

Minimierung der Produktionsdauer in Fabriken durch maschinelles Lernen

A Learning Beam Search for the No-Wait Flow Shop Scheduling Problem with Release Times

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Software Engineering & Internet Computing

eingereicht von

Jonas Mayerhofer, BSc

Matrikelnummer 01633065

an der Fakultät für Informatik

der Technischen Universität Wien

Betreuung: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Günther Raidl Mitwirkung: Projektass. Marc Huber, MSc

Wien, 20. April 2022

Jonas Mayerhofer

Günther Raidl



Minimizing Makespan in Flow Shops with a Reinforcement Learning like Approach

A Learning Beam Search for the No-Wait Flow Shop Scheduling Problem with Release Times

DIPLOMA THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Software Engineering & Internet Computing

by

Jonas Mayerhofer, BSc

Registration Number 01633065

to the Faculty of Informatics

at the TU Wien

Advisor: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Günther Raidl Assistance: Projektass. Marc Huber, MSc

Vienna, 20th April, 2022

Jonas Mayerhofer

Günther Raidl

Erklärung zur Verfassung der Arbeit

Jonas Mayerhofer, BSc

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 20. April 2022

Jonas Mayerhofer

Danksagung

Ich möchte mich bei allen Personen bedanken, welche mir beim Erstellen dieser Arbeit geholfen haben, die mir mit Rat und Tat zur Seite gestanden sind und welche mich über die Jahre hinweg in meiner universitären Laufbahn unterstützt haben.

Zuerst gilt mein Dank Herrn Prof. Günther Raidl, welcher meine Masterarbeit betreut und begutachtet hat. Danke für die hilfreichen Anmerkungen, die konstruktive Kritik und die investierte Zeit. Ebenso möchte ich mich bei Marc Huber, MSc bedanken. Danke für deine Mitwirkung bei der Betreuung dieser Arbeit und unsere Diskussionen.

Ich danke auch meiner Familie, meinen Studienkollegen und Freunden, die mich während der letzten Jahre begleitet und unterstützt haben. Danke, dass ihr mir ermöglicht und geholfen habt da zu sein, wo ich jetzt bin. Abschließend danke ich auch meiner Freundin. Danke für deine Geduld während des Schreibens an dieser Arbeit und deine Unterstützung durch deine LaTeX Kenntnisse, welche mir mehrmals Stunden mühsamen Suchens erspart haben.

Kurzfassung

Maschinenbelgungsplanungsprobleme (FSP) existieren in unterschiedlichen Varianten. Eine dieser Varianten ist das no-wait FSP with release times (NWFSP-RT). Es besteht aus einer Menge an Aufgaben und einer Menge an Maschinen. Beim NWFSP-RT müssen alle Aufgaben in derselben vordefinierten Reihenfolge auf allen Maschinen abgearbeitet werden. Zusätzlich dürfen Aufgaben auf der ersten Maschine erst nach einer gegebenen Release-Zeit beginnen. Anwendungen finden sich in der Stahl- und Lebensmittelproduktion zu finden. Hier können Wartezeiten zu einer Verschlechterung der Qualität führen.

Wir verwenden eine Beam Search (BS), d.h. eine Breitensuche mit beschränkter Breite, zum Lösen des NWFSP-RT. Auf jeder Suchebene behält BS nur die besten Knoten. Um zu entscheiden, welche Knoten behalten werden, verwendet BS eine heuristische Funktion (GF). In dieser Arbeit verwenden wir das Learning Beamsearch (LBS) Framework von Huber und Raidl [HR21] um GFs zu lernen. Bei der LBS werden in jeder Iteration Trainingsdaten unter Verwendung der aktuell gelernten GF erzeugt und diese anschließend verwendet um das neuronale Netzwerk (NN) zu trainieren.

Wir präsentieren zwei neuartige graphbasierte NN Typen inklusive Feature-Vektoren. Die NN Typen aggregieren Daten aller Aufgaben einer Instanz und der nähesten benachbarten Aufgaben jeder Aufgabe. Unter den Features ist eine neuartige Methode zur Berechnung unterer Schranken (ITLB) für das NWFSP-RT enthalten. Die beschriebenen Algorithmen und NN Typen wurden von uns implementiert und auf Vergleichsinstanzen. sowie zufälligen Instanzen evaluiert. Im Zuge der Evaluierung wurden auch statistische Signifikanztests durchgeführt. Die Ergebnisse zeigen, dass BS in Kombination mit den zwei neuartigen NN Typen in neun und zehn von 16 Konfigurationen signifikant bessere Ergebnisse erzielt als BS in Kombination mit ITLB verglichen auf Testinstanzen gleicher Größe, auf welcher die NN Typen trainiert wurden. Des Weitern generalisiert mindestens eine Konfiguration jedes der vier NN Typen gut über die Anzahl an Aufgaben, verglichen mit den besten bekannten Ergebnissen aus [Pou+20]. Während unserer Tests liefern die NN Typen für einzelne Instanzen bessere Ergebnisse als die besten bekannten Ergebnisse trotz der Einschränkung, dass die Tests mit einer kleineren maximalen Breite und ohne lokaler Suche durchgeführt wurden. Insgesamt war einer unserer Ansätze, evaluiert mit einer kleineren Breite als in [Pou+20], auf 11 von 46 getesteten Instanzklassen im Durchschnitt besser als die besten bekannten Ergebnisse und auf 43 von 46 Instanzklassen besser als BS ohne lokale Suche von [Pou+20].

Abstract

The flow shop problem (FSP) is a scheduling problem with many variants. One variant is the no-wait FSP with release times (NWFSP-RT). It consists of a set of jobs and a set of machines. It further imposes the constraints that jobs must pass all machines in a predefined order, the jobs are not allowed to wait on a machine until being processed, and jobs may only start processing on the first machine when their release time is exceeded. The goal is to find a schedule optimizing the desired objective, i.e., the makespan. The NWFSP-RT has applications in steel or food production where the product is not allowed to wait before being further processed to avoid degradation.

Beam search (BS), a limited width breadth-first search technique, has shown to be an effective heuristic in finding proper solutions to optimization problems within a limited time. Only the best nodes are kept and further branched on at every layer when the number of nodes exceeds a specific limit. To decide which nodes to keep, BS uses a guidance function (GF). We build upon the learning beam search (LBS) framework proposed by Huber and Raidl [HR21] to learn GFs. The LBS framework uses an iterative approach. In every iteration, training data is generated with a BS guided by the currently learned GF, and the neural network (NN) is trained to approximate the training data.

We propose two novel NN types, inspired by graph neural networks, that aggregate data over all individual jobs in a problem instance and their nearest neighbors. Further, we present feature sets for the NN types, including a novel lower bound, called ITLB, for the NWFSP-RT. We implement the algorithms, evaluate them over benchmark sets and random test instances, and perform statistical tests. The results show that a BS guided by two of our NN types produces significantly better results in nine and ten out of 16 configurations, respectively, than a BS guided by ITLB alone when run on similar instance sizes as the NNs were trained. The evaluation of the generalization abilities of the NNs shows that for each of the four NN types, at least one configuration generalizes well over the number of jobs compared with the best-known results. Our approaches frequently improve the state-of-the-art on even though running with a smaller beam width and without local search compared to the BS from Pourhejazy et al. [Pou+20] that represented the state-of-the-art so far. Overall, one of our approaches, evaluated with smaller beam width than in [Pou+20], was able to outperform the state-of-the-art on 11 out of 46 tested instance classes and to outperform the BS from [Pou+20] without local search on 43 out of 46 tested instance classes on average.

Contents

K	Kurzfassung iz					
A	Abstract					
Contents x						
1	Introduction					
2	Considered Problem2.1Scheduling Problem2.2The No-Wait Flow Shop Scheduling Problem with Release Times	5 5 6				
3	Related Work3.1The Flow Shop Scheduling Problem and its Variants3.2The Beam Search Framework3.3Machine Learning for Optimization Problems3.4The Learning Beam Search Framework	11 11 14 18 21				
4	Solving the NWFSP-RT with Beam Search4.1Beam Search for the NWFSP-RT4.2Lower Bounds for a Beam Search Node	23 23 25				
5	Learning Beam Search for the NWFSP-RT5.1General Details	 33 33 33 34 37 41 				
6	Results 6.1 Test Setup and Benchmark Instances	43 43 44 45 53				

	 6.5 Approximation Errors over Layers	$54\\58$	
7	Conclusion	73	
8	Future Work and Open Questions	75	
List of Algorithms			
Glossary			
Ac	Acronyms		
Bi	Bibliography		

CHAPTER 1

Introduction

Flow shops exist everywhere in today's world. Just think of a wooden furniture factory. To produce the furniture, first, the wood needs to be cut. After this, it needs to be ground and probably also painted before it is finally assembled. In a factory, not only one product is produced. Many kinds and variants of wooden furniture are created on-demand when ordered by the customer. Every product takes a different amount of time in each of the mentioned production steps. This yields idle times when the machines wait for the next product to be ready for processing. Nevertheless, most factory owners would like to increase the efficiency of the factory and lower the production costs. Therefore, their goal is to finish processing all products as early as possible. This problem maps directly to the Flow Shop Schedueling Problem (FSP). The FSPs' goal is to schedule a set of jobs on a set of machines. All machines must be passed in the same predefined order by each job. Furthermore, the jobs should be ordered in a way that minimizes some objective, like, e.g., the makespan. A variant of the FSP is the no-wait FSP with release times (NWFSP-RT), where jobs, i.e., products, are not allowed to wait before being processed on the next machine and only become available after a job-specific release time. An example for the NWFSP-RT would be a steel production factory, which produces different types of steel. First, the ore must be melted. Then it is processed in many intermediate steps until it is finally cast into a beam. During those steps, the melted product is not allowed to wait until the next machine is ready to process it, as this would mean that it cools down and degrades until it can not be processed further anymore. Now additionally assume that components needed during production are only available after a particular time, i.e., the release time, as they still need to be delivered to the factory.

The FSP as well as the NWFSP-RT are NP-hard for more than two machines [GJS76]. A non-deterministic Turing machine can solve problems in NP in polynomial time. Equivalently, a deterministic Turing machine can verify this class of problems in polynomial time.

1. INTRODUCTION

NP-hard implies that all other problems which are in NP can be reduced to NP-hard problems in polynomial time [KV18, Def. 15.40].

A heuristic approach to solve the NWFSP-RT is beam search (BS). Beam search is an incomplete breadth-first search technique where only a certain number of best nodes on every level are kept and branched further on while the others are discarded. A guidance function, also called a heuristic function, is used to decide which nodes on a level should be kept. In our setting, the guidance function should approximate the cost additionally introduced when completing the solution. Its result is added to the current objective value of the node's partial solution, and the best nodes according to this calculation are selected. However, finding proper guidance functions is a highly problem-specific task requiring expertise and testing.

To overcome the issue of finding proper guidance functions, Huber and Raidl [HR21] propose the learning beam search (LBS) framework for learning proper guidance functions for optimization problems in a reinforcement learning like setting. The framework iteratively generates random training samples. It labels them by using a BS run with the current learned guidance function and trains the used machine learning model on the generated training samples as well as training samples from previous iterations. In [HR21], it is shown that it is possible to use LBS to learn a proper guidance function for the Longest Common Subsequence Problem (LCSP). The LCSP asks to find the longest subsequence that occurs in all strings of the input set.

The LBS framework is inspired by the work of Silver et al. in 2017 [Sil+17]. The authors aimed to master the game of Go and reach super-human performance by letting the algorithm play games against itself [Sil+17]. In LBS the algorithm does not play against itself. However, it uses its acquired knowledge from the last iteration to compute better, w.l.o.g. smaller, labels than before and thus pushing itself to continue improving the learned guidance function.

In this work we aim to adapt LBS such that it can be applied to the NWFSP-RT. This is described in Chapter 5. In Chapter 4, we introduce a novel lower bound for the NWFSP-RT which can be computed in $O(n^2m)$ and call it improved Taillard lower bound (ITLB). It is inspired by the lower bound presented in [Tai93] and the asymmetric Traveling Salesperson Problem (ATSP) lower bound that takes the sum of the minimum incoming or outgoing arcs of every city [Chr72]. This lower bound is used as a feature in our machine learning models as well as a reference guidance function. A set of features for a single BS node is derived from the instance and the state of the search. This set of features is used in combination with a multilayer perceptron (MLP). Further, we present two novel machine learning models: aggregated jobs neural network (AJNN_{add}) and nearest neighbors neural network (NN_{nearest}). The models are based on the idea of graph neural networks (GNNs). Both AJNN_{add} and NN_{nearest} aim to use information about the structure of the instance and not only a global view of it. Thus, information about the structure is already incorporated into the feature set on a job-specific basis, considering information about close neighboring jobs. The distance measure used is defined later. In the AJNN_{add} a MLP is applied to each job's specific features, combined

with some global features, and the results are aggregated using a sum. The $NN_{nearest}$ uses a more elaborate structure that then the $AJNN_{add}$. Similarly, it applies to each job's specific features some learnable weights. Additionally, it contains links that pass aggregate information about the nearest neighbors of a job. The structure of the $NN_{nearest}$ also contains skip-connections and finally aggregates the results using the sum of all job-specific results. One can think of the $AJNN_{add}$ and the $NN_{nearest}$ in a way that they sum up the "contribution" of each job to the makespan increase.

We evaluate several configurations of the presented models on the same instance sizes they were trained on and perform significance tests. The results show that our novel neural network (NN) models trained with the LBS framework perform significantly better than a BS guided by ITLB for many configurations. Further, we evaluate how well the trained models generalize when evaluating instances with a different number of machines and/or jobs as they were trained. The evaluation shows that generalizing over the number of jobs works well, whereas generalizing over the number of machines does not. Finally, we compare the best configurations with the state-of-the-art results in [Pou+20]. Overall, one of our approaches, evaluated with smaller beam width than in [Pou+20], was able to outperform the state-of-the-art on 11 out of 46 tested instance classes and to outperform the BS from [Pou+20] without local search on 43 out of 46 tested instance classes on average.

CHAPTER 2

Considered Problem

This chapter describes and defines the necessary preliminaries and notations used in the following chapters. Section 2.1 defines the general concept of optimization for scheduling problems. A detailed description of the NWFSP-RT and its notation as well as an example follows in Section 2.2.

2.1 Scheduling Problem

In this work, we define an Scheduling Problem (SP) as the task of scheduling a set of jobs J on a set of resources M to optimize (i.e., minimizing or maximizing) an objective function f in a vast but finite search space S. Function f maps every solution $v \in S$ to a numerical value $f: S \to \mathbb{R}$. How a solution is represented in detail needs to be defined for the concrete scheduling problem at hand. In S we search for an optimal solution $v * \in S$ with $z^* := f(v^*)$ for which $\forall v \in S : f(v^*) \circ f(v)$ (with \circ one of $\{\geq, \leq\}$) holds. We assume w.l.o.g. for the whole work that a minimization problem is given and $\circ \equiv \leq$ and therefore:

$$z^* = \min_{v \in S} f(v) \tag{2.1}$$

$$v^* \in \operatorname*{arg\,min}_{v \in S} f(v) \tag{2.2}$$

Note that an SP may have multiple optimal solutions v^* . Throughout the work if it is stated that solution u is better than solution v, this refers to f(u) < f(v).

