A Competitive Co-Evolutionary Approach for the Nurse Scheduling Problem

Maizatul Farhana Mohamad Nazri[1], Zeratul Izzah Mohd Yusoh[2], Halizah Basiron[3], Azlina Daud[4]
Faculty of Information and Communication Technology, Universiti Teknikal Malaysia, Melaka, Malaysia[1]
Faculty of Artificial Intelligence and Cyber Security, Universiti Teknikal Malaysia, Melaka, Malaysia[2, 3]
Kulliyyah of Nursing, International Islamic University Malaysia, Pahang, Malaysia[4]

Abstract—**The Nurse Scheduling Problem (NSP) is a constrained combinatorial optimisation problem that plays a critical role in healthcare scheduling and constraint optimisation. Traditional evolutionary approaches often rely on static fitness evaluation, which struggles to balance feasibility and solution quality under complex real-world constraints. This study proposes a competitive co-evolutionary algorithm for the NSP that introduces adaptive adversarial evaluation, where candidate schedules are assessed under dynamic competitive pressure to expose structural weaknesses and guide evolution more effectively. The proposed competitive NSP is evaluated on a 20-nurse, one-week scheduling instance and compared against a classical Genetic Algorithm (GA) under identical conditions for 30 independent runs. Experimental results show that the competitive NSP achieves a mean best penalty of 447.28, compared to 651.30 for the classical GA, corresponding to an average improvement of approximately 31%. The competitive approach further exhibits smoother convergence behaviour across generations, indicating stronger optimisation dynamics and improved robustness. These findings demonstrate that competitive co-evolution provides an effective and practical alternative to static fitness-based evolutionary methods for nurse scheduling, with broader applicability to healthcare scheduling and constraint optimisation problems.**

Keywords*—*Nurse Scheduling Problem; competitive co-evolution; evolutionary algorithms; healthcare scheduling; constraint optimisation; adversarial evaluation

I. INTRODUCTION

Hospitals worldwide face increasing pressure to deliver high-quality healthcare services under growing demand, workforce shortages, and strict regulatory requirements. Population growth and demographic aging have intensified healthcare utilisation, particularly in inpatient and emergency settings, where nurses play a central role in daily clinical operations [1], [2]. Effective nurse scheduling is therefore essential for ensuring adequate staff coverage, maintaining care quality, and controlling operational costs, yet remains challenging due to the need to balance legal regulations, skill requirements, workload distribution, and individual preferences.

Within this context, the Nurse Scheduling Problem (NSP) is widely recognised as a highly constrained optimisation problem involving multiple conflicting objectives. Hard constraints, such as shift coverage, legal rest periods, and role-based rules, must be strictly satisfied to ensure feasibility, while soft constraints related to fairness, preferences, and workload distribution influence schedule quality and staff satisfaction [3],[4]. Studies have shown that poor balance between these objectives can lead to increased absenteeism, reduced job satisfaction, and higher operational costs, highlighting the importance of robust scheduling approaches [3]. As a result, NSP has been extensively studied over the past decades.

Consequently, numerous optimisation techniques have been proposed to address the NSP. Comprehensive survey and comparative studies consistently report difficulties when realistic constraints are considered, particularly in terms of scalability, robustness, and sensitivity to parameter tuning [5]. Despite these reported challenges, evolutionary approaches remain among the most widely explored solution paradigms for NSP due to their modelling flexibility and adaptability to complex constraints. Genetic algorithms have been widely adopted due to their flexibility in handling complex constraints [6], while memetic algorithms enhance search performance through hybridisation with local improvement strategies [7]. Harmony search and related metaheuristics have also been explored to address multi-constraint nurse rostering problems [8], alongside goal programming models that explicitly incorporate preference satisfaction [9]. Robust and scenario-based optimisation approaches have further been proposed to address uncertainty in nurse rostering, although these often increase computational complexity [10], while fuzzy optimisation models incorporate uncertainty directly into preference and constraint modelling [11]. A common limitation identified across these studies is the reliance on static evaluation functions, where schedules are assessed using fixed penalty weights throughout the search. Such static evaluation may fail to reflect changing problem difficulty, especially in highly constrained settings where different violations become dominant at different stages of optimisation.

In practical NSP settings, constraint interactions are inherently dynamic, where some violations are relatively easy to resolve during early search stages, while others become critical only after basic feasibility has been achieved. To address this challenge, co-evolutionary algorithms have been proposed as an extension of evolutionary optimisation, where multiple populations evolve simultaneously and influence one another during the search process [12]. Building on this foundation, competitive co-evolutionary algorithms introduce adaptive evaluation through direct competition rather than fixed fitness assessment. Rosin and Belew [13] first formalised this prototype by demonstrating how evolving opponents can expose weaknesses that static evaluation fails to capture, thereby sustaining selection pressure during optimisation. Subsequent studies extended this idea in adversarial optimisation and game theoretic contexts. Olsson [14] proposed a host-parasite genetic algorithm, showing how asymmetric competition can improve robustness by preventing dominance stagnation. Lehre [15] provided a runtime analysis of competitive co-evolution under maximin optimisation, highlighting its ability to overcome negative drift that commonly leads to premature convergence in static fitness landscapes. Fajardo et al. [16] further demonstrated that fitness aggregation mechanisms play a critical role in maintaining effective competition, while a follow-up study [18] showed that inappropriate aggregation can destabilise competitive dynamics and degrade optimisation performance. Harris et al. [17] complemented these findings by proposing opponent sampling strategies based on strength similarity to stabilise co-evolutionary interactions and preserve meaningful selection pressure.

Despite substantial progress in genetic algorithm-based nurse scheduling, most existing approaches optimise feasibility and schedule quality within a single evolutionary process. This coupled optimisation can limit further improvements once feasible solutions are reached, as enhancements in soft constraints often conflict with feasibility maintenance. In this study, a competitive nurse scheduling framework is proposed to explicitly model this conflict by introducing competitive evolutionary pressure between objectives, thereby enabling sustained optimisation beyond feasibility and addressing stagnation commonly observed in baseline genetic algorithms.

