

## Michael Hoffman

Department of Industrial and Manufacturing  
Engineering,  
Pennsylvania State University,  
University Park, PA 16802  
e-mail: hoffman@psu.edu

## Eunhye Song

Department of Industrial and Manufacturing  
Engineering,  
Pennsylvania State University,  
University Park, PA 16802  
e-mail: eus358@psu.edu

## Michael Brundage

National Institute of Standards and Technology,  
Gaithersburg, MD 20899  
e-mail: mpb1@nist.gov

## Soundar Kumara

Department of Industrial and Manufacturing  
Engineering,  
Pennsylvania State University,  
University Park, PA 16802  
e-mail: skumara@psu.edu

# Online Maintenance Prioritization Via Monte Carlo Tree Search and Case-Based Reasoning

*When maintenance resources in a manufacturing system are limited, a challenge arises in determining how to allocate these resources among multiple competing maintenance jobs. This work formulates an online prioritization problem to tackle this challenge using a Markov decision process (MDP) to model the system behavior and Monte Carlo tree search (MCTS) to seek optimal maintenance actions in various states of the system. Further, case-based reasoning (CBR) is adopted to retain and reuse search experience gathered from MCTS to reduce the computational effort needed over time and to improve decision-making efficiency. The proposed method results in increased system throughput when compared to existing methods of maintenance prioritization while also reducing the computation time needed to identify optimal maintenance actions as more information is gathered. This is especially beneficial in manufacturing settings where maintenance decisions must be made quickly to minimize the negative performance impact of machine downtime. [DOI: 10.1115/1.4053408]*

**Keywords:** artificial intelligence, case-based reasoning, machine learning for engineering applications, manufacturing planning, Monte Carlo tree search

## 1 Introduction

In this work, we examine the problem of online maintenance prioritization in a manufacturing setting with capacity-constrained maintenance resources. Under this setting, periodic maintenance conflicts may occur where the number of machines that require maintenance exceeds the currently available maintenance capacity. The goal is then to determine where to allocate the limited maintenance resources in these instances in order to maximize the expected system performance. This maintenance prioritization problem may arise in any system where maintenance resources are limited, regardless of the maintenance strategy that is in place.

When a machine fails in a manufacturing system, it is unable to continue production until it is restored to a healthy state by a maintenance action. Downtime of a machine results in lost production time and may also force other machines in the system to become starved or blocked [1]. In complex systems with many machines, there can be a large number of alternative maintenance actions to consider. Thoroughly evaluating the resulting performance of each alternative course of action is computationally demanding and may require a substantial amount of time. Meanwhile, as downtime accumulates in the system the negative impact to performance increases. Therefore, selecting proper maintenance actions is a crucial task that should be done as quickly as possible so as to minimize machine downtime.

We formulate the maintenance prioritization problem as a Markov decision process (MDP) and seek the best action using Monte Carlo tree search (MCTS). A search for the optimal decision in the form of which immediate maintenance action to perform is conducted each time a maintenance conflict occurs. Additionally, we retain the information from each search for the best action to learn an effective policy and improve the decision-making process over time. As more information is gathered, the reasoner

will be able to predict the best action in some states of the system, reducing the rate at which simulation effort is expended.

The following criteria are given by Ref. [2] to classify a maintenance scheduling problem as “online”: (1) jobs arrive as an input stream over time and scheduling decisions must be made without knowledge of future jobs, (2) duration of jobs are unknown until the job is completed, and (3) machine failure and maintenance intervals are unknown [2]. Since our problem setting of interest meets each of these criteria, we consider the proposed maintenance prioritization method to be online.

The rest of this document is organized as follows: Sec. 2 contains a review of relevant existing work in both maintenance prioritization and simulation-based planning. Section 3 gives a more detailed description of the problem setting and the assumed behavior of the target system. The methodology is described in Sec. 4, and experimental results of the proposed method are presented in Sec. 5. Lastly, conclusions are given in Sec. 6.

## 2 Background

Maintenance priority management addresses the problem of allocating limited resources among competing maintenance needs. These resources can include labor, time, material and spare parts, or funds. We focus mainly on a limited number of maintenance workers which imposes a limit on the number of maintenance jobs that may be carried out simultaneously.

A review of recent work in maintenance optimization is given by Ref. [3] which identifies several references that consider optimization under constrained maintenance capacity. However, in each case, there are additional simplifying assumptions such as negligible time to repair or identical machines. Such assumptions are often made for mathematical convenience and are not representative of many real-world manufacturing systems [4]. Maintenance prioritization subject to limited resources is also examined by Ref. [5] in a variety of industries, including manufacturing. They identify the following four prioritization methods as the most common among recent literature: analytical hierarchy process, priority criterion, matrix-based priority, and failure mode, and effect analysis.

Analytical hierarchy process (AHP) is a method of evaluating alternative decisions using pairwise comparisons of weighted

Contributed by the Computers and Information Division of ASME for publication in the JOURNAL OF COMPUTING AND INFORMATION SCIENCE IN ENGINEERING. Manuscript received August 9, 2021; final manuscript received December 30, 2021; published online February 7, 2022. Assoc. Editor: Matthew I. Campbell.

This work is in part a work of the U.S. Government. ASME disclaims all interest in the U.S. Government's contributions.

criteria and sub-criteria [6]. Each criterion is compared against all others and assigned a relative numeric importance. Then, each candidate decision is evaluated against the criteria and the alternative with the highest score is selected as the best. The criteria and weights are typically selected by expert opinion. References [7,8] apply AHP in a maintenance setting using production disruption, mean time between failure, mean time to repair, and resource availability as selection criteria. Each maintainable asset is evaluated using these criteria resulting in a fixed priority ranking of assets. This approach does not account for the changing dynamics of a production system, where the criteria score may change depending on the state of the system. For example, as shown in Refs. [9,10], the production disruption that results from machine downtime will depend on the distribution of buffer contents throughout the system which will change over time.

Similar to AHP, both priority criterion and matrix-based priority use expert opinion to establish relevant criteria and evaluate competing maintenance needs [11]. Several case studies using priority criterion for maintenance are presented by Ref. [12]. Priority criterion for both long-term strategic planning and short-term operational decision support is considered. In complex production systems, however, it is not feasible to establish maintenance priorities for all operating states of the system even if expert input is available.

Failure mode and effects analysis (FMEA) aims to identify potential machine failures and their impact on the system [13]. It involves measuring the risk priority number (RPN) of each possible failure which is based on the likelihood of occurrence of the failure, the severity of the failure, and the ability to detect the failure. Generally, maintenance of failure modes with a greater risk of occurrence and severe impact are prioritized over lower-risk failures. Maintenance of a manufacturing system using FMEA is presented by Ref. [14]. This approach again relies heavily on expert opinion to evaluate maintenance actions against the RPN criteria. Doing so can be difficult or impossible for complex systems where the consequence of downtime is not easily inferred.

In each method of maintenance priority management discussed so far, priority criteria are developed and evaluated for a specific instance of a system. If the system is modified, the criteria will need to be reevaluated within the context of the new system. This can be costly and time-consuming, especially if inputs from experts and stakeholders are required. Additionally, an expert may struggle to identify optimal maintenance actions in a complex system with a high degree of interactions among components. The system may also encounter rare or unforeseen conditions with which an expert has no experience.

