quad t = 1, \ldots, 2N; i = 1, \ldots, L \tag{2}$$

$$\sum_{i=1}^{L} x_{t,i} + y_{t,i} = 1, \quad t = 1, \ldots, 2N \tag{3}$$

$$e_{1,n} = c_n, \quad n = 1, \ldots, N \tag{4}$$

$$B_{1,i,j} = 0, \quad i = 1, \ldots, L; j = 1, \ldots, W \tag{5}$$

$$p_1 = 0 \tag{6}$$

$$\sum_{i=1}^{L} x_{t,i} = 1 \Rightarrow e_{t+1,n} = e_{t,n+1} \wedge e_{t+1,N} = 0, \quad t = 1, \ldots, 2N - 1; n = 1, \ldots, N - 1 \tag{7}$$

$$\sum_{i=1}^{L} x_{t,i} = 0 \Rightarrow e_{t+1,n} = e_{t,n} \quad t = 1, \ldots, 2N - 1; n = 1, \ldots, N \tag{8}$$

$$e_{t,1} = 0 \Rightarrow x_{t,i} = 0, \quad t = 1, \ldots, 2N; i = 1, \ldots, L \tag{9}$$

$$x_{t,i} + y_{t,i} = 0 \Rightarrow B_{t+1,i,j} = B_{t,i,j}, \quad t = 1, \ldots, 2N - 1; i = 1, \ldots, L; j = 1, \ldots, W \tag{10}$$

7

$$(x_{t,i} = 1 \land B_{t,i,W} = 0) \Rightarrow B_{t+1,i,W} = e_{t,1}, \quad t = 1, \ldots, 2N-1; i = 1, \ldots, L \tag{11}$$

$$(x_{t,i} = 1 \land B_{t,i,j} = 0 \land B_{t,i,j+1} \neq 0) \Rightarrow B_{t+1,i,j} = e_{t,1}, \tag{12}$$

$$(x_{t,i} = 1 \land B_{t,i,W} \neq 0) \Rightarrow B_{t+1,i,W} = B_{t+1,i,W}, \quad t = 1, \ldots, 2N-1; i = 1, \ldots, L \tag{13}$$

$$t = 1, \ldots, 2N-1; i = 1, \ldots, L; j = 1, \ldots, W-1$$

$$(x_{t,i} = 1 \land B_{t,i,j} \neq 0 \land B_{t,i,j+1} = 0) \Rightarrow B_{t+1,i,j} = B_{t,i,j}, \tag{14}$$

$$t = 1, \ldots, 2N-1; i = 1, \ldots, L; j = 1, \ldots, W-1$$

$$y_{t,i} = 1 \Rightarrow B_{t+1,i,1} = 0 \land B_{t+1,i,j+1} = B_{t,i,j}, \quad t = 1, \ldots, 2N-1, i = 1; \ldots, L; j = 1, \ldots, W \tag{15}$$

$$x_{t,i} = 1 \Rightarrow B_{t,i,1} = 0, \quad t = 1, \ldots, 2N; i = 1, \ldots, L; j = 1, \ldots, W \tag{16}$$

$$y_{t,i} = 1 \Rightarrow B_{t,i,W} \neq 0, \quad t = 1, \ldots, 2N; i = 1, \ldots, L; j = 1, \ldots, W \tag{17}$$

$$\sum_{i=1}^{L} y_{t,i} = 0 \Rightarrow p_t = p_{t-1}, \quad t = 2, \ldots, N \tag{18}$$

$$y_{t,i} = 1 \Rightarrow p_t = B_{t,i,W}, \quad t = 2, \ldots, N; i = 1, \ldots, L \tag{19}$$

$$|p_{t+1} - p_t| \leq C z_t, \quad t = 1, \ldots, 2N-1 \tag{20}$$

Our formalization extends the formalizations of the problem variants "store-then-retrieve" and "retrieval-only" studied by Ding & Sun (2004) and Spieckermann et al. (2004), respectively. The goal (1) is to minimize the total number of color changes. Constraint (2) ensures that the decision variables $x_{t,i}, y_{t,i}$ and dummy variables z_t are binary. Constraint (3) ensures that exactly one operation is performed at each time step. Constraints (4), (5) and (6) initialize the incoming sequence, the buffer, and the current color of the outgoing sequence at $t = 0$ with zeros. Constraints (7) and (8) prescribe updates to the incoming sequence: When no store operation is performed, the incoming sequence stays the same, otherwise it is moved by one position. Constraint (9) ensures that once the incoming sequence is empty, no further storing operations are performed. Constraints (10)-(17) describe how store and retrieve operations influence the buffer content. Constraint (10) ensures that if we neither store in a lane nor retrieve from it, all buffer variables for that lane stay the same. Constraints (11) and (12) ensure that if we store in a lane, the rightmost free buffer position in that lane is updated to the color of the current car from the incoming sequence. Constraints (13) and (14) ensure that all other elements of the lane stay unchanged. If we retrieve from a lane, we move all elements one spot to the right (15). Constraints (16) and (17) ensure that when we store in the buffer, the lane is not completely filled and if we retrieve, it is not completely empty. Finally, the last three constraints deal with the current colors and color changes. Constraint (18) ensures that if we do not retrieve the current color stays unchanged. Constraint (19) ensures that, if we retrieve a car, the current color is updated to the color of the corresponding buffer position. Lastly, constraint (20) defines that a color change occurs when the current color differs from the previous current color.

The implication constraints (7)–(19) can be converted into standard linear constraints (see e.g.,

Plastria, 2002).[3] For this purpose, an implication is converted into a constraint of the following normal form

$$a_1 = 0 \wedge \cdots \wedge a_n = 0 \Rightarrow b_1 = 0 \wedge \cdots \wedge b_m = 0. \tag{21}$$

Here, $a_1, \ldots, a_n$ are binary variables and $b_1, \ldots b_m$ are bounded integer variables with bounds $\mathcal{L}_j \leq b_j \leq \mathcal{U}_j$ for $j = 1, \ldots, m$. Such a normal form implication constraint is equivalent to the linear inequalities

$$\mathcal{L}_j \sum_{i=1}^{n} a_i \leq b_j \leq \mathcal{U}_j \sum_{i=1}^{n} a_i, \text{ for } j = 1, \ldots, m. \tag{22}$$

Each of the constraints (7)–(19) can be converted into the normal form (21) by introducing further dummy variables. For instance, for constraint (7) we define $b_2 = e_{t+1,N}$, $a_1 = 1 - \sum_{i=1}^{L} x_{t,i}$ and $b_1 = e_{t+1,n} - e_{t,n+1}$. The bounds are $\mathcal{L}_1 = -C, \mathcal{L}_2 = 0, \mathcal{U}_1 = \mathcal{U}_2 = C$. The implication constraint (7) is thus equivalent to the following two linear constraints

$$-C(1 - \sum_{i=1}^{L} x_{t,i}) \leq e_{t+1,n} - e_{t,n+1} \leq C(1 - \sum_{i=1}^{L} x_{t,i}), \quad t = 1, \ldots, 2N - 1; n = 1, \ldots, N - 1 \tag{23}$$

$$0 \leq e_{t+1,N} \leq C(1 - \sum_{i=1}^{L} x_{t,i}), \quad t = 1, \ldots, 2N - 1. \tag{24}$$

3.2. Solution quality

We now demonstrate that the solutions of the problem variant "flexible storage and retrieval" are superior to "store-then-retrieve" regarding the number of color changes. We first provide an example to illustrate the point. Subsequently, we formulate and prove the theorem stating that the solution quality of the variant "store-then-retrieve" can be arbitrarily worse compared to "flexible storage and retrieval".

Example 1. Consider a simple problem instance with two colors (coded as 1 and 2), a 2x2 buffer, and the incoming sequence $[2, 1, 1, 1, 2, 1, 1, 1]$. We assume that the sequence is processed in the two equal blocks $[2, 1, 1, 1]$ and $[2, 1, 1, 1]$. As a consequence, two iterations of "store-then-retrieve" will cause two color changes, one for each block since the last car of color 2 must also be stored in the buffer. However, if we instead generate a sequence with fully flexible store and retrieve operations, we can store and retrieve the first three 1s without any color change. Subsequently, we store the 2 in the first lane. We then store and retrieve the next three 1s through the second lane. Finally, we store the last 2 and retrieve it with the other 2, which results in only one color change.

Although this example seems small, it is actually generic in the sense that we can extend the problem instance to obtain an arbitrarily higher number of color changes. Also note that "flexible

---

[3]Note that Gurobi supports implication constraints, which are called indicator constraints in the Gurobi manual (Gurobi Optimization, LLC, 2023, p. 515).

storage and retrieval" always results in solutions which are as least as good as solutions generated by "store-then-retrieve".

Theorem 1. *Consider the paint shop problem with more than one buffer lane $L > 1$, more than one color $C > 1$, and lane width $W > 1$. For every number $n \in \mathbb{N}$ there is an incoming sequence for which "store-then-retrieve" causes at least n more color changes than "flexible storage and retrieval."*

PROOF. See Appendix B of the supplementary material.

4. Reinforcement learning approach

We now describe our reinforcement learning approach. For this purpose, we first formulate the paint shop problem as a reinforcement learning problem, consisting of an environment, state and action space as well as transition and reward function. Subsequently, we describe how action masking is used to incorporate prior human knowledge of optimal strategies. Finally, we explain the employed method (proximal policy optimization) for policy learning.

4.1. Environment

The environment models the whole paint shop, including store and retrieve operations which transform the incoming sequence into the outgoing sequence. The discrete time $t = 1, \ldots, 2n$ refers to the full sequence of store and retrieve operations. Each learning episode starts with an incoming sequence of length N and an empty buffer. At each time step, the environment provides the RL agent with the current state, which contains the buffer content, the most recent colors from the incoming sequence, and the current color of the last element in the outgoing sequence. The RL agent can then either store the current car from the incoming sequence in the rightmost empty position of a buffer lane or retrieve the car from the last position of a buffer lane which is then appended to the outgoing sequence. The environment simulates the performed action by updating the incoming sequence, buffer content, and outgoing sequence. A learning episode ends when the initial sequence and buffer are both empty. The environment also calculates the reward for the agent. We implement the environment in Python 3.8.0, using the "OpenAI Gym" package in version 0.21.0 (Brockman et al., 2016).