Motivated by this gap, this study introduces a competitive co-evolutionary optimisation approach for the Nurse Scheduling Problem. The proposed algorithm preserves the original NSP formulation and constraint structure while enhancing optimisation dynamics through adaptive adversarial evaluation, showing that improved performance can be achieved without problem reformulation or additional constraints. Comparative experiments against a classical Genetic Algorithm for 30 independent runs demonstrate improved solution quality, convergence stability, and robustness under realistic scheduling conditions. These findings indicate that competitive co-evolution is a viable optimisation paradigm for complex healthcare scheduling problems and can be generalised to other highly constrained settings where static fitness evaluation leads to premature convergence.

The remainder of this study is organised as follows: Section II reviews the literature on nurse scheduling and evolutionary optimisation. Section III presents the NSP formulation, including hard and soft constraints. Section IV describes the proposed competitive co-evolutionary methodology. Section V outlines the experimental setup and evaluation metrics. Section VI presents and discusses the results, followed by conclusions in Section VII.

II. LITERATURE REVIEW

A. Nurse Scheduling Problem and Conventional Solution Approaches

The Nurse Scheduling Problem (NSP) has been widely studied due to its significant impact on hospital operations, staff wellbeing, and patient safety. Classical NSP formulations distinguish between hard constraints, which ensure feasibility and legal compliance, and soft constraints, which influence schedule quality, fairness, and preference satisfaction [3],[4]. Realistic NSP models often include multiple nurse roles, skill requirements, shift-pattern rules, and workload balance, resulting in complex and highly constrained scheduling problems.

Early NSP solutions relied on exact mathematical programming and rule-based heuristics. However, these approaches scale poorly as problem size and constraint complexity increase. As a result, metaheuristic methods have become dominant in the literature. Techniques such as genetic algorithms and memetic algorithms have been widely applied with varying success [5],[6], while other metaheuristics, such as heuristic metaheuristic designs for hard and soft constraint nurse rostering, continue to be explored [7],[8]. While these methods are flexible and capable of handling complex constraints, many studies report sensitivity to fitness design, penalty calibration and stagnation during the search process [5],[6].

Several studies emphasise the importance of incorporating realistic operational constraints into NSP models. Azimi et al. showed that neglecting practical staffing rules can produce schedules that appear optimal numerically but are infeasible in practice [4]. Wright and Mahar further demonstrated that scheduling decisions directly influence both operational cost and nurse satisfaction, highlighting the need for balanced optimisation strategies [3]. Robust and scenario-based models have also been proposed to address uncertainty in staffing demand, although these often increase computational complexity [11]. More recently, exact and model-based approaches remain active, including mixed-integer programming nurse rostering models that incorporate preference and qualification structures [20] and practical MILP case studies that reduce workload imbalance [25].

Despite the diversity of solution methods, most NSP solvers rely on static evaluation functions with fixed penalty weights. While effective for benchmark problems, static evaluation can struggle when different constraint violations become dominant at different stages of optimisation, leading to premature convergence or extensive parameter tuning [5],[10].

B. Evolutionary and Co-Evolutionary Methods in Nurse Scheduling

Evolutionary algorithms are among the most commonly used techniques for the NSP due to their flexibility and ability to handle complex constraints. Genetic Algorithms (GAs) have been widely applied in nurse scheduling, while memetic algorithms, which combine evolutionary search with local improvement, have been shown to improve convergence speed and solution quality in some NSP variants [6]. Other approaches, such as bi-level heuristics and shift swapping strategies, further exploit problem structure to refine feasible schedules [9]. Recent evolutionary and hybrid systems continue to be proposed in practical healthcare settings, including Round-Robin GA-based scheduling systems that report improvements in fairness and execution time in real medical centre contexts [22], as well as healthcare workforce scheduling studies that formulate staffing as a multi-objective GA optimisation problem [24].

Co-evolutionary methods have received comparatively less attention in NSP research. Existing studies mainly focus on cooperative co-evolution, where multiple subpopulations contribute jointly to a single solution [12], [19]. While such approaches are effective for decomposable problems, they assume alignment between subcomponents as an assumption that does not always hold in NSP.

In practical nurse scheduling, improving one aspect of a roster, such as fairness, may negatively affect another aspect, such as coverage or legal compliance. Cooperative co-evolution does not explicitly model this tension, as subpopulations are designed to collaborate rather than challenge candidate schedules. As a result, cooperative approaches may lose selection pressure once near-feasible solutions are obtained.

While co-evolutionary approaches provide useful conceptual insights into evolutionary scheduling, this study focuses on the evaluation of a competitive nurse scheduling formulation in comparison with a baseline genetic algorithm.

C. Competitive Co-Evolution and its Relevance to NSP

Competitive co-evolutionary algorithms introduce adaptive selection pressure through direct competition rather than static fitness evaluation. Rosin and Belew [13] established the foundational concept of competitive co-evolution by demonstrating how evolving opponents can actively expose solution weaknesses that static evaluation fails to capture. Building on this idea, competitive co-evolution has been further explored in adversarial optimisation and robust learning contexts, where Olsson [14] showed that asymmetric host parasite interactions can improve robustness by preventing dominance stagnation.

More recent studies have examined the dynamics and stability of competitive co-evolution in greater depth. Lehre [15] provided a runtime analysis showing how competitive evaluation can overcome negative drift in maximizing optimisation problems. Fajardo *et al.* [16] demonstrated that fitness aggregation mechanisms play a critical role in sustaining effective competition, while a follow-up study by the same authors [18] showed that inappropriate aggregation can destabilise co-evolutionary dynamics and hinder optimisation progress. Complementing these findings, Harris *et al.* [17] proposed opponent sampling strategies based on strength similarity to stabilise competition and preserve meaningful selection pressure.

This competitive standard is particularly relevant to the NSP, where constraint violations can be viewed as challenges that reveal weaknesses in schedules. Certain shift patterns, role assignments or coverage configurations may consistently expose problematic structures. Competitive evaluation can adaptively focus on these weaknesses, enabling continued improvement even after basic feasibility has been achieved.