To alleviate the need for expert input when prioritizing maintenance, traditional queueing service disciplines have also been applied to real-time operational maintenance decision making. For instance, several common scheduling rules are compared by Ref. [15] including first-in, first-out (FIFO), shortest processing time first (SPTF), longest processing time first (LPTF), and an expert-derived static heuristic. Birnbaum importance is another heuristic metric that prioritizes machines according to their structural importance [16]. When used as a maintenance rule, Birnbaum importance prefers maintaining machines that have a greater likelihood of disrupting the system-level production, as demonstrated in Refs. [17,18]. While these scheduling rules may perform well in some scenarios, they do not consider the evolving state of the system. Reference [15] further proposes a dynamic priority that is found using an evolutionary algorithm for a single state of the system. Another dynamic scheduling priority measure based on the concept of a maintenance opportunity window [9] is proposed by Ref. [19]. The objective is to minimize the throughput disruption caused by machine downtime due to maintenance, although no substantial improvement is shown over a static FIFO rule.

While simple to implement, it is unlikely that a static priority will yield the optimal maintenance action in all states of a complex system. Monte Carlo planning algorithms aim to improve upon static rules for the action selection problem in systems with a

large state space. They use simulation to evaluate alternative sequences of actions in order to identify the optimal action in the current state. These algorithms typically require a generative model, or simulator, of the target system and strategically sample this model to estimate the expected performance of alternatives. The behavior of the generative model as well as the Markov decision process (MDP) model of the manufacturing system is formalized for our system of interest in Sec. 3.2.

Several methods exist to find optimal actions for an MDP. Value iteration, policy iteration, and  $Q$ -learning are common methods that involve exploring the state space and improving an agent's ability to choose rewarding actions as it gathers more experience [20]. However, the space complexity of each of these algorithms is a function of the system state space size, which limits their applicability to complex settings. To overcome this limitation, Kearns et al. propose a method of seeking optimal actions for large or infinite MDPs using an online search in the current state of the system. Their approach involves constructing a search tree by sampling each possible action a fixed number of times at each level of the tree until a specified depth is reached [21]. While this method provides useful theoretical guarantees on the optimality of the result, it is possible that significant simulation effort could be expended on suboptimal action trajectories. Reference [22] proposed the upper confidence bound for trees (UCT) algorithm to improve the efficiency at which simulation effort is expended during the tree search. UCT offers an action selection criterion that balances the exploitation of promising actions with the exploration of those that have been sampled less often. Monte Carlo tree search (MCTS) combines the UCT selection criteria with progressive tree building to effectively seek actions in large state space MDP settings and is the principal planning algorithm used throughout this work [23].

In Ref. [24], we have applied MCTS to the maintenance prioritization problem for a system subject to condition-based maintenance. Each time a maintenance conflict occurs, that is, in each instance, where the number of maintenance requests exceeds the currently available maintenance capacity, we apply MCTS to seek the optimal maintenance action. This approach has demonstrated an improvement in system throughput for a variety of system configurations. In this paper, we aim to improve upon this method by retaining and reusing experience that is gathered during each search.

Case-based reasoning (CBR) is a framework for using past experience to understand and solve newly encountered problems and is founded on the idea that similar problems have similar solutions [25–27]. CBR reflects the everyday human reasoning process that is often used when we attempt to solve problems. For example, from Ref. [25]:

When we order a meal in a restaurant, we often base decisions about what might be good on our other experiences in that restaurant and those like it. As we plan our household activities, we remember what worked and didn't work previously, and use that to create our new plans. A childcare provider mediating an argument between two children remembers what worked and didn't work previously in calming such situations, and bases her suggestion on that.

In a maintenance setting, when we encounter a state for which a maintenance decision is needed, we use CBR to identify past states that are sufficiently similar to the current state. Two system states might be considered similar if, for example, they share a similar distribution of buffer contents or if the same set of machines requires maintenance. We then examine the maintenance decisions that were made in those previously encountered similar states and determine if there is an action that can be retrieved and applied in the current state. Over time, as we gather more experience, we can retrieve existing actions more frequently which reduces the computational effort that is spent conducting an online search for the best action. A CBR agent, therefore, acts as a virtual expert that can provide the expert judgement needed in some maintenance priority management methods.

CBR has been applied to some areas of maintenance in past work. For example, in fault diagnosis for manufacturing equipment and

identification of the proper corrective action, as in Refs. [28–30]. Planning under limited maintenance resources is not considered, however. Reference [31] develops a maintenance decision support system that uses CBR but focuses on using cases to represent human expert knowledge rather than the results of an autonomous heuristic search.

CBR has also been effectively applied in some reinforcement learning (RL) settings which share a similar formulation as our maintenance planning problem. Many RL problems are also formulated as an MDP with the goal of choosing the best action in each state of the system. For example, Ref. [32] uses CBR to estimate the expected reward of actions in an episodic RL task. They address several challenges of using CBR for real-time decision support that are also important considerations in this work, such as storage and management of information in the case base. Reference [33] also examines the real-time decision support problem and proposes a general CBR framework for such settings which we adapt to our maintenance planning problem.

### 3 Problem Description

In this section, we elaborate on the manufacturing system behavior and the maintenance planning problem formulation.

**3.1 System Description.** The system studied in this work is a discrete manufacturing system consisting of multiple machines and buffers. While in a working state, each machine will retrieve a part from an upstream buffer (if one is available), process that part for some amount of time, and then place the completed part in a downstream buffer (if there is available space). The system has a single source of incoming unprocessed parts as well as a single sink to collect finished parts that leave the system. The overall system-level production is measured as the number of parts that traverse from source to sink over a specified time frame. The system can be viewed as a directed graph where each node represents either a machine or buffer and edges define possible routings through the system. Several real-world case studies that adopt this model of a manufacturing system are presented in Ref. [34].

Each machine is also subject to random failures, at which point the machine can no longer process parts and a corrective maintenance action is needed to restore the machine to a healthy state. It is assumed that failures are immediately observable and that a request for maintenance is generated as soon as a failure occurs.

The set of current maintenance requests forms a virtual “maintenance queue” of machines that are waiting to be repaired. When maintenance resources are limited, periodically situations may arise where the number of machines in the maintenance queue exceeds the available maintenance capacity. That is, at time  $t$  for the current maintenance queue  $L(t)$  and available maintenance capacity  $c_{\text{idle}}(t)$ , the system encounters a maintenance conflict if

$$1 \leq c_{\text{idle}}(t) < |L(t)| \quad (1)$$

In these instances, we must decide where to allocate the limited maintenance resources among the currently pending jobs. When a machine is chosen for maintenance, it is removed from the queue and seizes an available maintenance resource until the job is complete.

In the following, we consider only binary state machines (a machine is either working or failed) under a corrective maintenance strategy. This is to emphasize the contributions of the proposed planning methodology, which is easily extendable to any maintenance policy by using the maintenance queue formulation. The proposed method is agnostic to how or why maintenance jobs are placed in the queue so long as the other stated assumptions are met.

**3.2 Markov Decision Processes.** We model the underlying behavior of the manufacturing system as a Markov decision process (MDP), defined by the tuple  $M = \langle S, A, T, R, \gamma \rangle$  where

- $S$  is the state space
- $A$  is the action space and  $A(s)$  is the set of actions available in state  $s$
- $T$  is the transition probability function such that

$$T(s, a, s') = \Pr(s_{t+1} = s' | s_t = s, a_t = a)$$

- $R$  is the reward function where  $R(s, a, s')$  is the immediate reward obtained after transitioning from  $s$  to  $s'$  as a result of taking action  $a$
- $\gamma \in [0, 1]$  is the discount factor over the time horizon

MDP behavior is defined by a policy  $\pi$  that specifies the action to take in each state,  $\pi = \{\pi_t | \pi_t : S \mapsto A, t \geq 0\}$ . From an initial state  $s_0$ , the expected reward of following policy  $\pi$  at each step is given by

$$Q(s_0, \pi(s_0)) = \sum_{s' \in S} T(s_0, \pi(s_0), s') \cdot (R(s, \pi(s_0), s') + \gamma Q(s', \pi(s'))) \quad (2)$$

where  $Q(s, a)$  is the action-value function. It represents the expected discounted return when starting in a state  $s$  and following  $\pi$  indefinitely. The optimal policy  $\pi^*$  will maximize the reward in each state, defined as

$$\pi^* = \max_{\pi} Q(s, \pi(s)) \quad \forall s \in S \quad (3)$$

For our problem and many others,  $T$  and  $R$  are unknown but a generative model, or simulator,  $\mathcal{G}$  is available. When given a state and an action, the generative model is able to sample a resulting future state via  $\mathcal{G}(s, a) \rightarrow s'$ . The immediate reward is observable upon this transition to  $s'$ . This “single step” simulation can be chained together to obtain a trajectory of states and actions that emulates the system behavior over some (potentially infinite) time horizon. The problem then becomes to determine how to best allocate calls to the simulator in order to find a good sequence of actions.