Further, a partial solution v to an SP is a solution where not all jobs $j \in J$ are scheduled on all resources M. If f is applied to a partial solution, it returns the current objective value of the partial solution. We define the set F(v) to contain all complete solutions which can be constructed out of the partial solution v. Finally, we define f' to take a set of (partial) solutions $W \subseteq S$ and output the best objective value when applying f to all $v \in W$. We describe states of an (partial) solution in general for an SP. A state o(v) of the solution v is a mapping from a (partial) solution v to a value defined for the specific problem at hand. The representation of states is independent of f(v). Thus, it is possible that two solutions u, v with $u \neq v$ map to the same state o(u) = o(v) and still have different objective values $f(u) \neq f(v)$. The state o(v) needs to define all relevant parts of the problem such that o(v) can be used as starting point for constructing a solution. If w = v ||u| is the concatenation of solution v and a solution u for the instance defined by o(v), and || being the concatenation of vectors, then it must hold that f(w) = f(v) + f(u).

Finally, we define the concept of dominance between solutions, which is later used to filter out partial solutions which definitely do not yield better solutions than others.

Definition 1. (Dominance) A (partial) solution v dominates another (partial) solution u iff $f'(F(v)) \leq f'(F(u))$, and $v \neq u$.

2.2 The No-Wait Flow Shop Scheduling Problem with Release Times

The NWFSP-RT consists of n jobs contained in the set $J = \{1, \ldots, n\}$ and m machines contained in the set $M = \{1, \ldots, m\}$. Every job must pass all m machines in order from 1 to m. For each job j and every machine i the processing time is denoted by $p_{i,j} \ge 0$. The time when a job j is available on the first machine is denoted by r_j . It is prohibited to start with the processing of j before time r_j is exceeded. Jobs are not allowed to wait before being further processed after the processing on machine one is finished. Thus, if waiting is required, the job needs to be delayed as a whole and start its processing on machine one later, such that it can pass all machines without waiting. An early work that introduces the no-wait FSP (NWFSP) without release times is [Pie60]. A solution sto the NWFSP-RT is a permutation of all jobs in J, i.e., s = (3, 2, 1, 5, 4) for an instance with n = 5. Its length is denoted by |s|. The job on position $k \in \{1, \ldots, |s|\}$ of solution s is denoted by s_k . To compute the makespan of a solution, we compute its completion time matrix $C \in \mathbb{R}^{n \times m}$, which can be done in time O(nm). The indices of the C matrix start at one. An entry in the matrix denotes when the according job in the solution s is finished on the respective machine, i.e., $C_{k,i}(s)$ is the finishing time of job s_k on machine i. Note that the parenthesis and the denotation of the solution are omitted when clear from the context. If a solution v does not contain all jobs J, i.e. |J| > |v|, we call it a partial solution and denote the set of all unscheduled jobs by $U = J \setminus v$. The last scheduled job in a solution is denoted by a(v) and if v is empty, a(v) = 0. The makespan of a (partial) solution v is denoted by $C_{\max}(v) = C_{n,m}(v)$, which is the time when the last job of s finishes on the last machine m. Thus, the function $f(\cdot)$ as defined for SP yields for solution v: $f(v) = C_{\max}(v)$. To simplify the computation of the C matrix we define that all accesses to $C_{0,i} = 0$ for $i \in M$, as all machines are empty when starting the first job. The computation is done row-wise from top to bottom. To compute row k of the matrix for the (partial) solution v, a forward sweep and a backward sweep are

used. We treat the leftmost column separately:

$$C_{k,1} := \max(C_{k,1}, r_{s_k}) + p_{s_k,1} \tag{2.3}$$

For the forward sweep we do for all $2 \le i \le m$:

$$C_{k,i} := \max(C_{k,i-1}, C_{k-1,i}) + p_{s_k,i}$$
(2.4)

As we did not propagate the idle times introduced by taking the maximum back, only $C_{k,m}$ is computed correctly. To finish the computation, a backward sweep is needed, iterating over $m-1 \ge i \ge 1$:

$$C_{k,i} := C_{k,i+1} - p_{s_k,i+1} \tag{2.5}$$

Now the calculation of the completion times for the k^{th} row is finished and the $k + 1^{\text{th}}$ row can be computed, until all |v| rows are computed.

We define the start time of a job j on a machine i for a solution s analogous to the C matrix as $T_{j,i}(s)$, and compute it by:

$$T_{k,i} = C_{k,i} - p_{s_k,i} (2.6)$$

At any time during processing, before a machine has processed all jobs but is currently not processing any job, it incurs idle time. For a solution s the sum of all idle times on all machines is the idle time I(s). The idle time for solution s can be computed by:

$$I = \sum_{k \in \{1..|s|\}} \sum_{i \in M} T_{k,i} - C_{k-1,i}$$
(2.7)

In the following, we will provide an example solution for an instance with m = 4 and n = 5. The *p*-Matrix is:

$$p = \begin{bmatrix} 16 & 31 & 54 & 54 \\ 44 & 7 & 52 & 66 \\ 26 & 19 & 65 & 34 \\ 74 & 83 & 94 & 76 \\ 19 & 41 & 31 & 50 \end{bmatrix},$$

and thus, e.g., for job j = 5 and machine $i = 2 p_{5,2} = 41$. The release times are:

$$r = \begin{vmatrix} 60 & 180 & 33 & 17 & 95 \end{vmatrix}$$
.

Thus, one optimal solution to this problem is s = (3, 5, 1, 4, 2) with $C_{max} = 548$. The resulting matrix C(s) is:

$$C(s) = \begin{bmatrix} 59 & 78 & 143 & 177 \\ 114 & 155 & 186 & 236 \\ 155 & 186 & 240 & 294 \\ 229 & 312 & 406 & 482 \\ 423 & 430 & 482 & 548 \end{bmatrix}.$$

7

Solutions of an NWFSP-RT instance can be visualized as Gantt charts. A Gantt chart for s is provided in Figure 2.1. There, jobs are displayed as bars with different colors, and the release times are included as dotted vertical lines. It can be seen that between jobs three and five, idle time exists on every machine. This is since job five needs to start later than the minimum delay time between these jobs due to its release time of $r_5 = 95$. The total delay on machine one is d(s) = 244 and is, due to the release times, greater than if the sum of the according $d_{j,k}$ entries, described in Section 2.2.1, are taken, which yields 202.



Figure 2.1: Gantt chart of solution s = (3, 5, 1, 4, 2) of example sequence. Machines are on the vertical axis and jobs are displayed horizontally. Release times are included as vertical dotted lines in the color of the job they belong to.

2.2.1 Auxiliary Instance Data

During the search, when expanding partial solution v by appending some $j \in U(v)$, calculating the new row for j in the C matrix takes time O(m). We define auxiliary instance data to reduce the time needed when appending one job. The proposed alternative way to compute $f(\cdot)$ is provided in Chapter 4.1.2.

For each job $j \in J$, let $p_j^{\text{tot}} = \sum_{i \in M} p_{i,j}$ be its total processing time over all machines. Moreover, let $q_{i,j} = \sum_{i'=1}^{i} p_{i',j}$ be the aggregated processing time of job j over the first $i' \leq i$ machines. Note that $q_{m,j} = p_j^{\text{tot}}$ holds, and the job's processing on a machine i starts by $q_{i,j} - p_{i,j}$ after the job's overall start on the first machine. The q-matrix is stored within the instance as auxiliary data.

When a job $j' \in J$ is scheduled directly after a job $j \in J$, $j \neq j'$, and we disregard release times, their starting times (on the first machine) always differ by the (minimum) starting time differences:

$$\delta_{j,j'}^1 = \max_{i \in M} (q_{i,j} - q_{i,j'} + p_{i,j'}), \quad \forall j \in J, \forall j' \in J \setminus \{j\}.$$

$$(2.8)$$

Further, we define this difference in the starting times also on all other machines:

$$\delta_{j,j'}^{i} = \delta_{j,j'}^{1} + q_{i-1,j'} - q_{i-1,j}, \quad i \in \{2, \dots, m\}, \forall j \in J, \forall j' \in J \setminus \{j\}.$$
(2.9)

We further define that when speaking of *delay on machine i* we refer to $d_{j,j'}^i = \delta_{j,j'}^i - p_{i,j}$ and when speaking of *delay* in general of the delay on machine one.

Due to the property of the minimum starting time differences it follows that:

Lemma 1. For three different jobs $j, k, l \in J$ the triangle inequality $\delta^i_{j,k} \leq \delta^i_{j,l} + \delta^i_{l,k}$ holds for all $i \in M$.

I.e., it is impossible to reduce the minimum starting time differences between any two jobs j, k by scheduling some other job l in between. This yields for delays $d_{j,k}^i \leq d_{j,l}^i + d_{l,k}^i + p_{l,i}$.

Moreover, we define the total idle time over all machines arising when scheduling job j' as successor of job j in-between these jobs as

$$\omega_{j,j'} = \sum_{i \in M} (\delta^i_{j,j'} - p_{i,j}), \qquad (2.10)$$

where we again disregard release times.

Finally, we define how much job j' must end later on the last machine than job j such that j ends at an earlier or equal time on all machines:

$$\psi_{j,j'} = \max_{i \in M} (q_{i,j} - q_{i,j'}) + p_{j'}^{\text{tot}} - p_j^{\text{tot}}.$$
(2.11)

The above definition is needed to perform a dominance check in Section 4.1.3 efficiently.

2.2.2 Reduction of the NWFSP to the Asymmetric Traveling Salesperson Problem

Augmenting the jobs J with an artificial job 0 with $p_{i,0} = 0$, $i \in M$, we can reduce an NWFSP instance to an ATSP instance with the distance matrix δ^1 . A solution to this ATSP instance in the form of a permutation of the jobs $J \cup \{0\}$ is then transformed back to the respective NWFSP instance by rotating the solution, so that job 0 comes last and finally removing this job. Note that, in general, the minimum starting time differences on any machine δ^i could be used for the reduction. A reduction that uses the d^1 matrix, i.e., δ^1 without the processing times on the first machine, and instead uses two artificial jobs is presented in [Wis72].

Considering this, we also define $\delta_{j,0}^1$, $\delta_{0,j}^1$, $\omega_{0,j}$, and $\omega_{j,0}$, $j \in J$, respectively, in order to deal with the beginning and the end of a schedule. The remaining auxiliary data like δ^i and d^i is calculated accordingly. To simplify future computations we define $r_0 = 0$.

CHAPTER 3

Related Work

In this chapter, we present related work relevant to this work. First, we present the relevant history up to the state-of-the-art for the FSP and its relevant variants in Section 3.1. Section 3.2 explains how the BS framework works and Section 34 describes the BS how it was used in [Pou+20] to solve the NWFSP-RT. Further, Section 3.3 presents relevant work for Machine Learning (ML) methods in the context of combinatorial optimization and the FSP. Finally, the LBS framework is described in Section 3.4.

3.1 The Flow Shop Scheduling Problem and its Variants

The problem of scheduling a fixed set of n jobs J on a fixed set of m machines M, with a different execution time $p_{i,j} \ge 0$ for every job $j \in J$ on every machine $i \in M$, without preemption, and where each machine can only process one job at a time, is called the Job Shop Scheduling Problem (JSP). In the JSP, the order of machines a job is processed on is given by precedence constraints and may vary for each job. Therefore, the jobs can be ordered differently on every machine, and jobs can start on different machines. An early work analyzing the JSP and its anomalies is [Gra66]. There are different objectives for the JSP in the literature, such as minimizing the total tardiness or makespan. Nevertheless, in this work, for all scheduling problems, we aim to minimize the time needed until all jobs are finished on all machines, i.e., the makespan. Therefore, if not stated otherwise, we always refer to this objective.

If the same machine order for all jobs is fixed upfront, we get the Flow Shop Scheduling Problem (FSP). Still, the order on every machine can be different as jobs may overtake each other. The FSP was first introduced by [Joh54] in 1954, together with algorithms for solving the FSP with two machines and a special case with three machines. The Campbell-Dudek-Smith (CDS) algorithm [CDS70] extends the algorithm of [Joh54] to apply to mmachines by clustering the machines into two virtual machines and solving the generated two-machine problem by repeatedly using the algorithm of [Joh54]. Furthermore, [DJ64]

summarizes all assumptions implicitly used in the FSP and suggests to solve the FSP with an exact approach. Later, in 1976 [GJS76] showed that the FSP is NP-complete when minimizing the makespan for $m \geq 3$. For m = 2 there exist efficient algorithms [Joh54]. In the later years, many other variants of the FSP were introduced, and many solution heuristics for FSP and its variants were suggested. In the survey [Gra+79] the variants were first put in a common notation. Later, in [NEH83] the Nawaz-Enscore-Ham (NEH) algorithm is proposed to greedily construct a solution for the FSP. It iteratively tries to insert a job at the best position in a partial solution, starting with the jobs with higher total processing time. In [Tai90] NEH is refined to run in time $O(n^2m)$. Taillard et al. generated a random benchmark set for the FSP [Tai93]. They used different heuristic methods to generate upper bounds together with a self-developed lower bound to select the random instances with the biggest gap between the two bounds. The first part of the lower bound consists of multiple components. It computes for every machine the minimum time any job needs from the start up to this specific machine and similarly the minimum time any job needs from this machine up to finishing on all machines. Further, for every machine, the operations of all jobs on this machine are summed up. To compute the first part of the lower bound, these three components are summed up for every machine, and the maximum value over all machines is taken. The idea behind this computation is that a machine has to process all jobs operations for it and that by taking the minimum time until the computation can start on this machine and until it ends, a lower bound (LB) is derived. For the second part, the total processing time for every job is calculated by summing up the times for all its operations on every machine and taking the maximum of those. This yields a lower bound as every job must be executed until it is complete. Now the maximum of both parts is computed, which again yields a lower bound for the FSP. Since the paper of Johnson [Joh54], the FSP was intensively studied and became one of the most extensively investigated topics in literature [LH05]. Nevertheless, the FSP remained hard to solve, as up to the mid of the 1990s the best Branch and Bound (B&B) algorithms had problems solving instances with 15 jobs and 4 machines [And+97][p. 393]. Due to its NP-hard nature, the FSP was also tackled with many heuristic approaches like Genetic Algorithm (GA), Tabu Search (TS), B&B, and others.

A variant of the FSP is the Permutation Flow Shop Schedueling Problem (PFSP), which is similar but enforces the same processing order on every machine, and this order is subject to optimization. Starting with $m \ge 4$ there might exist non-permutation schedules (i.e. for the FSP, with a specific job order for every machine) which dominate permutation schedules for the PFSP [PSW91]. A comparison over seven lower bounds used by B&B approaches to solve the PFSP while minimizing the makespan is given in [LH05]. The summarized lower bounds reduce the PFSP to a one or two-machine problem, i.e. the constraint that a machine can only process one job at a time is relaxed for all machines except one or two. Release dates, time lags, and delivery dates are computed out of the other jobs' operations to yield good lower bounds. Most of the summarized lower bounds can be computed in polynomial time. Nevertheless, two of the proposed lower bounds can not be computed in polynomial time, as they use a B&B approach or an optimal algorithm to compute a lower bound for the PFSP by solving a relaxed problem.