Despite these advances, competitive co-evolution has received limited attention in nurse scheduling research. Existing co-evolutionary scheduling studies primarily adopt cooperative strategies, where multiple subpopulations work jointly toward a shared objective [12], [19]. While effective for certain decomposable problems, cooperative approaches do not explicitly model the inherent tension in NSP between feasibility and schedule quality, and may therefore struggle to maintain selection pressure once near feasible solutions are obtained. In parallel, recent non-coevolutionary NSP studies continue to explore alternative mechanisms for sustaining search pressure, including modern metaheuristics such as the Whale Optimization Algorithm adapted for nurse scheduling [23], highlighting the ongoing need for effective selection dynamics in realistic rostering problems.

In summary, although substantial progress has been made in evolutionary and hybrid optimisation methods for NSP, most existing approaches rely on static evaluation and cooperative or single-population search strategies. Competitive co-evolution, despite its ability to maintain adaptive selection pressure and robustness in other domains, has not been systematically explored for nurse scheduling. Existing co-evolutionary studies in scheduling primarily focus on cooperation rather than adversarial evaluation.

This gap motivates the development of a competitive co-evolutionary NSP algorithm that explicitly captures the adversarial nature of constraint satisfaction. By embedding adaptive competition into schedule evaluation, competitive co-evolution has the potential to address persistent challenges in NSP, including stagnation near feasibility, sensitivity to penalty tuning and limited robustness under complex real-world constraints. The proposed methodology, presented in the next section, adapts competitive co-evolutionary principles to the NSP context, while preserving practical feasibility and realistic scheduling requirements.

III. Problem Formulation

The Nurse Scheduling Problem (NSP) addressed in this study concerns the construction of a weekly duty roster that assigns nurses to shifts while satisfying operational feasibility and quality requirements. Let $N = \{1,2,\dots n\}$ denote the set of nurses and $D = \{1,2,\dots 7\}$ the set of planning days from Monday to Sunday. Each nurse is assigned exactly one shift per day from the set $S = \{1,2,3,4\}$ that represent Morning, Afternoon, Night, Post-Night and Off-day consecutively. Nurses belong to different roles, including sisters N_s, senior nurses N_{sen} and healthcare assistants N_{hca}, where role membership determines eligibility for specific shifts and Off-day rules. A schedule is represented by a discrete decision variable $x_{i,d} \in S$, where $x_{i,d}$ denotes the shift assigned to nurse i on day d. The feasibility and quality of a schedule are evaluated through a set of hard and soft constraints. Hard constraints capture mandatory operational requirements and include minimum coverage for working shifts, role-based Off-day restrictions and minimum weekly working hours. Any violation of these constraints is considered unacceptable from an operational standpoint and incurs a large penalty.

Soft constraints model preference and fairness considerations that influence schedule quality but do not invalidate feasibility. These include shift sequences such as the night shift must be followed by a post-night shift, balanced role coverage across shifts, acceptable ranges of weekly working hours and fair distribution of night shifts among eligible nurses. Soft constraint violations are penalised at a lower magnitude than hard constraints, reflecting their secondary priority. The overall penalty of a schedule x is computed using an additive formulation, as in Eq. (1):

$$F(x) = HC(x) + SC(x) \tag{1}$$

where, $HC(x)$ denotes the total hard constraint penalty and $SC(x)$ denotes the total soft constraint penalty. The classical NSP optimisation objective is therefore $min\ F(x)$, subject to the implicit domain restriction $x_{i,d} \in S$. This formulation aligns with conventional penalty based evolutionary approaches to nurse rostering and allows both feasibility and quality considerations to be captured within a single objective value.

However, because the penalty function aggregates violations across all nurses and all days, schedules that achieve low overall penalty may still contain localised structural weaknesses, such as fragile shift sequences or concentrated workload patterns affecting specific nurses. These weaknesses can remain hidden when evaluation relies solely on the global objective $F(x)$, particularly when penalty tuning balances competing constraints at an aggregate level. As a result, conventional optimisation may converge to schedules with low overall penalty values while still containing local weaknesses that are not explicitly captured by the aggregate evaluation.

To address this issue, the NSP formulation in this work is later extended with a competitive evaluation mechanism that preserves the same decision variables, constraints and penalty definitions, while altering the evaluation context under which schedules are compared. Rather than modifying the objective function itself, the competitive algorithm dynamically emphasises constraint sensitive positions during fitness assessment, enabling the optimisation process to distinguish between schedules that are merely low penalty and those that are robust under focused constraint stress. The competitive evaluation strategy and its integration into the evolutionary process are described in the following section.

Table I lists the hard and soft constraints used in the competitive NSP evaluation. The constraint set is consistent with conventional NSP formulations and is reused to isolate the effect of competitive fitness evaluation.

TABLE I. HARD AND SOFT CONSTRAINTS CONSIDERED IN COMPETITIVE NSP

ID	Constraint Description	Penalty Type
HC1	Minimum coverage for AM, PM, and Night shifts per day	Hard
HC2	Role based OFF-day rules (sisters: weekend OFF; others: exactly one OFF)	Hard
HC3	Minimum weekly working hours per nurse (Role based)	Hard
SC1	Legal shift sequence (Night must be followed by Post-Night)	Soft
SC2	Sisters assigned only AM/PM shifts on weekdays	Soft
SC3	Balanced role coverage across shifts (SEN and HCA presence)	Soft
SC4	Weekly working hours within preferred range	Soft
SC5	Fair distribution of night shifts among non-sisters	Soft

IV. METHODOLOGY

This study employs a competitive co-evolutionary optimisation algorithm, representing an advanced evolutionary optimisation model, to solve the Nurse Scheduling Problem (NSP) formulated in Section III. The underlying NSP model, including decision variables, constraints, and penalty definitions, remains unchanged. Unlike classical single population, genetic algorithms that rely on static fitness evaluation, the proposed approach introduces an adaptive competitive evaluation mechanism in which candidate schedules are assessed under adversarial selection pressure that varies dynamically throughout the evolutionary process.

A. Competitive Co-Evolutionary Algorithm

The proposed algorithm maintains two concurrently evolving populations with distinct roles. The first population consists of candidate nurse schedules, where each individual encodes a complete weekly roster represented as a nurse, day and shift assignment matrix. Each entry specifies the shift assigned to a nurse on a given day using discrete shift codes defined in the NSP formulation.

