

People-centric Evolutionary System for Dynamic Production Scheduling

Su Nguyen, *Member, IEEE* and Mengjie Zhang, *Fellow, IEEE* and Daminda Alahakoon,
and Kay Chen Tan, *Fellow, IEEE*

Abstract—Evolving production scheduling heuristics is a challenging task because of the dynamic and complex production environments and the interdependency of multiple scheduling decisions. Different genetic programming methods have been developed for this task and achieved very encouraging results. However, these methods usually have trouble in discovering powerful and compact heuristics, especially for difficult problems. Moreover, there is no systematic approach for the decision makers to intervene and embed their knowledge and preferences in the evolutionary process. This paper develops a novel people-centric evolutionary system for dynamic production scheduling. The two key components of the system are a new mapping technique to incrementally monitor the evolutionary process and a new adaptive surrogate model to improve the efficiency of genetic programming. The experimental results with dynamic flexible job shop scheduling show that the proposed system outperforms the existing algorithms for evolving scheduling heuristics in terms of scheduling performance and heuristic sizes. The new system also allows the decision makers to interact on the fly and guide the evolution towards desired solutions.

Index Terms—genetic programming, flexible job shop scheduling, diversity, visualisation.

I. INTRODUCTION

THERE have been an increasing number of studies on machine learning (ML) and artificial intelligence (AI) in the context of smart factories and Industry 4.0 [1]. Most of the studies emphasise the applications of ML and AI to advanced process control (APC) such as real-time equipment monitoring, fault detection and classification, quality control, and preventive maintenance. Scheduling is a critical function in APC to make sure that manufacturing resources are optimally utilised. However, the scheduling methods proposed for smart factories have been still based on traditional approaches (e.g. dispatching rules) with the support of computer simulation and real-time data [2]. There have been several studies on adaptive scheduling methods based on ML and online optimisation but these methods are restricted by many assumptions and not flexible enough to cope with the complexity of smart factories [3], [4]. Big data analytics solutions developed for scheduling purposes only provide real-time monitoring of orders and inventories, with only little emphases on making optimal planning and scheduling decisions.

Advances in computing power and optimisation have led to the developments of new solution methods for production scheduling. Meta-heuristics [5], [6], [7] is a popular approach to find near optimal solutions for a wide range of production scheduling problems. However, manually designing a good meta-heuristic algorithm is not a trivial task and it can be

very time consuming and requires a lot of problem domain knowledge. In addition, there is no guarantee that these algorithms may work in stochastic and dynamic production environments without major modifications. In recent years, genetic programming (GP) has been adopted to solve a wide range of production and supply chain management problems [8], [9]. Due to its flexibility and powerful search mechanisms, GP can be easily customised to deal with different planning and scheduling decisions in complex production and logistics systems, which cannot be easily handled by other ML and optimisation methods. One of the most popular applications of GP in this domain is the automated design of production scheduling heuristics [10]. Previous studies have shown that GP can successfully evolve superior scheduling heuristics compared with those manually designed by researchers and practitioners and other ML methods [11], [12], [13], [14]. Moreover, heuristics evolved by GP have some useful properties for practical applications such as simplicity, scalability, reactivity, and interpretability (in some cases), which are very difficult to achieve with conventional optimisation methods. As an evolutionary computation (EC) methods, GP can also take advantage of advanced techniques such as evolutionary multi-objective optimisation and surrogate models to improve its efficiency and to tackle real-world challenges. These properties make GP an attractive ML solution for dynamic and complex production systems such as smart factories.

Although the past experiments have demonstrated that GP can discover very powerful heuristics for different scheduling applications, searching for such heuristics is not trivial. When the production environments are complex and involve multiple planning and scheduling decisions [11], [13], [15], a large number of features (from jobs and machines) need to be taken into consideration, which results in a very large heuristic search space. Therefore, it is computationally expensive and difficult to find the most effective heuristics. To overcome this limitation, researchers have investigated and developed new representations [14], [16], surrogate models [17], [18], feature selection methods [19], and learning techniques [20], [21]. For instance, Hildebrandt and Branke [17] proposed a surrogate-assisted GP which relies on a simple surrogate model based on phenotypic characteristics of evolved scheduling heuristics for dynamic job shop scheduling. They showed that the surrogate model can improve GP convergence by efficiently generate new heuristics. Nguyen et al. [18] extended this surrogate model to deal with dynamic flexible job shop scheduling and made a modification to enhance population diversity. The experiments showed that surrogate-assisted GP is significantly

better than the simple GP method in terms of fitnesses. On the other hand, Mei et al. [19] proposed a new method to measure the importance of each feature and developed a two-phase learning algorithm for GP. They showed that it is possible to improve the effectiveness of GP if the optimal feature subset is selected. Some researchers focused on ensemble learning strategies [21] and they showed that combining multiple heuristics can improve scheduling decisions. Hart and Sim [20] developed a sophisticated algorithm based on ensemble learning to deal with a diverse set of problem instances and showed very promising results.

The above developments undoubtedly enhance the performance of evolved scheduling heuristics. However, there are key drawbacks with these existing methods. *First*, the existing methods cannot efficiently monitor what heuristics have been generated during the evolutionary process of GP. As a result, GP can waste its time examining bad heuristics or exploring the same areas of the search space without making significant improvements. In some extreme cases, GP can double the sizes of its heuristics just to achieve marginal improvements in terms of fitnesses. Given that the fitness evaluations can be very computationally expensive (e.g. using discrete-event simulation to measure heuristic performance), GP needs to be effective in generating potential heuristics when applying genetic operations.

Second, the sizes of evolved heuristics increase dramatically during the evolutionary process, especially with advanced GP methods mentioned above. As a consequence, it is very difficult to explain how the scheduling decisions are made. It is undesirable because interpretability is one of the attractive properties of GP. Although some techniques including simplification [22] and visualisation [14] have been proposed to gain more insights about the evolved heuristics, they are restricted to simple scheduling problems and a low number of features. Interpretable ML [23] and people-centric AI [24] are the current trend in the AI and ML research community to gain the values of these technologies by interpreting the representations learned and decisions made by these AI/ML models, and building new technologies to enhance human interaction with AI/ML.

Finally, GP evolution is currently a black-box optimisation process. In previous studies, researchers and decision makers can only provide some inputs (e.g. features, heuristic structures) and observe the final outputs. They have little or no knowledge about how or which heuristics are generated. Therefore, it is impossible to embed their insights and preferences to speed up the evolutionary process or to guide the search towards their most favourite solutions.

This paper aims to develop a new people-centric evolutionary system (PES) that can efficiently utilise the historical search data to adaptively and interactively guide the search towards effective and interpretable production heuristics. The main contributions of this paper are:

- 1) A new *mapping technique* that can efficiently capture the topological relations of evolved heuristics,
- 2) A new *adaptive surrogate model* that can integrate diversity control and bloat control to evolve powerful and compact scheduling heuristics,

- 3) A new *surrogate-assisted algorithm* that can evolve superior and compact production scheduling heuristics,
- 4) A new *human-computer interaction* component that allows decision makers to monitor and to guide the evolutionary process on the fly.

Different from existing approaches, the proposed system provides users and decision makers with a “*people-centric*” way to monitor the evolutionary process, control the evolutionary behaviours, integrate preferences, and generate compact solutions. To demonstrate the effectiveness of the new approach, the proposed PES algorithm is applied to evolve scheduling heuristics that simultaneously handle routing and dispatching decisions in dynamic flexible job shop scheduling problems. This problem is investigated as it is a good example of a complex production environment in which multiple scheduling decisions, uncertainty, and dynamic changes have to be taken into consideration. Extensive experiments using different simulated shop utilisations and flexibility are used to compare the performance and generalisation of PES and existing GP algorithms.

The rest of this paper is organised as follows. Section II describes the dynamic flexible job shop scheduling problems and scenarios to be investigated in this study. Section III presents the overview of the proposed algorithm and its key components. Section IV describes the experiment settings, and Section V presents the results in terms of test performance and sizes of evolved heuristic. Further analyses of the proposed algorithm are presented in Section VI. Finally, conclusions and future research directions are shown in Section VII.

II. DYNAMIC FLEXIBLE JOB SHOP SCHEDULING

Dynamic flexible job shop scheduling (DFJSS) is a variant of job shop scheduling problems in which a set of jobs \mathcal{J} must be allocated and processed by a set of machines \mathcal{M} . Each job $j \in \mathcal{J}$ has a number of predefined operations $O_j = \{o_{j1}, o_{j2}, \dots, o_{jN_j}\}$. Different from the traditional job shop scheduling problem, an operation o_{ji} can be processed by any machine in the subset $\mathcal{M}_{ji} \subset \mathcal{M}$. The processing time of the operation o_{ji} is $p(o_{ji}, m)$ if it is routed to a machine $m \in \mathcal{M}_{ji}$. Since routing decisions, i.e. specify which machine to process an operation, can greatly influence the scheduling performance, they need to be considered and optimised along with the sequencing decisions.

The flexible job shop in this study is similar to the simulation model used in [5] and [13], but we focus on the dynamic scheduling problems instead of static problem instances. The simulated shop will have $|\mathcal{M}| = 10$ machines and new jobs will arrive randomly over time. Each job j will have a number of operations N_j following a discrete uniform distribution $\mathcal{U}\{1, |\mathcal{M}|\}$, and each operation o_{ji} can be processed by a set of alternative machines \mathcal{M}_{ji} . The size of \mathcal{M}_{ji} is proportional to the flexibility $f\%$ [13]. If the flexibility $f\%$ is high, an operation is likely to be handled by more machines. In the simulated shop, if the flexibility $f\% = 50\%$, the size of \mathcal{M}_{ji} will follow the discrete uniform distribution $\mathcal{U}\{1, 5\}$ as $f \times |\mathcal{M}| = 5$. All processing times $p(o_{ji}, m)$ will follow an exponential distribution $Exp(1/\mu)$ with the mean processing time $1/\mu = 1$.

Similar to previous simulation studies of dynamic job shops, job arrivals will follow Poisson process with the arrival rate λ [25], [18] calculated as follow:

$$\lambda = \rho \frac{\mu |\mathcal{M}|}{\bar{n}} \quad (1)$$

where ρ is the utilisation of machines and \bar{n} is the average number of operations per job which is $\frac{1}{2}(1 + |\mathcal{M}|)$.

The due date d_j a new job j will be determined by the total work content (TWK) rule [26], i.e. $d_j = r_j + h \times \sum_{o_{ji} \in O_j} p_{avg}(o_{ji})$ where h is the allowance or tightness factor [5], [27] and $p_{avg}(o_{ji})$ is the average processing time of operation o_{ji} (across all alternative machines). For the weights w_j of jobs, the 4:2:1 rule is employed [5], [27].