4.2. State representation

The state contains the buffer content, the K next colors from the incoming sequence, and the current color of the last car in the outgoing sequence

$$s_t = (\underbrace{B_{t,1,1}, \ldots, B_{t,L,W}}_{\text{Buffer content}}, \underbrace{e_{t,1}, \ldots, e_{t,K}}_{\substack{\text{K next incoming cars} \\ \text{of input sequence}}}, \underbrace{p_t}_{\substack{\text{Current} \\ \text{painting color}}}). \tag{25}$$

Each value $(0, 1, \ldots, C)$ either denotes an empty position (encoded as 0) or a color (encoded as $1, \ldots, C$). We determine the value of the look-ahead length ($K = 5$) through a pre-study, see

10

Appendix C of the supplementary material. However, encoding a color by an integer number $\{0, \ldots C\}$ suggests that colors can be ordered in an ordinal way from low to high, although such an ordering is not present as colors reflect categorical values. Accordingly, we represent colors using a one-hot encoding, so that each color $c \in \{0, \ldots C\}$ is mapped to a vector $o(c) = (0, \ldots, 1, \ldots, 0)$ of length C, which has zeros throughout, except from a one at position c. The vector consisting of zeros only represents an empty buffer or sequence position. The one-hot state representation is thus given as

$$s_t^{\text{one-hot}} = (\underbrace{o(B_{t,1}), \ldots, o(B_{t,L,W})}_{\text{Buffer content}}, \underbrace{o(e_{t,1}), \ldots, o(e_{t,K})}_{\substack{K \text{ next incoming cars} \\ \text{of input sequence}}}, \underbrace{o(p_t)}_{\substack{\text{Current} \\ \text{painting color}}}). \tag{26}$$

We only consider the one-hot state representation in the remainder of our study.

4.3. Action space and transition function

At each timestep, the RL agent can either store the current car from the incoming sequence in a buffer lane or retrieve the rightmost car from a buffer lane and append it to the outgoing sequence. Therefore, the action space A consists of retrieve $(1, \ldots, L)$ and store actions $(L + 1, \ldots, 2L)$

$$A = \{\underbrace{1, \ldots, L}_{\substack{\text{Retrieve} \\ \text{actions}}}, \underbrace{L + 1, \ldots, 2L}_{\substack{\text{Store} \\ \text{actions}}}\}. \tag{27}$$

A retrieve action $a \in \{1, \ldots, L\}$ retrieves the rightmost car from lane a, while a store action $a \in \{L + 1, \ldots, 2L\}$ stores the current car from the incoming sequence in lane $a - L$.

Transition function

The transition function denotes how the paint shop evolves from state s_t to s_{t+1} if a store or retrieve action is performed. If an action is invalid, that is, attempting to store a car in a full lane, store once the incoming sequence is empty, or retrieving from an empty lane, the transition function returns the current state s_t.

Storage

A store action $a \in \{L + 1, \ldots, 2L\}$ stores the current car from the incoming sequence in the rightmost empty position in buffer lane $a - L$. All cars from the incoming sequence then move one position forward, while the color of the last car in the outgoing sequence does not change. The transition function for storage is formalized in Algorithm 1.

Retrieval

A retrieve action $a \in \{1, \ldots, L\}$ appends the rightmost car in lane a to the outgoing sequence. The current color is updated to the color of the retrieved car and all other cars of lane a move one position to the right. The incoming sequence does not change. This is formalized in Algorithm 2.

11

Algorithm 1 Transition function: Storage

1: **input:** state s_t, action $a \in \{L+1, \ldots, 2L\}$ // store car in lane $a-L$
2: **if** $B_{t,a-L,1} \neq 0$ **or** $e_{t,1} = 0$ **then**
3:     **output:** s_t // invalid action, return current state
4: $p_{t+1} = p_t$ // current color in outgoing sequence does not change
5: **for all** $n = 2, \ldots, N$ **do**
6:     $e_{t+1,n} = e_{t,n}$ // update incoming sequence
7: $e_{t+1,N} = 0$ // first position becomes empty
8: $j^* = \max_j$ so that $B_{t,a-L,j} = 0$
9: $B_{t+1,a-L,j^*} = e_{t,1}$ // car is placed in rightmost available space in lane $a-L$
10: **for all** $j = 1, \ldots, W$ with $j \neq j^*$ **do**
11:     $B_{t+1,a-L,j} = B_{t,a-L,j}$ // no changes to other cars in lane $a-L$
12: **for all** $i = 1, \ldots, L$ with $i \neq a - L$ **do**
13:     $B_{t+1,i,j} = B_{t,i,j}$ // no changes in other lanes
14: **output:** $s_{t+1} = (B_{t+1,1,1}, \ldots, B_{t+1,L,W}, e_{t+1,1}, \ldots, e_{t+1,K}, p_{t+1})$

Algorithm 2 Transition function: Retrieval

1: **input** state s_t, action $a \in \{1, \ldots, L\}$ // retrieve from lane a
2: **if** $B_{t,a,W} = 0$ **then**
3:     **output:** s_t // invalid action, return current state
4: $p_{t+1} = B_{t,a,W}$ // current color is updated
5: **for all** $n = 1, \ldots N$ **do**
6:     $e_{t+1,k} = e_{t,k}$ // no change in incoming sequence
7: $B_{t+1,a,1} = 0$ // left-most entry of lane a becomes empty
8: **for all** $j = 2, \ldots, W$ **do**
9:     $B_{t+1,a,j} = B_{t,a,j-1}$ // move cars to the right in lane a
10: **for all** $i \neq a$ and $j = 1, \ldots, W$ **do**
11:     $B_{t+1,i,j} = B_{t,i,j}$ // no changes in other lanes
12: **output:** $s_{t+1} = (B_{t+1,1,1}, \ldots, B_{t+1,L,W}, e_{t+1,1}, \ldots, e_{t+1,K}, p_{t+1})$

4.4. Reward function

The reward function should guide the RL agent towards performing valid actions, while minimizing the number of color changes. Accordingly, the reward function penalizes invalid actions and retrieve actions causing color changes. Invalid actions are penalized with a reward of -10. Valid store operations and retrieve actions causing a color change are assigned a zero reward. However, valid retrieve operations that do not cause a color change yield a positive reward of one. Taken together, the reward function is given as

$$
r(s_t, a) = \begin{cases} 0, & \text{if } a \leq L \text{ and } B_{t,a,W} \neq 0 \text{ and } B_{t,a,W} \neq p_t \text{ (retrieval with color change)} \\ 1, & \text{if } a \leq L \text{ and } B_{t,a,W} \neq 0 \text{ and } B_{t,a,W} = p_t \text{ (retrieval without color change)} \\ -10, & \text{if } a \leq L \text{ and } B_{t,a,W} = 0 \text{ (invalid retrieval action)} \\ 0, & \text{if } a > L \text{ and } B_{t,a-L,1} = 0 \text{ and } e_{t,1} \neq 0 \text{ (valid store action)} \\ -10, & \text{if } a > L \text{ and } B_{t,a-L,1} \neq 0 \text{ or } e_{t,1} = 0 \text{ (invalid store action)}. \end{cases}
\tag{28}
$$

4.5. Properties and action masking

The paint shop problem has several useful properties, which we can exploit using action masking. An action mask is a function $m: (s_t, a) \rightarrow \{0, 1\}$ that can be employed to reduce the set of admissible actions in a given state. An action a is admissible in state s_t if $m(s_t, a) = 1$. Specifically, it may be desired to reduce the set of admissible actions to (i) valid actions, (ii) provably optimal actions, or (iii) actions suggested by heuristics. We thus consider different action masks that enforce valid actions, as well as the provably optimal strategies "greedy retrieval" and so-called "fast-track actions." Furthermore, we consider an action mask that enforces greedy store actions. Thereby, we aim to increase the efficiency of the learning process by avoiding invalid actions and directly incorporating prior human knowledge about the problem. An illustration of all considered action masks is provided in Figure 3. Subsequently, we consider each action mask in detail.

Invalid actions

Our first action mask m^{INV} excludes invalid actions

$$
m^{\text{INV}}(s_t, a) = \begin{cases} 1, & \text{if } a \in \{1, \dots, L\} \text{ and } B_{t,a,W} \neq 0 \\ 1, & \text{if } a \in \{L+1, \dots, 2L\} \text{ and } B_{t,a-L,1} = 0 \text{ and } e_{t,1} \neq 0 \\ 0, & \text{else.} \end{cases} \tag{29}
$$

Adding this action mask hence waves the need to penalize invalid actions with negative rewards. Thus, the learning episodes can also be expected to be shorter, resulting in a more efficient learning process overall.

(a) Invalid action mask: Only storing in lane 1 or 3 and retrieving from lane 1 or 2 is allowed.

(b) Greedy retrieval action mask: Only retrieving from lane 1 or 2 is allowed.

(c) Fast track action mask: Only storing in lane 3 is allowed.