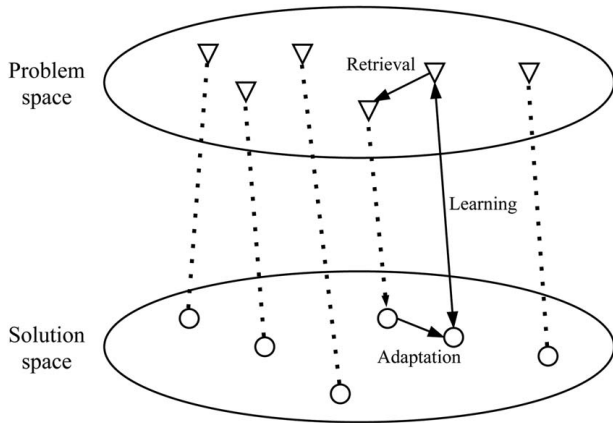
The online prioritization problem is formulated with  $s_0$  representing the current state of the system when Eq. (1) becomes true. We are interested in finding the action  $a$  that maximizes the expected return in the current state,  $Q(s_0, a)$ , given that our only knowledge of the system dynamics comes from  $\mathcal{G}$ .

The overall objective in selecting maintenance actions is to maximize the rate at which parts are produced by the system over an indefinite time horizon. We therefore model the reward function to represent the number of finished parts that leave the system. For example,  $R(s, a, s')$  would be equal to the difference in production in state  $s'$  and  $s$ . By choosing  $\gamma < 1$ , we place preference on obtaining rewards earlier instead of deferring rewards.

### 4 Methodology

The state space of the system grows exponentially with the number of machines and buffers, which precludes the application of classical dynamic programming algorithms for MDPs such as value iteration and policy iteration. The complexity of these algorithms is a function of the size of the state space, making them intractable for even moderately-sized systems. Instead, we turn to simulation-based decision tree search algorithms, namely Monte Carlo tree search (MCTS), to strategically sample the state space using the generative function  $\mathcal{G}$  and evaluate alternative actions.

In addition to using MCTS to seek an action from a given state, we would like to be able to learn from past experience to improve our decision making while reducing the computational effort needed over time. The proposed approach uses MCTS along with Case-based Reasoning (CBR) to retain and reuse the information that is gained over time. This approach is illustrated in Fig. 1, where the problem space consists of system states that are mapped to a point in the solution space representing the results of a search, including the estimated best action for that state. When



**Fig. 1 Experience storage and retrieval**

encountering a new state for which we need to determine the best action, we query the gathered information to determine if we have previously found a good solution for similar states and attempt to adapt that solution to the current state.

A high-level overview of the combined MCTS and CBR approach is as follows:

- (1) Operate the system until a state is encountered for which we want to find the optimal action.
- (2) Determine if the case base contains sufficient experience related to the current state. Sufficient experience is the ability to predict the optimal action in the current system state with an acceptable statistical confidence, as will be explained in Sec. 4.2.1.
- (3) If sufficient experience exists, retrieve the solution from the existing relevant cases, go to 4.
- (4) Formulate MCTS to seek a solution.
  - (a) Update the case base with experience gained during this search.
- (5) Return the estimated optimal action as the best action as for the current state, go to 0.

This procedure can be run indefinitely for a specified system so long as a simulator  $\mathcal{G}$  is available and accurately reflects the system's behavior. The following sections describe the steps of this approach in more detail.

**4.1 Monte Carlo Tree Search.** The general MCTS algorithm is initialized by creating a root node  $v_0$  representing the current state of the physical system,  $s_0$ . At each successive iteration, a leaf node  $v_\ell$  is selected according to a "tree policy." A simulation result is obtained from the state of this node,  $s(v_\ell)$ , using the "default policy" and then backed up through the tree from  $v_\ell$  to  $v_0$ . Once the specified search budget is exhausted, we choose the best estimated action from the current state.

Upper confidence bound for trees (UCT) [22] is the typical node selection criterion for the tree policy of MCTS. According to UCT, at each node  $v_i$ , the action  $a$  should be chosen that maximizes the quantity

$$UCT(v_i, a, C_p) = \hat{Q}_i(s(v_i), a) + 2C_p \sqrt{\frac{\ln \bar{n}}{N_{a,i}^{s(v_i)}}} \quad (4)$$

where  $\hat{Q}_i(s(v_i), a)$  is the estimated expected reward from choosing action  $a$  in state  $s(v_i)$  as defined by Eq. (2),  $C_p$  is the exploration constant,  $\bar{n}$  is the number of times node  $v_i$  has been visited so far, and  $N_{a,i}^{s(v_i)}$  is the number of times action  $a$  has been chosen from node  $v_i$ .

At termination, we choose the action that maximizes Eq. (4) with  $C_p = 0$  as the best action. For each candidate action, MCTS also

returns the mean reward, reward variance, and number of times that action was sampled from the initial state. These statistics are used during the CBR procedure to conduct a statistical comparison of action rewards.

**4.2 Case-Based Reasoning.** In general, experience is stored in a case base  $CB$  which is a set of cases describing problem instances and their associated solutions. Initially,  $CB = \emptyset$  and new cases are added over time as new problem instances are solved. As additional experience is gathered, the case base can be used to predict the optimal action in newly encountered states with greater confidence.

There are four main steps in the CBR process when a new problem instance is encountered:

- (1) Retrieve: Identify the existing cases most relevant to the target problem.
- (2) Reuse: Adapt a solution from the existing cases to be applied to the target problem.
- (3) Revise: Adjust the implemented solution based on observation of the results.
- (4) Retain: Add the new problem-solution pair to the case base as a new case.

We use the numeric system state vector described in Ref. [24] for problem descriptions in the case base under a manufacturing setting. To summarize, the state of a manufacturing system is determined by the level of each buffer in the system, the elapsed processing time of each machine that currently has a part in progress, the degradation state of each machine, and the elapsed repair time for machines currently under repair. We encode each of these attributes with numeric values to create a numeric state-space vector. The solutions are stored as the results of each search in the form of the mean, variance, and sample size of each candidate action in the initial state.

Given input as a vector of state variables representing the current system state, the case-based reasoner will return an output in the form of a prediction of the best action for that state. The following section will describe the formulation of this prediction problem in more detail.

**4.2.1 Action Label Prediction Task.** In this section, we formulate the problem of predicting the best action for a queried state given the set of experience that has been gathered so far. We assume that each state can be represented by a numeric vector of state variables.