The simulation starts with an empty shop and new jobs will arrive randomly over time. Upon their arrivals, new jobs will be immediately released to the shop. If the machine to process the current operation o_{ji} has not yet been specified, the routing rule is applied to calculate the priority $\mathcal{P}_r(m, j)$ for each machine $m \in \mathcal{M}_{ji}$ based on the information of the job j and the status of machine m . The machine m^* with the highest $\mathcal{P}_r(m, j)$ will process the operation o_{ji} . Similarly, the dispatching rule will be used to determine the priority $\mathcal{P}_d(j, \mathcal{M})$ for each job $j \in \mathcal{Q}(m)$ based on the attributes of the job j and the current status of related machines in the shop when the machine m become idle and its queue $\mathcal{Q}(m)$ is not empty. The objective of the dynamic FJSS problem here is to minimise the total weighted tardiness:

$$TWT = \sum_{j \in \mathbb{C}} w_j T_j \quad (2)$$

where $T_j = (d_j - C_j)$ is the tardiness, and C_j is the completion time of job j . In this equation, \mathbb{C} is the set of completed jobs recorded during the simulation. To measure the long-term performance of evolved scheduling heuristics, we set the warm-up time as the arrival time of the 500th job and the statistics from the next completed 5000 jobs will be recorded to calculate the objective value (i.e. $|\mathbb{C}| = 5000$).

III. PEOPLE-CENTRIC EVOLUTIONARY SYSTEM

Fig. 1 shows how the proposed PES works. The initial population $\mathbb{H} = \{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_N\}$ of scheduling heuristics are randomly generated. Then all heuristics in the population will be evaluated by using the simulator presented in Section II. The phenotypic characteristics $phenotype(\mathcal{H})$, the corresponding fitnesses $fitness(\mathcal{H})$, and the sizes of evaluated heuristics (i.e. number of nodes) are saved in an archive \mathcal{A} . Principal component analysis (PCA) is applied to transform the dataset X of phenotypic characteristics (with the dimension of D stored in \mathcal{A}) to X' with a lower dimensionality (by using the first K principal components where $K \ll D$). Then, the transformed dataset X' is fed into a modified version of growing neural gas (GNG) [28]. The network $\mathcal{N} = (V, E)$ obtained with the modified GNG (mGNG) contains a set of nodes V and a set of undirected edges E that represent the distributions of generated heuristics on the explored search space. After the network is learned, the user can optionally modify the network either by changing node positions or

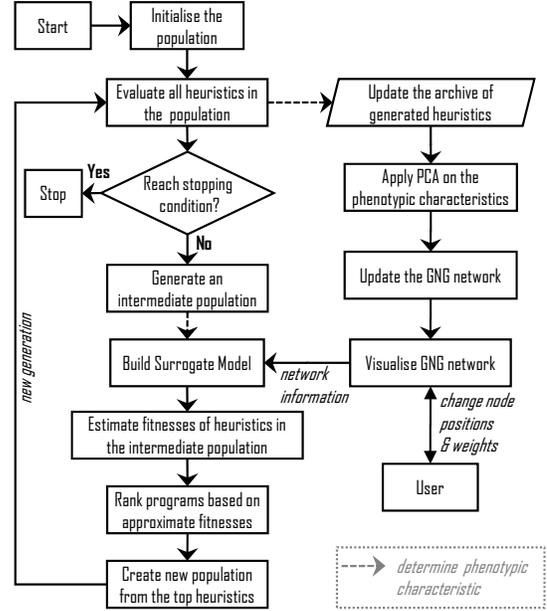


Fig. 1. People-centric evolutionary system.

selection pressures. PES will build a surrogate model based on the historical search data and the information from the obtained mGNG network. An intermediate population is then generated by applying genetic operations. Newly generated heuristics will be efficiently estimated by using the surrogate model with inputs from the phenotypic characteristics and sizes of those heuristics, archive \mathcal{A} , the network \mathcal{N} , and user preferences. The population for the next generation is the collection of the top heuristics in the intermediate population ranked by the estimated fitnesses. In the rest of this section, we will provide the details of each key component in this algorithm including the GP representation of scheduling heuristics, genetic operations, the mGNG algorithm, and the surrogate model.

A. Representation

The proposed PES utilises both genotype characteristics (heuristic sizes) and phenotype characteristics (rankings of jobs) of evolved heuristics to improve its efficiency.

1) *Genotype*: The standard GP tree representation is used for routing rules \mathcal{P}_r and dispatching rules \mathcal{P}_d . The terminal sets and functions sets to construct the two rules are presented in Table I and Table II [18]. In these tables, t is the decision moment or the current simulation time, mr_j is the time job j joins the queue of machine m , and o_{jk} is its current operation. The function $workload(Q(m))$ is the total processing time of all jobs waiting in the queue of machine m . These attributes are selected because they have been employed in previous studies on automated design of scheduling heuristics [29], [13], [17], and most attributes are used in existing rules in the literature. For the routing rules, the last two attributes are used as the rough estimate of the job priority when it is assigned to a specific machine.

2) *Phenotype*: This study used the phenotypic characteristics proposed by [17] to capture the behaviours of dispatching

TABLE I
ATTRIBUTES/FUNCTIONS OF ROUTING RULES $\mathcal{P}_r(j, m)$

Notation	Description
WLN	workload of machine $m = workload(Q(m))$
NON	number of operations waiting at $m = Q(m) $
rPT	ratio between workload of operations in $Q(m)$ with processing times larger than $p(o_{ji}, m)$.
$rSLACK$	ratio between workload of operations in $Q(m)$ with slack larger than the slack of job j
Function set	$+, -, \times$, protected division $\%$, min, max

TABLE II
ATTRIBUTES/FUNCTIONS OF DISPATCHING RULES $\mathcal{P}_d(j, \mathcal{M})$

Notation	Description
rrJ	time-in-system of job $(t - r_j)$
rRJ	job queuing time $(t - mr_j)$
RO	number of remaining operations $ O_j^r $.
RT	work remaining of the job $p(o_{jk}, m) + \sum_{o_{ji} \in O_j^r} p_{avg}(o_{ji})$
PT	operation processing time $p(o_{jk}, m)$
rDD	time to due date $= d_j - t$
SL	slack of the job $= d_j - (t + RT)$
W	weight of job w_j
$\#$	Random number from 0 to 1
NPT	average processing time of the next operation $p_{avg}(o_{ik+1})$
$WINQ$	average workload in the next queue $\frac{1}{ \mathcal{M}_{ik+1} } \sum_{m' \in \mathcal{M}_{ik+1}} workload(Q(m'))$
Function set	$+, -, \times$, protected division $\%$, min, max

rules. The phenotype of an evolved rule is a decision vector that captures how the rule ranks jobs in D specific decision situations (each situation has a set of predefined jobs). For each situation, the ranks of jobs based on a reference rule (smaller ranks for jobs with higher priorities) is first determined. Then, the rank (based on the reference rule) of the job with the highest priority given by an evolved rule for each decision situation is the decision value for that situation in the decision vector. For any pair of rules, Euclidean distance between their decision vectors is used to compute their similarity. As the routing rules evolved by PES are also priority functions, the same technique can be applied to determine the phenotypic characteristics of routing rules. The phenotype of a scheduling heuristic is simply a concatenation of the phenotypes of the two rules (routing and sequencing). An example of phenotypes for the considered scheduling heuristic can be found in [17] and [30].

B. Genetic operations

In the proposed algorithm, subtree crossover, subtree mutation [31], and subtree extraction [32] are used to produce new rules for the scheduling heuristic $\mathcal{H} = \{\mathcal{P}_r, \mathcal{P}_d\}$. To generate new heuristics, individuals in the population are randomly selected as parents. When a genetic operation is applied, it only focuses on one specific rule (routing or sequencing) to avoid generating random heuristics (i.e. too different from parents). The subtree crossover creates new rules by randomly recombining subtrees from two selected parent rules. The subtree mutation is performed by selecting a node of a chosen rule and replacing the subtree rooted by that node with a newly

randomly-generated subtree. Meanwhile, the subtree extraction simply replaces a rule with its random subtree so the size of the new rule is always smaller than the parent rule. To improve the diversity of the population, all the same duplicated heuristics (based on phenotypic similarity) will be eliminated before applying the proposed surrogate model [17].

C. Mapping evolutionary process

To efficiently explore competitive scheduling heuristics, it is crucial for PES to capture the topological relations and distributions of previously generated heuristics. Since the heuristics generated by genetic operations are complicated and usually contain redundant components, phenotypic characteristics are more useful than the genotypic characteristics to determine the similarity between generated heuristics. Also, phenotypic characteristics is a good predictor of program fitnesses [17], [18] and the phenotypic diversity has good correlations with the fitness improvements [33]. Therefore, we employ the phenotypic characteristics to map the evolutionary process of PES. There are two main steps in mapping evolutionary process with PES: (1) dimensionality reduction with PCA, and (2) topological learning with mGNG.

1) *Dimensionality reduction*: Directly using the high-dimensional dataset X to map the evolutionary process is not efficient and can greatly influence the scalability and the accuracy of the proposed algorithm. To overcome these challenges, a dimensionality reduction technique is needed to transform the input data in the high-dimensional space to a space of fewer dimensions. There are many dimensionality reduction methods designed to cope with different types of data. For the application investigated in this paper, PCA is used for dimensionality reduction because it is an efficient and well-established approach. Moreover, PCA is an appropriate method to deal with the multicollinearity issue in the decision vectors. It is noted that decision situations used to determine the decision vector are randomly generated. Thus, there is a good chance that similar situations (similar set of jobs) are created, which can lead to similar decision values when a specific rule is evaluated.