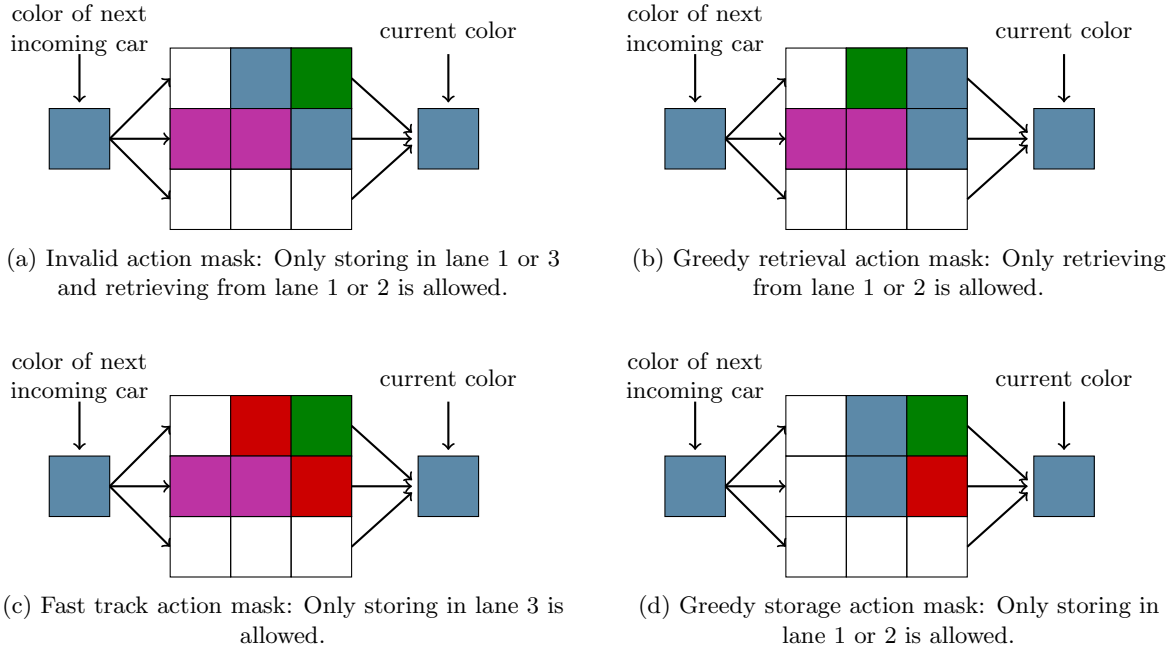(d) Greedy storage action mask: Only storing in lane 1 or 2 is allowed.

Figure 3: Illustration of the four considered action masks.

Greedy retrieval

An action is called a "greedy retrieve action" if it retrieves a car from the buffer that has the same color as the current color. To show that greedy retrieval is optimal, we first introduce the cumulative reward that arises when performing a sequence of actions $a_1, \ldots, a_n$ starting from a given state s_0

$$r(s_0, a_1, \ldots, a_n) = \sum_{t=0}^{t=n-1} r(s_t, a_{t+1}). \tag{30}$$

We call an action a^* *optimal* in state s_t if a sequence that maximizes the cumulative reward in s_t starts with a^*, formally

$$\max_{a_1, \ldots, a_n \in A} r(s_t, a^*, a_1, \ldots, a_n) \geq r(s_t, a', a'_1, \ldots, a'_n), \ \forall a', a'_1, \ldots, a'_n \in A. \tag{31}$$

Note that there can be multiple optimal actions in a given state. We define the optimality of a sequence of actions accordingly.

Theorem 2. *Greedy retrieve actions are optimal.*

PROOF. See Appendix B of the supplementary material.

Following Theorem 2, we consider the greedy retrieval action mask $m^{\text{GR}}(s_t, a)$

$$m^{\text{GR}}(s_t, a) = \begin{cases} 1, \text{ if } a \in \{1, \ldots, L\} \text{ and } p_t = B_{t,a,W} \\ 1, \text{ if } p_t \neq B_{t,i,W} \ \forall i = 1, \ldots, L \\ 0, \text{ else.} \end{cases} \tag{32}$$

The mask enforces greedy retrieval if it is possible. If greedy retrieval is not possible, the mask allows all actions.

Fast-track actions

We refer to a "fast-track action" as storing a car to an empty buffer lane in order to directly retrieve in the next step without causing a color change. Given optimality of greedy retrieval and the fact that store actions do not receive a reward, we can directly derive that fast-track actions are also optimal.

Corollary 2.1. *Fast-track actions are optimal.*

The fast-track action mask m^{FT} enforces store actions to empty lanes if a subsequent greedy retrieval from this lane is possible, consistent with Corollary 2.1. Conversely, the mask allows all actions if a fast-track action is not possible

$$m^{\text{FT}}(s_t, a) = \begin{cases} 1, \text{ if } a \in \{L+1, \ldots, 2L\} \text{ and } B_{t,a-L,1} = 0 \text{ and } p_t = e_{t,1} \\ 1, \text{ if } p_t \neq e_{t,1} \text{ or } B_{t,i,1} \neq 0, \ \forall i = 1, \ldots, L \\ 0, \text{ else.} \end{cases} \tag{33}$$

Greedy storage

Finally, we consider "greedy storage actions," i.e., storing the current car from the incoming sequence in a buffer lane, for which the leftmost car is of the same color (e.g., Leng et al., 2020, 2023). We hence evaluate a fourth action mask $m^{\mathrm{GS}}(s_t, a)$ given as

$$
m^{\mathrm{GS}}(s_t, a) = \begin{cases} 1, \text{ if } a \in \{L+1, \ldots, 2L\} \text{ and } \exists j \in \{2, \ldots, W\}: e_{t,1} = B_{t,a-L,j} \text{ and } B_{t,a-L,j-1} = 0 \\ 1, \text{ if } \nexists i \in \{1, \ldots, L\}, j \in \{2, \ldots, W\}: e_{t,1} = B_{t,i,j} \text{ and } B_{t,i,j-1} = 0 \\ 0, \text{ else.} \end{cases}
$$

$$(34)$$

Although greedy storage presents a reasonable policy, we cannot prove that it is indeed optimal.

We also want to note that employing the aforementioned action masks does not fully solve the paint shop problem. In fact, greedy retrieval, greedy storage, and fast-track actions are often not applicable. So far, we proposed several action masks independently. However, as detailed in the following, we also combine and prioritize multiple action masks.

Combining action masks

We consider two approaches of combining action masks. First, we define the combination $m_1 \oplus m_2$, which only allows an action a to be performed in state s_t, if both masks allow a

$$
(m_1 \oplus m_2)(s_t, a) = \begin{cases} 1, \text{ if } m_1(s_t, a) = 1 \wedge m_1(s_t, a) = 1 \\ 0, \text{ else.} \end{cases}
$$

$$(35)$$

The second combination $m_1 \oslash m_2$ applies two action masks in a sequential way, while assigning higher priority to m_1. This approach is particularly relevant when we want to prioritize one mask over another. We consider an action mask as applicable if it forbids at least one action. Hence, the combined mask $m_1 \oslash m_2$ applies m_1 if at least one action is not allowed by m_1, and m_2 otherwise

$$
(m_1 \oslash m_2)(s_t, a) = \begin{cases} m_1(s_t, a), \text{ if } \exists a: m_1(s_t, a) = 0 \\ m_2(s_t, a), \text{ else.} \end{cases}
$$

$$(36)$$

Following this notation, we evaluate the following combinations of action masks

- $m^{\mathrm{INV}} \oplus \left(\left(m^{\mathrm{GR}} \oslash m^{\mathrm{FT}} \right) \oslash m^{\mathrm{GS}} \right)$ (all action masks)

- $m^{\mathrm{INV}} \oplus \left(m^{\mathrm{GR}} \oslash m^{\mathrm{FT}} \right)$ (invalid + greedy retrieval + fast-track)

- $m^{\mathrm{INV}} \oplus m^{\mathrm{GR}}$ (invalid + greedy retrieval)

- m^{INV} (invalid)

4.6. Policy learning

Reinforcement learning alternates between applying the current policy $\pi_\theta(a|s_t)$ to generate experience in the form of trajectories $(s_1, a_1, r_1), \ldots, (s_H, a_H, r_H)$ and using the newly generated experience to perform policy updates. Here, H specifies the size of the experience set. We implement proximal policy optimization (PPO, Schulman et al., 2017) for policy learning as we found it to perform well for sequential decision problems in operations research (Brammer et al., 2022a,b, 2021), while being easy to implement and tune. In PPO, the policy is implemented as a neural network that takes the state s_t as input and outputs a probability distribution over the action space. The values $\pi_\theta(a|s_t)$ thus reflect the probabilities that the agent performs action a in state s_t. The policy parameters θ present the weights of the neural network.

We briefly describe how PPO works. All explanations are based on Schulman et al. (2017). Algorithm 3 provides a (simplified) pseudocode of policy learning with PPO. Besides the previously mentioned policy network $\pi_\theta(a|s_t)$, PPO also maintains an estimate of the value function $V^\theta(s_t)$. This value should be close to the expected reward that the agent receives when starting from state s_t and following the current policy π_θ. Given a trajectory $(s_1, a_1, r_1), \ldots, (s_H, a_H, r_H)$, PPO calculates the advantage $\hat{A}_t = R_t - V^\theta(s_t)$ with $R_t = \sum_{t=t'}^{H-1} r_{t'}$ that measures how much better (or worse) the given trajectory performed compared to the state-value estimate $V^\theta(s_t)$.[4] In summary, PPO updates its policy parameters by maximizing the following objective function $L(\theta)$

$$\max_\theta \; L(\theta) = \sum_{t=1}^{H} L_t^{CLIP}(\theta) - c_1 L_t^{VF}(\theta) + c_2 L_t^{H}(\theta). \tag{37}$$

$L(\theta)$ contains three terms. The purpose of the first term $L_t^{CLIP}(\theta)$ is to update the parameters θ such that actions with positive advantage $\hat{A}$ are assigned higher probabilities, while actions with negative advantage are assigned lower probabilities. This is the general idea of policy gradient methods. However, the novelty of PPO is to limit the extent of the policy updates by clipping the ratio between new and old probabilities $\frac{\pi_\theta(a|s_t)}{\pi_{\theta_{old}}(a|s_t)}$ to the range $[1 - \varepsilon, 1 + \varepsilon]$. The hyperparameter ε is set to 0.20 per default.

$$L_t^{CLIP}(\theta) = \min \left\{ \frac{\pi_\theta(a|s_t)}{\pi_{\theta_{old}}(a|s_t)} \hat{A}_t, \text{clip}\left(\frac{\pi_\theta(a|s_t)}{\pi_{\theta_{old}}(a|s_t)}, 1 - \varepsilon, 1 + \varepsilon \right) \hat{A}_t \right\} \tag{38}$$

The second term $L_t^{VF}(\theta) = (R_t - V^\theta(s_t))^2$ denotes the squared error in predicting the value of state s_t. Including this term with a negative sign ensures that updating the network weights θ also improves the value function estimate.

