

# A Deep Reinforcement Learning Approach to Supply Chain Inventory Management

**Francesco Stranieri**

Department of Informatics, Systems, and Communication  
University of Milano-Bicocca  
Viale Sarca 336, Italy

francesco.stranieri@unimib.it

**Fabio Stella**

Department of Informatics, Systems, and Communication  
University of Milano-Bicocca  
Viale Sarca 336, Italy

fabio.stella@unimib.it

## Abstract

This paper leverages recent developments in reinforcement learning and deep learning to solve the supply chain inventory management (SCIM) problem, a complex sequential decision-making problem consisting of determining the optimal quantity of products to produce and ship to different warehouses over a given time horizon. A mathematical formulation of the stochastic two-echelon supply chain environment is given, which allows an arbitrary number of warehouses and product types to be managed. Additionally, an open-source library that interfaces with deep reinforcement learning (DRL) algorithms is developed and made publicly available for solving the SCIM problem. Performances achieved by state-of-the-art DRL algorithms are compared through a rich set of numerical experiments on synthetically generated data. The experimental plan is designed and performed, including different structures, topologies, demands, capacities, and costs of the supply chain. Results show that the PPO algorithm adapts very well to different characteristics of the environment. The VPG algorithm almost always converges to a local maximum, even if it typically achieves an acceptable performance level. Finally, A3C is the fastest algorithm, but just like VPG, it never achieves the best performance when compared to PPO. In conclusion, numerical experiments show that DRL performs consistently better than standard reorder policies, such as the static ( $s, Q$ )-policy. Thus, it can be considered a practical and effective option for solving real-world instances of the stochastic two-echelon SCIM problem.

**Keywords:** machine learning, inventory management, deep reinforcement learning

## 1. Introduction

Supply chain inventory management (SCIM) is a complex *sequential decision-making problem* consisting of determining the optimal quantity of products to produce at the factory and to ship to different distribution warehouses over a given time horizon. Recently, deep reinforcement learning (DRL) has been applied to solve many challenging problems in various fields, including robotics, video games, medicine and healthcare, finance, transportation systems, and industry 4.0, to mention just a few (Li, 2017). However, as evidenced by the helpful roadmap of Boute et al. (2021), DRL algorithms are rarely applied to the SCIM field, although they can be used to develop near-optimal policies that are difficult, or impossible at worst, to achieve using traditional methods. Indeed, the uncertain and stochastic nature of products demand, as well as lead times, represent significant obstacles for mathematical programming approaches to be effective, with specific reference to those cases where the modeling of SCIM's entities is reasonable, for example, assuming a finite capacity of warehouses (de Kok et al., 2018).

Regarding the DRL algorithms that have been currently applied to tackle the SCIM problem, we found that they suffer the following *limitations*: i) given a supply chain *structure* (e.g., divergent<sup>1</sup> two-echelon<sup>2</sup>), no DRL algorithm has

1. In a *linear* supply chain, each participant has one predecessor and one successor; in a *divergent* supply chain, each has one predecessor but can have multiple successors, while the opposite is true in a *convergent* supply chain. Finally, in a *general* supply chain, each participant can have several predecessors and several successors.
2. A supply chain can include multiple stages, called formally *echelons*, through which the stocks are moved to reach the customer. When the number of echelons is greater than one, we refer to a *multi-echelon* supply chain.

been deeply tested with respect to different *topologies* (i.e., by changing the number of warehouses) as in Hubbs et al. (2020) and Gijbrecchts et al. (2022); ii) no extensive experiments have been performed on the same supply chain structure by varying different *configurations* (e.g., demands, capacities, and costs) as in Peng et al. (2019) and Alves and Mateus (2020); iii) no extension has been proposed for *comparing different DRL algorithms* and determining which one is more appropriate for a particular supply chain topology and configuration, as suggested by Alves and Mateus (2020) and Boute et al. (2021).

Furthermore, relevant aspects of the SCIM problem have not yet been addressed efficiently (Yan et al., 2022), for example: i) the *sequence of events* required to reproduce and validate a simulation model is not always well-defined or given. Hence, making available a consistent and universal open-source SCIM environment can improve reusability and reproducibility, especially if implemented with standard APIs (like those of OpenAI Gym (Brockman et al., 2016)). In this way, it is also possible to import DRL algorithms from reliable libraries and focus solely on their fine-tuning, instead of developing them from scratch; ii) DRL algorithms are typically compared with some standard *static reorder policies*. However, their performances are not always compared with those achieved by an oracle, i.e., a baseline who knows the optimal action to take a priori, thus making it difficult to evaluate the DRL effectiveness in real-world environments (the only paper in which an oracle is introduced is Hubbs et al. (2020)); iii) none of the DRL papers available in the specialized literature considers a *multi-product approach*, whereas it has been considered relating to other solution methods (Shervais et al., 2003; Sui et al., 2010; Cimen and Kirkbride, 2013). Considering more than one product type increases the dimensionality and complexity of the problem, consequently requiring an efficient implementation of the SCIM environment and DRL algorithms.

This paper makes the following *contributions* to the SCIM decision-making problem:

- design and formulation of a stochastic and divergent two-echelon SCIM environment under seasonal demand, which allows an arbitrary number of warehouses and product types to be managed;
- comparison of a set of state-of-the-art DRL algorithms in terms of their ability to find an optimal policy, i.e., a policy which maximizes the SCIM’s profit as achieved by an oracle;
- evaluation of performances achieved by state-of-the-art DRL algorithms and comparison to a static reorder policy, i.e., an  $(s, Q)$ -policy, whose optimal parameters have been set through a data-driven approach;
- design and run of a rich experimental plan involving different SCIM topologies and configurations as well as values of hyperparameters associated with DRL algorithms’;
- design and development of an open-source library for solving the SCIM problem<sup>3</sup>, thus embracing the open science principles and guaranteeing reproducible results.

The rest of the paper is organized as follows: Section 2 is devoted to introducing and providing main reinforcement learning (RL) definitions and notation, also highlighting how RL approaches have dealt with the SCIM problem. DRL is briefly introduced and described in Section 3; in this section, we also describe the state-of-the-art DRL algorithms and how they have been used to address the SCIM problem. Section 4 describes the main methodological contributions of this paper. The rich experimental plan is then reported in Section 5, while the results of numerical experiments are presented in Section 6. Lastly, discussions and conclusions are given in Section 7.

## 2. SCIM via Reinforcement Learning

RL has recently achieved remarkable results in the field of artificial intelligence, mainly when applied to video games and gaming in a more general sense (Mnih et al., 2015; Silver et al., 2017; Vinyals et al., 2019). Nevertheless, there are still few use-cases in *industrial applications*, even if RL proved to be effective in solving complex sequential decision-making problems.