The second population acts as an evaluation population and does not represent alternative schedules. Instead, individuals in this population encode binary evaluation masks defined over the same nurse-to-day grid. A value of 1 activates focused evaluation at a specific nurse-to-day position, while a value of 0 excludes it. This design allows the algorithm to dynamically intensify selection pressure on structurally weak regions of schedules without introducing new constraints or modifying feasibility rules.

Fig. 1 illustrates the overall workflow of the proposed competitive co-evolutionary nurse scheduling algorithm. The approach operates on the original NSP formulation and maintains two populations that evolve concurrently. The schedule population encodes complete weekly nurse rosters, while the evaluation population encodes binary evaluation masks that selectively emphasise constraint sensitive nurse to day assignments. During each generation, candidate schedules are assessed using a competitive fitness function that augments the base NSP penalty with a focused adversarial term. This mechanism introduces adaptive selection pressure without modifying the underlying constraints or objective definition. Evolutionary operators are applied independently to both populations, and the final solution is selected from the scheduled population based on the base NSP objective function.

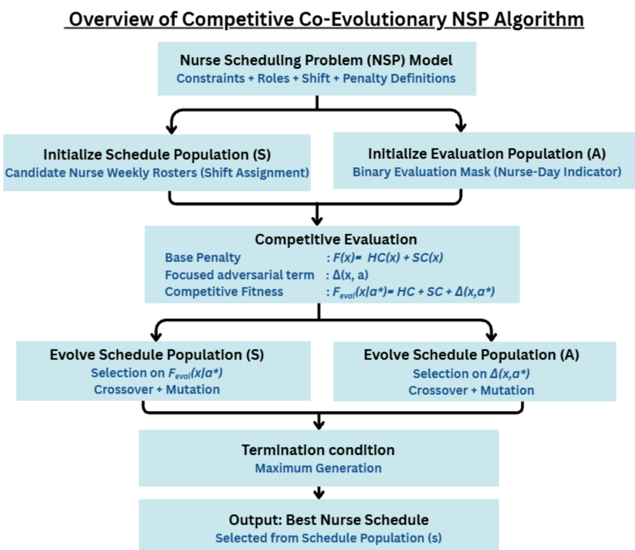

Fig. 1. Overview of the competitive co-evolutionary nurse scheduling algorithm.

B. Chromosome Representation

In the schedule population, each chromosome encodes a full weekly roster. For example, a row corresponding to Nurse i may be represented as:

Nurse i: [0 1 4 2 3 0 4]

where, the values denote Morning, Afternoon, Off-day, Night, Post-Night, Morning, and Off-day assignments respectively. In contrast, the evaluation population encodes binary masks over the same structure. A corresponding evaluation chromosome may be represented as:

Mask Nurse $\boldsymbol{i}$: [1 0 0 1 1 0 0]

Here, a value of 1 indicates that the corresponding nurse-to-day assignment is selected for focused evaluation, while a value of 0 indicates that it is ignored in the focused penalty computation.

C. Competitive Evaluation Mechanism

For each candidate schedule x, the base schedule quality is first evaluated using the standard NSP objective function as in Eq. (1). Competitive evaluation is introduced through an additional focused penalty term $\Delta(x, a)$ computed jointly from the schedule x and an evaluation mask a.

Here, $\Delta(x, a)$ denotes the focused adversarial penalty induced by evaluation mask a on schedule x. In the competitive process, $\Delta(x*, a)$ is used to evaluate adversaries against the current best schedule $x*$, while $\Delta(x, a*)$ is used to evaluate all candidate schedules under the selected strongest adversary $a*$. This distinction reflects the asymmetric roles of adversary selection and schedule evaluation in the competitive co-evolutionary algorithm.

The focused penalty re-evaluates existing constraint rules only at nurse-to-day positions activated by the mask. At these positions, local conditions such as role legality, Off-day requirements and night-to post-night sequencing are checked and penalised if violated. Positions not selected by the mask do not contribute to this term. Importantly, no new constraints are introduced, instead the mechanism amplifies the impact of existing violations at selected locations. The competitive fitness of a schedule is defined as Eq. (2):

$$F_{eval}(x|a^*) = HC(x) + SC(x) + \Delta(x, a^*) \qquad (2)$$

This formulation discourages schedules that achieve low aggregate penalties by relying on fragile local patterns, thereby promoting robustness under adversarial evaluation.

D. Adversarial Selection and Population Interaction

At each generation, a reference schedule x^* is identified as the individual with the lowest base penalty $F(x)$ in the schedule population. The evaluation population is then assessed against this reference schedule, and the strongest evaluator a^* is selected by maximising $x\ \Delta(x^*, a)$.

All candidate schedules are subsequently evaluated using the same evaluator a^*, ensuring a consistent competitive context within each generation. Evaluation individuals are selected based solely on their ability to expose weaknesses in the reference schedule, while schedule individuals are selected based on their competitive fitness. The two populations interact only through this adversarial evaluation mechanism, without chromosome exchange, cooperative pairing, or decomposition.

E. Evolutionary Operators and Termination

Both populations evolve using standard genetic operators, where the scheduled population applies selection, crossover, and mutation operators designed to preserve role eligibility and basic feasibility. The evaluation of the population evolves independently using analogous operators suitable for binary representations, encouraging diversity in evaluation focus and preventing stagnation.

The evolutionary process continues until a pre-defined termination condition is met. Throughout the search, solution quality is monitored using the base NSP objective $F(x)$ to ensure fair comparison with baseline algorithms.

The overall workflow of the proposed competitive co-evolutionary nurse scheduling algorithm is summarised in Pseudocode, illustrating the parallel evolution of schedule and evaluation populations and their interaction through adversarial fitness assessment.

Having described the proposed competitive co-evolutionary nurse scheduling algorithm and its adversarial evaluation mechanism combined, as in Eq. (2), the next section details the experimental setup used to assess its effectiveness. This includes the problem instances, parameter settings, baseline algorithms for comparison, and performance metrics employed to ensure a fair and reproducible evaluation of the proposed approach.