The third term $L_t^{H}(s_t)$ denotes the entropy (Shannon, 1948) of the policy in state s_t: $L_t^{H}(s_t) = -\sum_{a \in A} \pi_\theta(a|s_t) \log_2 \pi_\theta(a|s_t)$. Higher entropy values indicate that the probability distribution is less concentrated and more random in the sense that the agent can still explore different

---

[4]Note that there are different ways of estimating the advantage depending on the particular PPO implementation. Another common approach is the generalized advantage estimation (Schulman et al., 2015).

actions instead of solely relying on only the action, that is the best according to the current policy. All parameters except the number of time steps are set to their default values as stated in Appendix A of the supplementary material. We implement our RL approach in Python 3.8.0 using the PPO implementation from the RL framework "Stable Baselines 3" in version 1.5.0. We train all models for a total of 10 million timesteps.[5] Training a single policy on a machine with an AMD EPYC 7702 processor and 64 GB main memory takes approximately ten hours.

Algorithm 3 Policy learning with proximal policy optimization (Schulman et al., 2017)

1: **input:** policy π_θ, action mask m, discount parameter γ, size of the experience set H
2: initialize state value estimates $V^\theta(s_t)$, $\theta_{old} = \theta$
3: **while** total number of timesteps not reached **do**
4:     Use current policy π_θ and action mask m to generate trajectory $(s_1, a_1, r_1), \ldots, (s_H, a_H, r_H)$
5:     **for all** $t = 1, \ldots, H$ **do**
6:         $R_t = \sum_{t'=t}^{H-1} \gamma^{t'-t} r_{t'}$ // return from s_t
7:         $\hat{A}_t = R_t - V^\theta(s_t)$ // advantage estimate
8:     Update θ by maximizing $L(\theta)$ based on $(s_1, a_1, R_1, \hat{A}_1), \ldots, (s_H, a_H, R_H, \hat{A}_H)$
9:     $\theta_{old} = \theta$

---

5. Evaluation

5.1. Procedure

We perform a total of five analyses. First, the main analysis compares our RL approach against all baselines. The problem instances of all analyses (except the out-of-distribution instances) are generated from a distribution with equal probability for each color. The length of the incoming sequences is fixed to 100. We impose a cutoff time of five minutes for all approaches to model the necessity to quickly adapt to new incoming orders and short-term changes in production. Second, we compare our RL approach against multiple baselines with an extended cutoff time of ten hours. We expect Gurobi to solve small instances optimally, which allows us to assess the performance of our approach in comparison to the (near-)optimal Gurobi solutions. Third, we lift the assumption of square buffers and consider rectangular buffer sizes. Fourth, we consider problem instances with an initially filled buffer. While an initially empty buffer reflects the academic notion of the problem, real-world production never stands still. Hence, real-world instances of the paint shop problem also feature filled buffers. Fifth, we consider imbalanced (instead of balanced) color distributions. However, we deliberately do not retrain our RL approach to understand how it performs on out-of-distribution problems. Changes in the color distribution can, for instance, occur due to an unforeseen shortage of painting color and other supplies.

An overview of the considered problem instances is provided in Table 3. We consider 5–15 colors, while the buffer size (lanes x lane width) ranges from 2x2 to 8x10. For all instances, we require the number of buffer lanes to be smaller than the number of colors to ensure that the problem does not become trivial. Our problem parameters are similar to those of existing studies,

---

[5]We conducted a convergence analysis which suggests that no further improvement is achieved when training longer. The results of the convergence analysis are presented in Appendix G.

e.g., Leng et al. (2020) and Sun & Han (2017), who used 10, 14, and 20 colors and buffer sizes of 5x6, 7x8, and 10x10. Our problem instances are also publicly available, see Stappert et al. (2023).

Table 3: Problem parameters.

Analysis	Sequence length	Number of colors	Buffer size (lanes x lane width)
Main analysis	100	5, 10, 15	2x2, 3x3, 4x4, 5x5, 6x6, 7x7, 8x8
Evaluation with 10h cutoff time	100	5, 10, 15	2x2, 3x3, 4x4, 5x5, 6x6, 7x7, 8x8
Rectangular buffers	100	10	4x5, 4x6, 4x8, 5x6, 5x7, 5x8, 5x10, 6x4, 6x10, 7x8, 8x10
Initially filled buffer	100	5, 10, 15	2x2, 3x3, 4x4, 5x5, 6x6, 7x7, 8x8
Imbalanced color distribution	100	5, 10, 15	2x2, 3x3, 4x4, 5x5, 6x6, 7x7, 8x8

5.2. Baselines

We compare our RL approach against the following baselines: Gurobi based on our ILP formulation, Gurobi based on the store-then-retrieve formulation by Wu et al. (2021), the greedy heuristic, alternated greedy storage and RL for retrieval (Huang et al., 2024; Leng et al., 2020, 2023), multiple sequence alignment (MSA), and variants of the metaheuristic simulated annealing.

Gurobi

We implement our ILP formalization from Section 3 in Gurobi using "gurobipy" version 10.0.2. Gurobi is a state-of-the-art commercial ILP solver, free for academic use, that employs optimized branch-and-bound algorithms. Implication constraints are also supported by the "gurobipy" package, which allows us to enter the constraints without prior normalization.

Gurobi based on store-then-retrieve problem formulation (Wu et al., 2021)

We also evaluate Gurobi based on a modified version of the ILP formalization by Wu et al. (2021). We adapt their problem formulation to our problem in the following way: First, we remove car types to make all cars generic, i.e., of the same type. Second, we unify the costs of color changes to one as Wu et al. (2021) allow the costs of color changes to differ, e.g., changing from red to blue might be more costly than changing from red to green. Third, we exclude the assembly shop and focus only on the paint shop. Accordingly, we also place the buffer system in front of the paint shop instead. The adjusted ILP formulation is provided in Appendix H of the supplementary material. Wu et al. consider the paradigm "store-then-retrieve". Accordingly, we iteratively repeat the "store-then-retrieve" procedure until the entire incoming sequence is processed.

Greedy

We implement a greedy solution policy following the proposed action masks that enforce greedy retrieval, fast-track actions, and greedy storage. The greedy policy runs until the incoming sequence and buffer are both empty. Greedy alternates between storage and retrieval phases, while starting with the storage phase. Here, the heuristic performs greedy storage if possible. Otherwise, the

current car is stored in the least occupied lane. In case of ties, the lane with the lowest index is used. Once the buffer is filled, greedy changes to the retrieval phase. The heuristic performs greedy retrieval if possible and retrieves from the lane with the lowest index otherwise. Accordingly the heuristic is fully deterministic, which is preferable in regard to reproducibility. Note that fast-track actions are implicitly performed by the proposed greedy heuristic since cars are stored in an empty lane in the storage phase. The next retrieve operation in the retrieval phase will perform the greedy retrieval. The pseudocode of Greedy is provided in Appendix I of the supplementary material.

We evaluate another heuristic that changes between storage and retrieval phase based on the buffer fill rate. The fill rate is defined as the current number of cars in the buffer divided by the total size of the buffer (LW). Focusing on the fill rate is motivated by prior studies (e.g., Huang et al., 2024; Sun & Han, 2017) who aim to maintain certain buffer fill rates. The greedy approach based on buffer fill rate has two parameters $\mathcal{U}, \mathcal{L}$ denoting the upper and lower bounds of the fill rate. For our evaluation, we consider all combinations $\mathcal{U}, \mathcal{L} \in \{0, 0.01, 0.02, \ldots, 1\}$ and report the best respective result for each instance. An even more fine-grained parameter space does not yield better results as for the largest buffer size of our dataset, i.e., 8x10, any difference of fill rates is already greater than $1/(8 \cdot 10) > 0.01$. The pseudocode of "Greedy based on fill rate" is provided in Appendix I of the supplementary material.

Alternated greedy storage and reinforcement learning for retrieval

We consider another baseline that follows the RL approaches by Leng et al. (2020), Leng et al. (2023), and Huang et al. (2024). The buffer is first filled using greedy storage until it reaches the desired fill rate of 60% (Huang et al., 2024; Sun & Han, 2017). Subsequently, the baseline alternates between inserting one car into the buffer with greedy storage and then retrieving one car with RL. Note that these studies focused on a different problem variation that includes non-uniform costs of color changes and model types, whereas our problem has uniform costs of color changes and model types. Therefore, combining greedy storage with our own RL policy should yield a reasonable estimate of the performance of the existing RL approaches for the paint shop problem with multi-lane buffers.

Multiple sequence alignment

We implement multiple sequence alignment (MSA), a dynamic programming approach inspired by bioinformatics and introduced by Epping & Hochstättler (2003). MSA is a retrieval-only algorithm that works by first "merging" consecutive cars of the same color, which will be retrieved together. MSA then finds an optimal alignment of buffer columns, so that as many cars in a column have the same color, and then retrieves them all together. In this way, MSA forms color batches across lanes and columns simultaneously. Given a buffer state, the retrieval sequence that MSA produces is optimal as shown by Epping & Hochstättler (2003). Although MSA only provides a retrieval policy, we combine it with greedy storage, to obtain a full storage and retrieval method. MSA has a time complexity of $\mathcal{O}\left(2^L \prod_{i=1}^{L} W_i\right)$, where W_i is the number of cars in buffer lane i.

As a consequence, this approach requires longer than the cutoff time of five minutes to produce a solution for instances with 9x9 buffer. Multiple sequence alignment is implemented in C++14.

Simulated annealing

We implement the metaheuristic simulated annealing (see e.g., Gendreau & Potvin, 2018). Simulated annealing iteratively improves an initial action sequence by performing random permutations. If the resulting modified sequence reduces the number of color changes, the modified sequence is directly accepted as the new sequence. Otherwise, a modified sequence can still be accepted as the new sequence depending on the current temperature. The temperature is higher in early iterations to allow for exploration but it cools down after each iteration. Here, we implement geometric cooling ($T_i = \beta^i T_0$) as it is a frequently used cooling schedule (Goodson, 2015; Knopp et al., 2017) and logarithmic cooling ($T_i = \frac{T_0}{\log(1+i)}$) due to its favorable theoretical properties (Hajek, 1988). We determined the initial temperature T_0 and cooling factor β through a pre-study, see Appendix D of the supplementary material.