Another variant of the FSP, where jobs are not allowed to wait before being processed on the next machine, is called the no-wait FSP (NWFSP), or in the notation of [Gra+79]: Fm no-wait $|C_{max}$. This definition results in jobs not overtaking each other, i.e., the execution order of the jobs is the same on all machines as for the PFSP. Thus, every solution for the NWFSP is also a valid solution for the PFSP, with a different objective value though [HS96]. The NWFSP can be reduced to the ATSP [Pie60; RR72; Wis72] as already discussed in Section 2.2.2. The reduction, according to [Wis72], is to compute delays on the first machine between all job pairs and use them as distances in the ATSP together with two artificial jobs used as the first and the last job. As a reduction to the ATSP exists, all lower bounds designed for the ATSP might be used as well for the NWFSP. An example is the lower bound, where in the ATSP the sum of every city's cheapest in- and outgoing arc divided by two is taken, which is similar to the assignment problem described in [Chr72]. A lower bound for the NWFSP which can be calculated in time $O(\max(m, n^2 \log n)n)$ is provided in [KK07]. There the NWFSP is for every pair of machines h and i with $1 \le h \le i \le m$, reduced to a two machine problem, which is optimally solved with the Gilmore and Gomory algorithm (GG) [GG64], see also [Vai03]. If h = 1 or i = m then $O(\max(m, n \log n)n)$ and if h = 1 and i = m then $O(\max(m, \log n)n)$ time is needed. Due to the solving with the GG algorithm, also a solution sequence is returned, which might provide a good initial solution for the mmachine NWFSP to be further improved by metaheuristics. The authors of [KK07] present also a network representation for the NWFSP where the critical (longest) path of the network yields the minimum makespan of the instance.

An extensive survey about the FSP and its variants is provided in [All16], where the NWFSP is discussed in chapter 4.1. The NWFSP was tackled by greedy algorithms, TS and GA [All16] in the last years.

When jobs are only available for processing on the first machine after a job-specific release time r_j , fixed upfront, the problem is called NWFSP with release times (NWFSP-RT), or in the notation of [Gra+79]: (Fm|no-wait, $r_j|C_{max}$). This work focuses on the NWFSP-RT minimizing the makespan without preemption, i.e., processing can not be paused, and jobs must be immediately scheduled on the next machine after finishing on one machine. Setup times are included in the jobs processing time, and sequence-dependent setup times of machines are not used. The NWFSP-RT is a special case of the NWFSP. When adding release dates equal to 0, any NWFSP instance yields an NWFSP-RT instance.

Two lower bounds for the NWFSP-RT with sequence-dependent setup time are provided in [BDG99]. Note that when setting all setup times equal to zero, this problem is equivalent to the NWFSP-RT. The authors first provide a reduction to the ATSP with additional visiting time constraints (ATSP-RT) and a linear programming formulation to solve it. They then relax the formulation and iteratively compute a Lagrangian relaxation by adapting the input vector to the relaxation by a small factor towards the ascending direction until no further improvement happens. To calculate the second lower bound the ATSP-RT is taken again. First, the outgoing arcs of all jobs j are altered such that their weights are equal to the weight of the minimum outgoing arc of job j to any other job except itself. Then all jobs are sorted by increasing release dates, and the objective value of the solution is calculated and returned. This lower bound can be computed in time $O(n \log n)$.

We remark that we assume the jobs to be heterogeneous in the sense that typically they all have different processing times on the machines. The case that jobs are from a small set of different types and jobs of the same type have the same processing times on the machines is, for example, considered in [LS00], where a dynamic variant of the problem is approached in which jobs of predefined types become available only over time. Moreover, job types are considered in [EM95], where a reduction to the TSP is described.

3.2 The Beam Search Framework

The BS framework was originally proposed in the context of speech recognition [Low76]. Basically, BS is an incomplete breath first search utilizing a heuristic function for selecting promising nodes. It is frequently used to generate solutions for combinatorial optimization problems in a short time.

Some variants of the FSP were tackled with BS. The authors of [Pou+20] noticed the research gap for the NWFSP-RT and tackled it with a BS combined with a LS. The FSP with constraints on shared resources was solved by a combination of A* and BS, where only a fraction of all successor nodes is used for the A* search [Pas00]. A hybridization of Ant Colony Optimization (ACO) and BS was used in [Blu05] to tackle the open shop job problem. In [FVVF18] a BS approach was used to generate initial solutions for the PFSP with the objective to minimize the total tardiness. It is guided by a sophisticated heuristic function that weights the total tardiness, earliness, and idle time depending on the number of already scheduled jobs.

In BS a heuristic function $h(\cdot)$ is used to determine which partial solutions are promising and should be kept. In the context of our work, we aim to approximate the cost for completing the partial solution with $h(\cdot)$. The pseudo-code of the basic BS framework is given in Algorithm 3.1. In general, the heuristic function $h(\cdot)$ yields for every possible node a numerical value. During the search, nodes are used to keep track of the current partial solutions. A node v always has a parent node parent (v) and an action a(v) which, when applied to parent(v), results in node v. The beam B is a central element of the search, containing at most β nodes. The parameter β is a strategy parameter that determines how many partial solutions should be kept at most in every layer. Increasing β normally leads to an increase in solution quality but also a higher time and memory consumption. The search proceeds in layers and starts at the initial layer zero, whereas every layer L has its own beam B_L . The search starts from a root node r with parent(r) = \perp and $a(v) = \bot$, representing an empty solution, which is added to the initial layer $B_0 = \{r\}$. At every layer L all solutions in $v \in B_{L-1}$ are branched over all possible actions U(v)and added to the set V_{ext} . Further, V_{ext} is sorted by $h(\cdot)$ and the β -best nodes according to $f(\cdot) + h(\cdot)$ are added to B_L . The search ends when a terminal node t is reached. A



Figure 3.1: A beam search run with $\beta = 2$ for the example given in Section 2.2 with n = 5. The search starts at the root node r denoting the empty solution. It branches over all unscheduled jobs of every node in the beam. The arc labels denote the job which is added to the parents partial solution, i.e. scheduled next. The nodes values in the figure are the $f(\cdot)$ values of the partial solution represented by this node. We assume that $h(\cdot) = 0$ and thus the search greedily selects the two nodes with the best objective value in every layer. They are added to the beam and further branched on. The resulting solution is v = (3, 5, 1, 2, 4) with an objective value of 548.

criterion for determining terminal nodes is problem specific and needs to be defined at hand. In the context of this work, we define similarly to [SB99] that all nodes at layer $L_n = |U(\mathbf{r})| = n$ are terminal nodes, as all actions are taken and U(v) is empty. Every path $\pi_{\rm rt}$ leading from r to t corresponds to a feasible solution. To obtain the concrete solution sequence, one needs to iterate from t over the parent links up to r and concatenate all actions accordingly. As we only store one parent pointer in every node, a path $\pi_{\rm rv}$ of node v represents exactly one (partial) solution, and therefore, we use the term *node* also to denote a (partial) solution. Further, we define all functions defined for partial solutions equally for nodes. Algorithm 3.1 contains two optional parts marked with //optional. These parts are only applicable if a dominance definition for the problem at hand exists, which can be controlled by the parameter dominance enabled. Line 9 to 14 check if a certain state o(u) was already reached by any other node in V_{ext} or if the node u is dominated by any already added node. It follows directly from Definition 1 that no partial solution w which might yield a better overall solution is removed from $W = V_{ext} \cup u$, i.e. $f'(F(W)) = f'(F(W \setminus w))$ for any partial solution w removed, or not added, in lines 9 to 17. Note for clarity that the case that u dominates some node $w_1 \in V_{ext}$ and is itself dominated by some node $w_2 \in V_{ext}$ can not happen because Definition 1 is transitive. Thus w_1 would have already been removed earlier.

Line 23 defines a dominance check, which only adds v if it is not dominated by any $w \in B_i$ and in this case removes all dominated nodes from B_i . It follows again from Definition 1 that $f'(F(B_i \cup v))$ does not get worse by removing dominated solutions. The dominance check is not used in our implementation of BS for the NWFSP-RT, as we only define that nodes can dominate each other if they are in the same state, and thus dominance is already checked in lines 9 to 17.

Assuming that the reached and dominance check are disabled and that creation of one node takes the problem-specific time t_{node} , the BS takes time $O(n\beta t_{node})$. Further, we assume that checking the dominance between two nodes takes t_{dom} time. In theory with a naive implementation the reached check may take up to $O((\beta |U|)^2 t_{dom}) = O((\beta n)^2 t_{dom})$ time with $|U| \leq n$ and $|V_{ext}| \leq \beta |U|$ at most. When implemented with a hash map using the state as a key and a list of all nodes leading to this state as a value, finding duplicate states can be done in a much lower time in the average case, as only a few dominance checks need to be done per node. This is because, at the beginning of the search, many different states exist, and thus the number of nodes the dominance is checked with is small. Towards the end of the search, where the number of different states gets less, also the number of possible actions |U| gets less; thus, fewer than βn nodes are created, and again only some collisions of states happen. Therefore, we argue that the reached check needs significantly less time than $O((\beta n)^2 t_{dom})$ in the average case.

The dominance check checks, in the worst case, the dominance of all nodes in V_{ext} with all nodes in B. As $|V_{ext}| \leq \beta |U|$ and $|B| \leq \beta$, time $O(\beta^2 |U|)$ is needed.

Algorithm 3.1: General BS Algorithm							
Data: β (max. number of nodes in a layer), instance I, number of layers n,							
		heuristic function h , dominance_enabled denotes if dominance check is					
	enabled						
\mathbf{F}	Result: best terminal node						
1 initialize root r							
2 add r to B_0							
3 for $l \in \{1, \ldots, n\}$ do							
4	V_e	$_{xt} \leftarrow \emptyset$					
5	fo	reach node $v \in B_{l-1}$ do					
6		foreach valid action $a \in U(v)$ do					
7		create node u from v and a					
8		dominated = \perp					
9		if dominance_enabled; // optional					
10		then					
11		remove all nodes $w \in V_{ext}$ with $o(u) = o(w)$ or where u dominates					
12		if $\forall w \in V_{ext}$ with $o(u) = o(w) : \neg(w \text{ dominates } u) \land w \neq u$ then					
13		add u to V_{ext}					
14		end					
15		else					
16		add u to V_{ext}					
17		end					
18		end					
19	er	ıd					
20	so	rt V_{ext} according to $f(\cdot) + h(\cdot)$					
21	w	hile $ B_i < \beta$ do					
22		remove first node v of V_{ext}					
23		if dominance_enabled and $\forall w \in B_i : \neg(w \text{ dominates } v) ; // \text{ optional}$					
24		then					
25		remove all $w \in B_i$ dominated by v					
26		add v to B_i					
27		end					
28		$\mathbf{if} \neg \mathbf{dominance_enabled then}$					
29		\mid add v to B_i					
30		end					
31	er	ıd					
32	/ +	<pre> do postprocessing for LBS */ </pre>					
33 e	nd						
34 return best node from B_n (according to $f(\cdot)$)							

Reference Method from Pourhejazy et al. [Pou+20]

We take the work of Pourhejazy et al. [Pou+20] as a reference method. To the best of our knowledge, it is the first work aiming to solve the NWFSP-RT. It provides a mathematical formulation of the NWFSP-RT for solving it via MILP (mixed-integer linear programming) based on the completion time of each job in the solution sequence on each machine, i.e., the C matrix. Further, to solve the NWFSP-RT, it suggests a BS approach, which selects the ξ successors with the lowest idle time of each node of the beam and improves them via random local search (LS). The idle time for a partial solution s is calculated by the sum of the completion times of job s_{k-1} on machine i minus the starting times of job s_k on machine i, for all $i \in M$. Thus, the idle time after the last job in s is finished on machine i is not respected. However, idle time incurred before the first job starts on some machine i is respected.

These ξ nodes are improved by LS and then added to V_{ext} . They are sorted by a lower bound F_{pour} of the completion time. The β best nodes are then added to the beam. The lower bound $F_{\text{pour}}(v)$ of partial solution v is calculated by

$$F_{\text{pour}}(v) = \max\left(f(v), \min_{j \in U} r_j\right) + \sum_{j \in U} p_{m,j}.$$
(3.1)

First, the maximum of the current objective value and the minimum release time of all unscheduled jobs is taken. Then, the processing time of all unscheduled jobs on the last machine is added (later in this work denoted as P_m^U). Clearly this is a LB as $P_m^U = \sum_{j \in U} p_{m,j}$ time is taken in any completion of solution *s* additionally to f(v). Also, the minimum release time of all unscheduled jobs must be awaited before starting the next job. In Section 4.2 we present tighter LBs than this. The two LS procedures proposed are one that exchanges two jobs in a partial solution sequence and another which randomly removes one job and inserts it into all possible positions. They are randomly applied. It is not described in which order these LS procedures are used or how they are selected. The maximum number of LS executions is limited by *nmk* where *k* is a strategy parameter balancing between solution quality and time needed, which was set to k = 0.1 for $n \leq 400$ and k = 0.01 for $n \geq 500$. The number of successors per node in the beam ξ , in their work denoted by α , was evaluated and set to $\xi = 2$ by them.

3.3 Machine Learning for Optimization Problems

Scheduling problems, including the FSP, were also tackled with machine learning. An early work using NNs for solving scheduling problems is [PKL00], which uses a neural network to learn input parameters for some solution algorithm depending on the characteristics of an instance.

A learned weight matrix is used in [LS00] to solve a variant of the real-time FSP, where fixed types of jobs arrive over time. The weight matrix indicates the desirability of one job type following the other. The weight matrix is trained for a specific set of job types on training instances labeled with their optimal sequences.

In [ACE06] an adaptive learning strategy is used to solve the FSP. In every iteration of the algorithm, the jobs operation times are multiplied by a learned weight per operation, and a constructive heuristic, like, NEH or CDS, is executed with the product of each operations time with its corresponding weight. If the results on the test instances improve during training, the weights are reinforced else the weights are altered randomly. If no improvement happens after several rounds, backtracking to the best weights is done.

The authors of [Ram+11] used a back-propagation NN with two hidden layers to generate an initial solution for a Genetic Algorithm (GA) and a heuristic proposed by [Sul00] to solve the FSP to minimize the makespan. The NNs were trained on completely enumerated small instances (the number of jobs $n \in \{5, 6, 7\}$, the number of machines m = 5). The trained model is dependent on the number of machines but independent of the number of jobs. The input features for the NN contain the jobs processing times, the average processing times, and the standard deviation of the job lengths per machine, all of them on each machine. To solve an instance, the jobs are sequentially provided to the NN and sorted by the output values of the NN. Their approach outperformed two other established heuristics for generating the initial solution, namely: NEH [NEH83], and Suliman Heuristic [Sul00] initialized with CDS [CDS70].

The work of [GML20] compared different training algorithms for NNs for solving the FSP with the objective to minimize the makespan. The training setup was the same as in [Ram+11]. The learned networks were applied to instances with up to 100 jobs. The NNs trained with the Levenberg-Marquardt (L-M) algorithm approximated the objective best.