Pseudocode: Competitive Co-Evolutionary NSP

Initialize schedule population S
Initialize adversarial population A

Evaluate base penalty $F(x)$ **for** all $x \in S$
bestSoFar ← $min\, F(x)$
For generation = 1 **to** MAX_GENERATION **do**

  Identify reference schedule:
    $x* \leftarrow arg\ min\, F(x),\ x \in S$
  Identify strongest adversary:
    $a* \leftarrow arg\ max\, \Delta(x*, a),\ a \in A$

  For each schedule $x \in S$ **do**
    $F_{eval}(x) \leftarrow HC(x) + SC(x) + \Delta(x, a*)$
  End For

  For each adversary $a \in A$ **do**
    $G(a) \leftarrow \Delta(x*, a)$
  End For
  Evolve S using tournament selection on F_{eval}
  Evolve A using tournament selection on G

  If $min_{x \in S} F(x) < F(bestSoFar)$ then
    bestSoFar← $arg\ min\, F(x)_{x \in S}$
  End If
Update bestSoFar if a lower $F(x)$ is found
End For
Return bestSoFar // // best schedule from Schedule Population S
//selected by NSP population objective, $F(x)$

V. Dataset and Experimental Setup

The proposed competitive co-evolutionary approach was evaluated using benchmark nurse scheduling instances derived from realistic ward-level rostering scenarios. Each dataset represents a seven-day planning horizon from Monday to Sunday and includes multiple nurse roles, including sisters, senior nurses and healthcare assistants. The datasets incorporated role-based shift eligibility, coverage requirements, legal rest rules, and workload balance to reflect practical hospital scheduling conditions.

The experimental evaluation focuses on a 20-nurse, one-week scheduling instance, which was deliberately selected to provide a controlled and interpretable setting for comparative analysis. This scale enables clear observation of optimisation dynamics, convergence behaviour and robustness differences between evolutionary strategies without introducing confounding effects from instance size.

All compared methods were operated on identical problem instances and shared the same schedule representation, constraint definitions and penalty structure. This ensures that observed performance differences arise from optimisation strategy rather than differences in problem formulation.

The experimental evaluation considered two optimisation approaches, namely a classical genetic algorithm (GA) and the proposed competitive co-evolutionary NSP. The approaches differ only in their evolutionary interaction and fitness evaluation mechanisms, while genetic operators and representations remain consistent across both methods.

Algorithm parameters follow commonly adopted settings in evolutionary nurse scheduling studies and are held constant across methods unless otherwise stated. Each algorithm was executed for a fixed number of generations with identical population sizes and stopping criteria. To account for stochastic effects, 30 independent runs were performed for each dataset using different random seeds.

Performance was assessed using best and average penalty values across generations, together with convergence trends over time. All reported results were computed using the base NSP objective function to ensure that comparisons reflect genuine schedule quality rather than effects introduced by adversarial evaluation.

All experiments were conducted on the same computational platform using a consistent software environment. Execution time, convergence behaviour and solution quality were recorded for each run and form the basis of the comparative analysis presented in the following section.

While larger and more diverse nurse scheduling instances are important for evaluating scalability, the primary objective of this study is methodological validation of competitive co-evolutionary optimisation under controlled conditions. Extension of the proposed approach to larger scale and multiple ward nurse scheduling problems is therefore identified as a key direction for future work.

Table II summarises the dataset characteristics and scheduling requirements used to evaluate the proposed competitive co-evolutionary NSP. The problem definition follows standard nurse rostering practice and was adopted to enable fair comparison across optimisation strategies.

TABLE II. Problem Parameters for Competitive Nurse Scheduling Experiments

Parameter	Description	Value
Planning horizon	Number of days per scheduling period	7
Number of nurses	Total nurses per instance	20
Nurse roles	Sisters / Senior Nurses / HCAs	5/8/7
Shift types	AM, PM, Night, Post-Night, OFF	5
Shifts per nurse per day	Assignment constraint	Exactly one
Coverage requirement	Minimum staff per working shift	≥1 Nurse and ≥1 HCA (role-dependent)
Weekly OFF rules	Role-specific OFF constraints	Enforced
Working hours rule	Minimum weekly hours	Enforced
Shift sequence rules	Legal patterns (e.g., N → PN)	Enforced
Constraint handling	Hard vs soft constraints	Penalty-based
Objective definition	Schedule quality measure	$HC(x) + SC(x)$ (base NSP objective)

Table III reports the algorithmic parameters used in the experimental comparison of classical GA and the proposed competitive co-evolutionary approach. Both methods share identical operator settings unless stated otherwise.

TABLE III. ALGORITHM AND EXPERIMENTAL PARAMETERS

Parameter	Setting
Optimisation methods	Genetic Algorithm (GA) NSP, Competitive NSP
Population size	50 individuals per population
Number of populations	GA NSP = 1, Competitive NSP = 2
Parent selection	Tournament selection (size = 2)
Crossover operator	One-point crossover
Crossover probability	0.7
Mutation operator	Role-constrained shift reassignment
Mutation probability	GA NSP = 0.7, Competitive NSP = 0.05
Replacement strategy	Worst individual replacement based on the respective fitness definition
Maximum generations	15,000
Independent runs	30 runs per dataset
Fitness used for selection	GA NSP = Static base fitness, Competitive NSP = Adversarial competitive evaluation
Performance metric reported	Base NSP penalty (hard + soft constraints)

To ensure fair runtime measurement, execution time was recorded for the optimisation process only, excluding file input/output and result printing, which do not affect the optimisation logic.

VI. RESULTS AND DISCUSSION

This section presents the experimental evaluation of the proposed Competitive Co-Evolutionary Nurse Scheduling Problem (NSP) algorithm. The method was compared against the classical Genetic Algorithm (GA), where both algorithms shared identical problem representations, constraint formulations, and dataset parameter settings to ensure fair comparison. Performance was assessed using the total penalty value, defined as the sum of hard and soft constraint violations.

A. Performance Consistency for 30 Independent Runs

To evaluate robustness and repeatability, the Competitive NSP was executed for 30 independent runs using different random seeds. Across these runs, best penalty values ranged from 337.00 to 503.81, while average penalty values ranged from 339.72 to 590.16, indicating some variability inherent to stochastic optimisation.