This task can be viewed as a classification problem where our goal is to classify each state according to its optimal action. The action space  $A = \{a_1, a_2, \dots, a_m\}$  is the set of possible labels and states that share the same optimal action, therefore, belong to the same class. In order to train a classifier, we first need to gather experience by applying MCTS in various states to seek the best action. Using a traditional statistical classification approach, the best action as determined by MCTS is used as a label for the initial state of the search and this state-solution pair becomes a case in  $CB$  and one instance in the classifier training set. Each training instance, or case, for the classification problem is defined as

$$c^i = (s_i, a_i) \quad (5)$$

where  $s_i$  is the vector of state variables defining a particular state and  $a_i$  is the best action label assigned to that state. Once a sufficiently large training set is established, a classifier can be trained on these instances and used to generate predictions for newly observed states.

A significant limitation of this approach is that it does not account for the uncertainty of MCTS. In some cases, MCTS may be unable to distinguish which action is best if multiple actions from the current state yield a similar estimated reward. Conversely, MCTS may very easily distinguish an action that is clearly better than the others, so our predictions should ideally reflect these cases as well. To account for this uncertainty, we can instead aim to predict the likelihood that any particular action is optimal in the

current state. This requires assigning some probabilistic measure to the labels for each state. The cases are then represented by

$$c^i = (s_i, \langle a_1, q_1 \rangle, \dots, \langle a_m, q_m \rangle) \quad (6)$$

where  $q_i \in [0, 1]$  is a measure of our belief that action label  $a_i$  is optimal. Since we are now aiming to predict the continuous  $q_i$  values, we are faced with a regression problem. This problem setting is referred to by Ref. [35] as “learning classification with auxiliary probabilistic information.”

Predicting  $q_i$  for each action  $a_i$  at a queried state instance requires  $|A|=m$  regressors. To perform the regression, we use instance-based learning methods which avoids generalizing the training data until a query is made to the predictor. Instance-based learning algorithms (also sometimes called “lazy learning”) do not require a training period which is especially beneficial for case-based reasoning since we do not need to explicitly retrain each regressor when the experience stored in the case base changes [36]. In this work, we use  $k$ -nearest neighbor regression, which fits these criteria and has been successful in many previous applications of CBR.

Once the predicted action label confidence  $\hat{q}_i$  is calculated for each action of a queried state, we choose that with the greatest confidence exceeding some specified threshold  $\zeta$  as the predicted optimal action. If no predicted action label confidence exceeds this threshold, we formulate and solve MCTS to seek the best action. The experience obtained from this MCTS instance is then potentially added to the case base to aid future predictions.

**4.2.2 Measuring Action Label Confidence.** For a set of  $m$  alternative actions, we want to find the optimal action  $a_i^*$  where

$$a_i^* := \arg \max_{1 \leq i \leq m} E[X_i] \quad (7)$$

where  $X_i$  is the return of taking action  $a_i$  in the current state. However, with the uncertainty of MCTS we instead aim to determine the probability that an action  $a_i$  is optimal given the observations gathered from a search. This probability will serve as the auxiliary probabilistic information for the action labels of each state in the case base.

We introduce a stochastic process  $\eta_i$  to model the uncertainty regarding  $E[X_i]$ . By adopting a normal-normal model for  $\eta_i$  as described by Ref. [37], we assume

$$X_i \sim N(\eta_i, \lambda_i^2), \quad \eta_i \sim N(\mu_i, \sigma_i^2) \quad (8)$$

where  $\lambda_i^2$  is the simulation error variance, and  $\mu_i$  and  $\sigma_i^2$  are the mean and variance of the prior distribution of  $\eta_i$ , respectively.

Each instance of MCTS returns a set of statistics from sampling each candidate action  $a_i$  from the initial state including mean reward  $\bar{X}_i$ , sample size  $n_i$ , and sample variance  $S_i^2$ . We choose the noninformative prior  $\sigma_i^2 = \infty$  and use  $S_i^2$  as a plug-in estimate of  $\lambda_i^2$ . We then update the posterior mean and variance of  $\eta_i$  in

$$\mu_i = \bar{X}_i, \quad \sigma_i^2 = S_i^2/n_i \quad (9)$$

The probability that an action  $a_i$  is optimal in the current state is therefore given by

$$\begin{aligned} q_i &= \Pr\left(\eta_i \geq \max_{j \neq i} \eta_j\right) \\ &= \int \Pr(x \geq \max_{j \neq i} \eta_j) \cdot f_i(x) dx \\ &= \int \left[ \prod_{j \neq i} \Pr(\eta_j \leq x) \right] \cdot f_i(x) dx \end{aligned} \quad (10)$$

where  $f_i(\cdot)$  is the probability density function of the distribution  $N(\mu_i, \sigma_i^2)$  as defined in Eq. (8). We then set  $q_i$  to indicate our confidence that action  $a_i$  is optimal in the current state.

Additionally, we may be interested in finding the probability that the estimated reward of an action is within  $\delta$  of the best alternative for some user-specified  $\delta > 0$ . We refer to  $\delta$  as the optimality indifference parameter reflecting that the decision maker is indifferent about the optimality gap as long as it is within  $\delta$ . This probability is given by

$$\begin{aligned} q_i &= \Pr\left(|\eta_i - \max_{1 \leq j \leq m} \eta_j| \leq \delta\right) \\ &= \Pr\left(\eta_i + \delta \geq \max_{1 \leq j \leq m} \eta_j\right) \\ &= \Pr(i \text{ is best}) + \Pr\left(\eta_i + \delta \geq \max_{1 \leq j \leq m} \eta_j, i \text{ is not best}\right) \end{aligned} \quad (11)$$

The term  $\Pr(i \text{ is best})$  is given in Eq. (10). The second term of Eq. (11) is

$$\begin{aligned} \Pr\left(\eta_i \leq \max_{j \neq i} \eta_j \leq \eta_i + \delta\right) &= \int \Pr\left(x \leq \max_{j \neq i} \eta_j \leq x + \delta\right) \cdot f_i(x) dx \\ &= \int \left[ \sum_{\ell \neq i} b_\ell \right] \cdot f_i(x) dx \end{aligned} \quad (12)$$

where

$$\begin{aligned} b_\ell &= \Pr\left(x \leq \eta_\ell \leq x + \delta, \eta_\ell \geq \max_{j \neq i, j \neq \ell} \eta_j\right) \\ &= \int_x^{x+\delta} \left[ \prod_{j \neq i, j \neq \ell} \Pr(\eta_j \leq y) \right] \cdot f_\ell(y) dy \end{aligned} \quad (13)$$

We use Eq. (11) to calculate  $q_i$  for each alternative action whenever a new case is added to the case base or the statistics of an existing case are updated.

**4.2.3 CBR Procedure.** In this section, we describe each step of the CBR procedure as it applies to our problem.

**Retrieval.** The retrieval step involves identifying the set of cases that are most similar to the queried case. We use the Minkowski distance metric, defined as

$$D(\mathbf{x}_i, \mathbf{x}_j) = \left( \sum_{k=1}^n |x_{ik} - x_{jk}|^p \right)^{1/p} \quad (14)$$

for state vectors of variables scaled to the interval  $[0, 1]$ . We consider  $p=1$  (Manhattan distance) and  $p=2$  (Euclidean distance) in our experiments. A lesser distance between two states indicates a higher degree of similarity.

**Reuse.** Using a probabilistic classification approach, we attempt to reuse experience in the case base by predicting the likelihood that each available action is optimal in a particular state by training regression models using training examples of the form given by Eq. (6).

For each action and its associated regressor, we predict the action label confidence for the queried state. If the predicted confidence for action  $a_i$  is greater than or equal to the specified retrieval threshold  $\zeta$ , then this action is eligible for retrieval in the current state. If there is more than one action whose predicted confidence exceeds the threshold, we choose the action with the highest predicted value.