Essentially, RL adopts the Markov Decision Process (MDP) framework to represent the interactions between a learning agent and an environment. At each time step  $t$ , the agent observes the current state of the environment,  $S_t \in \mathcal{S}$ , chooses

3. The open-source library is available on <https://github.com/frenkowski/SCIMAI-Gym>.

an action,  $A_t \in \mathcal{A}(S_t)$ , and obtains a reward,  $R_{t+1} \in \mathcal{R} \subset \mathbb{R}$ ; then, the environment transitions into a new state,  $S_{t+1}$ . The goal of RL is thus to find an optimal policy,  $\pi_* : \mathcal{S} \rightarrow \mathcal{A}$ , that maximizes the *expected discounted return*,  $G_t = \sum_{k=t+1}^T \gamma^{k-t-1} R_k$ , where  $0 \leq \gamma \leq 1$  is a hyperparameter called *discount rate*.

One of the most common approaches for solving the SCIM problem through RL algorithms turns out to be Q-learning (Watkins, 1989). This approach is based on a tabular and temporal-difference (TD) algorithm that learns how to determine the *value* of an action  $A_t$  in a state  $S_t$ , referred to as the Q-value, in accordance with the following update rule:

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha \delta_t, \quad (1)$$

where  $\delta_t = [R_{t+1} + \gamma \max_a Q(S_{t+1}, a) - Q(S_t, A_t)]$  is the TD error, and  $0 \leq \alpha \leq 1$  is a hyperparameter called *learning rate*. Q-values of each state-action pair are stored in a table, known as Q-table, where each state is represented by a row and each action by a column. Through the Q-learning algorithm, Q-values associated with each state-action pair are estimated according to Eq. (1). Once convergence has been achieved (which is guaranteed under certain conditions (Tsitsiklis, 1994; Jaakkola et al., 1994)), an optimal policy can be easily obtained by identifying, for each state, the action with the highest Q-value, that is,  $\pi_*(S_t) = \arg \max_a Q(S_t, a)$ .

Chaharsooghi et al. (2008), authors of one of the most cited RL articles about SCIM, proposed an approach based on Q-learning to address (a *centralized* variant of) the beer game problem (Forrester, 1958), which consists of a linear supply chain with four participants (i.e., supplier, manufacturer, distributor, and retailer<sup>4</sup>) and is frequently discussed in academic contexts to demonstrate the *bullwhip effect* (Forrester, 1997; Lee et al., 1997), a phenomenon that demonstrates how small fluctuations in the final demand by customers can result in more significant fluctuations at previous stages of the supply chain, negatively impacting performances and costs. The beer game assumes no transportation costs but only storage and penalty costs, i.e., costs incurred by storing unsold inventories and when an order cannot be completely satisfied. Additionally, it also assumes deterministic lead times.

To implement the Q-learning algorithm, Chaharsooghi et al. (2008) defined the current system state as a vector consisting of the four inventory positions in terms of current stock levels. However, considering that inventory positions thus defined may take infinite values, applying this strategy appears unfeasible since the Q-table would be in turn infinite. Consequently, the authors *discretized* the state space into nine intervals. In this way, the possible state values amount to  $9^4$ . Regarding actions, their approach determines the number of products to order via the  $d+x$  policy<sup>5</sup>. The learning process's objective is hence to determine the value of the unknown variable  $x$  according to the given system state. For limiting the Q-table size,  $x$  was *constrained* by the authors to belong to  $[0, 3]$  so that the possible number of actions amounts to  $4^4$ . Lastly, the customer demand at the retailer was generated through a uniform distribution on  $[0, 15]$ , while the time horizon was set to 35 (weeks).

Obviously, by defining restricted state and action spaces, the resulting Q-table appears to be more manageable. However, analyzing various RL studies (Ravulapati et al., 2004; Sui et al., 2010; Mortazavi et al., 2015), it becomes evident that the Q-tables implemented are typically huge and, thus, *unscalable*. For example, the Q-table adopted by Chaharsooghi et al. (2008) has a number of cells equal to  $(9^4 \cdot 4^4 =) 1\,679\,616$ , equivalent to the number of states multiplied by the number of actions. Consequently, expanding the size of the state or action spaces might not be feasible, as the Q-tables can no longer be handled. Another critical point of reflection is given in Gevers (2020), where the author has shown that the approach proposed by Chaharsooghi et al. (2008) is not able to learn an effective policy since the impact of the actions (restricted to take values on the interval  $[0, 3]$ ) is so limited that even *random actions* provide basically the same results.

In conclusion, tabular RL methods can only be applied to discretized or constrained state and action spaces. However, discretization leads to a *loss of crucial information*, in addition to being unsuitable for real-world scenarios; thus, we need improved RL methods to address the SCIM problem effectively.

4. In a linear four-echelon supply chain, each of the participants is typically indicated with a *number* in ascending order, from the retailer to the supplier. Thus, for example, the distributor has to satisfy the orders made by the retailer (i.e., the downstream stage) while it can request stocks from the manufacturer (i.e., the upstream stage).

5. Precisely, if a participant in the previous time step received a request for  $d$  product units from the downstream stage, the  $d+x$  policy requires ordering  $d+x$  units to the upstream stage in the current time step.

### 3. SCIM via Deep Reinforcement Learning

DRL is a combination of RL with deep learning (DL) which promises to scale to previously intractable decision-making problems, i.e., environments with high dimensional state and action spaces. DL is rooted into artificial neural networks (ANNs) (LeCun et al., 2015), which are universal approximators capable of providing an optimal approximation of *highly nonlinear functions*. In practice, function parameters  $\theta$  are adjusted during the learning process in order to maximize the expected return (or, alternatively, to minimize the TD error).

The DRL algorithms we implemented belong to the policy-based methods, which can learn a *parameterized* and stochastic policy,  $\pi_\theta \approx \pi_*$  with  $\pi : \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ , to select actions directly (as opposed to the Q-learning algorithm, which is part of value-based methods (Sutton and Barto, 2018)). Inside them, policy gradient methods offer a considerable theoretical advantage through the *policy gradient theorem*, and the vanilla policy gradient (VPG) algorithm (Williams, 1992) is a natural result of this theorem; however, the *high variance* of gradient estimates usually results in policy update instabilities (Wu et al., 2018). Also to mitigate this issue, Schulman et al. (2015) proposed an actor-critic algorithm (which means that a policy and a value function are simultaneously learned) called trust region policy optimization (TRPO), which bounds the difference between the new and the old policy in a *trust region*. Proximal policy optimization (PPO) (Schulman et al., 2017) shares the same background as TRPO, but has demonstrated comparable or superior performance while being significantly *simpler* to implement and tune. Asynchronous advantage actor-critic (A3C) (Mnih et al., 2016) is also one of the available state-of-the-art actor-critic algorithms. Its core idea is to have different agents interacting with different representations of the environment, each with its parameters. Periodically (and asynchronously), they update a global ANN that incorporates *shared parameters*. For interested readers, an in-depth and more rigorous discussion on the various DRL algorithms can be found in François-Lavet et al. (2018).