We evaluate simulated annealing based on two initial action sequences. First, we consider the solution of the greedy heuristic as it creates a reasonable initial solution using several optimal heuristics. Second, we consider the simple action sequence $(L, 0, L, 0, \ldots, L, 0)$ that stores each car to the first lane and immediately retrieves it. This action sequence is less guided by human knowledge, which allows the metaheuristic to better explore the entire solution space. Simulated annealing is implemented in C++14.

5.3. Application of a trained reinforcement learning policy

The trained policy outputs a probability distribution over all actions for a given state. Based on this, we apply a trained policy in two ways. First, we apply the policy in a deterministic manner by always selecting the action with the highest probability. Second, we use the available cutoff time to continuously sample actions from the policy to search for the best solution. Thereby, we may find superior solutions than those of the deterministic application. In all analyses, we ignore the time needed for policy learning as this can be performed in advance for a given paint shop.

5.4. Performance metrics

Our main performance metric is given by the *Average Relative Percentage Deviation* (ARPD, Osman, 1993; Fernandez-Viagas et al., 2017; Abreu et al., 2020). To define ARPD, we first define the relative percentage deviation (RPD). Let $c^{i,A} = \sum_{t=1}^{2N-1} z_t$ denote the number of color changes of approach A on problem instance i. In addition, let c^{i,A^*} denote the total number of color changes of the best solution on instance i that was found by some algorithm A^*. $RPD^{i,A}$ is then defined as

$$\text{RPD}^{i,A} = \frac{c^{i,A} - c^{i,A^*}}{c^{i,A^*}} \cdot 100. \tag{39}$$

Note that $c^{i,A^*} \neq 0$ for every problem with more than one color. Now, let $\mathcal{I}$ denote the set of all considered problem instances. The ARPD of approach A is given as

$$\mathrm{ARPD}^A = \frac{\sum_{i \in \mathcal{I}} \mathrm{RPD}^{i,A}}{|\mathcal{I}|}. \tag{40}$$

If $\mathrm{ARPD}^A = 3.20$, the outgoing sequences generated by approach A have on average $3.20\,\%$ more color changes than the best solution that was found by some approach. Note that the best relative solution is not necessarily an optimal solution.

The metric ARPD is well suited to compare the relative performance of different approaches but it does not provide any information regarding the absolute number of color changes. Therefore, we also report the mean number of color changes caused by approach A

$$c^A = \frac{\sum_{i \in \mathcal{I}} \sum_{t=1}^{2N-1} z_t}{|\mathcal{I}|}. \tag{41}$$

6. Results

6.1. Main analysis

We first consider the results of our main analysis, which compares the performance of all approaches on problem instances with quadratic buffer and balanced color distribution. The results are shown in Table 4. Due to spatial limitations, we include only a subset of the results in Table 4 and exclude buffer sizes 3x3, 5x5, and 7x7 for 10 and 15 colors. The remaining results and computation times are provided in Appendices E and F of the supplementary material.

We find that our RL approach with all action masks and repeated sampling performs best on most problem instances. However, on the problems with 15 colors and 2x2 buffer, Gurobi achieves the best performance as they are comparably easier to solve due to the small buffer size. On the problems with 10 or 15 colors and 8x8 buffers, the Greedy heuristic based on buffer fill rate presents the best performing approach. A possible explanation could be that RL struggles to sufficiently explore the considerably larger state space, caused by the increased problem complexity. At the same time, the ARPD of RL with all action masks included and stochastic policy application is at most 15.0 for the problems with 8x8 buffer. Furthermore, we observe that repeated sampling from the RL policy always outperforms its respective counterparts of RL with deterministic policy application.

The best baseline is given by Greedy based on the buffer fill rate which achieves the lowest ARPD values for the problems with 10 and 15 colors. As mentioned before, it even outperforms all other approaches on the problems with 8x8 buffer. On the other problems with 10 and 15 colors and smaller buffers, it achieves ARPD values between 7.1 (6x6 buffer) and 22.1 (3x3 buffer). In particular, Greedy based on the buffer fill rate consistently outperforms Greedy following the paradigm "store-then-retrieve" which fills the buffer first using greedy storage to empty it later using greedy retrieval, if applicable. Multiple sequence alignment performs poorly on the problem

21

instances with five colors but its relative performance increases on larger problem instances with 10 and 15 colors. On the problems with 15 colors and 6x6, 7x7 buffer, it even presents the second best approach following our RL approach with all action masks and sampling.

Applying Gurobi to different ILP formulations does generally not lead to competitive results. Moreover, neither ILP problem formulation (ours vs Wu et al. (2021)), is strictly superior to the other. Gurobi based on our problem formulation performs well on problems with 2x2 buffer, whereas Gurobi based on Wu et al. (2021) achieves its lowest ARPD on problems with 15 colors and 4x4, 5x5 buffer.[6] Finally, we observe that following the paradigm of alternating between greedy storage and RL for retrieval at each timestep while maintaining a buffer fill rate of 60 % does not yield competitive results. However, decent ARPD values are achieved for the problem instances with 15 colors.

Table 4: Evaluation results (ARPD) for main analysis.

Colors	5			10				15			
Buffer size (lanes x lane width)	2x2	3x3	4x4	2x2	4x4	6x6	8x8	2x2	4x4	6x6	8x8
Reinforcement learning with stochastic policy application (sampling)											
All action masks	2.9	**0.6**	**0.0**	4.7	**0.0**	**0.3**	13.8	4.7	**0.5**	**2.5**	15.0
Invalid + greedy retrieval + fast-track	6.5	5.1	9.2	5.7	9.0	42.6	79.3	4.6	8.5	29.0	47.7
Invalid + greedy retrieval	5.2	9.8	22.0	7.9	16.6	43.1	99.3	5.0	14.6	37.5	62.6
Invalid	8.3	17.1	34.8	6.7	33.1	56.6	108.8	8.4	24.8	50.9	90.2
No action masks	**1.5**	20.7	22.5	**2.2**	25.6	57.2	107.8	2.7	24.9	84.1	151.9
Reinforcement learning with deterministic policy application											
All action masks	10.1	15.4	24.4	10.9	26.4	33.3	62.1	10.4	19.4	29.8	42.8
Invalid + greedy retrieval + fast-track	12.7	22.0	35.4	13.9	34.7	75.9	119.0	12.1	26.3	51.0	80.3
Invalid + greedy retrieval	15.6	30.8	40.6	15.2	36.5	73.2	150.6	12.5	34.6	60.8	93.0
Invalid	17.2	37.4	53.2	13.2	47.1	80.7	149.0	15.5	37.3	68.2	112.0
No action masks	9.6	39.6	49.0	12.4	40.5	93.3	156.8	12.1	43.5	102.6	176.9
Baselines											
Gurobi (our ILP formulation)	15.9	99.4	206.9	9.9	69.9	169.5	341.7	**2.5**	41.1	103.8	172.8
Gurobi (Wu et al., 2021)	32.2	32.4	28.5	19.3	9.3	75.5	233.8	12.9	4.8	42.8	111.4
Greedy	33.6	72.4	101.5	23.3	46.5	62.1	80.5	15.8	38.6	47.8	41.9
Greedy based on fill rate	25.0	39.4	35.9	16.3	19.7	7.5	**2.6**	12.9	12.5	10.2	**4.0**
Alternate greedy storage and RL	19.8	33.2	37.3	13.5	28.5	17.3	18.0	13.6	15.1	22.6	10.0
Multiple sequence alignment	40.6	54.3	67.8	25.7	24.6	14.7	13.8	19.8	14.8	7.9	7.1
Simulated annealing											
Greedy, logarithmic cooling	32.9	63.8	84.8	20.8	34.5	48.7	78.9	14.0	24.0	29.7	41.3
Simple, logarithmic cooling	49.7	121.3	180.5	29.7	54.7	89.1	154.6	19.9	28.9	43.8	70.8
Greedy, geometric cooling	32.3	66.1	94.5	20.6	39.8	52.1	78.0	13.4	31.0	36.8	38.2
Simple, geometric cooling	54.7	127.0	212.2	30.6	68.7	99.8	188.5	13.9	39.1	56.5	82.1

Figure 4 visualizes the mean number of color changes of RL with all action masks and policy sampling with several baselines. The values are averaged across ten instances per parameter setting. Evidently, the black line denoting RL with all action masks and policy sampling usually achieves

---

[6]See Appendix E of the supplementary material.

the lowest number of color changes. However, for problems with 8x8 buffer, RL is outperformed by Greedy based on buffer fill rate. In addition, we observe that the number of color changes decreases for larger buffers, which is expected. Yet, this finding does not hold for the Gurobi baselines, which can be explained by the fact that the problem complexity also increases for greater buffers.

We also analyze the influence of different combinations of action masks and policy applications as shown in Figure 5. Across all color settings, we find that including more action masks improves the overall performance for deterministic and repeated stochastic policy application. In addition, we observe that the improvement of action masks and repeated stochastic policy application depends on the buffer size. While the benefit is rather small for 2x2 buffers, the performance increase of action masking is maximized for large buffers. For 10 colors and 8x8 buffers, including action masks and sampling from the policy reduces color changes by approximately one half compared to deterministic policy application without action masking. This improvement is even greater for 15 colors and 8x8 buffers. Furthermore, we find that, from no action masks and deterministic policy application, the performance benefit of including all action masks dominates the effect of repeated policy sampling. Therefore, for all subsequent analyses, we only focus on the best RL approach, that is RL with all action masks and repeated policy sampling.

(a) Instances with 5 colors.   (b) Instances with 10 colors.   (c) Instances with 15 colors.