2) *Modified growing neural gas*: Algorithm 1 shows how mGNG can adapt the network to the input data. All data points in $X = \{phenotype(\mathcal{H}) | \mathcal{H} \in \mathcal{A}\}$ are transformed into X' by PCA. In each epoch, an input x is then sampled from the transformed data X' and mGNG determines, in terms of Euclidean distance, the nearest unit (s_1) and the second nearest unit (s_2) which are adapted towards the input data. An edge connecting these two units is also created (if it has not yet existed) and its age is set to zero. After each epoch, the algorithm will identify the unit with the maximum error and systematically insert a new unit nearby to reduce the total error of the network. Meanwhile, a unit can be removed if its utility is too low as compared to the maximum error. This mechanism is proposed in [28] to deal with non-stationary distributions. In PES, removing nodes with low utilities has two main advantages: (1) reducing the running times of mGNG, and (2) discarding noisy nodes (i.e. representing heuristics which are rarely generated in recent generations). The parameter θ is used to control the resolution of networks. If θ is high, nodes

Algorithm 1 Modified growing neural gas (mGNG)

Input: dataset X , current network \mathcal{N} (if exist)

Output: (updated) mGNG network \mathcal{N} , PCA model

- 1: $X' \leftarrow$ transform dataset X with PCA
 - 2: $epoch \leftarrow 0$
 - 3: randomly initialise \mathcal{N} with two nodes if \mathcal{N} is empty
 - 4: **repeat**
 - 5: randomly shuffle X'
 - 6: **for** each data point $x \in X'$ **do**
 - 7: $w_{s_1} \leftarrow$ the nearest (winning) node s_1 for input x
 - 8: update the error of s_1 : $E_{s_1} = E_{s_1} + \|x - w_{s_1}\|^2$
 - 9: updating unit s_1 : $w_{s_1} = w_{s_1} + \epsilon_b(x - w_{s_1})$
 - 10: updating neighbour n of s_1 : $w_n = w_n + \epsilon_n(x - w_n)$
 - 11: create an edge between s_1 and s_2 (the second-nearest unit to x) and set its age to zero if the edge does not exist; otherwise set the edge's age to zero
 - 12: increment the age of edges connecting s_1 and each neighbor $n \in \mathcal{N}$
 - 13: removing edges with age larger than $a_{max} = |X|$ and nodes with no emanating edges
 - 14: updating the utility of s_1 : $U_{s_1} = U_{s_1} + \|x - w_{s_2}\|^2 - \|x - w_{s_1}\|^2$
 - 15: let q is the unit with maximum error and l is the unit with minimum utility
 - 16: remove the unit l if $E_q/U_l > \theta$
 - 17: **for** each unit $c \in \mathcal{N}$ **do**
 - 18: $E_c = (1 - \beta)E_c$
 - 19: $U_c = (1 - \beta)U_c$
 - 20: insert a new unit to \mathcal{N} between q and its neighbor if the number of nodes in $\mathcal{N} < max_node$
 - 21: **until** $epoch = maximum_epoch$ or mGNG converges
-

will be unlikely to be removed and mGNG will try to fit the network to the whole dataset. If θ is low, mGNG will focus on more representative programs (frequently reproduced) and main evolutionary traits. The maximum age of edges is set to the size of the dataset to preserve most relations of programs. We also adopt a scheme like *k-means* to adapt the learning rate over time [34], i.e. $e_b = 1/n_{win}$ and $e_n = 100/n_{win}$ (see steps 9–10 in Algorithm 1) where n_{win} is the number of input signals for which the considered node has been a winner.

The mGNG network is updated after the whole population is evaluated and before new heuristics are produced. The obtained network \mathcal{N} captures distributions of generated heuristics and allows PES to monitor how GP explores the heuristic search space. mGNG has a number of advantages as compared to existing mapping techniques such as self-organizing map (SOM) [35] or MAP-Elites [36]. First, because of its growing nature, mGNG does not require any prior knowledge the input data, which is important for MAP-Elites or SOM to predefine the map size. Second, mGNG produces a network which allows PES to efficiently determine the topological relations of generated programs. Finally, mGNG can update the map (i.e. network) efficient during the evolutionary process, which makes mGNG a suitable approach to mapping the explored search space. In the next section, we show how the mGNG

network can be used as the input to help GP adaptively and effectively explore the search space.

D. Adaptive surrogate model

When a new heuristic $\mathcal{H} = (\mathcal{P}_r, \mathcal{P}_d)$ is generated by genetic operations, its corresponding phenotypic characteristic $phenotype(\mathcal{H}) = (phenotype(\mathcal{P}_r) \| phenotype(\mathcal{P}_d))$ is determined by using the technique discussed in section III-A2. In the previous study [17], the fitness was estimated by:

$$\hat{f}(\mathcal{H}) = fitness \left(\arg \min_{\mathcal{H}' \in \mathcal{A}} \| phenotype(\mathcal{H}) - phenotype(\mathcal{H}') \| \right) \quad (3)$$

In the proposed algorithm, we generalise this surrogate model by adding a *control factor* based on the information from the network \mathcal{N} . The new model will estimate the fitness as:

$$\hat{f}_a(\mathcal{H}) = \hat{f}(\mathcal{H}) \times control_factor(\mathcal{H}, \mathcal{N}) \quad (4)$$

in which:

$$control_factor(\mathcal{H}, \mathcal{N}) = df(\mathcal{H}, \mathcal{N}) \times \frac{bf(\mathcal{H}, \mathcal{N})}{pf(\mathcal{H}, \mathcal{N})} \quad (5)$$

where $df(\mathcal{H}, \mathcal{N})$, $bf(\mathcal{H}, \mathcal{N})$, and $pf(\mathcal{H}, \mathcal{N})$ are three key factors to control the behaviours of PES in terms of diversity, sizes of evolved heuristics, and user preferences. If $control_factor(\mathcal{H}, \mathcal{N}) = 1$, the proposed model will be the same as the surrogate model in [17] because $\hat{f}_a(\mathcal{H}) = \hat{f}(\mathcal{H})$. Heuristics with a higher $control_factor(\mathcal{H}, \mathcal{N})$ will be less likely to be selected.

1) *Diversity control*: PES uses a simple strategy to control the diversity of evolved heuristics, i.e. reduce the chance to select heuristics located in a well-explored area of the search space. Since the distribution of evolved heuristics has been captured by mGNG, we can easily determine the density of a certain area in the search space and use it to adjust the estimate fitness. Given a newly generated heuristic \mathcal{H} , a diversity factor is determined:

$$df(\mathcal{H}, \mathcal{N}) = \begin{cases} \frac{nd(\mathcal{H}, \mathcal{N})}{T} & nd(\mathcal{H}, \mathcal{N}) > T \\ 1 & otherwise \end{cases} \quad (6)$$

where $nd(\mathcal{H}, \mathcal{N})$ is the *neighbourhood density* of the heuristic \mathcal{H} given the network \mathcal{N} and T is a predefined diversity threshold ($0 < T \leq 1$). To calculate $nd(\mathcal{H}, \mathcal{N})$, $phenotype(\mathcal{H})$ must be transformed to lower dimensional space by the PCA model obtained from the Algorithm 1, and the node $n_{\mathcal{H}}^{\mathcal{N}*}$ in \mathcal{N} with the minimum distance to the transformed $phenotype(\mathcal{H})$ is determined. Then, the neighbourhood density is calculated as:

$$nd(\mathcal{H}, \mathcal{N}) = \frac{1}{|\mathcal{A}|} \left(freq(n_{\mathcal{H}}^{\mathcal{N}*}) + \sum_{n \in NB(n_{\mathcal{H}}^{\mathcal{N}*})} freq(n) \right) \quad (7)$$

where $freq(n)$ is the matching frequency of node $n \in \mathcal{N}$ determined by the number of entries in \mathcal{A} that best match node n , and $NB(n_{\mathcal{H}}^{\mathcal{N}*})$ is the set of neighbours of $n_{\mathcal{H}}^{\mathcal{N}*}$, i.e. basically any nodes in \mathcal{N} that are directly connected to $n_{\mathcal{H}}^{\mathcal{N}*}$ by an edge. A higher neighbourhood density indicates that the area covering the neighbourhood of the considered heuristic \mathcal{H} has been explored more frequently.

From equations (4)–(6), it is clear that the new surrogate model gives a heuristic \mathcal{H} a worse fitness (higher value) if it is located in the area with the neighbourhood density greater than T . This mechanism is created to discourage GP from visiting too crowded areas. If T is high, PES will focus on areas where good heuristics have been identified. If T is extremely small, PES will distribute its computational efforts to explore all areas in the search space without relying on the fitnesses. In other words, PES tends to explore more aggressively with a smaller T and the exploitation level of GP will be enhanced if a high T is used. In the proposed algorithm, T provides a convenient way for GP to adapt its search during the evolutionary process rather than predefining a fixed rule (e.g. with a fixed tournament size for tournament selection). In the early generations, T will only have a small influence on the approximate fitnesses because heuristics are randomly generated, i.e. the density of each node in mGNG network will be similarly low. In this case, PES will favour new heuristics with higher $\hat{f}(\cdot)$, i.e. more exploitation. However, when no improved heuristic is found in a few generation, the neighbourhood density of the areas where top evolved heuristics are located will increase due to the selection mechanism that favours high performing heuristics (see Fig. 1). When the neighbourhood density eventually rises above the threshold level T , PES will favour heuristics that have slightly worse fitnesses but are located in the less explored areas. Depending on the evolutionary progress, PES will adaptively adjust its exploration and exploitation intensity.

2) *Bloat control*: The most common approach to control bloat issue in GP is to predefine the maximum length, the maximum depth, or the penalty parameter to control sizes of evolved programs. However, without prior knowledge about the problem, it is difficult to determine good values for these parameters. To control bloat in PES, we further take advantage of the obtained network \mathcal{N} . Given a newly generated heuristic \mathcal{H} , a bloat factor is determined:

$$bf(\mathcal{H}, \mathcal{N}) = \begin{cases} 1 + \left(\frac{size(\mathcal{H})}{ms(\mathcal{H}, \mathcal{N})} - 1 \right)^B & \frac{size(\mathcal{H})}{ms(\mathcal{H}, \mathcal{N})} > 1 \\ 1 & otherwise \end{cases} \quad (8)$$

where $size(\mathcal{H})$ is the size of heuristic \mathcal{H} (total number of terminals and functions) and $ms(\mathcal{H}, \mathcal{N})$ is the maximum size of heuristics $\mathcal{H}' \in \mathcal{A}$ that best matched node $n_{\mathcal{H}}^{\mathcal{N}*}$ (defined in the previous section). If $size(\mathcal{H}) \leq ms(\mathcal{H}, \mathcal{N})$, the bloat factor will have no impact on the estimated fitness of \mathcal{H} . Otherwise, the parameter $B \geq 1$ governs the tradeoffs between the sizes of the evolved heuristics and their quality. The bloat factor curves with different B values are shown in Fig. 2. If B is higher, PES will put a higher priority on the fitnesses (a lower bloat factor) and a lower priority on the heuristic sizes. If B is low, PES will focus on generating compact rules. Similar to diversity control, this bloat control is also adaptive to the evolutionary progress, i.e. larger heuristics are only preferred when they have better fitnesses and there is no similar heuristic previously generated. It should be noted that B will have a reverse impact on the bloat factor is $size(\mathcal{H})/ms(\mathcal{H}, \mathcal{N}) > 2$, i.e. a higher B will control the sizes of newly generated heuristics more aggressively if their sizes $size(\mathcal{H})$ are double the maximum size of heuristics $ms(\mathcal{H}, \mathcal{N})$. This property

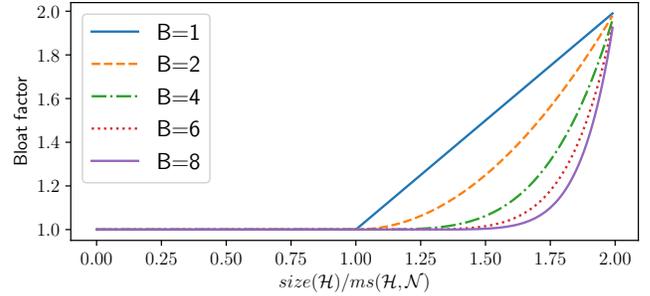


Fig. 2. Bloat factor curves.

helps PES avoid generating heuristics with too many redundant components in one generation which are likely to be included in the parent heuristics if a higher B value is used.