To the best of the authors' knowledge, only five papers have implemented DRL algorithms to solve the SCIM problem, despite some restrictions. More in detail, an *extension* of deep Q-network (DQN) (Mnih et al., 2015) has been proposed in Oroojlooyjadid et al. (2021) to solve (a *decentralized* variant of) the beer game problem. The authors revealed that a DQN agent, which basically involves an ANN instead of a Q-table to return the Q-value for a state-action pair, can learn a near-optimal policy when other supply chain participants follow a base-stock policy<sup>6</sup>. Because Q-learning requires a restricted action space cardinality, the authors performed numerical experiments using a  $d+x$  policy, with  $x$  constrained to one of the following intervals: [-2, +2], [-5, +5], and [-8, +8].

Alternatively, Peng et al. (2019) proposed the VPG algorithm to address a two-echelon supply chain with stochastic and seasonal demand. Due to storage capacity constraints, the authors designed a *dynamic action space*. As a result, the number of products to ship is determined also by considering the number of stocks actually present in the warehouses. To evaluate the VPG performance, three different numerical experiments are presented, and the results show that the VPG agent is able to outperform the  $(s, Q)$ -policy<sup>7</sup> employed as a baseline in all three experiments.

Using the same supply chain structure but with ten warehouses and a normal distribution, Gijsbrechts et al. (2022) applied and tuned the A3C algorithm for two different numerical experiments (taken from Van Roy et al. (1997)). The authors restricted the action space by implementing a *state-dependent* base-stock policy, and the results show that A3C can achieve performance comparable to state-of-the-art heuristics and approximate dynamic programming algorithms, despite its initial tuning remaining computationally intensive.

A SCIM problem with a linear four-echelon supply chain is considered in Hubbs et al. (2020), where different *operations research methods* are compared with the PPO algorithm in two different environments, i.e., without and with backlog (in case of backlog, unsatisfied orders will be fulfilled at a subsequent time step, while a penalty cost is applied). Numerical experiments show that PPO outperform the base-stock policy in both environments.

Finally, in the experimental scenario analyzed by Alves and Mateus (2020), a general four-echelon supply chain with two nodes per echelon is presented. The system state consists of product quantity currently available and in transit across the supply chain, plus *future* customer demands. To deal with the optimization problem, the authors proposed the PPO algorithm, while a deterministic linear programming agent (i.e., considering a deterministic demand) is employed as a baseline. Results of numerical experiments show that PPO still achieves satisfactory results.

6. Under a *base-stock policy*, each participant orders in each time step  $t$  a quantity to bring its stocks equal to a fixed number  $s$ , known as the base-stock level, to determine in an optimal way.

7. The  $(s, Q)$ -policy can be expressed by a rule: at each time step  $t$ , the current stock level is compared to the reorder point  $s$ . If the stock level falls below the reorder point  $s$ , then the  $(s, Q)$ -policy orders  $Q$  units of product; otherwise, it does not take any action. Also in this case, the parameters  $s$  and  $Q$  are to be determined optimally.

#### 4. Problem Definition

The SCIM environment we propose is primarily motivated by what was presented and discussed in [Kemmer et al. \(2018\)](#) and [Peng et al. \(2019\)](#). Inspired by these works, we designed a divergent two-echelon supply chain that includes a *factory*, a *factory warehouse*, and  $J$  *distribution warehouses*; an example of this structure is shown in Fig. 1.

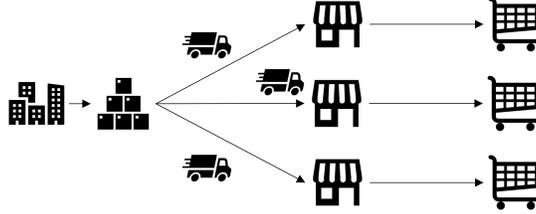


Figure 1: A divergent two-echelon supply chain consisting of a factory and its warehouse (first echelon), plus three distribution warehouses (second echelon). Shopping carts represent customers’ demands.

In our formulation, we assume that the factory produces  $I$  different product types. For each product type  $i$ , the factory decides, at every time step  $t$ , its respective production level  $a_{i,0,t}$  (we indicate  $j = 0$  for the factory and  $j > 0$  for the distribution warehouses), that is, how many units to produce, considering a fixed production cost of  $z_{i,0}$  per unit. Moreover, the factory warehouse is associated with a maximum capacity of  $c_{i,0}$  units for each product type  $i$  (this means that the overall capacity is given by  $\sum_{i=0}^I c_{i,0} = c_0$ ). The cost of storing one unit of product type  $i$  at the factory warehouse is  $z_{i,0}^S$  per time step, while the corresponding stock level at time  $t$  equals  $q_{i,0,t}$ . At every time step  $t$ ,  $a_{i,j,t}$  units of product type  $i$  are shipped from the factory warehouse to the distribution warehouse  $j$ , with an associated transportation cost of  $z_{i,j}^T$  per unit. For each product type  $i$ , each distribution warehouse  $j$  has a maximum capacity of  $c_{i,j}$  ( $\sum_{i=0}^I c_{i,j} = c_j$ ), a storage cost of  $z_{i,j}^S$  per unit, and a stock level at time  $t$  equal to  $q_{i,j,t}$ . The demand for product type  $i$  at distribution warehouse  $j$  for time step  $t$  is equivalent to  $d_{i,j,t}$  units, while each unit of product type  $i$  is sold to customers at sale price  $p_i$  (which is identical across all warehouses).

Products are non-perishable and provided in discrete quantities. Additionally, we assume that each warehouse is legally obligated to fulfill all the submitted orders. Consequently, if an order for a certain time step exceeds the corresponding stock level, a penalty cost per unsatisfied unit is applied (the penalty cost for product type  $i$  is obtained by multiplying the penalty coefficient  $z_i^P$  by the sale price value  $p_i$ ). Unsatisfied orders are maintained over time, and we design them as a negative stock level (which corresponds to *backordering*); this also implies that when the penalty coefficient is particularly high (e.g.,  $z_i^P \geq 1$ ), the agent may not be able to generate a positive profit if it causes backlog orders. Consequently, it should prefer a policy that leads to accumulating stocks in advance in order to pay storage costs rather than penalty costs. A summary of the *formulation parameters* is available through Table 1.