Nevertheless, having optimal results for the training data at hand is not always possible. Therefore, AlphaGo Zero, which is an agent excelling in the games of Go, chess, and shogi, [Sil+17] uses reinforcement learning via self-play to train the NN. It is based on a Monte Carlo tree search in which a deep NN is used to evaluate game states. Training data is continuously generated by simulating games against itself. The approach was able to achieve superhuman performance in the sense that it beat a previous approach, which was able to beat the best human Go players in numerous simulated games.

In recent works, the FSP has been solved via RL approaches. In [RYY21] the FSP is tackled by an RL approach in which a NN is trained, which should output for every machine the best next action given a state. In total, ten actions are designed, which use different heuristic rules for selecting the next job, such as shortest/longest processing time first, or a function computing a critical ratio. The actions even include a neighborhood search, i.e., an action which swaps two neighboring jobs if this reduces the idle time. Input features consist of the maximum, minimum, and mean of the makespan, remaining operations, and the load of machines. To provide immediate reward during training, idle time is used, as increased idle time leads to a worse solution, according to the authors. Their NN was trained for m = 5 and used 24 nodes, three hidden layers, and a logistic

activation function. The input layer has 45 nodes, nine nodes per machine, and the output layer has 50 nodes, ten nodes per machine, representing the probability of a specific action chosen on one machine. The NN was trained using training samples from instances with $n \in \{5, 10, 20, 30\}$ to increase adaptability, labeled with Branch and Bound or GA. It was applied to instances with m = 5 and $n \in \{5, 10, \ldots, 80\}$. One paper that uses machine learning also for bigger instances with hundreds of jobs and twenty machines is [Ni+21], which tackles the Hybrid Flow Shop Scheduling Problem (HFSP), in which the sequential machines of the FSP are replaced by sets (called stages) of parallel machines. The goal is still minimizing the total makespan. First, an initial solution is generated by some heuristic. The proposed framework then uses RL, i.e., a Markov-Decision-Process with search operators as actions, to train the model. It reformulates the problem state into a multi-graph data structure, in which the job to machine assignment of every stage is modeled as a graph. Therefore, a node for every job is created. All jobs assigned to the same machine in the respective layer are connected via directed edges in the executed order. Node features include idle time, waiting time, arriving time, starting time, completion time, and processing time. These graphs are then processed by a graph NN. To overcome the issue that stages might have different dimensions, an Attention-based Weighted Pooling approach is used, which uses a Multi-layer perceptron. The problem is then solved by unsupervised training of the framework, which should map graphs to specific actions by classifying them. Tests were performed on large real-world datasets, and, according to the authors, good-quality solutions were obtained within a limited time.

As the NWFSP can be reduced to ATSP, some machine learning approaches for the ATSP shall also be referred to here. The authors of [KHW19] propose an encoder/decoder approach to solve the TSP and related problems like the Vehicle Routing Problem (VRP). In this approach, an encoder produces an embedding of all input nodes of a single instance. These embeddings are then used by the decoder together with an input mask, masking cities that were already visited and context information, providing information about the first and the last city in the solution sequence. The decoder outputs probabilities about which city should be visited next. These probabilities are used to greedily construct a solution sequence by always visiting the city with the highest probability next. The decoder contains certain considerations for efficiency. The model is trained via selfcritical training using a greedy rollout as a baseline for estimating the total cost. This is similar to the self-play improvement of [Sil+17]. Also the authors of [Jos+21] use a encoder/decoder approach. Nevertheless, to speed up the computation, they explicitly sparsify the graph by considering only the k-nearest neighbors of each node. Further, beam search with different widths is compared with greedy search as well as supervised learning and reinforcement learning approaches. The ability of the models and approaches to generalize larger unseen instances than they have been trained on, is evaluated on TSP instances with up to 200 cities. The euclidean version of the TSP is tackled in [JLB19] by using a graph convolutional network. The authors propose to use the k-nearest neighbors graph during computation to reduce the computational complexity.

3.4 The Learning Beam Search Framework

The LBS framework was introduced by [HR21] and the description in this section follows their work. LBS is a general framework for learning a heuristic function to guide a BS. The framework uses a reinforcement learning (RL)-like approach where a BS with the currently learned heuristic function is used to label training data during self-learning. Further, it suggests an ML model and defines the interactions needed during training between the ML model, the BS and its other components like the instance generation and the replay buffer. The authors showed for the Longest Common Subsequence Problem (LCS), where the goal is to find the longest common supersequence of a set of strings, and related problems that the approach is highly competitive with the state-of-the-art of manually designed beam search heuristics. They could find several new best-known solutions so far. For solving the LCS, the NN gets the sorted remaining string lengths and the minimum number of remaining letters as input features.

The detailed description of the LBS Framework follows [HR21]. The pseudo-code of the algorithm for learning the heuristic function is provided in Algorithm 3.2. The training happens in iterations, whereas in each iteration, first, a random representative problem instance is generated, a BS with training data generation is performed, and finally the (re-)training of the ML model is performed. The total number of iterations z performed is given and used as a stopping criterion. Generating training instances is problem specific and needs to be defined for the problem at hand. Nevertheless, the generated instances should represent the problem instances that should be solved later. The BS calls generate in every iteration α training samples on average. To do so, nodes are selected as training data at an adaptive probability α/n_{nodes} . The number of training samples α , which should be generated in each iteration, is therefore divided by the number of generated non-dominated nodes n_{nodes} during one whole BS run. Initially, for the first BS run $n_{nodes} = 0$ and thus, no training data is generated, but the number of produced nodes is counted, and n_{nodes} is updated for successive iterations. In every following iteration, n_{nodes} is again updated to reflect the average number of all nodes produced over all so far performed LBS iterations. The nodes selected as training samples are labeled by a nested BS (NBS) call, again using the currently learned heuristic function as guidance. starting at the given node, i.e., using it as root node r. The NBS is performed with the beam width β' . The pseudo-code for data generation is shown in Algorithm 3.3. This algorithm is included into the general BS (Algorithm 3.1) at Line 32 when called with data generation. It is assumed that all additionally needed arguments are passed to the BS algorithm and are thus available and that Inst(v) creates an instance starting at node v. All labeled training samples are added to a replay buffer R, which can hold a given number of samples ρ . It is realized as a ring buffer. Thus if the number of training samples exceeds ρ , the oldest samples are overwritten. The training of the ML model starts when the buffer contains a minimum number of samples γ . As an ML model, we propose to use a NN, which we train by randomly selected mini-batches from R. In every iteration, one training step is performed, and all samples of R are included in randomly selected mini-batches. In general, any appropriate ML model can be used within the Algorithm 3.2: General LBS Algorithm [HR21]

Data: beam width β , beam width for NBS β' , number of layers n, number of training iterations z, size of the replay buffer ρ , buffer size with which training starts γ

Result: learned guidance function for BS

1 $h \leftarrow$ initialize guidance function

2 $R \leftarrow \emptyset$ // replay buffer: size ρ ring buffer

3 for z iterations do

4 $I \leftarrow$ randomly generated problem instance

- 5 BS with training data generation on instance I
- 6 | if $|R| \ge \gamma$ then
- 7 (re-)train h with data from R

8 end

9 end

```
10 return h
```

Algorithm 3.3: Data generation for LBS [HR21]. This piece of code is included in Algorithm 3.1 at line 32, if BS is called with data generation. For simplicity, it is assumed that all arguments additionally needed are passed to the BS algorithm.

Data: beam width for NBS β' , replay buffer R, number of samples to generate α , number of expected nodes n_{nodes} , guidance function for the BS h **1** $N = V_{ext} \cup B_i$ **2** for $v \in N$ do **3** | if BS with data generation $\wedge \operatorname{rand}() < \alpha/n_{nodes}$ then **4** | $t \leftarrow BS$ without data generation $(\beta', \operatorname{Inst}(v), n - |v|, h) //$ NBS call **5** | add training sample (v, f(t)) to R **6** | end **7** end

framework. The training approach must, of course, be adapted appropriately.

The computational complexity of one NBS call depends on the numbers of nodes that are expanded, which are at most β' , the problem-specific time for expanding and evaluating one node t_{node} and the height of the search tree, which is n. Thus, one NBS call requires time $O(\beta' n t_{\text{node}})$. The overall time complexity for one LBS run is dependent on the number of iterations z, their beam width β , and the number of NBS calls performed per iteration, which is equal to α on average. This yields an average time complexity of $O(z(\beta + \alpha \beta')nt_{\text{node}})$.
$_{\rm CHAPTER}$

Solving the NWFSP-RT with Beam Search

In this section we provide the missing parts to solve NWFSP-RT with the general BS algorithm from Section 3.2. Therefore, in Section 4.1 we describe states and nodes for the NWFSP-RT, the branching scheme, and give a dominance definition for the case when two solutions have an equal state. We present our novel lower bound together with three other lower bounds for the NWFSP-RT in Section 4.2.

4.1 Beam Search for the NWFSP-RT

To solve the NWFSP-RT with BS, we define all parts needed for the general BS framework presented in Section 3.2.

State: As *state* for a partial solution v of the NWFSP-RT we take the set of unscheduled jobs U(v) and the last scheduled job a(v), which yields o(v) = (a(v), U(v)). Thus, the state is independent of the concrete permutation of the jobs in v. Further, multiple solutions map to the same state if they already scheduled the same jobs and have the same time on all machines relative to the start time at the first machine, which is ensured by having a(v) in the state. Note that as we only use a(v) and not the completion times of a(v) relative to the completion time of a(v) on the first machine, it might happen that two jobs theoretically end in the same state, but are mapped to different ones in our framework. Nevertheless, this is unlikely, and our dominance filters such nodes. Therefore, we take the performance advantage of storing only a single value.

Node: A node v contains following necessary parts:

• The state o(v) of the current solution:

- $-a(v) \in J \cup 0$...action leading from the parent to node v, i.e., the job actually scheduled last. Job 0 is only used in the root node.
- $U(v)\subseteq J$. . . set of still unscheduled jobs. In the implementation stored as a bit vector.
- parent(v) ... a pointer to its parent node or \perp if it is the root node.
- g(v) = f(v) ... the costs of the respective partial solution so far, which is the makespan of the partial solution.
- h(v) ... a heuristic value that should estimate the cost-to-go, i.e., the expected further increase of the makespan for scheduling the remaining jobs in U(v) in an optimal fashion.

If it is clear from the context to which node we refer, we only write a, U, g, and h.

If two nodes are in an equal state, we keep the one with the smaller objective value $g(\cdot)$. This selection does not degrade the overall solution quality, as only the best node in each state is needed to reach the optimal solution. If both have an equal makespan $g(\cdot)$, an arbitrary one is kept.

For the root node r we set $U(\mathbf{r}) = J$, parent(r) = \bot , a(v) = 0, $g(\mathbf{r}) = 0$, and $h(\mathbf{r})$ calculated respectively.

A node is denoted a *terminal node* or *terminal*, if |U(v)| = 0, i.e., all jobs are scheduled. Terminal nodes only exist at layer n, where all jobs are scheduled. Layer n might have multiple terminal nodes, as not all nodes map to the same state due to the state containing the last scheduled job $a(\cdot)$.

An efficient method to calculate $f(\cdot)$ is provided in Section 4.1.2. The definition of the heuristic function $h(\cdot)$ remains open as we want to learn it with the LBS framework. Nevertheless, an example would be using a lower bound as done in [Pou+20].

4.1.1 Branching

In our implementation, we branch every node over its unscheduled jobs and append them to the end. Note that an alternative branching scheme suggested in [JSW06] would be to insert the next job out of an ordered list of unscheduled jobs into every possible position of the current partial solution sequence. Nevertheless, we argue that BS does not benefit from this alternative branching. This is because during a search, two jobs yielding a good objective value might reside next to each other, which might not occur in any optimal sequence next to each other. It would be even worse if two jobs are in the correct order for an optimal solution, but that are missing some other jobs in between them, are yielding a bad objective value. Thus, the BS heuristic function might have a hard time filtering out partial solutions which are locally optimal but yield a worse objective value when completed than other partial solutions. When branching over all unscheduled jobs and appending at the end, locally optimal solutions might also happen. However, it is at least guaranteed that these locally optimal parts must be included in the final solution of these nodes. An example, were we branch over all unscheduled jobs of every node in the beam is provided in Figure 3.1.

4.1.2 State Transition

To branch from node u to node v by scheduling job $j \in U(u)$ we set parent(v) = u and a(v) = j. In theory, these operations are enough to build up the search tree for the BS. Nevertheless, our nodes also store auxiliary data, so we need to update it depending on the parent node's values. Thus, we set $U(v) = U(u) \setminus j$. For calculating the objective value g(v) we use the following equation:

$$g(v) = g(u) - p_{a(u)}^{\text{tot}} + \max(\delta_{a(u),a(v)}^{1}, r_{a(v)} - g(u) + p_{a(u)}^{\text{tot}}) + p_{a(v)}^{\text{tot}}.$$
(4.1)

Note that this calculation can be performed in constant time O(1), other than calculating a new row for the C matrix that needs time O(m). The used heuristic function calculates the value for h(v).

If disregarding the calculation of h(v), the state transition is performed in time O(n), as U must be copied and has $|U| \leq n$ entries.

4.1.3 Dominance Definition for the NWFSP-RT

The dominance check needs to cover cases where nodes are on the same layer, i.e., the partial solutions have the same length. We define dominance for two nodes u and v of the NWFSP-RT only for the case where they have equal $U(\cdot)$.

Definition 2 (Dominance for partial solutions of the NWFSP-RT). Partial solution u dominates partial solution v iff $u \neq v \land U(u) = U(v) \land g(v) - g(u) \ge \psi_{a(u),a(v)}$.

This check can be done in O(1) time except for the comparison of U(u) and U(v). As we use bit vectors to represent the unscheduled jobs, we may assume that also the latter can be done efficiently in practice, although the time complexity is in O(n).

To perform the dominance check on a whole set of solutions V_{ext} obtained as extensions of the current beam, we store the nodes in a hash map with $U(\cdot)$ as keys. In this way, all groups of nodes with the same unscheduled nodes are determined efficiently, and O(n)time pairwise dominance checks must be done only within each group.

4.2 Lower Bounds for a Beam Search Node

This section describes lower bounds for the NWFSP-RT used in our implementation. In general, it needs to be said that every lower bound for the NWFSP and the FSP is also a lower bound for the NWFSP-RT, as those problems generally yield smaller objective

values for the same *p*-Matrix than the NWFSP-RT [LH05; PSW91]. Nevertheless, such LBs perform worse than bounds specifically designed for the NWFSP-RT as the second can use all constraints of the problem.

We will start with the LB used in [Pou+20] for the NWFSP-RT and continue by adapting the LB from [Tai93] used for generating hard FSP instances. Then we present an LB resulting from the idea of reducing the NWFSP to the ATSP, and finally, we combine the ideas of the former LBs into a new LB denoted as ITLB.