3) *User preferences*: PES allows the users and decision makers to interact with the evolutionary process by selecting the areas in the search space that they want PES to explore. This property is useful if there has been already an effective heuristic applied in the production system and the decision makers only want to make some improvements without dramatically changing the current scheduling practice. Another useful case is when the decision makers want to find a good tradeoff between program sizes and complexity on the fly rather than using a predefined objective function. With the obtained network \mathcal{N} , the decision makers can easily guide the search of PES simply by selecting a node or a set of nodes, corresponding to specific areas in the search space, to further explore. When the network \mathcal{N} is generated or updated, the default preference for each node n is $pref(n) = 1$. The decision makers can assign a particular preference value to each node within the range $[1, maxp]$. Given a newly generated heuristic \mathcal{H} , a preference factor is determined:

$$pf(\mathcal{H}, \mathcal{N}) = 1 + \frac{pref(n_{\mathcal{H}}^{\mathcal{N}*}) - 1}{maxp - 1} \quad (9)$$

From equation (5) and equation (9), it is clear that PES will select heuristics located near to nodes with high preference factors because of high preference values $pref(n_{\mathcal{H}}^{\mathcal{N}*})$. The users can select and increase the preference values of one or multiple nodes in \mathcal{N} and higher preference values should be assigned to more desirable nodes (e.g. similar to certain rules or explainable rules). The users may choose to interact with PES at the end of each generation or set up a periodic review depending on the speed of the evolutionary process.

In general, PES is developed to overcome a number of key challenges in existing GP algorithms. First, PES can efficiently capture the distributions of evolved heuristics across generations with the mGNG network. This network can provide useful aggregate information about the evolutionary progress. Second, PES provides a convenient way to adaptively control the diversity based on the evolutionary progress rather than simply applying random genetic operations, which is inefficient and sensitive to predefined parameters such as tournament size, crossover probability, and mutation probability. Third, PES can adaptively control the sizes of newly generated heuristics based on the sizes of their similar heuristics evolved in the previous generations instead of explicitly fixing the maximum

size of evolved heuristics. Finally, PES has a visualisation and interaction component that allows users to monitor and guide the evolutionary process. These properties help PES discover scheduling heuristics more efficiently, especially for highly complex environments with multiple decisions such as DFJSS.

IV. EXPERIMENT SETTINGS

This section describes the simulation configurations, parameter settings, and performance metrics used to evaluate the performance of the proposed PES. In our experiments, PES will be compared to a simple genetic programming (referred to as GP) algorithm [37] and the surrogate-assisted GP (referred to as SGP) [17].

A. Simulation configurations

Six different simulation scenarios based on two utilisation levels ($\rho = 85\%$ and 95%) and three levels of flexibility ($f\% = 30\%$, 50% , and 70%) are used to evaluate the performance of PES, GP, and SGP. During the training process, the fitnesses (i.e. total weighted tardiness in equation(2)) of evolved heuristics in a specific generation are calculated based on an independent simulation replication to improve the efficiency and effectiveness of the three algorithms [37], [38]. For testing, 50 independent simulation replications are used to evaluate the performance of the best heuristics evolved by the three algorithms. Common random number [25] is used as the variance reduction technique in our experiments for both DFJSS simulation and population initialisation. To reduce the computational cost, the simulation will terminate early if the number of jobs in the shop exceeds 500, i.e. the shop is overloaded [37].

B. Parameter settings

The default parameters for the three algorithms are shown in Table III. All algorithms use the same terminal set and function set in Table I and Table II. The maximum number of nodes in mGNG is 500. The utility factor θ is 50 (to ensure that representative programs generated during the evolution are captured) and the maximum number of epochs to learn the mGNG networks is 1000 or when the algorithm is converged, i.e. the total error is not improved in 5 epochs. These parameters are selected based on our preliminary experiments. It is noted that the mGNG algorithm is very efficient and does not show significant impacts on the running times of PES. As compared to the simulation time used to evaluate heuristics, the mapping time is negligible. The preference parameter $maxp$ for interactions is 4. For SGP and PES, the phenotype of each rule in an evolved heuristic is determined by 100 decision situations ($D = 2 \times 100$) and with a reference heuristic using *least waiting time* (LWT) [39] as the routing rule and $2PT + WINQ + NPT$ [29] as the dispatching rule.

C. Performance metrics

For comparisons, 30 independent runs of GP, SGP, and PES are performed and their best evolved heuristics are recorded. The Wilcoxon test with $\alpha = 0.05$ is used for statistical significance tests of the *test performance* and *sizes* of evolved heuristics. The test performance of an evolved heuristic is the

TABLE III
PARAMETER SETTINGS

Parameter	GP	SGP	PES
Population size		200	
Initialisation		ramp-half-and-half [31]	
Tournament size		5	
Maximum generation		50	
Maximum depth	17	17	–
Crossover rate	80%	80%	–
Mutation rate	15%	20%	–
Reproduction rate	5%	–	–
# of decision situations D	–	200	200
Size of intermediate population	–	2000	2000
# of principal components K	–	–	5
Diversity threshold T	–	–	0.05
Bloat control parameter B	–	–	1, 8

average TWT from 50 independent simulation replications. For these two performance metrics, the lower the better.

We also examine the behaviours of the three algorithms in terms of the *training fitnesses*, the *number of components* in the network \mathcal{N} , and the *number of nodes* in \mathcal{N} . It should be noted that the mGNG networks \mathcal{N} are only used to capture the evolutionary patterns of GP and SGP and do not influence their search mechanisms. The number of connected components shows the connectivity of evolved heuristics across generations. The connectivity can help explain the behaviours of each algorithm. Meanwhile, the number of nodes represents the diversity of evolved heuristics. If an algorithm converges to a certain area in the search space, the number of nodes will be low. If an algorithm continues to explore different areas in the search space, the number of nodes will be high.

V. COMPUTATIONAL RESULTS

This section presents the experimental results for PES, GP, and SGP. First, these three algorithms are compared in terms of the test performance and heuristic sizes. Then, the test results on unseen scenarios are provided to investigate their generalisation. Finally, we investigate the behaviours of the algorithms using visualisation and illustrate how users can intervene in the evolutionary process.

A. Test performance

The test performances of the three algorithms are shown in Fig. 3. PES-B1 and PES-B8 represent the PES algorithm using the default parameters in Table III with $B = 1$ and $B = 8$ respectively. For each PES variant, a notation $l|r$ is used to show the outputs of statistical tests. l and r are the algorithms, i.e. \diamond for GP and \star for SGP, which are significantly worse (or better) than a PES variant. In all scenarios, PES variants are significantly better than GP. Meanwhile, PES variants are either significantly better than or not significantly different from SGP in all scenarios. PES-B8 has a slightly lower variance as compared to PES-B1 and SGP. In general, PES shows consistent and superior performance in a wide range of scenarios with different levels of flexibility and utilisation.

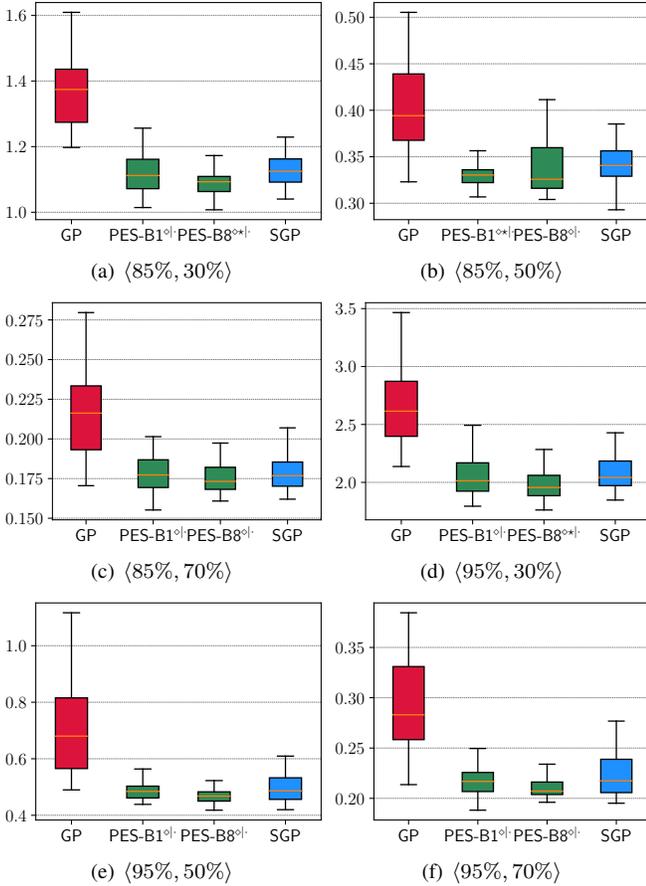


Fig. 3. Test performance (average TWT of evolved heuristics).

B. Program size

Fig. 4 shows the sizes of the best heuristics evolved by PES, GP, and SGP. PES-B1 obviously produces the most compact heuristics in all scenarios. Heuristics evolved by PES-B1 are roughly twice smaller than GP and three times smaller than SGP. In addition, the variances of heuristic sizes for PES-B1 are also much smaller than those for other algorithms. Given that PES-B1 is better than SGP in terms of the test performance, it is safe to conclude that PES-B1 is a superior algorithm to evolve heuristics for DFJSS. Although PES-B8 is slightly better than PES-B1 in terms of the test performance, PES-B8 evolves much larger heuristics than PES-B1.

C. Generalisation

To examine the generalisation of the three algorithms, we show the detailed results for unseen scenarios in Table IV. The results of PES are **bold** if they are significantly better than GP and *italic* if they are better than SGP (***bold-italic*** if PES is significantly better than both GP and SGP). In general, PES and SGP are better than GP in the scenarios used for training and are able to maintain their dominance in unseen scenarios. These results showed that PES, especially the PES-B8 variant, and SGP are not overfitted to the training scenarios. It means that loosely controlling bloat with a high B does not have a negative impact on the generalisation of PES.