TABLE IV. STATISTICAL SUMMARY OF COMPETITIVE NSP PERFORMANCE ACROSS 30 RUNS

Metric	Best Penalty	Average Penalty
Mean	447.28	520.14
Std. Dev.	41.63	60.27
Minimum	337	339.72
Maximum	503.81	590.16

The statistical summary of these runs is reported in Table IV. The method achieves a mean best penalty of 447.28 and a mean average penalty of 520.14, with relatively low standard deviations of 41.63 and 60.27, respectively. These results demonstrate stable convergence behaviour and confirm that the proposed competitive algorithm consistently produces high-quality nurse schedules across repeated executions rather than relying on isolated favourable runs.

B. Comparative Statistical Analysis

A comparative statistical summary for both methods was reported in Table V. The classical GA achieves a mean best penalty of 651.3 and a mean average penalty of 884.6. While the GA exhibits relatively low variance, its final penalty values remain substantially higher than those obtained by the Competitive NSP, suggesting premature convergence once basic feasibility is achieved.

TABLE V. COMPARATIVE PERFORMANCE SUMMARY OF GA AND COMPETITIVE NSP

Metric	Genetic Algorithm (GA) NSP	Competitive NSP
Mean Best Penalty	651.30	447.28
Mean Average Penalty	884.60	520.14
Std. Dev. (Best)	37.90	41.63
Std. Dev. (Average)	40.80	60.27
Minimum Penalty	581.70	337.00
Maximum Penalty	703.60	503.81

In contrast, the Competitive NSP demonstrates markedly improved performance. The mean best penalty is reduced to 447.28, representing a 31.3% reduction compared to the baseline GA, while the mean average penalty decreases from 884.6 to 520.14, corresponding to a 41.2% reduction. These results indicate that introducing competitive evaluation significantly enhances optimisation effectiveness without modifying the underlying NSP formulation.

TABLE VI. EXECUTION TIME COMPARISON OF GA AND COMPETITIVE NSP

Metric	Genetic Algorithm (GA) NSP	Competitive NSP
Mean (sec)	20.53	12.67
Std. Dev (sec)	0.17	0.38
Minimum (sec)	20.22	12.12
Maximum (sec)	20.73	13.41
Median (sec)	20.54	12.48

Table VI reports the execution time statistics for both optimisation methods over 30 independent runs under identical computational settings. The baseline Genetic Algorithm (GA) records the higher mean execution time, with an average runtime of 20.53 seconds, reflecting its reliance on a single population search and static fitness evaluation, which typically requires additional generations to refine feasible schedules.

In contrast, the Competitive NSP achieves a substantially lower mean execution time of 12.67 seconds, representing a reduction of approximately 38.3% compared to the baseline GA. Although the Competitive NSP introduces additional evaluation mechanisms to model competition between objectives, its runtime remains efficient and well within practical limits for offline nurse scheduling. This improvement in computational efficiency, together with the significant gains in solution quality reported earlier, demonstrates the effectiveness of competitive evaluation without incurring prohibitive computational cost.

C. Convergence Behaviour Analysis

The convergence behaviour of both methods is illustrated in Fig. 2 and Fig. 3, corresponding to GA NSP and Competitive NSP, respectively.

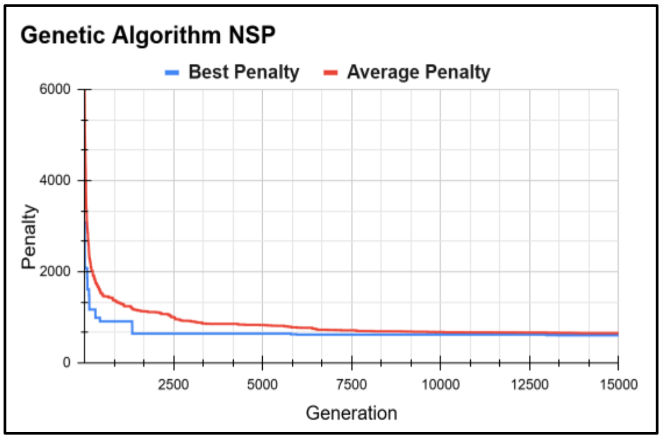

Fig. 2. Convergence of best and average penalty values for the genetic algorithm in NSP.

As shown in Fig. 2, the GA exhibits rapid initial improvement during the early generations, with penalty values decreasing sharply as feasible schedules are identified. However, convergence stagnates shortly thereafter, indicating premature convergence. This behaviour can be attributed to the reliance on static fitness evaluation within a single population, which limits further exploration once basic feasibility is achieved.

In contrast, Fig. 3 demonstrates that the Competitive NSP maintains a smooth and sustained reduction in penalty values throughout the evolutionary process. While its early convergence trend is visually similar to that of the GA, the Competitive NSP continues to refine solutions beyond the initial feasibility phase. This is evidenced by the progressively narrowing gap between the best-of-generation and average penalty values, indicating stronger population-wide improvement.

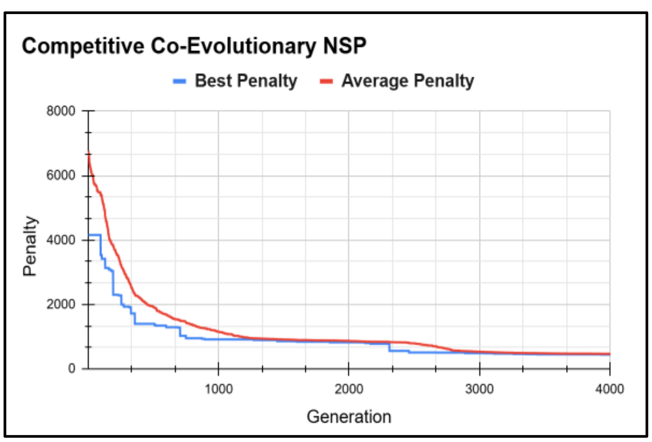

Fig. 3. Convergence of best and average penalty values for the competitive co-evolutionary NSP.