4.2.1 LB from Pourhejazy et al. [Pou+20] (PLB)

This section refers to the LB on the completion time presented in [Pou+20] formula 12. The provided formula is translated into our notation. The idea of this LB is to take the minimum release time and add up the time all unscheduled jobs take on the last machine. The formula is presented in Equation 3.1. It can be calculated in time $O(|U|) \leq O(n)$.

4.2.2 LB from Taillard (TLB) [Tai93]

The LB proposed in [Tai93] was used for generating random test instances for the PFSP which have large gaps between this LB and the best-known solution determined with different approaches. In this section, we adapt it to the NWFSP-RT and refine it by incorporating release times and, to a small extent, also delays, but do not consider the no-wait aspect yet. The bound consists of two parts that both yield valid LBs, from which we take the maximum. The first part calculates a bound for a shortest possible schedule. More specifically, let η_i be the minimum time needed before machine *i* starts to work relative to $g - p_a^{\text{tot}} + p_{1,a}$, i.e., from the time when *a* finished on machine one, and λ_i the minimum amount of time that machine *i* remains inactive after it finished all jobs. Note that $q_{1,j} = p_{1,j}, \forall j \in J \cup \{0\}$. Before we start describing the bound we define how to calculate the job-specific release time relative to the starting time of job a(v) of solution v:

$$r_j^{\text{rel}}(v) = \max(0, \ r_j - (g(v) - p_{a(v)}^{\text{tot}}))$$
(4.2)

To calculate η_i for node v we use:

$$\eta_i = \max(\min_{j \in U} (q_{i,j} - p_{i,j} + \max(0, r_j^{\text{rel}} - p_{1,a})), \ q_{i,a} - p_{1,a})$$
(4.3)

The above formula's maximum consists of two parts. In the first part, we aim to find the shortest possible time any job $j \in U$ needs until it starts on machine *i*. If release times are ignored, this is $q_{i,j} - p_{i,j}$. To incorporate release times we use $\max(0, r_j^{\text{rel}} - p_{1,a})$. In this formula the release time has to be shifted such that it is also relative to $g - p_a^{\text{tot}} + p_{1,a}$, i.e., does not include $p_{1,a}$. The maximum with zero is taken for the cases where the release time has already passed. The second part calculates the time when machine *i* becomes available in the current state, i.e., when job *a* finishes on machine *i*. As we defined to calculate the lower bound from the time where *a* finishes on machine one, we

have to subtract the time a takes on the first machine $p_{1,a}$ from $q_{i,a}$ as it is included there.

The minimum amount of time that machine i remains inactive after it finished all jobs is calculated for partial solution v according to the following formula:

$$\lambda_i = \min_{j \in U} p_j^{\text{tot}} - q_{i,j} \tag{4.4}$$

The calculation of λ_i uses the pre-computed values of $p_j^{\text{tot}} - q_{i,j}$ which denote the time machine *i* is idle after executing job *j* if job *j* is the last job in the solution sequence.

To get the first part of the TLB we take the maximum of the per machine sum of $\eta_i + \lambda_i$ together with the total processing time $P_i^U = \sum_{j \in U} p_{i,j}$ on machine *i* for all unscheduled jobs. We further need to correct LB₁ by $-(p_a^{\text{tot}} - p_{1,a})$ such that it can be directly added to g(v) for solution v:

$$LB_{1} = \max(0, \max_{i \in M} (\eta_{i} + P_{i}^{U} + \lambda_{i}) - (p_{a}^{tot} - p_{1,a}))$$
(4.5)

This is a valid LB as all jobs need to be processed on any machine *i*. Further, also the minimum times to start computation on machine *i* and minimum idle time afterward, and the minimum release time are always implied. The correction by $-(p_a^{\text{tot}} - p_{1,a})$ is necessary as we need to account for the fact that η_i starts after the finishing of *a* on machine one. Through this correction, LB₁ can be directly added to *g*. Lower bounds are generally defined to yield values ≥ 0 , so we take the maximum with zero.

The second part of the LB calculates the sum of the total processing times, and release times and takes the maximum over all unscheduled jobs:

$$LB_2 = \max(0, \ \max_{j \in U} (p_j^{tot} + \max(0, r_j^{rel} - p_{1,a})) - (p_a^{tot} - p_{1,a}))$$
(4.6)

Again, this is a valid LB as every job's operations need to be processed sequentially. One operation cannot start before the preceding one is finished, and every job can only be scheduled when its release time is exceeded. As above, the relative release time needs to be corrected, and the result corrected such that we can directly add it to g.

An improvement of LB₂ follows from Lemma 1. We can take the maximum of r_j^{rel} and $\delta_{a,j}^1$, as any job $j \in U$ must always follow job a somewhere later in the sequence:

$$LB'_{2} = \max(0, \max_{j \in U}(p_{j}^{\text{tot}} + \max(r_{j}^{\text{rel}}, \delta_{a,j}^{1}) - p_{1,a}) - (p_{a}^{\text{tot}} - p_{1,a}))$$
(4.7)

Because $\delta_{a,j}^1$ always includes $p_{1,a}$, which we defined not to be contained in the bound, we have to subtract $p_{1,a}$ from the dominating factor of release time and delay.

Finally, by combining these two bounds we get

$$TLB = \max(LB_1, LB_2'). \tag{4.8}$$

27

Note that for the root node r, all p_a^{tot} and $p_a^{\text{tot}} - q_{1,a}$ are zero as we defined $a(\mathbf{r}) = 0$ and 0 denotes the artificial job. Further, note that all processing times of the artificial job 0 are zero. For improvements upon this LB, see ITLB in Section 4.2.4

Runtime Analysis: Calculating η_i and λ_i is in $O(|U|) \leq O(n)$ as the minimum of all $j \in U$ is taken. Thus calculating LB₁ takes time O(nm) as it iterates over all machines. Calculating LB₂' takes $O(|U|) \leq O(n)$ as iterating over all jobs is necessary. Note that calculating P^U takes time O(nm), which could be reduced by storing it in the node and calculating it in the state transition incrementally in time O(m). Thus, the total time needed is O(nm).

4.2.3 Delay-Based LB (DLB)

The idea for DLB stems from the reduction of NWFSP to the TSP [Chr72]. The delays and the artificial job play a central role in this reduction. Looking at the reduction to the ATSP designed by us earlier, we aim to sum up the minimum cost of all incoming arcs for every $j \in U \cup \{0\}$ and all outgoing arcs $j \in U \cup \{a\}$. When adding release times to this bound, one must consider carefully that if the minimum starting time difference $\delta_{a,j}^1$ is $\geq r_j$, no additional idle time is introduced by scheduling job j next. In theory, any machine's minimum starting time differences could be taken, and the calculation altered accordingly. More specifically, for the NWFSP the makespan of solution s represented by node v is calculated as $\sum_{k=1}^{|s(v)|-1} \delta_{s(v)_k, s(v)_{k+1}}^1 + \delta_{a,0}^1$.

Regarding the case when v = r, we highlight that a(r) = 0. We define the temporary distance matrix δ' to use release times when they exceed the delays and subtract $p_{1,j}$ as we want to consider the operation times separately by adding P_i^U later:

$$\delta_{j,k}' = \begin{cases} \max(\delta_{j,k}^{1}, r_{k}^{\mathrm{rel}}) - p_{1,j} & \text{if } a = j \\ \delta_{j,k}^{1} - p_{1,j} & \text{otherwise} \end{cases}, \forall j, k \in J \cup \{0\}$$
(4.9)

A bound for the NWFSP-RT considering incoming arcs which can be directly added to g would be:

$$LB_{in} = \max(0, \sum_{j \in U \cup \{0\}} \min_{k \in U \cup \{a\} \setminus \{j\}} \delta'_{k,j} + P_1^U - (p_a^{tot} - q_{1,a}))$$
(4.10)

Subtracting $p_a^{\text{tot}} - q_{1,a}$ in the above formula is necessary as $\delta'_{j,k}$ is calculated on machine one and starts at the time when *a* finishes on machine one. If doing this calculation e.g. on machine *m* the subtraction is not needed as it would result in $-p_a^{\text{tot}} + q_{m,a}$ with $q_{m,a} = p_a^{\text{tot}}$.

Further, we define a second bound based on the outgoing arcs when thinking in terms of an ATSP:

$$DLB = \max(0, \sum_{j \in U \cup \{a\}} \min_{k \in U \cup \{0\} \setminus \{j\}} \delta'_{j,k} + P_1^U - p_a^{\text{tot}} + p_{1,a})$$
(4.11)

28

This time, all $j \in U \cup \{a\}$ have outgoing arcs which need to be considered and which might lead to unscheduled jobs, including the artificial job but excluding self-loops $k \in U \cup \{0\} \setminus \{j\}$.

To compute the final bound we take the maximum of the above parts:

$$LB_{delay} = \max(LB_{in}, LB_{out}) \tag{4.12}$$

The lower bound LB_{delay} can be directly added to g.

Note that this bound can be further improved by calculating it on all machines $i \in M$. Additional considerations for the release times must be taken. Further, $q_{i,a}$ must be substracted instead of $q_{1,a}$ or $p_{1,a}$, and the maximum over the resulting bounds on all machines must be taken.

Runtime Analysis: To calculate LB_{in} and LB_{out} we iterate over all combinations of unscheduled jobs including the artificial job. Therefore, calculating LB_{delay} takes time $O(|U|^2) \leq O(n^2)$.

4.2.4 Improved Taillard Lower Bound (ITLB)

This section aims to improve TLB by incorporating a bound on the minimum delays on each machine. The ITLB can be seen as using DLB (with adaptions) within the calculation of TLB once for every machine. The bound is calculated from the end of job *a* on machine one until the end of all jobs on machine *m*. Thus, we correct it by subtracting $(p_a^{\text{tot}} - q_{1,a})$ such that it can be directly added to *g*. Note that the artificial job 0 is not used in calculating this bound, as we consider the processing times for finishing the last job by λ_i .

We adapt the calculation of the TLB for η_i accordingly for partial solution v to not add up the whole release time but only the part of the release time that exceeds the necessary delay. Further, we add the minimum delay on this machine calculated by d(v, i, j) stated later. In the second part of the maximum function, we calculate the minimum delay over all jobs, as it is not given which job will be the next one scheduled. All other parts of the formula remain the same as for the TLB.

$$\eta_{i} = \max(\min_{j \in U} (q_{i,j} - p_{i,j} + \max(\delta_{a,j}^{1}, r_{j}^{\text{rel}}) - p_{1,a} + d(v, i, j)),$$

$$q_{i,a} - p_{1,a} + \min_{j \in U} d(v, i, j)) \quad (4.13)$$

The calculation of λ_i is not altered, i.e.,

$$\lambda_i = \min_{j \in U} p_j^{\text{tot}} - q_{i,j}. \tag{4.14}$$

To calculate d(v, i, j) we use the following formula to first calculate a delay bound $d^{in}(v, i, j)$ on the incoming arcs. This bound only considers the additional time needed (delay) on the respective machine but ignores the processing times as they are added

later. Thus, the processing times must always be subtracted from the minimum starting time differences. We assume that every unscheduled job has an incoming arc, where we consider job j separately as it must follow job a:

$$d^{\text{in}}(v,i,j) = \sum_{j' \in U \setminus \{j\}} \min_{k \in U \setminus \{j'\}} \delta^{i}_{k,j'} - p_{i,k}$$
(4.15)

To calculate a bound considering only outgoing arcs $d^{out}(v, i, j)$, we assume that every unscheduled job has one outgoing arc, except the last job in the sequence. To correct the delay the last job in the sequence introduces, we subtract the maximum of the delays summed up before. Again, the processing times are subtracted from the delays:

$$d^{\text{out}}(v,i,j) = \sum_{j' \in U} \min_{k \in U \setminus \{j,j'\}} (\delta^{i}_{j',k} - p_{i,j'}) - \max_{j' \in U} \min_{k \in U \setminus \{j,j'\}} (\delta^{i}_{j',k} - p_{i,j'})$$
(4.16)

In the above formula it is again clear that job j must follow job a and j is therefore considered separately. The second part of the formula calculates the sum over all outgoing arcs for all unscheduled jobs. Note that we now added one outgoing edge too much. To correct this we remove the most expensive outgoing edge counted before by subtracting it.

Our final delay bound is thus:

$$d(v, i, j) = \max(d^{\text{in}}(v, i, j), d^{\text{out}}(v, i, j))$$
(4.17)

To get the first part of the LB we proceed as for the TLB and derive a bound which can be directly added to g:

$$LB_{1} = \max(0, \max_{i \in M} \eta_{i} + P_{i}^{U} + \lambda_{i} - p_{a}^{\text{tot}} + p_{1,a})$$
(4.18)

The second part of the TLB, LB'_2 , was already improved before and therefore stays the same.

Finally, we get:

$$ITLB = \max(LB_1, LB_2') \tag{4.19}$$

Note that in the case of the root node $p_{a(\mathbf{r})}^{\text{tot}}$ and $q_{1,a}$ evaluate to zero as the duration of the artificial job which comes first is zero on all machines.

This bound clearly dominates TLB as it contains all its components. It also dominates DLB, as it contains all delays. It was also tested on 100 randomly generated instances (n=10, m=20) with a BS where TLB, DLB, and ITLB were calculated for every node that ITLB outperforms the others.

To speed up the calculation of d(v, i, j) and calculate it for all $j \in U$ in time $O(n^2)$ together, we do the following:

- For calculating $d^{\text{in}}(v, i, j)$ for all $j \in U$ at once: First, we compute the minimum $x_k = \delta^i_{k,j'} p_{i,k}, \forall k \in U$ while discarding self loops (i.e. $\delta^i_{j,j} p_{i,j}$) and then take the sum of it which we denote here by sum. To compensate for counting job j too much in the sum we now calculate sum $-x_j, \forall j \in U$. The whole computation happens in time $O(n^2)$ for all $j \in U$.
- For calculating $d^{\text{out}}(v, i, j)$ for all $j \in U$ at once: We compute the two minimum delays $x_{j'} = \delta^i_{j',k} p_{i,j'}, \forall k \in U$ and store them together with their indices denoted by index_{j'}. The minimum is denoted by $x_{j',1}$ and the second smallest value by $x_{j',2}$, similarly for indices. We then take the sum sum $= \sum_{j \in U} x_{k,1}$ and store it in a repeated way as vector sum_j, $\forall j \in U$. We further copy $y = x_{:,1}$. Now we have to account for the cases where the job j is also the minimum object and thus contained in sum_j. Therefore, we do sum_{index_{j,1} = sum_{index_{j,1} $x_{j,1} + x_{j,2}$ for all $j \in U$. To account for the subtraction of the maximum values we compute $y_{\text{index}_{j,1}} = \max y_{\text{index}_{j,1}}, x_{j,2}$. This calculation happens in time $O(n^2)$ for all $j \in U$.}}

Note that due to performance reasons and to stay within reasonable time limits, we will only compute d^{in} during evaluation. Even though the theoretical time complexity of d^{in} and d^{out} is the same, it takes a significantly higher time to compute d^{out} in our implementation.

Runtime Analysis: Calculating λ_i is in $O(|U|) \leq O(n)$ as the minimum of all $j \in U$ is taken as for the TLB. Calculating d(v, i, j) for all $j \in U$ at once is in $O(|U|^2) \leq O(n^2)$ and thus calculating η_i takes time $O(|U|^2)$. Further, all computations must be done for every machine. This finally yields that calculating ITLB takes time $O(m|U|^2) \leq O(mn^2)$.