Parameter	Explanation	Parameter	Explanation
$I$	Number of Products Types	$q_{i,j,t}$	Stock Level (units)
$J$	Number of Warehouses	$c_{i,j}$	Storage Capacity (units)
$T$	Episode Length	$z_{i,j}^S$	Storage Cost (per unit)
$a_{i,j,t}$	Production and Shipping Level (units)	$z_i^P$	Penalty Coefficient
$z_{i,0}$	Production Cost (per unit)	$p_i$	Sale Price (per unit)
$z_{i,j}^T$	Transportation Cost (per unit)	$d_{i,j,t}$	Demand (units)

Table 1: The considered SCIM parameters with relative explanation (and units of measure). All these parameters are integrated and editable within our open-source library; therefore, it can be seen as a solid foundation for further research (e.g., for benchmarking new state-of-the-art DRL algorithms and deriving promising conclusions on the basis of different topologies and configurations of the SCIM environment).

#### 4.1 Environment Formulation

In this subsection, we formalize the RL problem as an MDP. More precisely, we introduce and define the *main components* of the SCIM environment that we propose: the state vector, the action vector, and the reward function.

The *state vector* includes all current stock levels for each warehouse and product type, plus the last  $\tau$  demand values, and is defined as follows:

$$s_t = (q_{0,0,t}, \dots, q_{I,J,t}, d_{t-\tau}, \dots, d_{t-1}),$$

where  $d_{t-1} = (d_{0,1,t-1}, \dots, d_{I,J,t-1})$ . It is worth noticing that the actual demand  $d_t$  for the current time step  $t$  will not be known until the next time step  $t + 1$ . This implementation choice ensures that the agent may benefit from learning the demand pattern so as to integrate a sort of *demand forecasting* directly into the policy. Additionally, we include the last demand values in order to enable the agent to have *limited knowledge* about the demand history and, consequently, to gain a basic comprehension of its fluctuations (similar to what was made originally by Kemmer et al. (2018)). In our SCIM implementation, the agent can access the demand values of the last five time steps, even if preliminary results suggest that comparable performances are obtained by accessing the last three or four time steps.

Regarding the *action vector*, we chose to implement a *continuous action space* (i.e., the ANN generates the action value directly) consisting, for each product type  $i$ , of the number of units to produce at the factory and of the number of units to ship to each distribution warehouse  $j$ :

$$a_t = (a_{0,0,t}, \dots, a_{I,J,t}). \quad (2)$$

According to the literature, a relatively small and *identical upper bound* is typically adopted for all the action values to reduce the computational effort. However, the drawback is that this might lead to a significant drop in terms of performance. Indeed, if the upper bound is set too small, the agent may select an inefficient action given that the optimal one is outside the admissible range. Otherwise, if the upper bound is set too high, the agent may repeatedly choose an incoherent action, i.e., one that falls within the admissible range but exceeds a specified maximum capacity.

Our implementation thus provides a continuous action space with an *independent upper bound* for each action value, in order to find a trade-off between efficiency and performance. In practical terms, the lower bound for each value is simply zero. In fact, it would be illogical to produce or ship negative quantities of products. Conversely, the upper bound for each distribution warehouse corresponds to its maximum capacity with respect to each product type (by referring to Eq. (2),  $0 \leq a_{i,j,t} \leq c_{i,j}$ ). To guarantee that the factory can adequately handle the various demands, its upper bound amounts to the sum of all warehouses' capacities with regard to each product type ( $0 \leq a_{i,0,t} \leq \sum_{j=0}^J c_{i,j}$ ). Accordingly, we now have a well-defined and coherent action space. We expect to improve both efficiency and performance with this intuition, as the action space is bounded (and hence restricted) but contains only coherent (and possibly optimal) actions. We specify that available stocks are not explicitly considered when the agent chooses an action. However, producing or shipping a number of stocks that it is not possible to store leads to a cost and, therefore, an *implicit penalty* for the agent. An alternative (and possibly better) formulation would be to consider a dynamic action space (e.g.,  $0 \leq a_{i,j,t} \leq c_{i,j} - q_{i,j,t}$ ), as in Peng et al. (2019). For the sake of simplicity, we also assume that there are no lead times both for production and transportation (or to refer to the literature, we assume *constant lead times equal to 0*).

To evaluate the performance of the DRL agents, we simulate a seasonal behavior by representing the *demand* as a co-sinusoidal function with a stochastic component, defined according to the following equation:

$$d_{i,j,t} = \left\lfloor \frac{d_{max_i}}{2} \left( 1 + \cos \left( \frac{4\pi(2ij + t)}{T} \right) \right) + \mathcal{U}(0, d_{var_i}) \right\rfloor, \quad (3)$$

where  $\lfloor \cdot \rfloor$  is the floor function,  $d_{max_i}$  is the maximum demand value for each product type  $i$ ,  $\mathcal{U}$  is a random variable uniformly distributed on the support  $(0, d_{var_i})$  representing the demand variations (i.e., the *uncertainty*), and  $T$  is the final time step of the episode. At each time step  $t$ , the demand may vary for each warehouse  $j$  and product type  $i$  while maintaining the same behavior, as can be seen in Fig. 2.

The *reward function* for each time step  $t$  is then defined as follows:

$$r_t = \sum_{j=1}^J \sum_{i=0}^I p_i \cdot d_{i,j,t} - \sum_{i=0}^I z_{i,0} \cdot a_{i,0,t} - \sum_{j=1}^J \sum_{i=0}^I z_{i,j}^T \cdot a_{i,j,t} - \sum_{j=0}^J \sum_{i=0}^I z_{i,j}^S \cdot \max(q_{i,j,t}, 0) + \sum_{j=0}^J \sum_{i=0}^I z_i^P \cdot p_i \cdot \min(q_{i,j,t}, 0). \quad (4)$$

The first term represents revenues, the second one production costs, while the third one transportation costs. The fourth term is the overall storage costs. The function  $\max$  is implemented to avoid negative inventories (i.e., backlog orders) from being counted. The last term denotes the penalty costs, which is introduced with a plus sign because stock levels would already be negative in the eventuality of unsatisfied orders. The DRL agents' goal is thus to *maximize the*