This convergence behaviour reflects more persistent selection pressure introduced by competitive evaluation, enabling continued optimisation and reducing the likelihood of stagnation. These results highlight the effectiveness of the Competitive NSP in sustaining meaningful optimisation progress compared to the baseline GA.

D. Discussion of Competitive Advantage

The superior performance of the Competitive NSP can be attributed to its relative evaluation mechanism, which fundamentally alters how selection pressure is applied during the search process. Unlike classical Genetic Algorithms, which optimise schedules against static evaluation criteria, the competitive algorithm continuously challenges candidate schedules under evolving evaluation contexts. This adaptive pressure discourages overfitting to fixed penalty weightings and prevents premature stagnation, allowing the algorithm to refine solution quality even after feasible schedules have been identified. Similar limitations associated with static fitness evaluation have been reported in recent nurse scheduling studies [21] employing hybrid genetic algorithms and multi-objective formulations, where convergence tends to plateau once basic feasibility is achieved due to fixed objective definitions [22],[24]. In contrast, the competitive evaluation strategy sustains selection pressure beyond feasibility, enabling continued improvement in solution quality.

From an optimisation perspective, the competitive evaluation mechanism effectively exposes localised structural weaknesses that may remain hidden under aggregate penalty evaluation. Conventional NSP solvers assess schedules primarily through global penalty accumulation, which can mask fragile shift sequences or concentrated workload patterns affecting individual nurses. By selectively amplifying such weaknesses, adversarial evaluation forces the evolutionary process to address them explicitly. This behaviour is consistent with observations from competitive co-evolutionary research in other optimisation domains, where adaptive opponent-based evaluation has been shown to maintain meaningful selection pressure and mitigate stagnation caused by static fitness landscapes [15],[16]. The present results extend these findings by demonstrating that similar advantages can be realised within a practical nurse scheduling setting.

Importantly, both methods evaluated in this study optimise the same base objective function and share identical constraint definitions. The observed performance gains, therefore, stem from differences in evolutionary interaction and selection dynamics rather than changes to the NSP formulation itself. This characteristic differentiates the proposed approach from recent integer and mixed-integer programming-based nurse rostering models, which typically rely on fixed optimisation objectives and predefined constraint weighting schemes to achieve solution quality [20],[25]. While such exact and hybrid models offer strong feasibility guarantees, their optimisation behaviour remains tightly coupled to the chosen formulation and does not adapt dynamically during the search.

The competitive co-evolutionary algorithm also provides a distinct alternative to recent nature-inspired metaheuristics proposed for NSP, such as whale optimisation and related population-based methods, which primarily enhance exploration through novel operators while retaining static fitness evaluation [23]. Rather than increasing operator complexity, the proposed method introduces adaptivity directly at the evaluation level, resulting in improved robustness without additional tuning parameters. This property is particularly relevant for real-world nurse scheduling environments, where constraint interactions are dynamic and difficult to capture through fixed penalty calibration.

Overall, these findings indicate that competitive co-evolution improves optimisation effectiveness through adaptive search dynamics rather than problem reformulation or operator augmentation. The results position competitive evaluation as a promising direction for healthcare scheduling and constraint optimisation, especially for highly constrained NSP instances where sustained adaptability and robustness are essential for producing high-quality schedules.

VII. Conclusion

This study examined the effectiveness of a competitive nurse scheduling formulation for solving the Nurse Scheduling Problem (NSP) and compared its performance with a classical Genetic Algorithm (GA). Unlike conventional evolutionary approaches that rely on static fitness evaluation within a single population, the proposed method introduces competitive evaluation into the evolutionary process, enabling adaptive and continuously informative selection pressure throughout the search.

Experimental results across 30 independent runs demonstrate clear differences in convergence behaviour between the two methods. The classical GA achieves rapid early improvement, but exhibits premature stagnation once near-feasible solutions are obtained, leading to higher final penalty values. In contrast, the Competitive NSP maintains sustained and smooth improvement across generations, achieving lower best and average penalties with a progressively narrowing gap between them. This behaviour indicates stronger population-wide refinement and improved robustness of the resulting schedules.

Notably, these performance gains are achieved without altering the underlying NSP formulation, constraint definitions, or penalty structure. The observed improvements arise solely from the introduction of competitive evaluation and its influence on evolutionary dynamics, rather than from problem reformulation or additional model complexity. This demonstrates the effectiveness of competitive evolutionary pressure in addressing stagnation commonly observed in GA-based nurse scheduling.

Future work will focus on extending the proposed competitive NSP to larger-scale nurse scheduling instances and integrated healthcare scheduling problems involving operating theatres and patient admissions. Further investigation into incorporating domain-specific knowledge within the competitive evaluation process is also planned. In addition, cooperative co-evolutionary algorithms may be explored as a complementary framework to examine how cooperative and competitive evolutionary pressures interact in complex and highly constrained healthcare scheduling environments.

Acknowledgment

The authors gratefully acknowledge the support of Universiti Teknikal Malaysia, Melaka and the Ministry of Higher Education, Malaysia under the Fundamental Research Grant Scheme (FRGS/1/2023/FTMK/F00546). This research was also supported by the University's Kesidang Scholarship, which funded the first author's postgraduate studies.

Declaration on the Use of AI-Assisting Tools

The authors declare that generative AI tools were used solely for language refinement and editorial assistance. All scientific content, including problem formulation, methodology, experiments, and analysis, was developed and verified by the authors, who take full responsibility for the accuracy and integrity of this work.