We use  $k$ -nearest neighbor (kNN) regression to predict the confidence for each action in the queried state. Under this method, we identify the  $k$  cases in  $CB$  with the minimum distance to the queried case according to Eq. (14). The predicted confidence for action  $a_i$  is then the average value  $q_i$  among the  $k$  neighbors. We also consider the distance-weighted average where neighbors closer to the queried point have a greater influence on the predicted value than those further away.

*Revision.* Once a solution for the queried state is identified, we may choose to revise that solution based on the observed outcome of implementing that solution. Necessary revisions are applied to the solution before storing this new problem-solution pair in the case base.

In our setting, however, it is difficult to retroactively determine if good (or poor) system performance was due to the action selected or the random behavior of the system. Newly encountered cases for which we retrieve an existing solution are therefore not stored in the case base. Experience is stored only if it is obtained as a result of MCTS.

*Retention.* After each instance of MCTS, if the initial state is represented in the case base, we update the case with the results from the new search. These include the mean reward, variance, and sample size of each candidate action in that state. Each statistic is updated incrementally, so that the statistics in the case base reflect the results of every search conducted from that state. As the number of samples for each action increases, we become more confident in the estimation of that action's reward.

If the initial state of an MCTS instance is not already in the case base, we create a new case from the initial state vector and the search statistics and add it to  $CB$ . Limiting the number of cases in  $CB$  is necessary to limit the time needed to query the case base and to avoid exceeding memory constraints. Once the case base exceeds its maximum specified size, we must determine which cases to keep and which ones to discard.

Three factors are identified by Ref. [32] to determine the overall "value" of each case stored in the case base:

- the age of the case,
- the density of cases in the neighborhood in which the case resides, and
- the accuracy of predictions generated by the case.

Given that we are interested in studying a system in its steady-state over an indefinite time horizon, the age of a case is less relevant than the other criteria. We therefore only consider the case neighborhood score and the regress error score when evaluating each case.

The first term, the case neighborhood score, determines the density of cases surrounding  $c^i$  and is given by

$$S_n(c^i) = \varphi(c^i) \cdot \sigma(c^i) \quad (15)$$

where

$$\varphi(c^i) = \frac{1}{k} \sum_{c^j \in NN_k(c^i)} sim(c^i, c^j) \quad (16)$$

$$\sigma(c^i) = \frac{1}{k} \sum_{c^j \in NN_k(c^i)} 1 - \kappa(c^i, c^j) \quad (17)$$

$NN_k(c^i)$  is the set of  $k$  nearest neighbors of  $c^i$ , and  $\kappa(c^i, c^j)$  is the proportion of actions in  $A(s_i) \cap A(s_j)$  for which  $c^i$  and  $c^j$  agree on whether or not confidence in the action should be above or below the retrieval threshold  $\zeta$ . If  $A(s_i) \cap A(s_j) = \emptyset$  then  $\kappa(c_i, c_j) = 0$ . A greater value of  $\varphi(c^i)$  indicates that a case is very similar to its neighbors on average, while a greater  $\sigma(c^i)$  indicates greater rate of disagreement (and therefore greater variability) among cases in the neighborhood.

The second term, the regress error score, is given by

$$S_e(c^i) = \sum_{c^j \in NN_k(c^i)} sim(c^i, c^j) \cdot (1 - \kappa(c^i, \hat{c}^j)) \quad (18)$$

where  $\hat{c}^j$  is the prediction of  $c^j$  using  $CB \setminus c^j$ . A lower value of  $S_e(c^i)$  indicates that  $c^i$  is more useful for predicting the state value of its neighbors.

These components are then combined into a single heuristic:

$$S(c^i) = S_n(c^i) + S_e(c^i) \quad (19)$$

A higher score relative to other cases indicates that a case is "worse" with regards to each of these criteria. We remove the case with the greatest score until the number of cases is within the limit.

## 5 Results

To demonstrate the proposed methodology, we examine the performance of several maintainers, or "agents," by evaluating (1) the average production (reward) rate obtained by the agent and (2) the simulation effort that is expended over time. If the CBR approach is effective, an agent should be able to maintain a relatively constant reward rate while reducing its simulation effort. This would indicate that an agent is effectively learning from its experience as it interacts with the environment. Throughout our experiments, we use the PYTHON discrete event simulation package Simantha to model and simulate the system of interest [38]. This package was developed by the authors to support maintenance optimization via simulation for complex manufacturing systems.

We apply the following procedure to evaluate each CBR maintainer:

- (1) Simulate each system while applying MCTS at each decision point to gather initial experience.
- (2) Tune CBR retrieval regressor parameters using the initial experience.
- (3) Continue simulating the system using the proposed MCTS + CBR method.

The structure of the system studied in this example is shown in Fig. 2. We set the maintenance capacity in each experiment to 1. This system includes a complex arrangement of machines for which effective maintenance planning is not easily derived by inspection or achieved by analytical methods.

To model machine degradation we use the notion of Bernoulli machines as described by Ref. [34]. Under this reliability model, each machine has some probability of failure at the beginning of each discrete time-step. The time to failure for each machine is therefore geometrically distributed. If a failure occurs, then that machine can no longer function until it receives maintenance that restores it to a healthy state. We can therefore represent the health state of each machine with a binary state variable indicating whether or not the machine is failed. Similarly, the time to repair is geometrically distributed and the repair state of a machine is also a binary variable indicating whether or not the machine is currently under repair.

The proposed method may be used with any discrete distribution of cycle time, time to failure, and time to repair as long as the state variable is chosen such that it is adequate for representing a machine's condition. Choosing a geometric distribution for each of these quantities and the use of binary state variables allows us to reduce the size of the state space. MCTS is therefore able to converge to an optimal solution more quickly which is convenient for our illustrative example.

The production status of each machine must also be captured in the state representation. Assuming geometrically distributed cycle times for each machine again allows us to use binary state variables




Fig. 2 Nine-machine complex production line

to indicate whether a machine has a part, and also whether or not that part has finished processing. Considering each of these machine state variables, there are five unique machine states as summarized in Table 1. Our method also allows for a more complex system state space, such as that examined by Ref. [24] where state variables are integer-valued rather than binary. If machines do not abide by the Benoulli reliability model, for example, the degradation state of each machine can be represented by a variable indicating the time since the last repair.

Table 2 shows the cycle time distribution for each machine that is used throughout the following example. Each buffer has a maximum capacity of five units. The time to repair distribution for machines  $M_1$ ,  $M_6$ ,  $M_7$ ,  $M_8$ , and  $M_9$  is  $Geom(1/10)$  minutes. For machines  $M_2$ ,  $M_3$ ,  $M_4$ ,  $M_5$ , the time to repair distribution is  $Geom(1/20)$  minutes. The time to failure for each machine is distributed  $Geom(1/100)$  minutes.

We evaluate FIFO, Birnbaum importance (BI), and the expert heuristic (EH) suggested by Ref. [15] as static heuristics for maintenance prioritization in this system. The Birnbaum importance measure for each machine in this arrangement is given in Table 3 where a greater value indicates a higher priority for maintenance. We defer to Ref. [16] for details of this calculation and resolve ties using FIFO.

The expert heuristic gives the following list of rules that are used to determine which machine should be prioritized for maintenance:

- (1) A machine with a longer cycle time should be prioritized over a machine with a shorter cycle time.
- (2) Machines in a serial section of the line should be prioritized over machines in a parallel section.
- (3) A machine in a parallel station containing a fewer number of machines should be prioritized over machines in parallel stations with a greater number of machines.