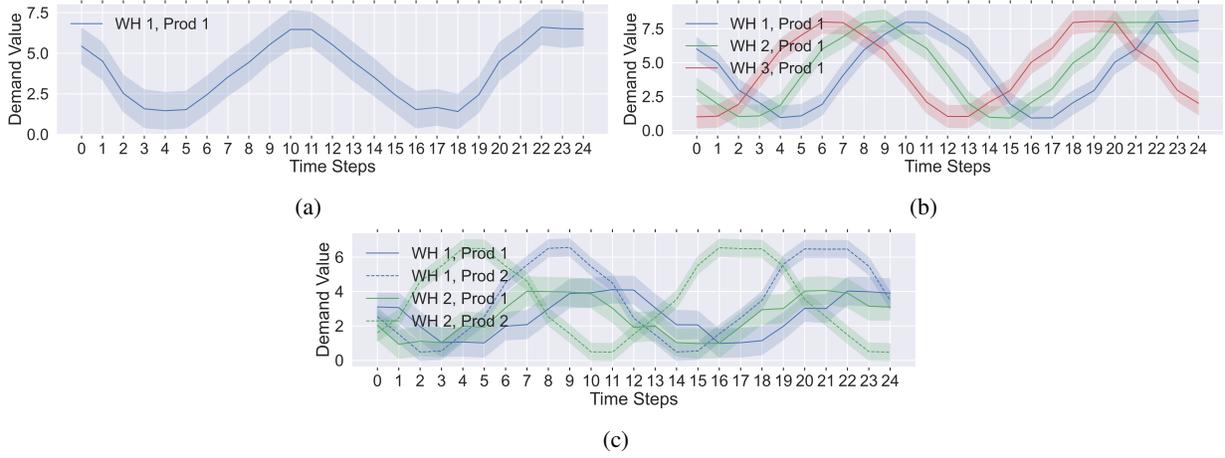


Figure 2: Some instances of different demands behavior generated according to Eq. (3) for different topologies and configurations of the SCIM problem: (a) one product type and one distribution warehouse with  $d_{max} = 5$  and  $d_{var} = 3$ , (b) one product type and three distribution warehouses with  $d_{max} = 7$  and  $d_{var} = 2$ , and (c) two product types and two distribution warehouses with  $d_{max} = (3, 6)$  and  $d_{var} = (2, 1)$  (referring to the values in round brackets, the first denotes the first product type, whereas the second indicates the second product type).

*supply chain profit* as defined in the reward function. By design, revenues are always calculated regardless of whether the demand is effectively satisfied; however, in the event of unsatisfied orders, the penalty costs will impact the actual return for each time step in which backlog orders are counted (in the amount of the penalty coefficient).

Finally, the *state's updating rule* is defined as follows:

$$s_{t+1} = (\min[(q_{0,0,t} + a_{0,0,t} - \sum_{j=1}^J a_{0,j,t}), c_{0,0}], \dots, \min[(q_{I,J,t} + a_{I,J,t} - d_{I,J,t}), c_{I,J}], d_{t+1-\tau}, \dots, d_t).$$

This implies that, at the beginning of the next time step, the factory's stocks are equal to the initial stocks, plus the units produced, minus the stocks shipped. Similarly, the warehouses' stocks are equal to the initial stocks, plus the units received, minus the current demand. When surplus stocks are generated, a storage cost is imposed; otherwise, a penalty cost is considered. Lastly, the demand values included in the state vector are also updated, discarding the oldest value and concatenating the most recent one.

## 5. Numerical Experiments

Once the environment has been specified, we implemented the agents according to three different DRL algorithms: A3C, PPO, and VPG, which have been briefly introduced in Section 3. In this respect, we relied on the implementations made available by Ray (Moritz et al., 2018), an open-source Python framework that is bundled with RLib, a scalable RL library, and Tune, a scalable hyperparameter tuning library. An advantage of Ray is that it natively supports OpenAI Gym. As a result, we exploited the OpenAI Gym APIs to develop the *simulator* representative of the environment and used for the agents' training process.

To assess and compare performances achieved by the adopted DRL algorithms, we also implemented a static reorder policy known in the specialized literature as the  $(s, Q)$ -policy. In our SCIM implementation, we opted to make reordering decisions independently; this means that the  $(s, Q)$ -policy parameters,  $s_{i,j}$  and  $Q_{i,j}$ , can differ for each warehouse and product type (this policy is still defined *static* because these parameters do not change over time). To find the best possible parameters that maximize Eq. (4), we developed a *data-driven approach* based on Bayesian optimization (BO). In this way, the solution method does not require making any assumptions or simplifications, and hence it is no longer problem-dependent; therefore, it can be applied to any SCIM topology and configuration just as it happens for DRL algorithms (they share, in fact, the same identical simulator).

To compare DRL and BO approaches, we also implemented an *oracle*, that is, a baseline that knows the real demand value for each product type and warehouse in advance and can accordingly select the optimal action to take a priori.

### 5.1 Scenarios Considered

A rich set of numerical experiments have been designed and performed to compare the performances of DRL algorithms and BO under *three different scenarios*. Each scenario is associated with different demand patterns with respect to each product type and warehouse (i.e., seasonal and stochastic fluctuations). Furthermore, each scenario has different capacities and costs for evaluating in-depth the adaptability and robustness of DRL algorithms.

Under the *one product type one distribution warehouse* (1P1W) scenario, the supply chain is set to manage just one product type. Accordingly, it consists of one factory, a factory warehouse, and one distribution warehouse; thus the input dimension of the ANN (representing the state vector) is equal to 7, given by the number of warehouses (i.e., 2, including the factory warehouse) times the number of product types (i.e., 1), plus the last demand values for each distribution warehouse and product type (i.e., 5), while the output dimension of the ANN (expressing the action vector) is 2, equivalent to the number of warehouses (including the factory warehouse) multiplied by the number of product types. Under the 1P1W scenario, which consists of five experiments as summarized in Table 5 of Appendix A, sale prices and costs are manipulated so as to increase or decrease revenues and, consequently, the margin of return. Moreover, in the first experiment, we bound the warehouses' capacities in such a way that they are smaller than the maximum demand value (also considering the stochastic demand variation). This decision is made to study whether DRL algorithms are able to learn an efficient strategy, i.e., a strategy capable of predicting a *growing demand* and thus saving and shipping stocks in advance. Finally, we generate multiple penalty coefficients to determine whether a hefty punishment forces DRL algorithms to be more or less effective, with particular attention to the more challenging experiments where low revenues and high costs are considered.

The *one product type three distribution warehouses* (1P3W) scenario concerns a more complex supply chain configuration, consisting of a factory, a factory warehouse, and three distribution warehouses. Even in this case, the supply chain still manages a single product type, while the input and output dimensions of the ANN are equal to 9 and 4, respectively; hence, the difficulty of the problem is increased because there is a higher number of both ANN parameters to be optimized and actions to be determined. The design of the five experiments follows that of the previous 1P1W scenario. However, a remarkable difference is found in storage capacities and costs, as depicted in Table 6 of Appendix A. In fact, we set warehouses' costs to be *directly proportional* to their corresponding capacities, that is, the less storage space we have, the more expensive it is to store a product. This scenario is also designed to investigate the DRL algorithms strategy when capacities increase, given that the search space of optimal actions grows accordingly. Furthermore, we are interested in studying how DRL algorithms react when demand, with the associated costs (i.e., production and transportation), becomes greater than actual capacities, considering that the supply chain now consists of three distribution warehouses and, consequently, the SCIM problem becomes more challenging to be tackled.