—o— RL (all action masks, sampling)   —o— Gurobi   —o— Greedy based on fill rate
—o— Multiple sequence alignment   —o— Gurobi (Wu et al., 2021)   —o— Alternate greedy storage and RL
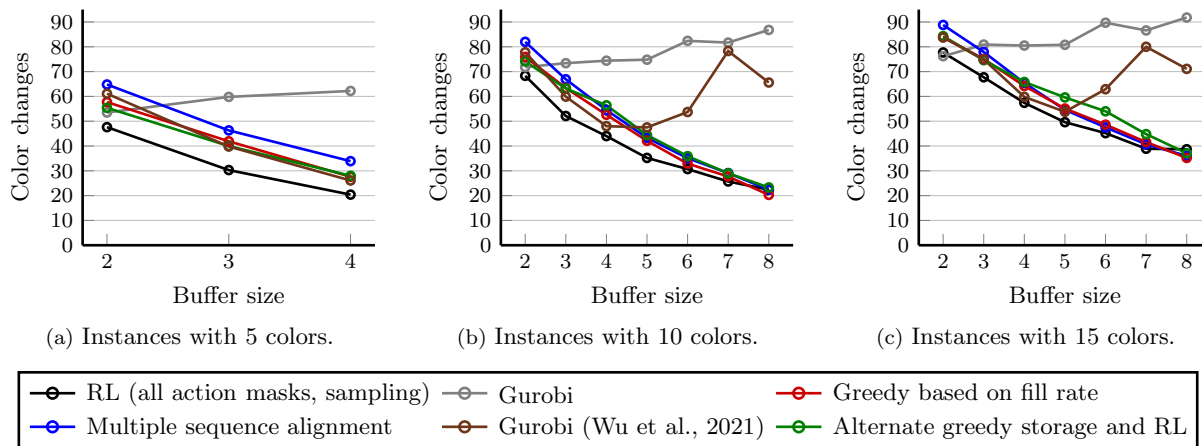
Figure 4: Evaluation results (color changes) of main analysis for best RL approach and baselines.

| (a) Instances with 5 colors. | (b) Instances with 10 colors. | (c) Instances with 15 colors. |

RL (all action masks, sampling) —— RL (all action masks, deterministic)
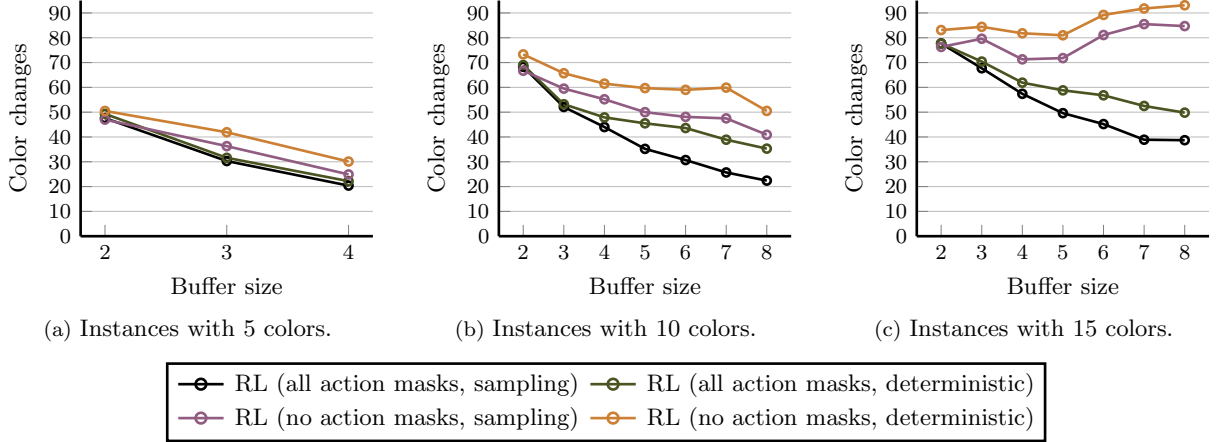RL (no action masks, sampling) —— RL (no action masks, deterministic)

Figure 5: Evaluation results (color changes) of main analysis for different RL policy applications.

6.2. Benchmarks with ten hours time limit

We now compare the solutions of our RL approach with the possible (near-)optimal solutions found by several competing approaches with a time limit of ten hours. Here, we evaluate Gurobi based on our problem formulation, Gurobi based on the "store-then-retrieve" problem studied by Wu et al. (2021), and simulated annealing. Given the high time limit, Gurobi should be able to solve at least several of the small instances optimally. To account for the increased time limit of simulated annealing, we adjust the cooling schedule to ensure that the final temperature for ten hours equals the final temperature after 5 minutes. We achieve this by rescaling the time index for the ten hour computation with the factor $\frac{10h}{5min} = 120$. We specifically keep the sampling time of RL at five minutes like in our main analysis.

The results are provided in Table 5. As expected, we find that Gurobi (based on our ILP formulation) is dominant for the small problem instances with 2x2 buffer. Here, the number of color changes of the RL solutions differ on average by 5.3–11.9% from the near-optimal Gurobi solutions. However, even with a time limit of ten hours, Gurobi never managed to find an optimal solution of the 2x2 problems. For the larger problems with 3x3 buffer and larger, Gurobi (10h) is no longer competitive to RL with all action masks. Nevertheless, the extended cutoff time for Gurobi still leads to considerable reductions in ARPD compared to five minute cutoff time.

Table 5: Evaluation results (ARPD) with time limit 10h.

Colors	5			10				15			
Buffer size (lanes x lane width)	2x2	3x3	4x4	2x2	3x3	4x4	8x8	2x2	3x3	4x4	8x8
RL all action masks (sampling, 5min)	5.3	**0.0**	**0.0**	11.0	**0.0**	**0.0**	**0.0**	11.9	**0.6**	**0.5**	**1.0**
Gurobi (our ILP formulation, 5min)	18.1	98.1	206.9	16.5	41.6	69.9	289.6	9.5	20.3	41.1	139.9
Gurobi (our ILP formulation, 10h)	**0.2**	52.1	107.0	**0.0**	30.4	51.9	184.8	**0.0**	9.6	23.7	91.4
Gurobi (Wu et al., 2021, 5min)	35.3	31.6	28.5	26.5	15.1	9.3	194.2	20.7	11.4	4.8	85.8
Gurobi (Wu et al., 2021, 10h)	35.3	31.6	28.0	26.5	15.1	8.1	45.1	20.7	11.4	4.8	10.6
Simulated annealing (5min)											
Greedy, logarithmic cooling	35.8	62.8	84.8	28.1	34.0	34.5	58.1	21.8	17.4	24.0	24.9
Simple, logarithmic cooling	53.1	119.9	180.5	37.5	47.8	54.7	124.5	28.1	20.2	28.9	49.8
Greedy, geometric cooling	35.2	65.1	94.5	28.0	37.0	39.8	57.3	21.1	18.9	31.0	22.1
Simple, geometric cooling	58.3	125.6	212.2	38.7	53.8	68.7	155.0	21.7	27.0	39.1	60.1
Simulated annealing (10h)											
Greedy, logarithmic cooling	27.6	61.9	82.8	22.3	26.1	25.8	42.7	15.9	10.8	8.6	5.7
Simple, logarithmic cooling	44.4	85.3	99.3	25.8	30.6	27.7	44.0	19.0	10.3	7.8	7.8
Greedy, geometric cooling	27.4	57.2	87.7	22.1	30.0	33.4	54.0	15.9	14.3	21.0	20.5
Simple, geometric cooling	43.3	95.3	134.0	24.1	38.6	44.2	93.8	18.1	17.4	23.0	26.5

Gurobi based on the formulation by Wu et al. (2021) does not benefit from an increased time limit for smaller instances as the 5min cutoff time is already sufficient. However, we observe considerable improvements for larger buffers. In fact, Gurobi (Wu et al., 2021) performs best for problems with 10 colors and 5x5 buffer, as well as 15 colors and 5x5, 6x6 buffer, see Appendix E of the supplementary material. Concerning simulated annealing, we also observe strong improvements of up to 42 percentage points when increasing the cutoff time from 5 minutes to ten hours. Yet, apart from the problems with 15 colors and 4x4 or 8x8 buffers, simulated annealing never achieves a competitive performance compared to RL with all action masks.

6.3. Rectangular buffers

Next, we consider rectangular buffers instead of quadratic buffers. For this analysis, we fix the number of colors to 10 and choose a selection of rectangular buffer sizes that were considered in prior studies (Spieckermann et al., 2004; Lin et al., 2011; Sun & Han, 2017; Huang et al., 2024). The results are shown in Table 7. We find that RL still yields the best results out of all approaches for most buffer sizes. However, the Greedy heuristic based on the buffer fill rate is the best approach for problems with 7x8 and 8x10 buffers. In summary, the results suggest that the performance of our RL approach is not limited to quadratic buffers; instead, it also persists on rectangular buffers.

Table 6: Evaluation results (ARPD) for rectangular buffers.

Buffer size (lanes x lane width)	4x5	4x6	4x8	5x6	5x7	5x8	5x10	6x4	6x10	7x8	8x10
RL all action masks (sampling)	**0.0**	**0.0**	**0.0**	**0.0**	**0.0**	**0.6**	**0.3**	**0.0**	**3.9**	3.7	15.9
<u>Baselines</u>											
Gurobi (our ILP formulation)	79.2	100.4	106.8	127.0	139.3	151.0	144.2	119.2	226.6	209.3	387.6
Gurobi (Wu et al., 2021)	12.8	43.2	73.2	81.7	84.5	110.0	143.8	43.7	167.6	177.3	332.3
Greedy	48.7	52.5	52.8	68.0	73.4	60.1	63.1	70.1	61.2	70.7	59.9
Greedy based on fill rate	17.4	12.7	12.9	16.7	14.5	13.0	10.7	19.2	4.1	**2.8**	**2.1**
Alternate greedy storage and RL	20.1	30.9	23.0	29.5	19.7	27.4	29.5	17.3	19.4	16.3	17.0
Multiple sequence alignment	22.0	20.9	15.1	26.4	15.1	19.5	6.7	28.6	12.2	7.7	26.9
<u>Simulated annealing</u>											
Greedy, logarithmic cooling	36.6	36.2	34.4	49.7	48.1	40.7	50.3	52.7	58.2	58.7	59.9
Simple, logarithmic cooling	50.0	49.3	46.9	72.8	74.0	75.6	74.0	83.5	111.3	112.3	180.7
Greedy, geometric cooling	42.2	44.7	44.8	57.1	62.3	54.4	54.6	62.1	56.6	62.2	57.4
Simple, geometric cooling	68.6	69.6	72.5	92.6	91.5	91.5	97.0	108.5	126.2	131.5	215.9

6.4. Buffer initially filled

We also consider problem instances, where the buffer is not empty but initially filled. Thereby, the problem setting is shifted from the academic notion towards a real-world scenario as buffers are unlikely to be empty due to a halt of production. We thus create novel problem instances with a fully occupied buffer with random content. Note that these instances have a higher complexity as the number of necessary retrieve operations increases from 100 to $100 + L \times W$. We do not retrain the RL policies based on these instances to ensure a fair comparison.