**Table 1 Machine state summary**

Has part	Processing	Failed	Under repair	State description
0	0	0	0	Starved
1	1	0	0	Processing
1	0	0	0	Blocked
0	0	1	0	Awaiting repair
0	0	1	1	Under repair

**Table 2 Machine cycle time distribution in minutes**

Machine	Cycle time
$M_1, M_9$	$Geom(1/5)$
$M_2$	$Geom(1/30)$
$M_3, M_4, M_5$	$Geom(1/90)$
$M_6, M_7$	$Geom(1/20)$
$M_8$	$Geom(1/10)$

**Table 3 Static Birnbaum importance (BI) and expert heuristic (EH) priorities for the nine-machine line where a higher priority value indicates a machine should be repaired earlier**

Machine	Priority	
	BI	EH
$M_1, M_9$	5	1
$M_2$	2	4
$M_3, M_4, M_5$	1	5
$M_6, M_7$	3	3
$M_8$	4	2

To apply this heuristic, we first use rule (1) to find the machine in the maintenance queue with the greatest average cycle time. If there is a tie, we then consider rule (2), and finally rule (3). If a tie remains after applying each rule, we defer to FIFO. The priorities of each machine under this method are given in Table 3, where again a greater value indicates a higher priority. Lastly, we also include a random selection policy for reference.

We consider two factors when creating maintainers using the proposed MCTS+CBR method: the retrieval threshold  $\zeta$  and the maximum size of the case base. We use three levels of  $\zeta$  in our experiments, 0.5, 0.9, and 0.95. A smaller  $\zeta$  increases the likelihood that we will be able to retrieve an action from the case base, although we run the risk of retrieving an action that is suboptimal. On the other hand, a greater  $\zeta$  will force the reasoner to be more selective in the actions it retrieves. We also consider the case with no limit on the maximum case base size as well as a limit of 100 cases. In our experiments, instances with a limited case base size are denoted  $kNN_{100}$ . Three levels of  $\zeta$  and two levels of case base capacity results in six total maintainers that we evaluate using our approach. For each case, we set the optimality indifference ( $\delta$  in Eq. (11)) to be 10% of the maximum average observed reward. The confidence label of each action, therefore, represents the probability that an action is within 10% of the best given the observations so far.

Lastly, we also evaluate a maintainer that resolves every maintenance conflict using MCTS and does not retain any experience from the search. For each instance of MCTS, including the searches conducted by the CBR maintainers, we use a search budget of 200 iterations, discount factor  $\gamma=0.95$ , and exploration constant  $C_p=10$ . These settings performed well throughout our experimentation.

**5.1 Case-Based Reasoning Maintainer Tuning.** We simulate for an initial period of 12 weeks to obtain experience that is used to tune the regression parameters of the CBR maintainers. After 12 weeks, 3268 unique states were encountered where MCTS was conducted to determine a maintenance action. When tuning the parameters of the action label confidence regression models, we aim to maximize the probabilistic classification accuracy of each model. For a given set of case instances with known action confidence levels, the probabilistic classification accuracy is the proportion of predictions that are correctly above or below the retrieval threshold  $\zeta$ .

Since we have nine possible maintenance actions corresponding to the nine machines in the system, we tune nine regressors independently so that each predict the likelihood of a particular action is optimal in each state. We perform a grid search of parameters for the  $kNN$  regression model and use  $K$ -fold validation with  $K=5$  to find the best parameters. Tables 4–6 show the resulting best parameters for  $\zeta=0.50, 0.90$ , and  $0.95$ , respectively. In each of these tables, the parameter  $k$  is the number of neighbors to consider for  $kNN$  regression and  $p$  is the Minkowski distance metric

**Table 4 kNN regressor tuning results for  $\zeta=0.50$**

	Training instances	Best model			
		$k$	$p$	Weights	Accuracy
$M_1$	1119	33	1	Distance	76.85%
$M_2$	1357	25	2	Distance	74.73%
$M_3$	1411	33	2	Distance	74.34%
$M_4$	1358	25	1	Distance	74.96%
$M_5$	1413	29	1	Distance	75.31%
$M_6$	1148	13	1	Distance	74.82%
$M_7$	1103	33	1	Distance	72.26%
$M_8$	637	5	1	Uniform	87.28%
$M_9$	628	9	2	Uniform	85.51%

Note: Overall accuracy: 76.21%.

**Table 5 kNN regressor tuning results for  $\zeta = 0.90$**

Training instances	Best model				
	$k$	$p$	Weights	Accuracy	
$M_1$	1119	5	2	Distance	66.93%
$M_2$	1357	5	1	Distance	78.78%
$M_3$	1411	9	2	Distance	79.73%
$M_4$	1358	9	2	Distance	80.48%
$M_5$	1413	9	1	Uniform	81.31%
$M_6$	1148	5	2	Distance	67.77%
$M_7$	1103	5	1	Distance	69.45%
$M_8$	637	1	1	Uniform	71.27%
$M_9$	628	1	2	Uniform	71.18%

Note: Overall accuracy: 74.99%.

**Table 6 kNN regressor tuning results for  $\zeta = 0.95$**

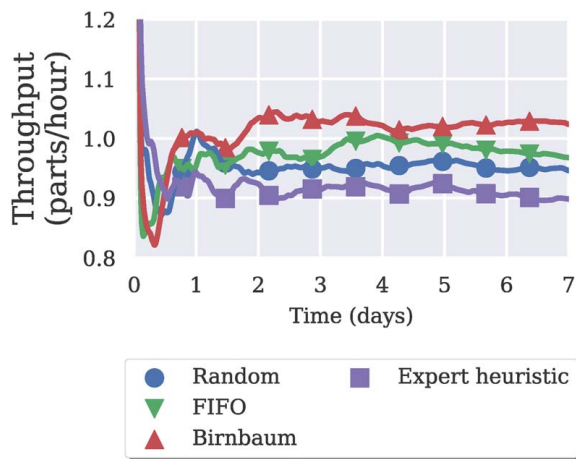
Training instances	Best model				
	$k$	$p$	Weights	Accuracy	
$M_1$	1119	5	2	Distance	66.93%
$M_2$	1357	5	1	Distance	78.78%
$M_3$	1411	9	2	Distance	79.73%
$M_4$	1358	9	2	Distance	80.48%
$M_5$	1413	9	1	Uniform	81.31%
$M_6$	1148	5	2	Distance	67.77%
$M_7$	1103	5	1	Distance	69.45%
$M_8$	637	1	1	Uniform	71.27%
$M_9$	628	1	2	Uniform	71.18%

Note: Overall accuracy: 80.01%.

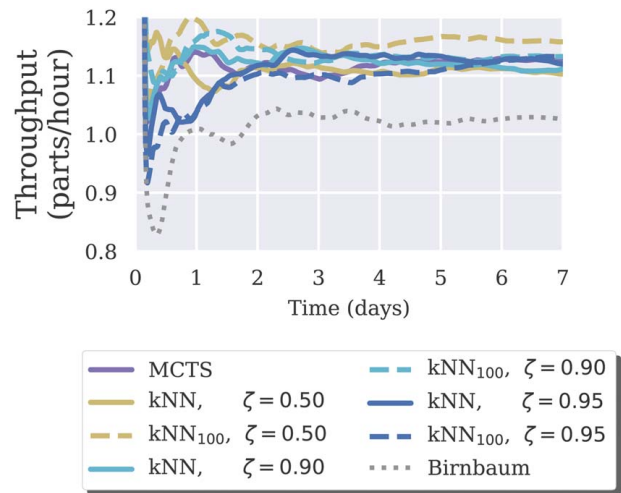
parameter.  $p = 1$  is equivalent to the Manhattan distance while  $p = 2$  is Euclidean distance.

**5.2 Maintainer Comparison Results.** The system is simulated under each maintainer for a period of one week for 48 independent replications. The resulting average throughput for each of the static heuristic maintenance priority rules is shown in Fig. 3. The Birnbaum importance heuristic yields the greatest average throughput with 1.0225 parts per hour.