In terms of training scenarios, heuristics evolved with $\langle 85\%, 30\% \rangle$ are most effective across scenarios. The main

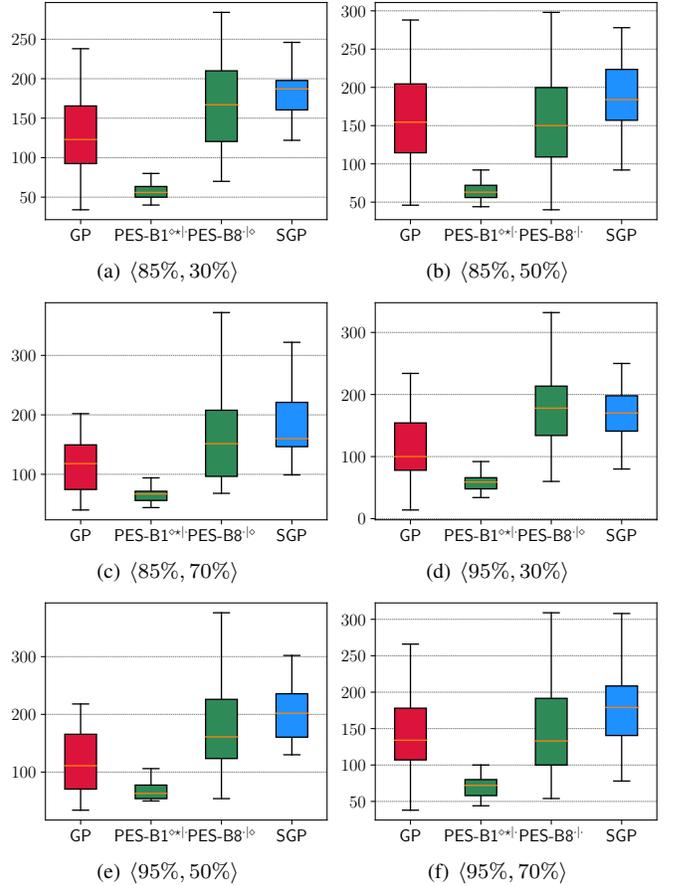


Fig. 4. Sizes of evolved heuristics.

reason is that a low flexibility and a low utilisation force the algorithms to evolve dispatching rules and routing routes that assigned accurate priorities to effectively utilise machines. For a high flexibility, an operation will have many options and it is easy to find a reasonably good routing rule. Similarly, for a high utilisation, simple time-based dispatching rules such as shortest processing time (SPT) and 2PT+WINQ+NPT perform quite well and their variances can also easily be discovered. Therefore, heuristics evolved with a high utilisation or a high flexibility may not generalise well to unseen scenarios.

D. Visualisation

With the mNGG, PES has a built-in ability to visualise the evolutionary process. Fig. 5 shows the visualisation of PES with $B = 8$ and $K = 2$, GP, and SGP through generations of a particular run for the scenario $\langle 95\%, 30\% \rangle$. It is noted that a higher K value can be used but may lead to a complex network which is hard to understand. Each node in a network visualised in Fig. 5 is a node in the obtained mNGG network at the end of each generation. The position of a node is determined using the weights of that node. The node colour (from red to green) and node size represent the average fitness and the number of all heuristics $\mathcal{H} \in \mathcal{A}$ that are nearest to that node. Red (green) nodes represent the area with bad (good) heuristics.

After initialisation, all algorithms have the same population (the same network), in which heuristics are scattered randomly. After a few generations, we start to see the difference in

TABLE IV
TEST PERFORMANCE OF EVOLVED SCHEDULING HEURISTICS (RESULTS IN THE FORM OF MEDIAN [MINIMUM, MAXIMUM])

Test Scenario	Trained with ⟨85%, 30%⟩				Trained with ⟨95%, 30%⟩			
	GP	PES-B1	PES-B8	SGP	GP	PES-B1	PES-B8	SGP
⟨0.85%, 30%⟩	1.37 [1.20,1.78]	1.11 [1.01,1.26]	1.09 [1.01,1.21]	1.12 [1.04,1.32]	1.51 [1.24,2.33]	1.16 [1.03,1.43]	1.13 [1.03,1.28]	1.19 [1.06,1.38]
⟨0.95%, 30%⟩	2.45 [2.17,3.32]	2.02 [1.79,2.35]	2.01 [1.75,2.19]	2.04 [1.86,2.53]	2.61 [2.14,3.47]	2.01 [1.79,2.49]	1.96 [1.76,2.33]	2.04 [1.85,2.43]
⟨0.85%, 50%⟩	0.46 [0.37,0.67]	0.34 [0.30,0.40]	0.33 [0.30,0.42]	0.35 [0.31,0.46]	0.58 [0.40,1.22]	0.39 [0.32,0.49]	0.37 [0.32,0.47]	0.41 [0.32,0.53]
⟨0.95%, 50%⟩	0.75 [0.54,1.18]	0.50 [0.42,0.62]	0.48 [0.42,0.63]	0.51 [0.45,0.73]	0.90 [0.60,1.55]	0.57 [0.46,0.77]	0.52 [0.44,0.70]	0.61 [0.46,0.81]
⟨0.85%, 70%⟩	0.27 [0.20,0.41]	0.20 [0.16,0.25]	0.19 [0.17,0.26]	0.20 [0.17,0.27]	0.36 [0.23,0.89]	0.22 [0.18,0.28]	0.22 [0.17,0.29]	0.24 [0.18,0.32]
⟨0.95%, 70%⟩	0.38 [0.27,0.65]	0.25 [0.19,0.31]	0.23 [0.19,0.34]	0.25 [0.21,0.38]	0.51 [0.30,1.04]	0.29 [0.22,0.40]	0.28 [0.21,0.40]	0.32 [0.22,0.46]
Test Scenario	Trained with ⟨85%, 50%⟩				Trained with ⟨95%, 50%⟩			
	GP	PES-B1	PES-B8	SGP	GP	PES-B1	PES-B8	SGP
⟨0.85%, 30%⟩	1.43 [1.16,2.66]	1.26 [1.11,1.47]	1.28 [1.05,1.57]	1.27 [1.11,1.62]	1.44 [1.23,1.81]	1.19 [1.09,1.59]	1.17 [1.04,1.30]	1.22 [1.06,1.37]
⟨0.95%, 30%⟩	2.98 [2.25,8.29]	2.58 [2.16,3.49]	2.60 [1.98,3.54]	2.67 [2.14,3.46]	2.81 [2.43,3.81]	2.34 [2.06,3.20]	2.26 [1.84,2.70]	2.33 [2.05,3.06]
⟨0.85%, 50%⟩	0.39 [0.32,1.00]	0.33 [0.30,0.38]	0.33 [0.30,0.41]	0.34 [0.29,0.43]	0.43 [0.34,0.65]	0.33 [0.31,0.42]	0.32 [0.30,0.37]	0.33 [0.28,0.44]
⟨0.95%, 50%⟩	0.63 [0.49,2.78]	0.50 [0.45,0.60]	0.51 [0.45,0.66]	0.53 [0.43,0.71]	0.68 [0.49,1.12]	0.48 [0.44,0.64]	0.47 [0.42,0.57]	0.49 [0.42,0.65]
⟨0.85%, 70%⟩	0.21 [0.18,0.58]	0.18 [0.15,0.21]	0.18 [0.16,0.22]	0.18 [0.16,0.23]	0.24 [0.18,0.38]	0.18 [0.17,0.22]	0.18 [0.16,0.20]	0.18 [0.16,0.25]
⟨0.95%, 70%⟩	0.28 [0.22,1.41]	0.22 [0.19,0.26]	0.22 [0.20,0.29]	0.23 [0.20,0.32]	0.32 [0.22,0.60]	0.22 [0.20,0.29]	0.22 [0.19,0.26]	0.22 [0.18,0.33]
Test Scenario	Trained with ⟨85%, 70%⟩				Trained with ⟨95%, 70%⟩			
	GP	PES-B1	PES-B8	SGP	GP	PES-B1	PES-B8	SGP
⟨0.85%, 30%⟩	1.64 [1.30,2.63]	1.43 [1.10,1.77]	1.38 [1.19,1.79]	1.42 [1.17,1.87]	1.52 [1.26,2.22]	1.35 [1.17,1.65]	1.30 [1.19,1.80]	1.30 [1.14,1.59]
⟨0.95%, 30%⟩	3.59 [2.65,6.28]	3.16 [2.12,4.23]	3.06 [2.30,4.33]	3.18 [2.19,5.06]	3.07 [2.50,5.64]	2.79 [2.19,5.82]	2.81 [2.40,4.09]	2.65 [2.16,3.55]
⟨0.85%, 50%⟩	0.43 [0.33,0.69]	0.35 [0.30,0.44]	0.33 [0.31,0.41]	0.35 [0.31,0.43]	0.40 [0.32,0.60]	0.33 [0.30,0.41]	0.32 [0.30,0.44]	0.33 [0.29,0.39]
⟨0.95%, 50%⟩	0.73 [0.52,1.32]	0.57 [0.44,0.74]	0.54 [0.46,0.72]	0.57 [0.47,0.78]	0.67 [0.46,1.04]	0.52 [0.44,0.66]	0.50 [0.45,0.73]	0.50 [0.45,0.65]
⟨0.85%, 70%⟩	0.22 [0.17,0.33]	0.18 [0.16,0.20]	0.17 [0.16,0.24]	0.18 [0.16,0.21]	0.21 [0.17,0.34]	0.17 [0.15,0.22]	0.17 [0.16,0.21]	0.18 [0.16,0.20]
⟨0.95%, 70%⟩	0.30 [0.21,0.54]	0.23 [0.19,0.29]	0.22 [0.19,0.32]	0.23 [0.20,0.30]	0.28 [0.21,0.52]	0.22 [0.19,0.29]	0.21 [0.20,0.27]	0.22 [0.20,0.28]

their evolutionary processes. At generation 10, while GP still randomly explores the search space (small nodes scattered across the network), PES and SGP have a focus on a certain area and the diversity in PES is slightly higher than that of SGP. At generation 20, GP starts to narrow down its search but it evolves a mix of both good and bad heuristics. PES also narrows its search but in a more efficient way, i.e. focusing on good heuristics (a green area on the left of the network). SGP does not change significantly. At generation 30, GP shows a tendency to converge to the left area. SGP still focus on one area but slightly expand the search. PES, on the other hand, starts to form two clusters in the network. At generation 40, GP and SGP converge to a particular in their networks but PES converges to two distinct areas represented by the two large clusters. This visualisation suggests that SGP and PES are much more efficient than GP as they can identify and evolve good heuristics rather than random heuristics. PES, with the ability to maintain its diversity and efficiently explore the search space, shows a better adaptive ability than SGP which highly restricts its exploration and loses its diversity quickly in later generations (the number of nodes is reduced).

E. Human-guided evolutionary process

To illustrate how a decision maker can intervene the evolutionary process, we conduct a small experiment in which the decision maker can interact with the mGNG network and decide which nodes or areas that PES should explore. To make the visualisation easier to interpret, we use four benchmark scheduling heuristics: (1) FIFO-LWT (first-in-first-out for the dispatching rule and least-waiting time for the routing rule), (2) WSPT-LWT (weighted shortest processing time for the

dispatching rule and LWT), (3) EDD-LWT (earliest due date for the dispatching rule and LWT), and (4) SL-LWT (minimum slack for the dispatching rule and LWT). Their phenotypic characteristics are obtained to identify their position in the network (gray crosses). In this case, we assume that the production system is currently employing FIFO-LWT because of its simplicity, and the decision maker wants to make some improvements without significantly changing the scheduling behaviours. To achieve this goal, the decision maker will interact with the mGNG network at the end of each generation and select the nodes nearest to FIFO-LWT and adjust their preference values (see Section III-D3) for PES to explore.