The results are shown in in Table 7. We find that several approaches can no longer solve larger instances, which we denote with "–". We find that, on average, RL yields the best solutions for all buffer sizes. Gurobi reaches the smallest ARPD among all baselines for a buffer size of 2x2, whereas the Greedy based on fill rate and multiple sequence alignment come closest to RL for large buffers.

Table 7: Evaluation results (ARPD) for initially filled buffer.

Colors	5			10				15			
Buffer size (lanes x lane width)	2x2	3x3	4x4	2x2	4x4	6x6	8x8	2x2	4x4	6x6	8x8
RL all action masks (sampling)	**0.0**	**0.0**	**0.0**	**1.8**	**0.0**	**0.6**	**0.6**	**1.9**	**0.0**	**1.5**	**0.0**
<u>Baselines</u>											
Gurobi (our ILP formulation)	10.7	93.1	158.1	5.2	71.3	122.8	—[a]	4.6	40.2	78.1	96.0
Gurobi (Wu et al., 2021)	22.4	36.3	55.5	15.0	20.4	99.3	162.6	9.9	10.6	57.9	92.3
Greedy	26.2	58.5	77.4	18.6	44.0	61.0	69.7	11.4	28.6	48.1	44.6
Greedy based on fill rate	18.3	31.6	26.1	15.0	19.6	18.6	13.4	7.6	11.5	13.8	7.9
Alternate greedy storage and RL	13.9	41.6	60.6	13.5	32.5	49.5	74.2	7.5	20.3	28.9	35.7
Multiple sequence alignment	30.4	49.2	41.5	17.9	19.9	5.7	—[b]	11.6	12.5	2.1	—[b]
<u>Simulated annealing</u>											
Greedy, logarithmic cooling	22.6	49.7	68.4	17.2	31.5	51.9	57.3	9.5	17.6	32.3	28.3
Simple, logarithmic cooling	39.7	106.1	137.4	25.1	51.8	74.9	87.3	13.7	24.0	35.1	37.2
Greedy, geometric cooling	25.5	56.6	76.6	18.6	40.1	58.5	66.5	10.4	25.7	41.4	38.2
Simple, geometric cooling	38.9	117.0	167.6	25.1	64.1	94.8	104.6	12.5	36.0	52.3	48.9

[a] one out of ten instances not solved in cutoff time, [b] no instances solved in cutoff time

6.5. Out-of-distribution instances

We finally assess whether the performance of our RL approach remains robust against out-of-distribution problem instances. In our main analysis, the colors were sampled from a balanced distribution with equal probability for each color $(p_1, \ldots, p_C = 1/C)$. We now consider imbalanced color distributions where the probabilities decrease exponentially or linearly. The exponentially decreasing probabilities decay by a factor of 0.8. Hence, we have $p_1 = \frac{1}{\sum_{i=0}^{C-1} 0.8^i} = \frac{0.2}{1-0.8^C}$ and $p_i = 0.8\,p_{i-1}$, for $i = 2, \ldots, C$.

Table 8: Evaluation results (ARPD) for imbalanced distribution with exponentially decreasing probabilities.

Colors	5			10				15			
Buffer size (lanes x lane width)	2x2	3x3	4x4	2x2	4x4	6x6	8x8	2x2	4x4	6x6	8x8
RL all action masks (sampling)	**1.1**	**0.0**	**0.0**	**1.5**	**0.0**	**0.8**	14.3	**3.5**	**0.2**	**2.6**	5.5
<u>Baselines</u>											
Gurobi (our ILP formulation)	9.2	105.6	190.4	4.7	81.6	197.3	372.2	6.8	59.8	116.4	192.0
Gurobi (Wu et al., 2021)	22.7	35.8	26.4	14.5	8.7	101.2	251.0	14.8	6.0	68.1	135.8
Greedy	22.7	64.5	76.8	9.7	40.1	64.9	65.7	11.9	26.4	39.3	39.4
Greedy based on fill rate	11.9	28.3	30.0	4.3	10.0	12.1	**1.7**	6.0	7.0	3.2	**2.0**
Alternate greedy storage and RL	15.3	32.8	31.9	10.1	18.1	19.1	12.9	14.2	11.8	14.9	16.9
Multiple sequence alignment	31.4	58.6	61.1	17.9	22.9	14.5	17.9	17.3	20.5	12.5	8.7
<u>Simulated annealing</u>											
Greedy, logarithmic cooling	20.9	57.4	62.3	8.5	27.7	56.0	65.1	10.6	18.2	31.6	37.8
Simple, logarithmic cooling	49.7	121.2	164.8	25.1	60.1	104.1	159.5	19.8	41.6	58.3	84.4
Greedy, geometric cooling	20.5	60.7	72.5	8.2	33.4	59.2	64.5	10.3	22.7	35.4	37.7
Simple, geometric cooling	52.9	125.3	192.4	23.4	76.8	125.2	195.7	15.1	53.8	73.3	98.8

For the linearly decreasing probabilities, we define $p_1 = \frac{1}{C} + d(C)\frac{C-1}{2}$ and $p_i - p_{i-1} = d(C)$ for $i = 2, \ldots, C$ with slope $d(C)$ depending on the number of colors. We choose $d(5) = 0.08$,

$d(10) = 0.02$ and $d(15) = 0.008$. Note that the slope $d(C)$ needs to account for normalization and positivity so that $d(C) < \frac{2}{C(C-1)}$.[7]

For all analyses, we specifically do not retrain the RL policies. Instead, the trained policies from the main analysis based on balanced color distributions are directly applied to the out-of-distribution problem instances.

Table 9: Evaluation results (ARPD) for imbalanced distribution with linearly decreasing probabilities.

Colors	5			10				15			
Buffer size (lanes x lane width)	2x2	3x3	4x4	2x2	4x4	6x6	8x8	2x2	4x4	6x6	8x8
RL all action masks (sampling)	**0.5**	**0.0**	**0.0**	**3.8**	**0.3**	**0.8**	10.7	**5.5**	**0.5**	**1.2**	7.9
Baselines											
Gurobi (our ILP formulation)	10.9	96.6	215.0	5.1	85.5	201.4	351.8	5.9	39.9	108.5	183.7
Gurobi (Wu et al., 2021)	33.9	43.5	32.8	19.8	8.9	118.3	257.9	12.5	2.4	52.9	117.4
Greedy	25.6	57.5	80.5	18.7	43.2	70.4	65.7	13.4	24.6	39.3	41.1
Greedy based on fill rate	17.8	29.6	16.6	12.4	12.5	9.6	**1.7**	10.0	5.6	4.5	**3.3**
Alternate greedy storage and RL	18.4	40.3	36.3	13.9	25.2	20.3	15.8	13.0	13.6	14.4	15.9
Multiple sequence alignment	42.0	66.6	68.0	23.5	29.5	19.8	11.1	15.2	15.5	5.6	8.1
Simulated annealing											
Greedy, logarithmic cooling	25.4	54.9	71.2	17.1	35.4	49.7	65.7	12.3	18.2	26.1	40.5
Simple, logarithmic cooling	56.8	143.9	179.4	26.8	63.7	110.3	175.8	17.8	28.4	46.2	69.9
Greedy, geometric cooling	24.5	55.0	76.9	17.2	38.7	63.3	65.7	11.8	21.6	34.7	38.4
Simple, geometric cooling	62.5	145.4	216.2	26.1	86.3	136.1	191.8	19.2	40.5	61.0	87.6

The results are presented in Table 8 for the exponentially decreasing probabilities and in Table 9 for the linearly decreasing probabilities.

Evidently, we find that RL produces the best solutions for most situations. Similarly to our previous analyses, Gurobi provides the best results from all baselines for 2x2 buffers and the Greedy heuristic based on fill rate outperforms RL for 8x8 buffers. These results point towards the robustness of our RL approach as it manages to successfully solve out-of-distribution instances.

7. Conclusion

We proposed a reinforcement learning approach for the paint shop problem with multi-lane buffers that allows fully flexibility in performing store and retrieve operations. To the best of our knowledge, our study is the first to provide an ILP formalization of this problem variant. We showed formally that the less flexible problem variant "store-then-variant" can lead to solutions that are arbitrarily worse than our flexible problem variant allowing store and retrieve operations to be performed in an arbitrary order. We then present a RL approach based on proximal policy

---

[7]We need $\sum_{i=1}^{C} p_i = 1$ by definition of probabilities. Accordingly, we get $\sum_{i=1}^{C} p_i = \sum_{i=1}^{C} (p_1 - (i-1)d(C)) = Cp_1 - d(C) \sum_{i=1}^{C} (i-1) = Cp_1 - d(C)\frac{C(C-1)}{2}$. Thus, we have $p_1 = \frac{1}{C} + d(C)\frac{C-1}{2}$. All probabilities must be positive, in particular, we need $p_C > 0$ for the color with the smallest probability. Hence, we require $p_C = p_1 - (C-1)d(C) = \frac{1}{C} + d(C)\frac{C-1}{2} - (C-1)d(C) = \frac{1}{C} - d(C)\frac{C-1}{2} > 0$, which implies $d(C) < \frac{2}{C(C-1)}$.

optimization that is trained to minimize color changes. In contrast to prior works that proposed RL approaches for the paint shop problem with multi-lane buffers (Huang et al., 2024; Leng et al., 2020, 2023), we employ RL to perform both, store and retrieve operations. The one-hot encoded state representation describes the buffer content, the colors of the most recent cars from the incoming sequence, and the color of the next car in the outgoing sequence. After showing that greedy retrieval and fast-track actions are optimal, we employed action masking to incorporate this prior knowledge into policy learning and application. Our evaluation based on 170 problem instances suggests that the proposed RL approach reduces color changes for most considered problem instances by considerable margins. In addition, we demonstrated the robustness of our approach towards rectangular buffer sizes, initially filled buffers, and out-of-distribution instances.