Finally, in the *two product types two distribution warehouses* (2P2W) scenario, the supply chain consists of two product types, a factory with its warehouse, and two distribution warehouses. With this design, the number of parameters to optimize is still higher, considering that the ANN input dimension is equal to 26, while the ANN output dimension is 6. Due to computational time, we performed just three experiments under this scenario, as reported in Table 7 of Appendix A. Regarding the demand, we explore demand variations which can be different or equal, according to the specific experiment. Additionally, we thought of something different concerning storage capacities and, consequently, the search space of optimal actions. Indeed, in the last experiment, warehouses' capacities for the first product type are designed in descending order, while for the second product type in ascending order; this implies that, for example, the second distribution warehouse can store the minimum amount of stocks for the first product type and the maximum amount for the second product type. We expect that this *imbalance*, especially when combined with greater uncertainty, makes the SCIM problem more unexpected and, thus, more difficult to be effectively solved.

## 6. Results

To compare the performances between DRL algorithms, BO, and oracle, we simulated, for each scenario and experiment, 200 different episodes. Each episode consists of 25 time steps, and we reported the *average cumulative profit* achieved, i.e., the sum of the per-step profit at the last time step  $T$ . All experiments were run on a machine equipped

with an Intel® Xeon® Platinum 8272CL CPU at 2.6 GHz and 16 GB of RAM. The hyperparameters of DRL algorithms selected for tuning have been chosen following what is presented in the Ray documentation and discussed in the papers of Alves and Mateus (2020) and Gijsbrechts et al. (2022); they are reported in Table 8 of Appendix A, along with their corresponding values. To early stop training instances associated with *bad hyperparameters configurations*, we also implemented, through Ray, the asynchronous successive halving (ASHA) scheduling algorithm (Li et al., 2018). It is important to note that the simulation results presented and commented in this section have been obtained by selecting, for each algorithm and experiment, the respective *best training instance*<sup>8</sup>.

Results of numerical experiments under the 1P1W scenario are summarized in Table 2. BO and PPO achieve a near-optimal profit in the first experiment where the demand is greater than warehouses’ capacities, whereas A3C and VPG perform slightly worse. All DRL algorithms achieve comparable results in the second and simpler experiment, with higher revenues but lower transportation and penalty costs. In the third and more complex experiment, which, on the contrary, involves lower revenues and higher transportation costs and penalties, the optimal profit is relatively small, but PPO tends to behave better than other DRL algorithms. BO, PPO, and A3C obtain satisfactory profits in the fourth and more balanced experiment, with increasing revenues and maximum demand value but reducing uncertainty, while VPG seems to perform poorly. The main difficulty here is represented by a wider search space (caused by greater storage capacities) and higher storage costs, especially for the factory. In the fifth and last experiment, the demand uncertainty increases, the penalty costs decrease, and it is more expensive to maintain stocks at the distribution warehouse rather than at the factory, but all DRL algorithms achieve comparable and near-optimal results.

	A3C	PPO	VPG	BO	Oracle
Exp 1	870 ± 67	1213 ± 68	885 ± 66	<b>1226 ± 71</b>	1474 ± 45
Exp 2	1066 ± 94	1163 ± 66	1100 ± 77	<b>1224 ± 60</b>	1289 ± 68
Exp 3	-36 ± 74	<b>195 ± 43</b>	12 ± 61	101 ± 50	345 ± 18
Exp 4	1317 ± 60	1600 ± 62	883 ± 95	<b>1633 ± 39</b>	2046 ± 37
Exp 5	736 ± 45	838 ± 58	789 ± 51	<b>870 ± 67</b>	966 ± 55

Table 2: Results covering the 1P1W scenario. It is possible to note how in general BO and PPO obtain near-optimal profits, while A3C and VPG seem more distant in terms of performance.

Table 3 summarizes the results for the 1P3W scenario, which in design is similar to the 1P1W scenario. The first experiment is characterized by a high maximum demand value, especially if compared with the capacities of the factory and of the first distribution warehouse; with this setting, BO performs worse than DRL algorithms. However, as PPO, it obtains a nearly optimal profit in the second experiment, where a simpler configuration is investigated. In the third and more challenging experiment, none of the algorithms achieves a profit greater than zero, with PPO achieving the worst one. Still, PPO outperforms A3C and VPG in the fourth and more balanced experiment, characterized by an increased search space and higher storage costs. Finally, BO and PPO achieve the best profits in the fifth experiment, where uncertainty and search space are increased, but fewer penalties are considered.

	A3C	PPO	VPG	BO	Oracle
Exp 1	1606 ± 139	<b>2319 ± 122</b>	803 ± 154	486 ± 330	3211 ± 60
Exp 2	2196 ± 104	<b>3461 ± 120</b>	2568 ± 112	3193 ± 101	3848 ± 95
Exp 3	-2142 ± 128	-4337 ± 216	-2638 ± 121	<b>-1682 ± 196</b>	772 ± 21
Exp 4	-561 ± 237	<b>2945 ± 135</b>	656 ± 140	1256 ± 170	4389 ± 64
Exp 5	1799 ± 306	<b>2353 ± 131</b>	1341 ± 79	2203 ± 152	2783 ± 91

Table 3: Results regarding the 1P3W scenario. PPO performs better than BO and other DRL algorithms on average, except in the third and more challenging experiment.

To conclude, Table 4 summarizes performances under the 2P2W scenario. The first experiment provides a balanced configuration, with maximum demand values and variations that change according to the specific product type, storage costs at the factory greater than those at the two distribution warehouses, and revenues particularly high for the first product type. Under such a mix, PPO achieves a good profit, as it also does A3C, which overcomes BO. For the second

8. All the figures regarding the three scenario and related to the convergence and the behavior of DRL algorithms and BO are available on <https://github.com/frenkowski/SCIMAI-Gym>.

experiment, sales prices for the second product type are increased and, accordingly, the associated revenues grow as well. Even storage and transportation costs are decreased, while penalties increase. With this configuration, PPO still obtains a nearly optimal result, and the same happens for VPG, while BO also behaves well. In the third experiment, capacities are increased, and we design alternating storage costs; this means, for example, that maintaining stocks of the first product type at the factory warehouse is the most inexpensive option while maintaining stocks of the second product type is the most expensive. The results allow us to conclude that PPO, followed by VPG, continues to perform successfully, whereas BO seems to suffer the most.