Fig. 6 shows the outcomes of the interaction. With the inputs of the decision makers, PES narrows its search to areas close to FIFO-LWT. In the caption of each sub figure, we also show the best-so-far scheduling heuristics. In the initial generation, the best heuristic is quite random. At generation 3, after two rounds of interaction, the best heuristic has changed to $\mathcal{P}_d = W$ and $\mathcal{P}_r = (rPT - WLN)$. This heuristic is actually a variant of FIFO-LWT. With $\mathcal{P}_d = W$, jobs are basically prioritised as FIFO if they have the same weight. Similarly, $\mathcal{P}_r = (rPT - WLN)$ will behave like LWT if the term rPT is removed. Having rPT will help the routing rule break the tie when the considered machines have the same workload. At generation 6, a better heuristic is identified. The new heuristic has the same routing rule as one found at generation 3 but its dispatching rule is a variant of WSPT. It is reasonable as PES has explored a new good area between FIFO-LWT and WSPT-LWT. This experiment demonstrates that it is possible to include decision makers' inputs on the fly by using the proposed PES. This property is very useful in the case that

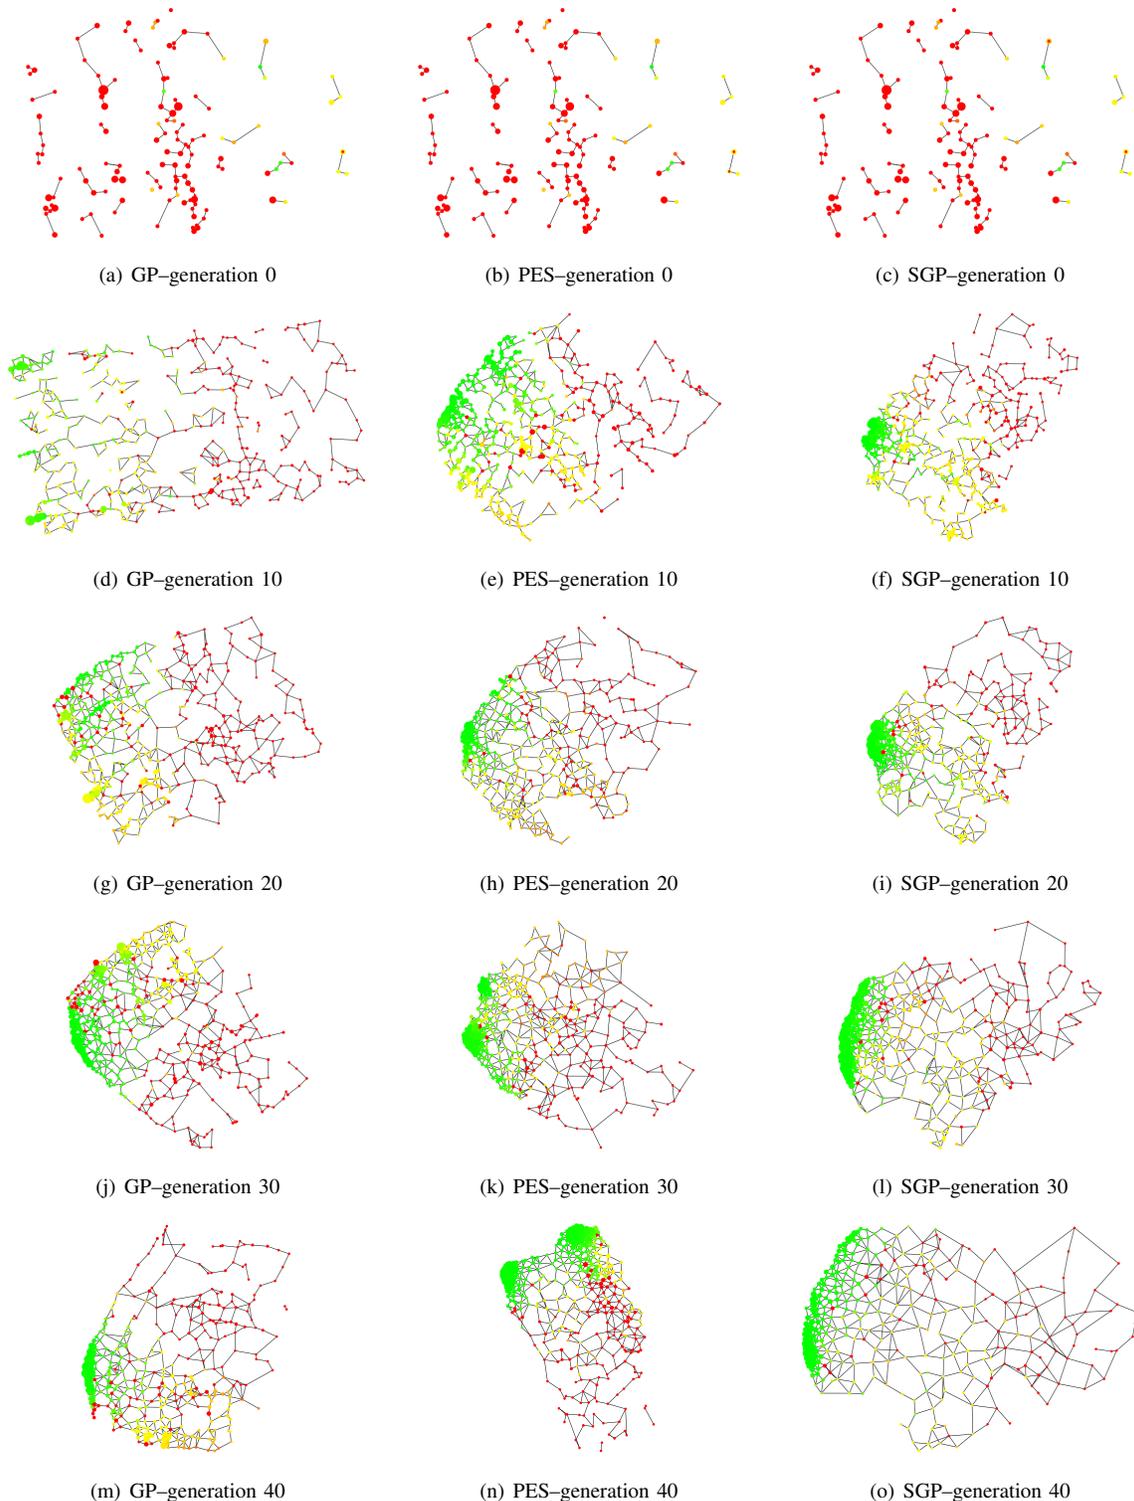


Fig. 5. Visualisation of the evolutionary process. The node colour and node size show the average fitness and the density for each corresponding node. The network rotation through generation indicates a genetic drift, i.e. the changes in the distributions of generated heuristics.

the user or decision maker wants to navigate the evolutionary process to find the solutions of interest without hardcoding the search algorithm or the fitness function.

VI. FURTHER ANALYSES

The previous section has shown effectiveness and capability of PES. In this section, we further analyse the key parameters

of PES to understand how they influence the evolutionary process of PES. Because of the space limitation, we only present the results for the scenario $\langle 95\%, 30\% \rangle$. Similar patterns are also observed for other scenarios.

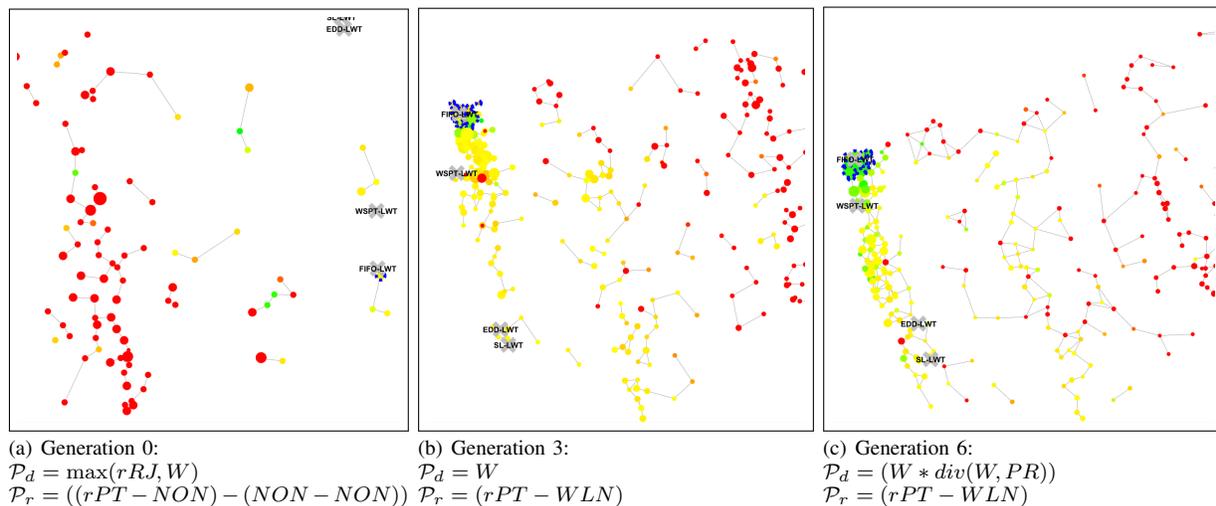


Fig. 6. Visualisation of the evolutionary process. For better presentation, only the areas with high fitnesses in the mGNG network are shown. The gray crosses represent benchmark rules. The blue boundaries indicates that the nodes are preferred. $\text{div}(\cdot, \cdot)$ is the protected division.

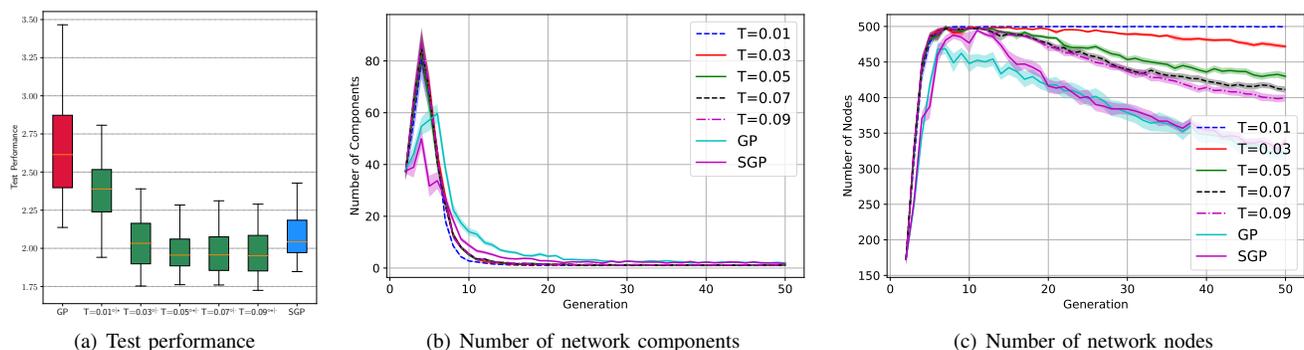


Fig. 7. Influence of diversity control.