Our evaluation suggests several guidelines for the implementation of RL approaches for the paint shop problem with multi-lane buffers. First, researchers should include action masking in policy learning as it leads to a more efficient learning process by directly enforcing valid and provably optimal actions. We presented several action masks and showed how they can be combined into a single action mask while prioritizing one mask over another. Second, our results showed that a RL policy trained on problem instances with balanced color distributions also performs well on instances with imbalanced color distributions. More balanced color distributions make the problem harder to solve as uncertainty over the next color is maximized. For instance, if one color occurs much more frequently than others, the algorithm could simply learn to reserve one lane for the most frequent color, instead of learning to perform smart store and retrieve operations using all lanes. Third, we found that repeated sampling from the trained policy results in considerably higher solution quality than deterministic policy application. Given that a trained RL policy is generally very fast in generating solutions, the limited cutoff time can efficiently be used to continuously sample and evaluate solutions from the policy.

Our study also provides several opportunities for future research. First, one could consider more complex problem variants, e.g., with more than one painting line. If multiple painting lines are employed, the outgoing sequence can be split further into multiple outgoing sequences, which reduces the number of color changes. Second, other performance metrics, like makespan of individual cars could be considered in addition to color changes. So far, we have allowed each car to remain in the buffer for an arbitrary amount of time without considering any constraints regarding makespan. The model would then have to learn to account for constrained makespans by performing timely retrieve actions. Third, future research could attempt to integrate the paint shop problem into other scheduling problems like mixed-model sequencing or flow shop problems. Here, researchers first need to weigh the performance metrics of the main scheduling problem against the number of color changes. Another challenge is to quickly estimate the number of color changes of a sequence under a (near-)optimal action sequence in the paint shop.

References

Abreu, L. R., Cunha, J. O., Prata, B. A., & Framinan, J. M. (2020). A genetic algorithm for scheduling open shops with sequence-dependent setup times. *Computers & Operations Research*, *113*, 104793.

Bengio, Y., Lodi, A., & Prouvost, A. (2021). Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research*, *290*(2), 405–421.

Boysen, N., Scholl, A., & Wopperer, N. (2012). Resequencing of mixed-model assembly lines: Survey and research agenda. *European Journal of Operational Research*, *216*(3), 594–604.

Boysen, N., Schulze, P., & Scholl, A. (2022). Assembly line balancing: What happened in the last fifteen years? *European Journal of Operational Research*, *301*(3), 797–814.

Brammer, J., Lutz, B., & Neumann, D. (2021). Solving the mixed model sequencing problem with reinforcement learning and metaheuristics. *Computers & Industrial Engineering*, *162*, 107704.

Brammer, J., Lutz, B., & Neumann, D. (2022a). Permutation flow shop scheduling with multiple lines and demand plans using reinforcement learning. *European Journal of Operational Research*, *299*(1), 75–86.

Brammer, J., Lutz, B., & Neumann, D. (2022b). Stochastic mixed model sequencing with multiple stations using reinforcement learning and probability quantiles. *OR Spectrum*, *44*(1), 29–56.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). OpenAI Gym. Available at `https://arxiv.org/abs/1606.01540`.

Bysko, S., & Krystek, J. (2019). A game theory approach for solving the new concept of car sequencing problem. In *2019 5th International Conference on Control, Automation and Robotics (ICCAR)* (pp. 531–535). IEEE.

Bysko, S., Krystek, J., & Bysko, S. (2020). Automotive paint shop 4.0. *Computers & Industrial Engineering*, *139*, 105546.

Ding, F.-Y., & Sun, H. (2004). Sequence alteration and restoration related to sequenced parts delivery on an automobile mixed-model assembly line with multiple departments. *International Journal of Production Research*, *42*(8), 1525–1543.

Epping, T., & Hochstättler, W. (2003). Sorting with line storage systems. In *Operations Research Proceedings 2002* (pp. 235–240). Berlin: Springer.

Epping, T., Hochstättler, W., & Oertel, P. (2004). Complexity results on a paint shop problem. *Discrete Applied Mathematics*, *136*(2-3), 217–226.

Fernandez-Viagas, V., Ruiz, R., & Framinan, J. M. (2017). A new vision of approximate methods for the permutation flowshop to minimise makespan: State-of-the-art and computational evaluation. *European Journal of Operational Research*, *257*(3), 707–721.

Gendreau, M., & Potvin, J.-Y. (2018). *Handbook of Metaheuristics* volume 272. New York: Springer.

Goodson, J. C. (2015). A priori policy evaluation and cyclic-order-based simulated annealing for the multi-compartment vehicle routing problem with stochastic demands. *European Journal of Operational Research*, *241*(2), 361–369.

Gurobi Optimization, LLC (2023). Gurobi Optimizer Reference Manual. Available at `https://www.gurobi.com`.

Hajek, B. (1988). Cooling schedules for optimal annealing. *Mathematics of Operations Research*, *13*(2), 311–329.

Hong, S., Han, J., Choi, J. Y., & Lee, K. (2018). Accelerated dynamic programming algorithms for a car resequencing problem in automotive paint shops. *Applied Mathematical Modelling*, *64*, 285–297.

Huang, Y., Fu, G., Sheng, B., Lu, Y., Yu, J., & Yin, X. (2024). Deep reinforcement learning for solving car resequencing with selectivity banks in automotive assembly shops. *International Journal of Production Research*, *Forthcoming*, 1–22.

Inman, R. R., & Schmeling, D. (2003). Algorithm for agile assembling-to-order in the automotive industry. *International Journal of Production Research*, *41*(16), 3831–3848.

Knopp, S., Dauzère-Pérès, S., & Yugma, C. (2017). A batch-oblivious approach for complex job-shop scheduling problems. *European Journal of Operational Research*, *263*(1), 50–61.

Ko, S.-S., Han, Y.-H., & Choi, J. Y. (2016). Paint batching problem on M-to-1 conveyor systems. *Computers & Operations Research*, *74*, 118–126.

Lahmar, M., Ergan, H., & Benjaafar, S. (2003). Resequencing and feature assignment on an automated assembly line. *IEEE Transactions on Robotics and Automation*, *19*(1), 89–102.

Leng, J., Jin, C., Vogl, A., & Liu, H. (2020). Deep reinforcement learning for a color-batching resequencing problem. *Journal of Manufacturing Systems*, *56*, 175–187.

Leng, J., Wang, X., Wu, S., Jin, C., Tang, M., Liu, R., Vogl, A., & Liu, H. (2023). A multi-objective reinforcement learning approach for resequencing scheduling problems in automotive manufacturing systems. *International Journal of Production Research*, *61*(15), 5156–5175.

Lin, L., Sun, H., & Xu, Y.-q. (2011). A research of resequencing problem in automobile paint shops using selectivity banks. In *2011 IEEE 18th International Conference on Industrial Engineering and Engineering Management* (pp. 658–662). IEEE.

Meunier, F., & Neveu, B. (2012). Computing solutions of the paintshop–necklace problem. *Computers & Operations Research*, *39*(11), 2666–2678.

Mosadegh, H., Fatemi Ghomi, S., & Süer, G. (2020). Stochastic mixed-model assembly line sequencing problem: Mathematical modeling and Q-learning based simulated annealing hyper-heuristics. *European Journal of Operational Research*, *282*(2), 530–544.

Neufeld, J. S., Schulz, S., & Buscher, U. (2023). A systematic review of multi-objective hybrid flow shop scheduling. *European Journal of Operational Research*, *309*(1), 1–23.

Osman, I. H. (1993). Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Annals of Operations Research*, *41*, 421–451.

Plastria, F. (2002). Formulating logical implications in combinatorial optimisation. *European Journal of Operational Research*, *140*(2), 338–353.

Schulman, J., Moritz, P., Levine, S., Jordan, M., & Abbeel, P. (2015). High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, . Available at `https://arxiv.org/abs/1506.02438`.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. Available at `https://arxiv.org/abs/1707.06347`.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*(3), 379–423.

Spieckermann, S., Gutenschwager, K., & Voß, S. (2004). A sequential ordering problem in automotive paint shops. *International Journal of Production Research*, *42*(9), 1865–1878.

Stappert, M., Lutz, B., & Neumann, D. (2023). Data for: Solving the paint shop problem using reinforcement learning. Dataset on Mendeley. Available at `https://doi.org/10.17632/zbg64f6vb3.1`.

Sun, H., Fan, S., Shao, X., & Zhou, J. (2015). A colour-batching problem using selectivity banks in automobile paint shops. *International Journal of Production Research*, *53*(4), 1124–1142.

Sun, H., & Han, J. (2017). A study on implementing color-batching with selectivity banks in automotive paint shops. *Journal of Manufacturing Systems*, *44*, 42–52.

Taube, F., & Minner, S. (2018). Resequencing mixed-model assembly lines with restoration to customer orders. *Omega*, *78*, 99–111.

Winter, F., Musliu, N., Demirović, E., & Mrkvicka, C. (2019). Solution approaches for an automotive paint shop scheduling problem. In *Proceedings of the International Conference on Automated Planning and Scheduling* (pp. 573–581). AAAI Press volume 29.

Wu, J., Ding, Y., & Shi, L. (2021). Mathematical modeling and heuristic approaches for a multi-stage car sequencing problem. *Computers & Industrial Engineering*, *152*, 107008.