	A3C	PPO	VPG	BO	Oracle
<b>Exp 1</b>	2227 ± 178	<b>2783 ± 139</b>	1585 ± 184	2086 ± 173	3787 ± 102
<b>Exp 2</b>	1751 ± 83	<b>2867 ± 90</b>	2329 ± 98	2246 ± 114	3488 ± 63
<b>Exp 3</b>	1414 ± 128	<b>2630 ± 138</b>	2434 ± 156	552 ± 268	3549 ± 103

Table 4: Results concerning the 2P2W scenario. Results suggest that PPO behaves well generally, whereas BO seems slightly inferior compared to the other DRL algorithms.

## 7. Discussions and Conclusions

Results of numerical experiments demonstrated that the SCIM environment we propose is *effective* in representing states, actions, and rewards; indeed, the DRL algorithms we implemented have achieved *nearly optimal solutions* in all three investigated scenarios. In detail, PPO is the one that better adapts to different topologies and configurations of the SCIM environment achieving higher profits than other algorithms on average, although it fails to reach a positive profit in the most challenging experiment of the 1P3W scenario. VPG frequently appears to converge to a local maximum that seems slightly distant from PPO, especially when the number of warehouses increases, but it still obtains acceptable results. Finally, A3C is constantly the fastest DRL algorithm, as reported in Table 9 of Appendix B; however, just like VPG, it is never the best-performing one, perhaps because of its heightened sensitivity to the complex task of hyperparameter tuning.

It is worthwhile to mention that the BO approach also shows remarkable results, especially when the search space of optimal actions is limited, as in the 1P1W scenario. When compared to DRL algorithms, the BO approach seems to suffer more when there are two product types or when the demand exceeds the capacities. This is mainly due to the static and non-dynamic nature of the  $(s, Q)$ -policy, which does not allow developing an effective strategy, for example, for saving stocks in advance, but, conversely, culminates in a *myopic behavior*. Nevertheless, BO usually achieves respectable results, and the absence of hyperparameters to be tuned offers a considerable advantage.

### 7.1 Future Research

This paper can be extended and improved in many directions as:

- develop a *more comprehensive SCIM environment*, for example, by considering additional configurations mentioned in de Kok et al. (2018) (e.g., different demand distributions or different customers’ reactions to backordering). This allows us to evaluate the robustness of the proposed approaches under more realistic experiments;
- take into account the *non-linearity of transportation costs* (e.g., introducing a fixed cost independent of the number of stocks shipped effectively), as well as *non-zero leading times*. In this respect, a significant issue would be how to represent the system state as an MDP without violating the Markov property;
- *real-world data* could also be used to validate DRL algorithms and check whether they improve the performances of currently used SCIM systems in practice.

Lastly, even the BO approach could be *extended* to other standard static reorder policies, such as the base-stock policy, which has exactly half of the  $(s, Q)$ -policy parameters and can therefore enable faster convergence times.

## References

- Júlio César Alves and Geraldo Robson Mateus. Deep reinforcement learning and optimization approach for multi-echelon supply chain with uncertain demands. In *International Conference on Computational Logistics*, pages 584–599. Springer, 2020.
- Robert N Boute, Joren Gijsbrechts, Willem van Jaarsveld, and Nathalie Vanvuchelen. Deep reinforcement learning for inventory control: A roadmap. *European Journal of Operational Research*, 2021.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- S Kamal Chaharsooghi, Jafar Heydari, and S Hessameddin Zegordi. A reinforcement learning model for supply chain ordering management: An application to the beer game. *Decision Support Systems*, 45(4):949–959, 2008.
- Mustafa Cimen and Christopher Kirkbride. Approximate dynamic programming algorithms for multidimensional inventory optimization problems. *IFAC Proceedings Volumes*, 46(9):2015–2020, 2013.
- Ton de Kok, Christopher Grob, Marco Laumanns, Stefan Minner, Jörg Rambau, and Konrad Schade. A typology and literature review on stochastic multi-echelon inventory models. *European Journal of Operational Research*, 269(3):955–983, 2018.
- Jay W Forrester. Industrial dynamics. a major breakthrough for decision makers. *Harvard business review*, 36(4):37–66, 1958.
- Jay Wright Forrester. Industrial dynamics. *Journal of the Operational Research Society*, 48(10):1037–1041, 1997.
- Vincent François-Lavet, Peter Henderson, Riashat Islam, Marc G Bellemare, and Joelle Pineau. An introduction to deep reinforcement learning. *arXiv preprint arXiv:1811.12560*, 2018.
- Kevin Geevers. Deep reinforcement learning in inventory management. Master’s thesis, University of Twente, 2020.
- Joren Gijsbrechts, Robert N Boute, Jan A Van Mieghem, and Dennis J Zhang. Can deep reinforcement learning improve inventory management? performance on lost sales, dual-sourcing, and multi-echelon problems. *Manufacturing & Service Operations Management*, 2022.
- Christian D Hubbs, Hector D Perez, Owais Sarwar, Nikolaos V Sahinidis, Ignacio E Grossmann, and John M Wassick. Or-gym: A reinforcement learning library for operations research problems. *arXiv preprint arXiv:2008.06319*, 2020.
- Tommi Jaakkola, Michael I Jordan, and Satinder P Singh. On the convergence of stochastic iterative dynamic programming algorithms. *Neural computation*, 6(6):1185–1201, 1994.
- Lukas Kemmer, Henrik von Kleist, Diego de Rochebouët, Nikolaos Tziortziotis, and Jesse Read. Reinforcement learning for supply chain optimization. In *European Workshop on Reinforcement Learning*, volume 14, 2018.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- Hau L Lee, Venkata Padmanabhan, and Seungjin Whang. Information distortion in a supply chain: The bullwhip effect. *Management science*, 43(4):546–558, 1997.
- Liam Li, Kevin Jamieson, Afshin Rostamizadeh, Ekaterina Gonina, Moritz Hardt, Benjamin Recht, and Ameet Talwalkar. A system for massively parallel hyperparameter tuning. *arXiv preprint arXiv:1810.05934*, 2018.
- Yuxi Li. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*, 2017.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.
- Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. Ray: A distributed framework for emerging {AI} applications. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pages 561–577, 2018.
- Ahmad Mortazavi, Alireza Arshadi Khamseh, and Parham Azimi. Designing of an intelligent self-adaptive model for supply chain ordering management system. *Engineering Applications of Artificial Intelligence*, 37:207–220, 2015.
- Afshin Oroojlooyjadid, MohammadReza Nazari, Lawrence V Snyder, and Martin Takáč. A deep q-network for the beer game: Deep reinforcement learning for inventory optimization. *Manufacturing & Service Operations Management*, 2021.
- Zedong Peng, Yi Zhang, Yiping Feng, Tuchao Zhang, Zhengguang Wu, and Hongye Su. Deep reinforcement learning approach for capacitated supply chain optimization under demand uncertainty. In *2019 Chinese Automation Congress (CAC)*, pages 3512–3517. IEEE, 2019.
- Kiran Kumar Ravulapati, Jaideep Rao, and Tapas K Das. A reinforcement learning approach to stochastic business games. *IIE Transactions*, 36(4):373–385, 2004.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Stephen Shervais, Thaddeus T Shannon, and George G Lendaris. Intelligent supply chain management using adaptive critic learning. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 33(2):235–244, 2003.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Zheng Sui, Abhijit Gosavi, and Li Lin. A reinforcement learning approach for inventory replenishment in vendor-managed inventory systems with consignment inventory. *Engineering Management Journal*, 22(4):44–53, 2010.
- Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.
- John N Tsitsiklis. Asynchronous stochastic approximation and q-learning. *Machine learning*, 16(3):185–202, 1994.
- Benjamin Van Roy, Dimitri P Bertsekas, Yuchun Lee, and John N Tsitsiklis. A neuro-dynamic programming approach to retailer inventory management. In *Proceedings of the 36th IEEE Conference on Decision and Control*, volume 4, pages 4052–4057. IEEE, 1997.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Christopher John Cornish Hellaby Watkins. *Learning from Delayed Rewards*. PhD thesis, King’s College, Cambridge United Kingdom, 1989.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.