CHAPTER 5

Learning Beam Search for the NWFSP-RT

To apply the LBS framework to the NWFSP-RT, we define and describe all missing parts of the general LBS framework as described in Section 3.4. We aim to learn a heuristic function $h(\cdot)$, approximating the minimum makespan increase when completing a partial solution. Therefore, we first describe some general details in Section 5.1. Then, the training data generation is described in Section 5.2. Section 5.3 proposes three NN types used, and Section 5.4 presents the used features and states the idea behind each. Those features are combined to observations in Sections 5.5.

5.1 General Details

The description of the time complexity of one LBS run in Section 3.4 leaves the description of the time expanding and evaluating one node takes t_{node} open as it is problem specific. This time is dependent on the time the state transition described in Section 4.1.2 takes which is O(n) as well as the time needed for evaluating it t_{eval} . Further, each node might have up to n children and thus $t_{node} = O(n^2) + t_{eval}$. The evaluation time t_{eval} is dependent on the time needed to evaluate $h(\cdot)$, and the time needed to compute the used features which are further described in Sections 5.3, and 5.4.

5.2 Training Instance Generation

We generate the training instances to match the characteristics of the benchmark instances of [VRF15]. We note again that we assume the jobs to be different from each other such that they can not be clustered into types of jobs. This assumption was experimentally verified by giving a tolerance area around each processing time on each machine and trying to find jobs for which all their processing times on all machines fit into the tolerance area. It was found that the assumption also holds if the criterion is ignored for up to two machines using a tolerance area of 40. Thus, we assume that the times of all operations of all jobs $p_{i,j}$ are equally distributed over [0, 99] and independently chosen. Further, the release times r_j for an randomly created instance are calculated by the formula of [ZLC15]:

$$r_j = 0.05 \sum_{j \in J} p_{j,1}.$$
(5.1)

Note, that Formula 5.1 was also used to adapt the benchmark sets we use from the FSP to the NWFSP-RT in [Pou+20].

5.3 Neural Networks

We propose different types of NNs to be used within the LBS framework. The inputs used for these NNs are described in Section 5.4 and combined to observations that are later used during evaluation in Section 5.5.

5.3.1 Multi-Layer Perceptron

The use of a MLP within LBS for approximating the costs for completing a partial solution is proposed in [HR21]. We use a dense network of neurons with two hidden layers and ReLu activation functions in our work. The output layer consists of a single neuron without an activation function. The number of nodes in the hidden layers is set to 20 as it was done in [HR21] and as preliminary tests with {10, 20, 30} showed no significant difference. Preliminary tests showed no significant difference in the results when adapting this value. When ignoring the weights between the input layer and the first hidden layer, this network has 420 learnable weights. It is trained using the mean squared error (MSE) as a loss function.

5.3.2 Aggregated Job Neural Network (AJNN_{add})

The AJNN_{add} consists of an internal MLP, further denoted as M_1 . An overview of its structure is given in Figure 5.1. The observation vector consists of a global part x^{glob} and job $j \in U$ specific features x_j^{job} . The job specific part x_j^{job} is then split further into $x_j^{\text{job}} = (x_j^{\text{job,glob}}, x_j^{\text{job,in}}, x_j^{\text{job,out}})$. The first part $x_j^{\text{job,glob}}$ contains specific features for job j. The other parts $x_j^{\text{job,in}}, x_j^{\text{job,out}}$ each contain κ features dependent on the nearest pread succeeding jobs of j.

The nearest jobs are determined by the ω matrix, containing idle time. For every job j in |U| the κ nearest jobs must be selected, which takes time $O(|U|\min(\kappa, \log |U|))$ for one job. Note that when κ gets big enough, it is more efficient to sort all jobs in U instead of iteratively selecting the nearest one. Thus selecting the κ nearest pre- and succeeding jobs for all $j \in U$ takes time $O(|U|^2 \min(\kappa, \log |U|)) \leq O(n^2 \log n)$. These 2κ nearest jobs are stored in vectors of sorted lists $\mathcal{N}_j^{\text{red,in}}$ and $\mathcal{N}_j^{\text{red,out}}$ for each job j. In

the implementation, the values in $\mathcal{N}^{\text{red,in}}$ and $\mathcal{N}^{\text{red,out}}$ refer to the jobs' indices when taking only jobs from U, sorting them by their id, and numbering them with 1 to |U|.

If there are too few jobs left to fill up $x_j^{\text{job,in}}$ and $x_j^{\text{job,out}}$, they are filled up with the last job's values. We do similarly for $\mathcal{N}_j^{\text{red,in}}$ and $\mathcal{N}_j^{\text{red,out}}$ by filling the values up with the last job. When only one job remains, we use it in the feature vectors $x_j^{\text{job,in}}$, and $x_j^{\text{job,out}}$ and fill $\mathcal{N}_j^{\text{red,in}}$, and $\mathcal{N}_j^{\text{red,out}}$ with ones as the index of the job left is one.

Take a sample instance with n = 5, $U = \{1, 2, 4\}$, $\kappa = 1$ and four job specific features $|x^{\text{job}}|$. Further, assume that $\mathcal{N}_1^{\text{red,in}} = [2]$, $\mathcal{N}_1^{\text{red,out}} = [2]$, $\mathcal{N}_2^{\text{red,in}} = [4]$, $\mathcal{N}_2^{\text{red,out}} = [1]$, $\mathcal{N}_4^{\text{red,in}} = [2]$ and $\mathcal{N}_4^{\text{red,out}} = [1]$. We also assume that $x^{\text{glob}} = (10, 20, 30)$. Further, we assume $x_{1,2}^{\text{job,in}} = (8)$, $x_{1,2}^{\text{job,out}} = (9)$, $x_{2,4}^{\text{job,in}} = (10)$, $x_{2,1}^{\text{job,out}} = (11)$, $x_{4,3}^{\text{job,in}} = (12)$, and $x_{4,1}^{\text{job,out}} = (13)$. Finally, assume $x_1^{\text{job,glob}} = (22, 33)$, $x_2^{\text{job,glob}} = (44, 55)$, and $x_4^{\text{job,glob}} = (66, 77)$. Then $x^{\text{job}} = (22, 33, 8, 9) ||(44, 55, 10, 11)||(66, 77, 12, 13)$, where || denotes the concatenation of vectors. The nearest job vectors are $\mathcal{N}^{\text{red,in}} = (2, 4, 2)$ and $\mathcal{N}^{\text{red,out}} = (2, 1, 1)$. In our implementation, we map those vectors to (2, 3, 2) and (2, 1, 1) respectively, representing the jobs' indices in U when sorted by their ids. These values can be used as indices in the feature vector x^{job} in the NN_{nearest} in Section 5.3.3.

If the $AJNN_{add}$ needs to generalize to an instance with more jobs than it was trained on, M_1 is called more often but always with the same number of input features, and all its results are summed up.

The $AJNN_{add}$ aims to approximate the increase of the makespan when completing the solution. The $AJNN_{add}$ is trained the same way as the MLP in Section 5.3.1 is, using the mean squared error (MSE) as a loss function. The loss is calculated between the result value of the $AJNN_{add}$ and the label. We use for M_1 the same setup as in Section 5.3.1, a dense network of neurons with two hidden layers and ReLu activation functions. The output layer consists of a single neuron without an activation function.



Figure 5.1: Structure Diagram of the AJNN_{add}. The shared MLP is denoted as M_1 and can itself have multiple layers like a NN.

5.3.3 Nearest Neighbors Neural Network

In this section, we describe the $NN_{nearest}$. This network takes a similar observation vector as $AJNN_{add}$. The parameter κ used by the $NN_{nearest}$ denotes how many preceding and succeeding neighbors should be included in the observation, similar as for the $AJNN_{add}$. The structure of the features passed to the NN is the same as for the $AJNN_{add}$. Therefore the description is omitted here.

For our NN model we consider the input vectors x_j^{glob} , x_j^{job} , $\mathcal{N}_j^{\text{red,in}}$ and $\mathcal{N}_j^{\text{red,out}}$ for $j \in U$. A graphical representation giving an overview is provided in Figure 5.2. Parameters are denoted by W and are learnable weight matrices. The learnable bias vectors are denoted by b, the internal states by h, and the final output value by o. The job specific features x_j^{job} are first preprocessed by following formula:

$$h_j^0 = \operatorname{ReLu}(W^0 x_j^{job} + b^0), \forall j \in U.$$
(5.2)

The internal node state h_i^0 is a vector of dimension d_h .

The feature vectors of the κ nearest jobs are combined in the next layer. The κ nearest pre- and succeeding jobs for job j are denoted by $\mathcal{N}_{j}^{\text{red,in}}$ and $\mathcal{N}_{j}^{\text{red,out}}$, and thus we get as input vector:

$$h_{j}^{0'} = h_{\mathcal{N}_{j,1}^{\text{red,in}}}^{0} || \dots || h_{\mathcal{N}_{j,\kappa}^{\text{red,in}}}^{0} || h_{j}^{0} || h_{\mathcal{N}_{j,1}^{\text{red,out}}}^{0} || \dots || h_{\mathcal{N}_{j,\kappa}^{\text{red,out}}}^{0}, \forall j \in U.$$
(5.3)

The combined internal state vectors $h_j^{0'}$ are then processed by a layer containing a ReLu activation function, biases and skip connections:

$$h_j^1 = \text{ReLu}(W^1 h_j^{0'} + b^1) + h_j^0, \forall j \in U.$$
(5.4)

The skip connection is implemented by adding h_j^0 at the end of the formula. Formula 5.4 yields an output of dimension d_h .

The global features are preprocessed by a feed-forward network with one hidden layer, including ReLu activation functions, biases, and d_h output nodes. The number of internal nodes is 20. The output vector of this network is denoted as h^{glob} .

Next, we combine the job specific features with the global features, by using a dense layer with ReLu activation functions, biases, and a skip connection and another layer with biases reducing the results to dimension one:

$$h_j^{1'} = \left[h^{\text{glob}}||h_j^1\right]h_j^3 = W^3\left(\text{ReLu}\left(W^2h_j^{1'} + b^2\right) + h_j^{1'}\right) + b^3, \forall j \in U.$$
(5.5)

Note that the addition of $h_j^{1'}$ results from using a skip connection The output h_j^3 is defined as a single value. To achieve this, we use an additional weight matrix W^3 . In the above equation, W^3 is designed to be a learnable weight matrix. Nevertheless, it might be enough if W^3 is designed as a static matrix with all entries equal to 1 only used for reducing dimensionality.

These single per job values are finally combined with each other and the result is a single value which is the output of the $NN_{nearest}$:

$$o = \sum_{j \in U} h_j^3 \tag{5.6}$$

We do not use weights when summing up the results of all job-individual features, as the jobs do not have an intrinsic order and thus, using weights is not applicable.

The size of the mentioned dimensions d_h and the number of internal nodes for preprocessing the global features is a parameter that influences the number of overall weights in the network and is subject to optimization. In our experiments we use $d_h = 10$ and 20 internal nodes in the hidden layers as done for the MLP in Section 5.3.1. The selection of $d_h = 10$ follows preliminary tests where $d_h \in \{10, 20, 30\}$ were compared on the VRF-small test instance set, see Section 6.1, and $d_h = 10$ showed to work best.

5.4 Feature Sets

This section describes the features used later in the observations. Each feature is motivated, and the time complexity for calculating it is provided. All features are calculated depending on the state and the instance. In Section 5.4.1 global features are presented. These are features that represent the whole instance. The features presented in Section 5.4.2 are dependent on a individual unscheduled job j.

5.4.1 Global Features

The global features of a partial solution v are denoted by $x^{\text{glob}}(v)$. Note that we omit the explicit mentioning of v if it is clear from the context. Following global features are used within the work:

1. number of unscheduled jobs |U|.

Motivation: This feature should provide the NN with information about how many jobs are unscheduled such that the output can be adapted accordingly. Time complexity: O(n) (Can be reduced to O(1) when storing the current layer in the state.)

2. minimum sum delay incoming $\delta_{\min}^{\text{in},m}$: $\sum_{j \in U} \min_{k \in U \cup \{a\} \setminus \{j\}} \delta_{k,j}^m - p_{m,k}$. **Motivation:** This feature should provide the NN with the minimum delay expected

Motivation: This feature should provide the NN with the minimum delay expected when scheduling any job $j \in U$ and might be used to calculate a (bad) lower bound. Nevertheless, as one delay matrix is enough to make a reduction to the ATSP, we take the delays on machine m here. Time complexity: $O(n^2)$

3. minimum sum delay outgoing $\delta_{\min}^{\text{out,m}}$: $\sum_{j \in U \cup \{a\}} \min_{k \in U \setminus \{j\} \cup \{0\}} \delta_{j,k}^m - p_{m,j}$. **Motivation:** See the motivation for the incoming delay. Time complexity $O(n^2)$

37



Figure 5.2: Graphical representation of the internal structure of the $NN_{nearest}$. Arrows denote data flow, and boxes denote components like layers and aggregation functions. Details about the structure are provided in Section 5.3.3.

4. maximum total job time: $\max_{j \in U} p_j^{\text{tot}}$

Motivation: The maximum total job time is part of a lower bound to the resulting solution. Even if it can not be added directly to g, it might be used to identify solutions still missing long jobs to be scheduled. Time complexity: O(n) when precomputing p^{tot} .

5. minimum release time left: $\min_{j \in U} \max(0, r_j^{\text{rel}} - p_{1,a})$

Motivation: This feature should provide the NN with information about the release time relative to the earliest scheduling time on machine one. The minimum is provided as guidance if there are jobs that can be scheduled soon. The maximum to provide a lower bound about how much additional time is at least needed. Time complexity: O(n)

6. maximum release time left: $\max_{j \in U} \max(0, r_j^{\text{rel}} - p_{1,a})$

Motivation: See the motivation for the minimum release time left.

7. minimum sum idle time incoming per machine $\omega_{\min}^{\text{in}}$: $\sum_{j \in U} \min_{k \in U \cup \{a\} \setminus \{j\}} \omega_{k,j}/m$.