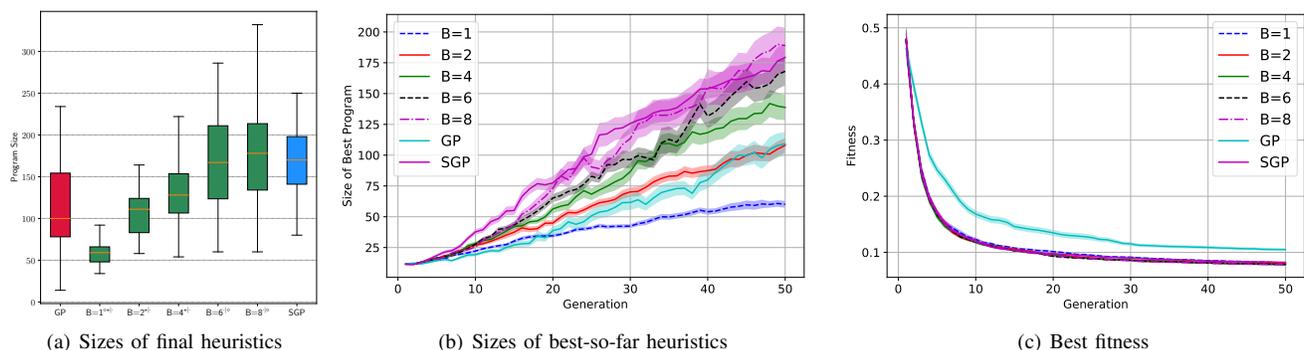


Fig. 8. Influence of bloat control.

A. Influence of diversity control

The threshold T is the key parameter of PES to control the diversity of PES across generations. Fig. 7(a) shows the test performance of evolved heuristics when T increases from 0.01 to 0.09. It is easy to see that setting T too low has a negative effect on the test performance as PES tries too hard to maintain the diversity rather than boosting the quality of evolved heuristics. However, with the lowest $T = 0.01$, PES is still better than GP. The T values from 0.05 to 0.09 seems to be consistent. As SGP is a special case of PES with $T = 1$ (without bloat control and user preferences), it is reasonable

to say that using a high T value is also undesirable.

Fig. 7(b) and Fig. 7(c) show the number of connected components and the number of nodes in the obtained mGNG network. It is easy to see that PES has a different behaviour as compared to GP and SGP. A common observation is that the numbers of connected components increase quickly in the first few generations and then converge to 1 (i.e. each node has at least one connection). However, GP and SGP are slower than PES in terms of the increasing speed and the convergence speed. There are no significant differences in terms of the number of connected components when different T values are

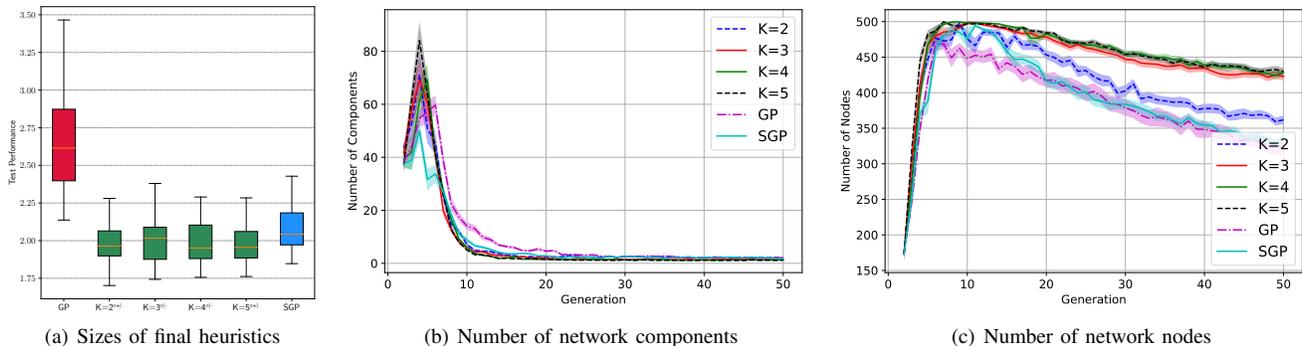


Fig. 9. Influence of dimensionality reduction.

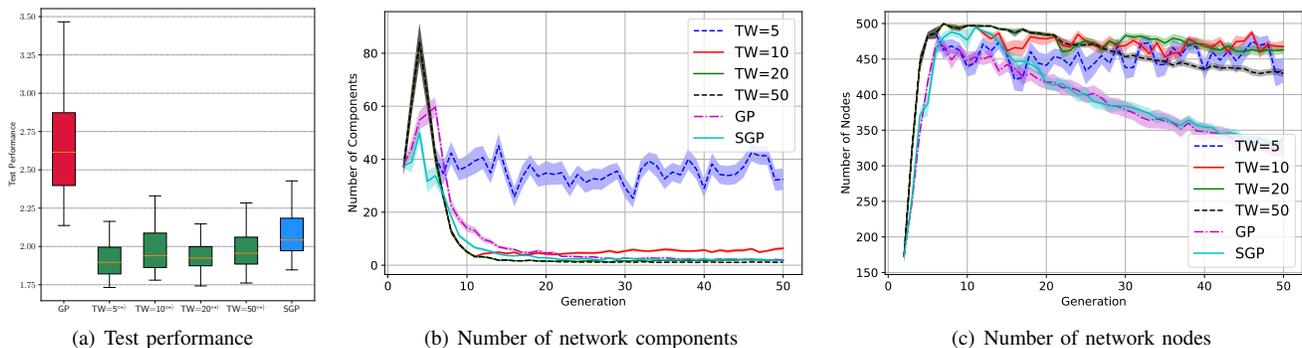


Fig. 10. Influence of time window.

applied. The differences in terms of the number of nodes are much more obvious. All algorithms search unexplored areas very quickly and reach the maximum number of nodes (500) in the first few generations. Then most algorithms except for PES with $T = 0.01$ starts to lose their diversity. As expected, a low T helps PES maintain a high diversity across generations. GP and SGP lose their diversity very quickly as compared to PES. These results confirm that T is a very useful and predictable parameter to control the population diversity.

B. Influence of bloat control

The ability to produce compact heuristic is an attractive property of PES. Fig. 8(a) shows that the sizes of evolved heuristics increase very quickly when B increases from 1 to 8. When $B \geq 4$, PES behaves like SGP. This can be further observed in Fig. 8(b). Clearly, PES with $B = 1$ is best at control the heuristic sizes when the average sizes of its best heuristics only increase from 10 to more than 50. Meanwhile, the sizes of final heuristics evolved by PES with $B = 8$ and SGP are about 20 times larger than their initial heuristics. PES with $B = 2$ is similar to GP in terms of heuristic sizes. It should be noted that the best fitnesses of all algorithms except for GP are not significantly different as shown in Fig. 8(c) although the heuristic sizes are very different.

C. Influence of dimensionality reduction

As dimensionality reduction plays an important role to improve the efficiency of mGNG, it is important to see if they have a negative impact on PES. Fig. 9(a) show the

test performance of PES when different numbers of principal components K are applied. The results show that there is no significant difference in terms of the test performance when different K values are used. K also does not influence the number of connected components as shown in Fig. 9(b). However, the diversity seems to be reduced when a smaller K is used as presented in Fig. 9(c). One explanation for this phenomenon is that a lower K will reduce the discrimination ability when mGNG tries to learn the topological relations of evolved heuristics. Therefore, PES cannot effectively see the difference in terms of phenotypic characteristics of evolved heuristics to accurately determine their estimate fitnesses.

D. Influence of time window

The archive \mathcal{A} is needed for mGNG to map the explore areas in the search space. However, keeping all evolved heuristics in \mathcal{A} may reduce the efficiency of mGNG and PES. In this section, we will examine the recency of the archive by restricting the number of recorded heuristics only to the most recent generation, i.e. time window TW . Fig. 10(a) shows the test performance when different TW s are used (e.g. $TW = 5$ means that only heuristics evolved in the last 5 generations are kept in \mathcal{A}). The results show that PES is not sensitive to TW . In Fig. 10(b), it is interesting to see that the number of components are very high for $TW = 5$ and this number is significantly reduced when $TW = 10$. The reason is that the number of evolved heuristics in \mathcal{A} is not large enough for mGNG to form a well connected network. Fig. 10(c) shows that the number of nodes with $TW = 5$ is still higher than those of GP and SGP although their variances are much higher

(because there are significant changes in \mathcal{A} in each generation). These results suggest that TW can be significantly reduced without influencing the performance of PES.

VII. CONCLUSIONS

This paper presents a new evolutionary system for learning powerful and compact scheduling heuristics for a complex production environment. The proposed system is a systematic combination of conventional genetic operations, surrogate-assisted modelling, topological data analysis, and visualisation. The core principle of PES is to monitor the evolutionary progress by incrementally learning topological relations of evolved heuristics and using that knowledge to guide the evolution. The experiments have shown that PES can outperform the state-of-the-art SGP in terms of heuristic performance. Moreover, PES has overcome a major limitation of SGP by successfully controlling the sizes of evolved heuristics without deteriorating its performance. On average, heuristics evolved by PES are twice smaller than the simple GP and three times smaller than SGP. Smaller evolved heuristics will make it much easier to analyse and interpret how scheduling decisions are made. Mapping the evolutionary progress also provides PES with the visualisation and interaction abilities that allow the users and decision makers to integrate their inputs on the fly. The extensive analyses presented in this paper also demonstrate that the key parameters are intuitive and can be used to control many aspects of the evolution.

Future studies will focus on improving the efficiency of PES by investigating other machine learning methods for topological data analysis such as growing self-organizing maps and self-organizing incremental neural networks. As finding the right parameters for PES for an arbitrary problem is not straightforward, an important task is to reduce the number of parameters and make the algorithm more adaptive. It would be interesting to also examine how PES can be used for multi-objective optimisation and other machine learning tasks such as classification, symbolic regression, and computer vision.