Cathy Wu, Aravind Rajeswaran, Yan Duan, Vikash Kumar, Alexandre M Bayen, Sham Kakade, Igor Mordatch, and Pieter Abbeel. Variance reduction for policy gradient with action-dependent factorized baselines. *arXiv preprint arXiv:1803.07246*, 2018.

Yimo Yan, Andy HF Chow, Chin Pang Ho, Yong-Hong Kuo, Qihao Wu, and Chengshuo Ying. Reinforcement learning for logistics and supply chain management: Methodologies, state of the art, and future opportunities. *Transportation Research Part E: Logistics and Transportation Review*, 162:102712, 2022.

## Appendix A. Experimental Plan

	Exp 1	Exp 2	Exp 3	Exp 4	Exp 5
Max Demand Value	10	5	5	10	5
Max Demand Variation	2	2	2	1	3
Sale Price	15	20	15	20	15
Production Cost	5	5	10	5	5
Storage Capacities	5, 10	5, 10	5, 10	10, 15	5, 10
Storage Costs	2, 1	2, 1	2, 1	4, 2	1, 2
Transportation Cost	0.25	0.05	1	0.25	0.25
Penalty Coefficient	1.5	0.1	2	1.5	0.1

Table 5: Experiments concerning the 1P1W scenario. When two values are present, the first one refers to the factory, while the second one refers to the first (and only) distribution warehouse.

	Exp 1	Exp 2	Exp 3	Exp 4	Exp 5
Max Demand Value	7	5	5	7	5
Max Demand Variation	2	2	2	1	3
Sale Price	15	20	15	20	15
Production Cost	5	5	10	5	5
Storage Capacities	3, 6, 9, 12	3, 6, 9, 12	3, 6, 9, 12	4, 8, 12, 16	4, 8, 12, 16
Storage Costs	4, 3, 2, 1	4, 3, 2, 1	4, 3, 2, 1	8, 6, 4, 2	4, 3, 2, 1
Transportation Costs	0.3, 0.6, 0.9	0.03, 0.06, 0.09	3, 2, 1	0.3, 0.6, 0.9	0.3, 0.6, 0.9
Penalty Coefficient	1.5	0.1	2	1.5	0.1

Table 6: Experiments regarding the 1P3W scenario. In this case, when there are four values, the first relates to the factory, while the remaining to the first, the second, and the third distribution warehouse, respectively.

	Exp 1	Exp 2	Exp 3
Max Demand Values	3, 6	3, 6	4, 2
Max Demand Variations	2, 1	2, 1	2, 2
Sale Prices	20, 10	10, 15	20, 10
Production Costs	2, 1	2, 1	2, 1
Storage Capacities	(3, 4), (6, 8), (9, 12)	(3, 4), (6, 8), (9, 12)	(9, 4), (6, 8), (3, 12)
Storage Costs	(6, 3), (4, 2), (2, 1)	(0.5, 0.3), (1.0, 0.6), (1.5, 0.9)	(1, 3), (2, 2), (3, 1)
Transportation Costs	(0.1, 0.3), (0.2, 0.6)	(0.01, 0.025), (0.02, 0.050)	(0.1, 0.3), (0.2, 0.6)
Penalty Coefficient	0.5	1.5	0.5

Table 7: Experiments covering the 2P2W scenario. Referring to the round brackets, the first value denotes the first product type, whereas the second value indicates the second product type.

	A3C	PPO	VPG
Hidden Layers	{(64, 64), (128, 128)}	{(64, 64), (128, 128)}	{(64, 64), (128, 128)}
Learning Rate	{1e-4, 1e-3}	{5e-4, 5e-3}	{4e-4, 4e-3}
Rollout Fragment Length	{10, 100}	{20, 200}	{10, 100}
Train Batch Size	{200, 2000}	{400, 4000}	{200, 2000}
Grad Clip	{20, 40}	{0, 20}	-
SGD Mini-Batch Size	-	{128, 256}	-
SGD Iterations	-	{15, 30}	-

Table 8: The hyperparameters of DRL algorithms selected for tuning. Through a grid search, we instantiated 880 DRL algorithm instances for the 1P1W scenario, 880 for the 1P3W, and 528 for the 2P2W, for a total of 2 288 instances. Each instance is trained for a given number of episodes: 15 000 episodes for the 1P1W scenario and 50 000 for the 1P3W and 2P2W scenarios.

**Appendix B. Training Times**

	<b>A3C</b>	<b>PPO</b>	<b>VPG</b>
<b>1P1W</b>	<b>5 ± 1</b>	13 ± 3	5 ± 0
<b>1P3W</b>	<b>21 ± 1</b>	47 ± 9	24 ± 2
<b>2P2W</b>	<b>24 ± 2</b>	53 ± 2	25 ± 0

Table 9: Average training times (in minutes) of the best DRL algorithm instances for the three scenarios considered. A3C proves to be constantly the fastest DRL algorithm, also due to the values of the hyperparameters related to the best training instances.