Motivation: This feature should provide the NN with the anyway minimum incurred idle time per machine. The sum of the minimum idle times is divided by m. This division makes the feature independent of a change in the number of machines. Time complexity: $O(n^2)$

- 8. minimum sum idle time outgoing $\omega_{\min}^{\text{out}}$: $\sum_{j \in U \cup a} \min_{k \in U \setminus \{j\} \cup \{0\}} \omega_{j,k}/m$. Motivation: See the motivation for the incoming idle times.
- 9. total job time on machine m: P_m^U **Motivation:** The total execution time of all unscheduled jobs on machine m yields a lower bound and is always contained in the makespan. Time complexity: O(n)
- 10. Lower Bound: LB = ITLB (without the calculations for d_out due to performance reasons)

Motivation: A (tight) LB should help the NN to not underestimate the makespan. If the lower bound is stable in the sense that all its results are off by the same factor independent if a partial solution leading to a better or worse final state is evaluated, it might give good guidance. Time complexity: $O(mn^2)$

5.4.2 Job-Individual Features

The job-individual features of a partial solution are denoted by $x^{job}(v)$, whereas the specific features for each job $j \in U$ are denoted by $x_j^{job}(v)$. Note that we do not mention partial solution v if it is clear from the context. The job-individual features for job j consist of a global part $x_j^{job,glob}$ only dependent on job j and features dependent on the nearest neighbors $x_j^{job,in}$ and $x_j^{job,out}$. The features $x_j^{job,in}$ and $x_j^{job,out}$ are only provided when $\kappa > 0$. Following global job-individual features $x_j^{job,glob}$ are used:

- 1. minimum delay incoming $\min_{k \in U \cup \{a\} \setminus \{j\}} \delta_{k,j}^m p_{m,k}$ **Motivation:** This feature should provide the NN with information how much additional time is at least needed in scheduling this job. Time complexity: O(n)
- 2. minimum delay outgoing $\min_{k \in U \cup \{0\} \setminus \{j\}} \delta_{j,k}^m p_{m,j}$

Motivation: This feature should provide the NN with information how much additional time is at least needed in scheduling this job. Time complexity: O(n)

3. minimum idle time incoming per machine: $\min_{k \in U \cup \{a\} \setminus \{j\}} \omega_{k,j}/m.$

Motivation: The minimum idle time for incoming nodes gives a measure how well job j fits after the other unscheduled jobs and the last scheduled job a. Time complexity: O(n)

4. minimum idle time outgoing per machine: $\min_{k \in U \cup \{0\} \setminus \{j\}} \omega_{j,k}/m.$

Motivation: The minimum idle time for outgoing nodes gives a measure how well job j fits before the other unscheduled jobs. Time complexity: O(n)

- 5. release time left: $\max(0, r_j^{\text{rel}} p_{1,a})$ **Motivation:** When scheduling job j next the release time must be respected, thus it is also provided to the NN. Time complexity: O(1)
- 6. the job duration on machine $m: p_{m,j}$

Motivation: We use the time on machine m, as it can be directly added to the total makespan, yielding a LB. This does not hold for processing times on the other machines. Time complexity: O(1)

Following job-individual features depending on the nearest preceding and succeeding jobs $x_j^{\text{job,in}}$, $x_j^{\text{job,out}}$ are included if $\kappa > 0$:

- 1. for each $k = 1, \ldots, \kappa$ features $x_{j,k}^{\text{job,in}}$ in respect to $\mathcal{N}_j^{\text{red,in}}[k]$:
 - Delay: $\delta_{k,j}^m p_{m,k}$

Motivation: This feature gives the delay when job k is scheduled directly before job j, which is reasonable as k is one of the nearest preceding jobs of j. Time complexity: O(1)

• Idle Time: $\omega_{k,j}/m$

Motivation: This feature gives the idle time introduced when job k is scheduled directly before job j, which is reasonable as k is one of the nearest preceding jobs of j. The idle time is divided by m to be independent of the number of machines m. Time complexity: O(1)

If $|\mathcal{N}_{j}^{\mathrm{red,in}}| < \kappa$, we fill the features up to size κ by copying the last ones. Thus, we always have the same number of features for each job j independent of the layer of the BS the observation is made at.

- 2. for each $k = 1, ..., \kappa$ features $x_{j,k}^{\text{job,out}}$ in respect to $\mathcal{N}_j^{\text{red,out}}[k]$:
 - Delay: $\delta_{j,k}^m p_{m,j}$

Motivation: This feature gives the delay when job k is scheduled directly after job j, which is reasonable as k is one of the nearest succeeding jobs of j. Time complexity: O(1)

• Idle Time: $\omega_{j,k}/m$

Motivation: This feature gives the idle time introduced when job k is scheduled directly after job j, which is reasonable as k is one of the nearest succeeding jobs of j. The idle time is divided by m to be independent of the number of machines m. Time complexity: O(1)

If $|\mathcal{N}_j^{\text{red,out}}| < \kappa$, we fill the list up to size κ by copying the last job. Thus, we always have the same number of jobs and features for each job j independent of the layer of the BS the observation is made at.

5.5 Observations

In the following sections, the observations evaluated later are described. We provide one observation for the $AJNN_{add}$ and the $NN_{nearest}$ and two similar observations for the MLP.

5.5.1 Observation O_{glob}

This observation is designed for use with the MLP. It uses the global features described in Section 5.4.1 with the exception that ITLB is omitted. The observation can be computed in time $O(n^2)$.

5.5.2 Observation $O_{\text{glob,itlb}}$

Similar as O_{glob} , observation $O_{\text{glob,itlb}}$ uses the global features described in Section 5.4.1 and is designed for the MLP. Nevertheless, this time also ITLB is included in the features. Thus calculating this observation takes time $O(mn^2)$.

5.5.3 Observation $O_{AJNN,add}$

This observation is designed for the AJNN_{add}. It consists of all mentioned global features presented in Section 5.4.1 and all job-individual features presented in Section 5.4.2. Note that when setting $\kappa = 0$, the feature sets $x_j^{\text{job,in}}$ and $x_j^{\text{job,out}}$ would not be included. Nevertheless, as it is mentioned in Section 6.1, we set $\kappa = 3$ for all of our tests. Thus $x_j^{\text{job,in}}$ and $x_j^{\text{job,out}}$ are always included. The observation can be computed in time $O(n^2 \max(m, \log n))$, where the additional time compared to $O_{\text{glob,itlb}}$ results from calculating $\mathcal{N}^{\text{red,in}}$ and $\mathcal{N}^{\text{red,out}}$.

5.5.4 Observation O_{nearest}

The observation O_{nearest} is designed for the NN_{nearest} and is similar to $O_{\text{AJNN,add}}$ as it contains all global features of Section 5.4.1 and all job-individual features presented in Section 5.4.2. Due to the specific structure of the AJNN_{add}, we require this observation to have $\kappa > 0$ and thus always include the feature sets $x_j^{\text{job,in}}$ and $x_j^{\text{job,out}}$. Further, this observation includes $\mathcal{N}_j^{\text{red,in}}$ and $\mathcal{N}_j^{\text{red,out}}$ itself for all $j \in U$ as a feature. Note that due to simplicity during the evaluation of the AJNN_{add} the job IDs in $\mathcal{N}_j^{\text{red,in}}$ and $\mathcal{N}_j^{\text{red,out}}$ refer to the positions in the unscheduled jobs, i.e., the unscheduled jobs are sorted by their original IDs and then their indices in the resulting sequence starting at one to |U|are used. This observation can be computed in time $O(n^2 \max(m, \log n))$.

CHAPTER 6

Results

In this chapter, we present the results of our computational study. Before doing so, we describe the test setup and the benchmark instances used in Section 6.1. After that, in Section 6.2, we conduct tests on the same instance classes the NNs are trained, and perform statistical tests to check whether the NNs perform better than the reference approach ITLB as heuristic function (h_{itlb}). Section 6.3 experimentally evaluates how well the NNs generalize over m and Section 6.4 evaluates the generalization over n alone and both n and m. In Section 6.5 the approximation errors of the NNs over the layers of a BS are evaluated. Finally, we compare the best approaches with the state-of-the-art method from Pourhejazy et al. [Pou+20] in Section 6.6.

6.1 Test Setup and Benchmark Instances

All benchmark tests are run with Julia 1.7.0 on a computing cluster on Intel Xeon E5540, 2.53 GHz Quad-Core nodes in single-threaded mode. Each neural network is trained ten times independently of the others but with the same settings. The tests are performed with all these ten trained models, and results are reported as boxplots if not described differently. Deterministic methods like when using ITLB as guidance function for the BS are only run once. In our work, we use four different benchmark sets, all adapted for the NWFSP-RT in [Pou+20]. The first benchmark set is denoted by VFR and can be split into small and large instances. The set of small instances VFR-small consists of instance-classes of all combinations of $n \in \{10, 20, \ldots, 60\}$ and $m \in \{5, 10, 15, 20\}$ and VFR-large of $n \in \{100, 200, \ldots, 800\}$ and $m \in \{20, 40, 60\}$. Each instance class consists of ten individual instances. Thus, there are 240 VFR-small and 240 VFR-big instances. Note that when evaluating a NN on a per instance-class basis, the boxplots contain 100 values as we perform ten runs per instance. A smaller test set called REC consists of three instances of sizes $(n = 20, m \in \{5, 10, 15\}), (n = 30, m \in \{10, 15\}), (n = 50, m \in \{10\}), and (n = 75, m \in \{20\}), thus, in total 21 instances. The TA-small$

instance set consists of ten instances of $n \in \{20, 50, 100\}$ and $m \in \{5, 10, 20\}$ and bigger instances of $(n = 200, m \in \{10, 20\})$ and (n = 500, m = 20). Thus, it has 120 instances. We describe the instance set TA-large for completeness even if we did not test on it due to its big instance sizes. The set TA-large consists out of ten instances of sizes $n \in \{1000, 1500, 2000\}$ and m = 20, thus, in total 30 instances.

As baseline we use the results provided in [Pou+20] for their BS with local search from [Pou+20] (BS_{LS,pour}) approach which are also the so far best known results per instance. If not mentioned differently, we report the relative percentage difference (RPD) to the currently best known results BS_{LS,pour}. The RPD is calculated by the following formula, in which x denotes the length of solution u on a specific instance and x_{best} the result of BS_{LS,pour} on the same instance:

$$RPD = (x - x_{best}) * 100/x_{best}$$
(6.1)

We experimentally evaluate four configurations in the later sections. These are one configuration per observation presented in Section 5.5:

- MLP_{basic}: The MLP as described in Section 5.3.1 together with Observation O_{glob} .
- MLP_{ITLB}: The MLP as described in Section 5.3.1 together with Observation $O_{\text{glob,itlb}}$.
- AJNN_{add}: The aggregated job neural network as presented in Section 5.3.2 with $\kappa = 3$ and Observation $O_{AJNN,add}$.
- NN_{nearest}: The nearest Neighbors network as presented in Section 5.3.3 with $\kappa = 3$ and Observation O_{nearest} .

All NNs are trained on instances $n \in \{10, 20\}$ and $m \in \{10, 20\}$, in total on four different instance sizes. The LBS was performed with the following settings: number of iterations z = 3000, minimum number of observation to start learning $\gamma = 3000$, size of the replay buffer $\rho = 5000$, beam width of the NBS calls $\beta' = 50$, expected number of samples $\alpha = 60$. Resulting from the settings approximately $180 \cdot 10^3$ training samples are created and the loss function is called approximately $14.5 \cdot 10^6$ times. This is far more than the number of weights contained in our proposed networks.

6.2 Comparison on Same Instance Sizes as Trained

In this section we compare all trained NNs on separate test instance with the same size m and n as the NNs were trained on. This is done to evaluate whether the NNs are able to improve the results compared with h_{itlb} (using ITLB as guidance function). As beam width we use $\beta \in \{1, 10, 20, 50\}$. Due to the lack of being able to reproduce the results of [Pou+20] we can not run their BS with smaller beam widths and thus only

have results for $\beta = 600$. Therefore, we use h_{itlb} as reference method here. Figures 6.1,6.2, 6.3, and 6.4 show the results as boxplots for the instances $n \in \{10, 20\}$ and $m = \{10, 20\}$. To check whether the results provided by the NNs are significantly better than h_{itlb} we perform a Wilcoxon rank-sum test with a significance level of $\alpha = 0.05$ and the alternative hypotheses that the compared method is better, i.e. yields smaller results, than h_{itlb} . The results of the statistical tests are provided in Table 6.1.

The Figures and the significance tests show that the MLP_{basic} does not outperform h_{itlb}. For the MLP_{ITLB} the results are mostly quite similar to h_{itlb} and only the configuration trained on n = 20 and m = 10 was able to significantly outperform h_{itlb} when run with $\beta \geq 10$. The AJNN_{add} and NN_{nearest} are able to outperform h_{itlb} significantly on ten and nine of sixteen tested configurations and performed best overall. Also note that some of the tested methods yield results better than the best known results on individual instances even if run with small beam widths of at most $\beta = 50$ compared to $\beta = 600$ which was used for BS_{LS,pour}. Especially for the instance class n = 10, m = 10 the tested methods outperformed BS_{LS,pour} on average. When visually comparing the results of BS from [Pou+20] (BS_{pour}) shown as dashed lines in the figures that were evaluated with $\beta = 600$ with the trained models and h_{itlb} on $\beta = 50$ it shows that all methods outperform BS_{pour} on average. Thus, we conclude that LBS is highly effective for learning guidance functions for BS as some models outperformed the well working guidance function h_{itlb} and MLP_{basic} which does not contain ITLB is comparably or better on all instances than BS_{pour}. For a comparison of the average runtimes see Figure 6.11 in Section 6.4.

6.3 Generalization over m

This section evaluates whether the trained NNs generalize well over the number of machines m. Therefore, all NNs are evaluated on instance classes with the same number of jobs as they are trained on and $m \in \{5, 10, 15, 20\}$ of the VFR-small set. The results are shown in Figures 6.5, and 6.6. It shows that the AJNN_{add} and the NN_{nearest} generalize poorly over m. Especially the NNs trained on m = 20 perform worse than the NNs m = 10 on instances with m = 5. For instances with n = 20 and m = 15 the AJNN_{add} and the NN_{nearest} trained on m = 20 yield better results when visually compared. For instances with n = 10 and m = 15 this does not hold when visually compared. The average RPD of the MLP_{basic} gets worse when generalizing. The MLP_{ITLB} has a good generalization performance as the average RPD does not change much. The Wilcoxon rank-sum test performed to check whether the method trained with m = 10 or m = 20 is better than the other is given in Table 6.1 and shows that there are no significant differences on a level of $\alpha = 0.05$.



Figure 6.1: The figure shows the results of the NNs on the same instance sizes for n and m as they were trained on and the results for ITLB relative to $BS_{LS,pour}$ for n = 10 and m = 10. Each boxplot of the NNs represents results over ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots indicate the median and the triangles the mean results. Instances from the VFR-small set are used. The dashed line shows the average RPD of BS_{pour} .



Figure 6.2: The figure shows the results of the NNs on the same instance sizes for n and m as they were trained on and the results for ITLB relative to $BS_{LS,pour}$ for n = 10 and m = 20. Each boxplot of the NNs represents results over ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots indicate the median and the triangles the mean results. Instances from the VFR-small set are used. The dashed line shows the average RPD of BS_{pour} .



Figure 6.3: The figure shows the results of the NNs on the same instance sizes for n and m as they were trained on and the results for ITLB relative to $BS_{LS,pour}$ for n = 20 and m = 10. Each boxplot of the NNs represents results over ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots indicate the median and the triangles the mean results. Instances from the VFR-small set are used. The dashed line shows the average RPD of BS_{pour} .



Figure 6.4: The figure shows the results of the NNs on the same instance sizes for n and m as they were trained on and the results for ITLB relative to $BS_{LS,pour}$ for n = 20 and m = 20. Each boxplot of the NNs represents results over ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots indicate the median and the triangles the mean results. Instances from the VFR-small set are used. The dashed line shows the average RPD of BS_{pour} .