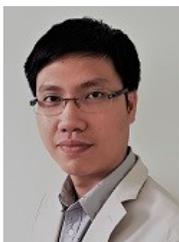
ACKNOWLEDGMENT

This work was partially supported by David Myers Research Fellowship from La Trobe University, the Marsden Fund of New Zealand Government (VUW1509), the National Natural Science Foundation of China (Grant No. 61876162/F060604), and the Shenzhen Scientific Research and Development Funding Program (Grant No. JCYJ20180307123637294)

REFERENCES

- [1] J. Moyné and J. Iskandar, "Big data analytics for smart manufacturing: Case studies in semiconductor manufacturing," *Processes*, vol. 5, no. 3, 2017.
- [2] D. Ivanov, A. Dolgui, B. Sokolov, F. Werner, and M. Ivanova, "A dynamic model and an algorithm for short-term supply chain scheduling in the smart factory industry 4.0," *International Journal of Production Research*, vol. 54, no. 2, pp. 386–402, Jan. 2016.
- [3] Y.-C. Wang and J. M. Usher, "Application of reinforcement learning for agent-based production scheduling," *Engineering Applications of Artificial Intelligence*, vol. 18, no. 1, pp. 73–82, Feb. 2005.
- [4] H. Mao, M. Alizadeh, I. Menache, and S. Kandula, "Resource Management with Deep Reinforcement Learning," in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, ser. HotNets '16. New York, NY, USA: ACM, 2016, pp. 50–56.
- [5] X. N. Shen and X. Yao, "Mathematical modeling and multi-objective evolutionary algorithms applied to dynamic flexible job shop scheduling problems," *Information Sciences*, vol. 298, pp. 198–224, Mar. 2015.
- [6] J. Li, Q. Pan, and P. Duan, "An improved artificial bee colony algorithm for solving hybrid flexible flowshop with dynamic operation skipping," *IEEE Transactions on Cybernetics*, vol. 46, no. 6, pp. 1311–1324, 2016.
- [7] Y. Han, D. Gong, Y. Jin, and Q. Pan, "Evolutionary multiobjective blocking lot-streaming flow shop scheduling with machine breakdowns," *IEEE Transactions on Cybernetics*, vol. 49, no. 1, pp. 184–197, 2019.
- [8] S. Nguyen, D. Thiruvady, A. Ernst, and D. Alahakoon, "Genetic programming approach to learning multi-pass heuristics for resource constrained job scheduling," in *Proceedings of the Genetic and Evolutionary Computation Conference*, ser. GECCO '18. New York, NY, USA: ACM, 2018, pp. 1167–1174.
- [9] S. Chand, Q. Huynh, H. Singh, T. Ray, and M. Wagner, "On the use of genetic programming to evolve priority rules for resource constrained project scheduling problems," *Information Sciences*, vol. 432, pp. 146–163, 2018.
- [10] J. Branke, S. Nguyen, C. W. Pickardt, and M. Zhang, "Automated design of production scheduling heuristics: A review," *IEEE Transactions on Evolutionary Computation*, vol. 20, no. 1, pp. 110–124, Feb. 2016.
- [11] C. W. Pickardt, T. Hildebrandt, J. Branke, J. Heger, and B. Scholz-Reiter, "Evolutionary generation of dispatching rule sets for complex dynamic scheduling problems," *International Journal of Production Economics*, vol. 145, no. 1, pp. 67–77, 2013.
- [12] S. Nguyen, M. Zhang, M. Johnston, and K. C. Tan, "Dynamic multi-objective job shop scheduling: A genetic programming approach," in *Automated Scheduling and Planning*. Springer Berlin Heidelberg, 2013, no. 505, pp. 251–282.
- [13] L. Nie, L. Gao, P. Li, and X. Li, "A GEP-based policies constructing approach for dynamic flexible job shop scheduling problem with job release dates," *Journal of Intelligent Manufacturing*, vol. 24, no. 4, pp. 763–774, Aug. 2013.
- [14] J. Branke, T. Hildebrandt, and B. Scholz-Reiter, "Hyper-heuristic evolution of dispatching rules: A comparison of rule representations," *Evolutionary Computation*, vol. 23, no. 2, pp. 249–277, Jun. 2015.
- [15] S. Nguyen, M. Zhang, M. Johnston, and K. C. Tan, "Automatic design of scheduling policies for dynamic multi-objective job shop scheduling via cooperative coevolution genetic programming," *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 2, pp. 193–208, Apr. 2014.
- [16] M. Durasevic, D. Jakobovic, and K. Knezevic, "Adaptive scheduling on unrelated machines with genetic programming," *Applied Soft Computing*, vol. 48, pp. 419–430, Nov. 2016.
- [17] T. Hildebrandt and J. Branke, "On using surrogates with genetic programming," *Evolutionary Computation*, vol. 23, no. 3, pp. 343–367, Sep. 2015.
- [18] S. Nguyen, M. Zhang, and K. C. Tan, "Adaptive charting genetic programming for dynamic flexible job shop scheduling," in *Proceedings of the Genetic and Evolutionary Computation Conference*, ser. GECCO '18. New York, NY, USA: ACM, 2018, pp. 1159–1166.
- [19] Y. Mei, S. Nguyen, B. Xue, and M. Zhang, "An efficient feature selection algorithm for evolving job shop scheduling rules with genetic programming," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 1, no. 5, pp. 339–353, Oct. 2017.
- [20] E. Hart and K. Sim, "A hyper-heuristic ensemble method for static job-shop scheduling," *Evolutionary Computation*, vol. 24, no. 4, pp. 609–635, Dec. 2016.
- [21] M. Durasević and D. Jakobović, "Comparison of ensemble learning methods for creating ensembles of dispatching rules for the unrelated machines environment," *Genetic Programming and Evolvable Machines*, vol. 19, no. 1, pp. 53–92, Jun. 2018.
- [22] S. Nguyen, M. Zhang, and K. C. Tan, "Surrogate-assisted genetic programming with simplified models for automated design of dispatching rules," *IEEE Transactions on Cybernetics*, vol. 47, no. 9, pp. 2951–2965, Sep. 2017.
- [23] F. Doshi-Velez and B. Kim, "A roadmap for a rigorous science of interpretability," *CoRR*, vol. abs/1702.08608, 2017. [Online]. Available: <http://arxiv.org/abs/1702.08608>
- [24] Google AI, "People + AI research," in *People + AI Research Initiative Symposium*, Sep. 2017. [Online]. Available: <https://ai.google/research/teams/brain/pair>
- [25] A. M. Law and D. M. Kelton, *Simulation Modeling and Analysis*. McGraw-Hill Higher Education, 1999.
- [26] S. Kreipl, "A large step random walk for minimizing total weighted tardiness in a job shop," *Journal of Scheduling*, vol. 3, pp. 125–138, May 2000.

- [27] S. Nguyen, M. Zhang, M. Johnston, and K. C. Tan, "A computational study of representations in genetic programming to evolve dispatching rules for the job shop scheduling problem," *IEEE Transactions on Evolutionary Computation*, vol. 17, no. 5, pp. 621–639, Oct. 2013.
- [28] B. Fritzke, "A self-organizing network that can follow non-stationary distributions," in *Proceedings of International Conference on Artificial Neural Networks: ICANN'97*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1997, pp. 613–618.
- [29] V. Sels, N. Gheysen, and M. Vanhoucke, "A comparison of priority rules for the job shop scheduling problem under different flow time- and tardiness-related objective functions," *International Journal of Production Research*, vol. 50, no. 15, pp. 4255–4270, Sep. 2011.
- [30] S. Nguyen, M. Zhang, D. Alahakoon, and K. C. Tan, "Visualizing the evolution of computer programs for genetic programming [research frontier]," *IEEE Computational Intelligence Magazine*, vol. 13, no. 4, pp. 77–94, Nov 2018.
- [31] J. R. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press, 1992.
- [32] S. Nguyen, M. Zhang, M. Johnston, and K. Tan, "Automatic Programming via Iterated Local Search for Dynamic Job Shop Scheduling," *IEEE Transactions on Cybernetics*, vol. 45, no. 1, pp. 1–14, 2015.
- [33] E. K. Burke, S. Gustafson, and G. Kendall, "Diversity in genetic programming: An analysis of measures and correlation with fitness," *IEEE Transactions on Evolutionary Computation*, vol. 8, no. 1, pp. 47–62, Feb. 2004.
- [34] S. Furao and O. Hasegawa, "An incremental network for on-line unsupervised classification and topology learning," *Neural Networks*, vol. 19, no. 1, pp. 90–106, 2006.
- [35] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological Cybernetics*, vol. 43, no. 1, pp. 59–69, Jan. 1982.
- [36] J. Mouret and J. Clune, "Illuminating search spaces by mapping elites," *CoRR*, vol. abs/1504.04909, 2015.
- [37] T. Hildebrandt, J. Heger, and B. Scholz-Reiter, "Towards improved dispatching rules for complex shop floor scenarios – A genetic programming approach," in *Proceedings of 2010 Genetic and Evolutionary Computation Conference*, ser. GECCO'10. Portland, Oregon, USA: ACM Press, 2010, pp. 257–264.
- [38] S. Nguyen, M. Zhang, M. Johnston, and K. C. Tan, "Selection Schemes in Surrogate-Assisted Genetic Programming for Job Shop Scheduling," in *SEAL'14: Simulated Evolution and Learning*. Springer International Publishing, 2014, pp. 656–667.
- [39] J. C. Tay and N. B. Ho, "Evolving dispatching rules using genetic programming for solving multi-objective flexible job-shop problems," *Computers & Industrial Engineering*, vol. 54, no. 3, pp. 453–473, Apr. 2008.



Su Nguyen (M'13) received the Ph.D. degree in artificial intelligence and data analytics from the Victoria University of Wellington, New Zealand, in 2013. He is currently a David Myers Research Fellow with the Centre for Research in Data Analytics and Cognition, La Trobe University, Australia. His primary research interests include computational intelligence, optimization, data analytics, large-scale simulation, and their applications in operations management and social media.



Mengjie Zhang (M04-SM10-F19) received the B.E. and M.E. degrees from Artificial Intelligence Research Center, Agricultural University of Hebei, Hebei, China, and the Ph.D. degree in computer science from RMIT University, Melbourne, VIC, Australia, in 1989, 1992, and 2000, respectively. He is currently Professor of Computer Science, Head of the Evolutionary Computation Research Group, and the Associate Dean (Research and Innovation) in the Faculty of Engineering. His current research interests include evolutionary computation with application areas of image analysis, multi-objective optimization, feature selection and reduction, job shop scheduling, and transfer learning. He has published over 500 research papers in refereed international journals and conferences. Prof. Zhang is a Fellow of Royal Society of New Zealand and has been a Panel member of the Marsden Fund (New Zealand Government Funding), a Fellow of IEEE, and a member of ACM.

He was the chair of the IEEE CIS Intelligent Systems and Applications Technical Committee, and chair for the IEEE CIS Emergent Technologies Technical Committee and the Evolutionary Computation Technical Committee, and a member of the IEEE CIS Award Committee. He is a vice-chair of the IEEE CIS Task Force on Evolutionary Feature Selection and Construction, a vice-chair of the Task Force on Evolutionary Computer Vision and Image Processing, and the founding chair of the IEEE Computational Intelligence Chapter in New Zealand. He is also a committee member of the IEEE NZ Central Section.



Daminda Alahakoon received the Ph.D. degree from Monash University, Australia, in 2002. He is currently a Professor of business analytics with the La Trobe University Business School, Melbourne, Australia, and the Director of the Research Centre for Data Analytics and Cognition. He has published over 100 research articles in data mining, clustering, neural networks, machine learning, and cognitive systems. He received the Monash Artificial Intelligence Prize from Monash University.



Kay Chen Tan (SM08-F14) is a full Professor with the Department of Computer Science, City University of Hong Kong. He is the Editor-in-Chief of *IEEE Transactions on Evolutionary Computation*, was the EiC of *IEEE Computational Intelligence Magazine* (2010-2013), and currently serves on the Editorial Board member of 10+ journals. He is an elected member of IEEE CIS AdCom (2017-2019). He has published 300+ refereed articles and 6 books. He is a Fellow of IEEE.