References

[1] MacroTrends, World Population Growth Rate 1950–2024. Available: https://www.macrotrends.net/global-metrics/countries/WLD/world/population-growth-rate

[2] Z. A. Bdalkareem, A. Amir, M. A. Al-Betar, P. Ekhan, and A. I. Hammouri, "Healthcare scheduling in optimization context: A review," Health and Technology, vol. 11, no. 3, pp. 445–469, 2021.DOI: https://doi.org/10.1007/s12553-021-00547-5

[3] P. D. Wright and S. Mahar, "Centralized nurse scheduling to simultaneously improve schedule cost and nurse satisfaction," Omega, vol. 41, no. 6, pp. 1042–1052, 2013. DOI: https://doi.org/10.1016/j.omega.2012.08.004

[4] S. Azimi, M. M. Sepehri, and M. Etemadian, "A nurse scheduling model under real life constraints," International Journal of Hospital Research, vol. 4, no. 1, pp. 1–8, 2015. Available: https://doi.org/10.1007/s10916-014-0160-8

[5] C. M. Ngoo, S. L. Goh, S. N. Sze, N. R. Sabar, S. Abdullah, and G. Kendall, "A survey of the nurse rostering solution methodologies: The state-of-the-art and emerging trends," IEEE Access, vol. 10, pp. 56504–56524, 2022. DOI: https://doi.org/10.1109/ACCESS.2022.3177280

[6] N. Nico, N. Charibaldi, and Y. Fauziah, "Comparison of memetic algorithm and genetic algorithm on nurse picket scheduling at public health center," International Journal of Artificial Intelligence and Robotics, vol. 4, no. 1, pp. 9–23, 2022. DOI: https://doi.org/10.25139/ijair.v4i1.4323

[7] P.-S. Chen and Z.-Y. Zeng, "Developing heuristic algorithms with metaheuristics to solve nurse rostering problems with hard and soft constraints," Applied Soft Computing, vol. 93, 2020. DOI: https://doi.org/10.1016/j.asoc.2020.106336

[8] E. Yağmur and A. Sarucan, "Nurse scheduling with opposition-based parallel harmony search algorithm," Journal of Intelligent Systems, vol. 28, no. 4, 2019. DOI: https://doi.org/10.1515/jisys-2017-0150

[9] A. B. Youssef and S. Senbel, "A bi-level heuristic solution for the nurse scheduling problem based on shift-swapping," in Proc. IEEE Conf., 2018. DOI: https://doi.org/10.1109/CCWC.2018.8301623

[10] M. R. Hassani and J. Behnamian, "A scenario-based robust optimization with a pessimistic approach for nurse rostering problem," Journal of Combinatorial Optimization, 2021. DOI: https://doi.org/10.1007/s10878-020-00667-0

[11] P. Rerkjirattikal, V. Huynh, S. Olapiriyakul, and T. Supnithi, "A goal programming approach to nurse scheduling with individual preference satisfaction," Mathematical Problems in Engineering, 2020. DOI: https://doi.org/10.1155/2020/2379091

[12] X. Ma, X. Li, Q. Zhang, K. Tang, Z. Liang, and Z. Zhu, "A survey on cooperative co-evolutionary algorithms," IEEE Transactions on Evolutionary Computation, vol. 23, no. 3, pp. 421–442, 2019. DOI: https://doi.org/10.1109/TEVC.2018.2868770

[13] C. D. Rosin and R. K. Belew, "Methods for competitive co-evolution: Finding opponents worth beating," in Proc. International Conference on Genetic Algorithms, 1995.

[14] B. Olsson, “A host–parasite genetic algorithm for asymmetric tasks,” in European Conference on Machine Learning, 1998. DOI: https://doi.org/10.1007/BFB0026705

[15] P. K. Lehre, “Runtime analysis of competitive co-evolutionary algorithms for maximin optimisation of a bilinear function,” in Proc. GECCO, 2022. DOI: https://doi.org/10.1145/3512290.3528853

[16] M. A. Fajardo, P. K. Lehre, and D. Sudholt, “Runtime analysis of a co-evolutionary algorithm: Overcoming negative drift in maximin optimisation,” in Proc. GECCO, 2023. DOI: https://doi.org/10.1145/3583133.3590701

[17] M. A. Fajardo, P. K. Lehre, and D. Sudholt, “How fitness aggregation methods affect the performance of competitive CoEAs on bilinear problems,” in Proc. GECCO, 2023. DOI: https://doi.org/10.1145/3583131.3590506

[18] C. Harris, S. Didi, and A. Bucci, “Elo-based similar-strength opponent sampling for multiobjective competitive coevolution,” in Proc. GECCO, 2021. DOI: https://doi.org/10.1145/3449726.3459559

[19] L. Sun, L. Lin, H. Li, and M. Gen, “Cooperative co-evolution algorithm with an MRF-based decomposition strategy for stochastic flexible job shop scheduling,” Mathematics, vol. 7, no. 4, 2019. DOI: https://doi.org/10.3390/math7040318

[20] E. Ouda, A. Sleptchenko, and M. C. E. Simsekler, “Nurse Rostering via Mixed-Integer Programming,” in Industrial Engineering and Applications, Advances in Transdisciplinary Engineering, vol. 35, pp. 815–823, 2023. DOI: https://doi.org/10.3233/ATDE230110

[21] Z. Zhang, “Nurse Scheduling Algorithms based on Different Scenarios,” Frontiers in Computing and Intelligent Systems, vol. 4, no. 2, pp. 12–16, 2023. DOI: https://doi.org/10.54097/fcis.v4i2.9700

[22] D. Vallejos, K. Villalobos, J. L. Castillo Sequera, and L. Wong, “Intelligent System Based on Round Robin and Genetic Algorithm for Managing Nurse Schedules in Health Centres in Peru,” International Journal of Online and Biomedical Engineering (iJOE), vol. 20, no. 14, pp. 116–138, 2024. DOI: https://doi.org/10.3991/ijoe.v20i14.50571

[23] M. Sadeghilalimi, M. Mouhoub, and A. Ben Said, “Solving the Nurse Scheduling Problem Using the Whale Optimization Algorithm,” in Optimization and Learning (OLA 2023), Communications in Computer and Information Science, vol. 1824, pp. 62–73, 2023. DOI: https://doi.org/10.1007/978-3-031-34020-8_5

[24] V. Patel, A. Deodhar, and D. Birru, “A Multi-Objective Genetic Algorithm for Healthcare Workforce Scheduling,” in Proc. MODeM 2025 (ECAI 2025 Workshop), 2025. DOI: https://doi.org/10.48550/arXiv.2508.20953

[25] K. Thongsopa and U. Janjarassuk, “Optimization of Nurse Scheduling Problem: A hospital case study,” Engineering and Technology Horizons, vol. 38, no. 4, 2021. DOI: https://doi.org/10.1088/1757-899X/536/1/012131