Figure 4 compares the throughput of each of the online prioritization maintainers. The average throughput of the best static heuristic maintainer is also included for reference. In each case, online prioritization via MCTS and CBR outperforms the static priority



**Fig. 3 Complex system throughput comparison for static scheduling rules**



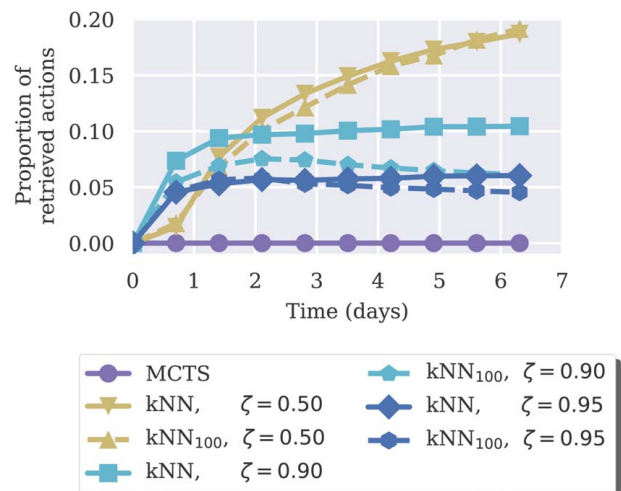
**Fig. 4 Complex system throughput comparison for CBR scheduling**

scheduling rules. The average throughput of each maintainer is given in Table 7.

We also consider the rate at which simulation effort is expended by examining the frequency at which MCTS is conducted over time as shown in Fig. 5. The vertical axis gives the average proportion of maintenance actions that are retrieved from the case base at the

**Table 7 Maintainer throughput comparison results**

Maintainer	$\zeta$	Throughput (parts/hour)		Retrieval rate
		Mean	SE (%)	
Random	—	0.9455	1.6611	—
FIFO	—	0.9667	1.9717	—
BI	—	1.0225	1.7977	—
EH	—	0.8977	1.9370	—
MCTS	—	1.1266	1.1145	0.0000
kNN	0.50	1.1031	1.0497	0.1942
kNN <sub>100</sub>	0.50	1.1584	1.1568	0.2017
kNN	0.90	1.1132	1.3950	0.1063
kNN <sub>100</sub>	0.90	1.1334	1.3738	0.0587
kNN	0.95	1.1192	1.1574	0.0614
kNN <sub>100</sub>	0.95	1.1323	1.4562	0.0442



**Fig. 5 CBR retrieval rate results**



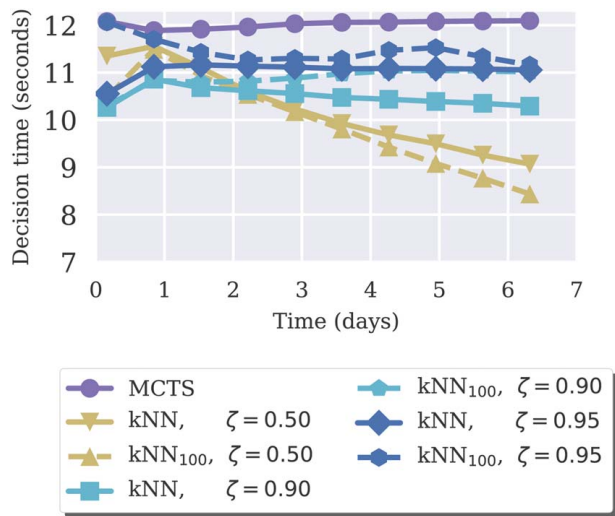


Fig. 6 CBR decision time results

occurrence of a maintenance conflict. For each CBR maintainer, the rate at which we are able to retrieve actions increases as we gather more experience. When the threshold for retrieval  $\zeta$  is lower, we retrieve actions more often. Each CBR maintainer provides a relatively constant throughput improvement over the baseline static heuristic policies, indicating that there is no loss in performance when reusing solutions from their experience.

The reduction in simulation effort is also demonstrated in Fig. 6, which shows the average duration of time in seconds needed to make a decision when a maintenance conflict occurs. Since the computational effort needed to retrieve an existing action from the case base is much lower than that needed to conduct MCTS, the maintainers with a higher proportion of retrieved actions are able to make decisions more quickly on average. This is particularly valuable in manufacturing settings where critical maintenance decisions must be made as soon as possible. Reference [39] identifies this time between machine failure and the start of its repair, or “time to dispatch”, as a key performance indicator (KPI) in maintenance. Reducing this time by making maintenance decisions more quickly can reduce the overall downtime of a machine and increase the productivity of the system.

## 6 Conclusions

In this work, we have demonstrated a method of online maintenance prioritization in a dynamic manufacturing setting by using strategic sampling of a generative simulation model via MCTS. Further, we have used CBR to retain and reuse search experience to infer optimal maintenance actions in cases where the current state is similar to those previously encountered. The proposed methods offer improved system-wide performance compared to static heuristic scheduling rules without the need for expert judgement.

A natural extension of this work is to apply the proposed method of maintenance prioritization to more advanced maintenance strategies such as condition-based or predictive maintenance. This would mainly affect the system state representation since we may no longer be able to represent the health state of a machine as a binary variable. Under condition-based maintenance, for example, a machine can be in several possible states of degradation beyond “healthy” or “failed” as we assume in this work. Considering these additional health states will greatly increase the size of the state space and present additional computational challenges that will be addressed in future research.

Future work also includes examining alternative state similarity metrics beyond Minkowski distance. The chosen metric has a

significant impact on nearest-neighbor methods, and a learned distance metric may result in better performance of the CBR model. Additionally, an active learning approach to case base management may further improve CBR performance. Such an approach would involve identifying prototypical cases whose best action estimates would provide the most useful information to surrounding states. A search for the best action in prototype states can also be conducted offline, so that the case-based reasoner does not have to wait for a maintenance conflict to occur before gathering additional useful experience.

## Acknowledgment

Michael Hoffman would like to acknowledge the National Institute of Standards and Technology (NIST) for their support of this work through the Graduate Measurement Science and Engineering Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this dissertation are my own and do not necessarily reflect the views of NIST. Soundar Kumara would like to acknowledge the Clean Energy Smart Manufacturing Innovation Institute (CESMII) for their support of this work. Any opinions, findings, and conclusions or recommendations expressed in this dissertation are the authors’ own and do not necessarily reflect the views of CESMII.

## Conflict of Interest

There are no conflicts of interest.

## Data Availability Statement

The datasets generated and supporting the findings of this article are obtainable from the corresponding author upon reasonable request. The data and information that support the findings of this article are freely available.

## References

- [1] Chang, H. S., Fu, M. C., Hu, J., and Marcus, S. I., 2005, “An Adaptive Sampling Algorithm for Solving Markov Decision Processes,” *Oper. Res.*, **53**(1), pp. 126–139.
- [2] Albers, S., 2009, “Online Scheduling.” Introduction to Scheduling, Y. Robert and F. Vivien, eds. CRC Press, ch. 3.
- [3] Keizer, M. C. O., Flapper, S. D. P., and Teunter, R. H., 2017, “Condition-Based Maintenance Policies for Systems With Multiple Dependent Components: A Review,” *Eur. J. Oper. Res.*, **261**(2), pp. 405–420.
- [4] Alrabghi, A., and Tiwari, A., 2015, “State of the Art in Simulation-Based Optimisation for Maintenance Systems,” *Comput. Ind. Eng.*, **82**, pp. 167–182.
- [5] Chong, A. K. W., Mohammed, A. H., Abdullah, M. N., and Rahman, M. S. A., 2019, “Maintenance Prioritization—a Review on Factors and Methods,” *J. Facil. Manage.*, **17**(1), pp. 18–39.
- [6] Saaty, T. L., 2008, “Decision Making with the Analytic Hierarchy Process,” *Int. J. Serv. Sci.*, **1**(1), pp. 83–98.
- [7] Sharma, S., and Sisodia, A., 2016, “Prioritization of Tools in Joint Production-Maintenance Environment of Auto Component Manufacturer Using AHP-Fuzzy-TOPSIS,” *Intell. Ind. Syst.*, **2**(1), pp. 73–84.
- [8] Khanlari, A., Mohammadi, K., and Sohrabi, B., 2008, “Prioritizing Equipments for Preventive Maintenance (PM) Activities Using Fuzzy Rules,” *Comput. Ind. Eng.*, **54**(2), pp. 169–184.
- [9] Chang, Q., Xiao, G., Biller, S., and Li, L., 2013, “Energy Saving Opportunity Analysis of Automotive Serial Production Systems,” *IEEE Trans. Autom. Sci. Eng.*, **10**(2), pp. 334–342.
- [10] Gu, X., Jin, X., and Ni, J., 2015, “Prediction of Passive Maintenance Opportunity Windows on Bottleneck Machines in Complex Manufacturing Systems,” *ASME J. Manuf. Sci. Eng.*, **137**(3), p. 031017.
- [11] Dekker, R., 1995, “Integrating Optimisation, Priority Setting, Planning and Combining of Maintenance Activities,” *Eur. J. Oper. Res.*, **82**(2), pp. 225–240.
- [12] Dekker, R., and Scarf, P. A., 1998, “On the Impact of Optimisation Models in Maintenance Decision Making: the State of the Art,” *Reliab. Eng. Syst. Saf.*, **60**(2), pp. 111–119.
- [13] Teng, S.-H. G., and Ho, S.-Y. M., 1996, “Failure Mode and Effects Analysis,” *Int. J. Qual. Reliab. Manage.*, **13**(5), pp. 8–26.
- [14] Ding, S.-H., Kamaruddin, S., and Azid, I. A., 2014, “Maintenance Policy Selection Model—a Case Study in the Palm Oil Industry,” *J. Manuf. Technol. Manage.*, **25**(3), pp. 415–435.

- [15] Yang, Z., Chang, Q., Djurdjanovic, D., Ni, J., and Lee, J., 2007, "Maintenance Priority Assignment Utilizing on-line Production Information," *ASME J. Manuf. Sci. Eng.*, **129**(2), pp. 435–446.
- [16] Birnbaum, Z. W., 1968, On The Importance of Different Components in A Multicomponent System. Technical Report, Washington University Seattle Lab of Statistical Research.
- [17] Nguyen, K.-A., Do, P., and Grall, A., 2014, "Condition-Based Maintenance for Multi-Component Systems Using Importance Measure and Predictive Information," *Int. J. Syst. Sci.: Oper. Logist.*, **1**(4), pp. 228–245.
- [18] Si, S., Liu, M., Jiang, Z., Jin, T., and Cai, Z., 2019, "System Reliability Allocation and Optimization Based on Generalized Birnbaum Importance Measure," *IEEE Trans. Reliab.*, **68**(3), pp. 831–843.
- [19] Hoffman, M., Song, E., Brundage, M., and Kumara, S., 2018, "Condition-Based Maintenance Policy Optimization Using Genetic Algorithms and Gaussian Markov Improvement Algorithm," Annual Conference of the PHM Society, Vol. 10, Philadelphia, PA, Prognostics and Health Management Society.
- [20] Sutton, R. S., and Barto, A. G., 2018, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA.
- [21] Kearns, M., Mansour, Y., and Ng, A. Y., 2002, "A Sparse Sampling Algorithm for Near-Optimal Planning in Large Markov Decision Processes," *Mach. Learn.*, **49**(2), pp. 193–208.
- [22] Kocsis, L., and Szepesvári, C., 2006, "Bandit Based Monte-Carlo Planning," European Conference on Machine Learning, Berlin, Germany, Springer, pp. 282–293.
- [23] Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., and Colton, S., 2012, "A Survey of Monte Carlo Tree Search Methods," *IEEE Trans. Comput. Intell. AI in Games*, **4**(1), pp. 1–43.
- [24] Hoffman, M., Song, E., Brundage, M., and Kumara, S., 2021, "Online Improvement of Condition-Based Maintenance Policy Via Monte Carlo Tree Search," *IEEE Trans. Autom. Sci. Eng.*, pp. 1–12.
- [25] Kolodner, J. L., 1992, "An Introduction to Case-Based Reasoning," *Artif. Intell. Rev.*, **6**(1), pp. 3–34.
- [26] Aamodt, A., and Plaza, E., 1994, "Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches," *AI Commun.*, **7**(1), pp. 39–59.
- [27] Bergmann, R., and Wilke, W., 1996, "On the Role of Abstraction in Case-Based Reasoning," European Workshop on Advances in Case-Based Reasoning, Lausanne, Switzerland, Springer, pp. 28–43.
- [28] Bengtsson, M., Olsson, E., Funk, P., and Jackson, M., 2004, "Technical Design of Condition Based Maintenance System-A Case Study Using Sound Analysis and Case-Based Reasoning," 8th International Conference of Maintenance and Reliability, Knoxville, TN.
- [29] Tsai, Y., 2009, "Applying a Case-Based Reasoning Method for Fault Diagnosis During Maintenance," *Proc. Inst. Mech. Eng., Part C: J. Mech. Eng. Sci.*, **223**(10), pp. 2431–2441.
- [30] Wan, S., Li, D., Gao, J., and Li, J., 2019, "A Knowledge Based Machine Tool Maintenance Planning System Using Case-Based Reasoning Techniques," *Rob. Comput.-Integr. Manuf.*, **58**, pp. 80–96.
- [31] Yu, R., Iung, B., and Panetto, H., 2003, "A Multi-Agents Based E-Maintenance System with Case-Based Reasoning Decision Support," *Eng. Appl. Artif. Intell.*, **16**(4), pp. 321–333.
- [32] Gabel, T., and Riedmiller, M., 2005, "CBR For State Value Function Approximation in Reinforcement Learning," International Conference on Case-Based Reasoning, Chicago, IL, Springer, pp. 206–221.
- [33] Varshavskii, P., and Eremeev, A., 2010, "Modeling of Case-Based Reasoning in Intelligent Decision Support Systems," *Sci. Techn. Inf. Process.*, **37**(5), pp. 336–345.
- [34] Li, J., and Meerkov, S. M., 2008, *Production Systems Engineering*, Springer Science & Business Media, New York.
- [35] Nguyen, Q., Valizadegan, H., and Hauskrecht, M., 2011, "Learning Classification With Auxiliary Probabilistic Information," 2011 IEEE 11th International Conference on Data Mining, Vancouver, Canada, IEEE, pp. 477–486.
- [36] Aha, D. W., 2013, *Lazy Learning*, Springer Science & Business Media, New York.
- [37] Kim, K.-K., Tacho, K., and Song, E., 2021, "Selection of The Most Probable Best Under Input Uncertainty," 2021 Winter Simulation Conference, Phoenix, AZ, IEEE.
- [38] Hoffman, M., 2021, Simantha, <https://github.com/m-hoff/simantha>
- [39] Brundage, M. P., Morris, K., Sexton, T., Mocozet, S., and Hoffman, M., 2018, "Developing Maintenance Key Performance Indicators From Maintenance Work Order Data," International Manufacturing Science and Engineering Conference, College Station, TX, Vol. 51371, American Society of Mechanical Engineers, p. V003T02A027.