Table 6.1: The table shows the results of Wilcoxon rank-sum tests when testing whether h_{itlb} is worse in terms of the resulting objective value than the NNs when run on the same instances as trained. As an alternative hypothesis, we use that the compared method is better than h_{itlb} . Further, $\alpha = 0.05$ is used. The tests are performed on runs with the same β values. Note that the results for h_{itlb} are ten times the same per instance as it is a deterministic method. Significant results where H0 can be rejected are written in bold and marked with a *.

Methods	$\beta = 1$	$\beta = 10$	$\beta = 20$	$\beta = 50$
MLP	1.00	1.00	1.00	1.00
MLP , $n = 10, \ m = 20$	1.00	1.00	0.96	0.86
MLP , $n = 20, m = 10$	1.00	1.00	1.00	1.00
MLP , $n = 20, m = 20$	1.00	0.99	0.95	0.99
$\overline{\text{MLP}_{\text{ITLB}}, n = 10, m = 10}$	0.28	0.08	0.20	0.75
MLP_{ITLB} , $n = 10$, $m = 20$	0.21	0.42	0.62	0.32
MLP_{ITLB} , $n = 20$, $m = 10$	0.90	*0.02	*0.00	*0.01
MLP_{ITLB} , $n = 20$, $m = 20$	*0.01	0.12	0.17	0.24
AJNN _{add} , $n = 10, m = 10$	*0.00	*0.00	0.13	0.32
$AJNN_{add}$, $n = 10, m = 20$	*0.00	0.28	*0.04	0.57
$AJNN_{add}$, $n = 20, m = 10$	0.79	*0.01	*0.00	*0.02
$AJNN_{add}$, $n = 20, m = 20$	*0.00	*0.03	*0.03	0.07
$NN_{nearest}$, $n = 10, m = 10$	*0.00	*0.02	*0.03	0.77
$NN_{nearest}$, $n = 10$, $m = 20$	*0.00	0.30	*0.02	0.24
$NN_{nearest}$, $n = 20, m = 10$	0.76	*0.00	*0.00	*0.00
$NN_{nearest}$, $n = 20$, $m = 20$	*0.01	0.17	0.13	0.10

Table 6.2: The table shows the results of Wilcoxon rank sum tests, when testing for two models trained on the same n but different m the one is yields better results than the other. In the table e.g. m = 10 < m = 20 denotes the results when testing whether the model trained with m = 10 is better than the model trained with m = 20. The table shows that there are no significant differences on a level of $\alpha = 0.05$.

	n = 10		n = 20	
Methods	m = 10 < m = 20	m = 20 < m = 10	m = 10 < m = 20	m = 20 < m = 10
MLP	0.11	0.89	0.72	0.28
MLP_{ITLB}	0.58	0.42	0.47	0.53
$\mathrm{AJNN}_{\mathrm{add}}$	0.16	0.84	0.89	0.11
$\mathrm{NN}_{\mathrm{nearest}}$	0.40	0.60	0.65	0.35



Figure 6.5: The figure shows the results of the NNs on instances with the same number of jobs as they were trained on n = 10 and $m \in \{5, 10, 15, 20\}$. The results for ITLB are included as reference. For evaluation $\beta = 10$ is used. All results are relative to BS_{pour}. Each boxplot of the NNs contains results for ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots show the median and the triangles the mean results. Instances from the VFR-small set are used. The dashed line shows the average RPD of BS_{pour}.



Figure 6.6: The figure shows the results of the NNs on instances with the same number of jobs as they were trained on n = 20 and $m \in \{5, 10, 15, 20\}$. The results for ITLB are included as reference. For evaluation $\beta = 10$ is used. All results are relative to BS_{pour}. Each boxplot of the NNs contains results for ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots show the median and the triangles the mean results. Instances from the VFR-small set are used. The dashed line shows the average RPD of BS_{pour}.

6.4 Generalization over n and m

In this section we evaluate whether the NNs generalize well over n and also over n and m at the same time. We evaluate all trained NNs on all instances in VFR-small with $\beta = 100$. The results for instances with $m \in \{5, 10, 15, 20\}$ are reported in Figures 6.7, 6.8, 6.9, and 6.10.

Generalization over n with m = 5:

The results in Figure 6.7 show that MLP_{ITLB} and h_{itlb} generalize well and are able to constantly improve the best known results for some instances, even if only run with $\beta = 100$. The MLP_{basic} performs second best on average and shows only a moderate growth of its average RPD with increasing n

Generalization over n with m = 10:

The results in Figure 6.8 show that MLP_{ITLB} generalizes again well over increasing n and looks slightly better than h_{itlb} . Nevertheless, the statistical test results in Table 6.3 show that this difference is not significant. Further, AJNN_{add} and NN_{nearest} trained on n = 20, m = 10 generalize well over increasing n as they only have a moderate growth of their average RPD. Surprisingly this does not hold for their configurations trained on n = 10.

Generalization over n with m = 15:

The results in Figure 6.9 show that MLP_{ITLB} and h_{itlb} again perform well and also the results of MLP_{basic} have a moderate RPD growth.

Generalization over n with m = 20:

The results in Figure 6.10 show that again MLP_{ITLB} and h_{itlb} perform best. Further, MLP_{basic} performs good. For this instances AJNN_{add} and NN_{nearest} trained on n = 20, m = 20 generalize well.

Overall Results:

Overall MLP_{ITLB} and h_{itlb} show excellent performance in generalization over n and also good performance when generalizing over n and m at the same time. We observe that training on bigger instances regarding n yields better results in generalizing to bigger instances. Further, MLP_{basic} shows well but not excellent performance. Nevertheless, we want to highlight that this method is the fastest and thus could in theory use a higher β and terminate in the same time as the methods containing ITLB. The average runtimes over instances of increasing n and m = 20 are shown in Figure 6.11. Figures for $m \in \{5, 10, 15\}$ are omitted as they show the same trend with smaller runtimes. The figure shows that AJNN_{add} has the highest average runtime, followed by NN_{nearest}. Both have a steep increase in runtime. The average runtimes of MLP_{ITLB} and h_{itlb} are similar, and also have a steep increase. The average runtimes of MLP_{Dasic} increase only slowly and are the lowest. For AJNN_{add} and NN_{nearest} we observed that they benefit from training on bigger instances, i.e., n = 20. As already observed in Section 6.3 these network types do not perform well in generalizing over m. The AJNN_{add} and NN_{nearest} trained with n = 20 show a well generalization behavior when evaluated on instances with the same



Figure 6.7: The figure shows the results of the NNs when generalizing over m and n. The results for ITLB are included as a reference. For evaluation $\beta = 100$ is used. All results are relative to BS_{pour}. Instances from the VFR-small set with m = 5 are used. The dashed line shows the average RPD of BS_{pour}.

number of machines m they were trained with. Nevertheless, neither configuration is statistically significantly better than h_{itlb} , when compared over all instances or instances with the same m used for training. The results of the statistical test are provided in Tables 6.4, and 6.3.

6.5 Approximation Errors over Layers

This section evaluates how well the NNs approximate the remaining cost on the different layers in the BS. The evaluation is done for all four NNs types trained on n = 20, m = 20 as these configurations scaled well over the number of jobs n.

Results are shown in Figures 6.12, 6.13, 6.14, 6.15, and 6.16 for $n \in \{10, 20, \dots, 50\}$.



Figure 6.8: The figure shows the results of the NNs when generalizing over m and n. The results for ITLB are included as a reference. For evaluation $\beta = 100$ is used. All results are relative to BS_{pour}. Instances from the VFR-small set with m = 10 are used. The dashed line shows the average RPD of BS_{pour}.

Table 6.3: The table shows the results of the Wilcoxon rank-sum test when testing whether h_{itlb} is worse in terms of the resulting objective value than the NNs when evaluated over all instances of the VFR-small set which have the same m as the model was trained with and $n \in \{10, 20, \ldots, 60\}$. As an alternative hypothesis, we use that the compared method is better than h_{itlb} . In the table, one row represents the statistical test results for one NN type, and each column represents the instance class it was trained. The table shows that none of the NNs is significantly better than h_{itlb} .

Methods	n = 10, m = 10	n = 10, m = 20	n = 20, m = 10	n = 20, m = 20
MLP	0.96	0.98	0.62	0.55
$\mathrm{MLP}_{\mathrm{ITLB}}$	0.36	0.52	0.46	0.16
$\mathrm{AJNN}_{\mathrm{add}}$	1.00	0.95	0.69	0.60
$\mathrm{NN}_{\mathrm{nearest}}$	1.00	0.92	0.71	0.52



Figure 6.9: The figure shows the results of the NNs when generalizing over m and n. The results for ITLB are included as a reference. For evaluation $\beta = 100$ is used. All results are relative to BS_{pour}. Instances from the VFR-small set with m = 15 are used. The dashed line shows the average RPD of BS_{pour}.

Table 6.4: The table shows the results of the Wilcoxon rank-sum test when testing whether h_{itlb} is worse in terms of the resulting objective value than the NNs when evaluated over all instances of the VFR-small set. As an alternative hypothesis, we use that the compared method is better than h_{itlb} . In the table, one row represents the statistical test results for one NN type, and each column represents the instance class it was trained. The table shows that none of the NNs is significantly better than h_{itlb} .

Methods	n = 10, m = 10	n = 10, m = 20	n = 20, m = 10	n = 20, m = 20
MLP	0.79	0.94	0.78	0.95
MLP_{ITLB}	0.36	0.37	0.23	0.67
$\mathrm{AJNN}_{\mathrm{add}}$	1.00	1.00	1.00	1.00
$\mathrm{NN}_{\mathrm{nearest}}$	1.00	1.00	0.99	0.98



Figure 6.10: The figure shows the results of the NNs when generalizing over m and n. The results for ITLB are included as a reference. For evaluation $\beta = 100$ is used. All results are relative to BS_{pour}. Instances from the VFR-small set with m = 20 are used. The dashed line shows the average RPD of BS_{pour}.

Note that layer c yields instances of size n - c, which the NNs must approximate. A NN trained on n = 20 therefore must generalize if instances of sizes $n - c \ge 20$ are inputted. For AJNN_{add} and NN_{nearest} it shows that for $n \in \{10, 20, 30\}$ the boxplots without outliers stay into a +/-25%RPD range and also most outliers stay within this range. For n = 40, still, most of the boxplots stay within the +/-25%RPD range. However, the layers one to ten outliers are approximately within a +/-50%RPD. The observed behavior seems reasonable, as the models were not trained on sizes of instances of $n \in \{31, 32, \ldots, 40\}$. Starting from layer 11, the figure is similar to the figure of n = 30. For n = 40 AJNN_{add} seems to underestimate the results, while NN_{nearest} is closer at zero. For n = 50 the trend of the evaluation on n = 40 continues. Again, the approximation results for layers where the size of the instances is bigger than the models are trained, i.e., layers one to 20, have a higher approximation error than on the sizes where they



Figure 6.11: The figure shows average runtime in seconds over instances with increasing n and m = 20. Even though the data is not continuous, we connect the data points of the instance classes to visualize the trend.

were trained. The range of the boxplots and outliers is approximately within [-75, 100]. On average, the AJNN_{add} again underestimates the results.

The figures show for the MLP_{basic} and MLP_{ITLB} that they approximate the results well for layers yielding instance sizes where they are not trained. Nevertheless, starting with n = 40, they tend to underestimate the results for lower layers. For higher layers, the approximation errors get higher for both MLP_{basic} and MLP_{ITLB}. For MLP_{basic} the approximation errors on the last layers in every figure become almost infinitely high, which is not shown in the figures.

6.6 Comparison with State-of-the-Art

In this section we compare the NNs with $BS_{LS,pour}$ and BS_{pour} . The comparison with $BS_{LS,pour}$ is important as $BS_{LS,pour}$ represents the current state-of-the-art. However, BS_{pour} uses the same guidance function as $BS_{LS,pour}$ but without local search and will therefore be our primary reference when evaluating the quality of the learned functions.


Figure 6.12: The figure shows the resulting approximation errors made by the machine learning models for n = 10 and m = 20. In total, 1000 random samples on different layers of the BS were created for each method independently by using the implementation of our replay buffer together with a fully trained NN. The beam width during the generation used for the NBS calls is $\beta' = 50$, samples are taken from BS runs with $\beta = 100$. For each configuration, the training samples and labels were generated once using the configuration itself for labeling and then evaluated once with every trained model, i.e., ten times. Each boxplot consists of a different number of samples due to the randomness of the generation. In total, there are $10 \cdot 10^3$ samples per configuration. On the vertical axis, the charts show the RPD to the labels. Thus, if values greater than zero are reported, the NN reported a higher objective value than it was, i.e., it overestimated the result. Values smaller than zero denote that the values reported by the NN are lower than the labels. The horizontal axes are fixed to a range from 100 to -100 and thus might not show all data.

Note that we frequently compare our models evaluated with $\beta = 100$ with BS_{LS,pour} and BS_{pour} that were evaluated with $\beta = 600$. This is done to limit the runtime our approaches need.

The comparison on the VFR-small test instances in Table 6.5 and 6.6 show that all NNs could outperform BS_{pour} on at least eight out of 24 instance classes on average with $\beta = 100$. The best performance is shown by $MLP_{ITLB} m \in 10, 20$ which outperformed BS_{pour} on all instance classes on average and $BS_{LS,pour}$ on seven and five instance classes. Equally well performance is also shown by h_{itlb} which also outperformed BS_{pour} on average on all instance classes. Observe that $AJNN_{add}$ and $NN_{nearest}$ outperform BS_{pour} on nine and ten instance classes on average and are able to outperform $BS_{LS,pour}$ on two to five instance classes. This behavior seems reasonable, as there are in total 12 instances with $n \leq 30$ and thus in a range where the $AJNN_{add}$ and $NN_{nearest}$ trained on



Figure 6.13: The figure shows the resulting approximation errors made by the machine learning models for n = 20 and m = 20. In total, 1000 random samples on different layers of the BS were created for each method independently by using the implementation of our replay buffer together with a fully trained NN. The beam width during the generation used for the NBS calls is $\beta' = 50$, samples are taken from BS runs with $\beta = 100$. For each configuration, the training samples and labels were generated once using the configuration itself for labeling and then evaluated once with every trained model, i.e., ten times. Each boxplot consists of a different number of samples due to the randomness of the generation. In total, there are $10 \cdot 10^3$ samples per configuration. On the vertical axis, the charts show the RPD to the labels. Thus, if values greater than zero are reported, the NN reported a higher objective value than it was, i.e., it overestimated the result. Values smaller than zero denote that the values reported by the NN are lower than the labels. The horizontal axes are fixed to a range from 100 to -100 and thus might not show all data.

n = 20 perform well. Even though MLP_{basic} is not among the top performers in this comparison, it yields staple results by outperforming BS_{pour} on eight and 12 instance classes on average while having the lowest runtime with less than 23 seconds on average for n = 60 m = 20.

The results for TA-small in Table 6.7 and Figure 6.17 and the results for REC in Table 6.8 and Figure 6.18 show similar results as already observed for the VFR-small data set. The MLP_{ITLB} trained on n = 20 and m = 20 outperformed BS_{pour} on nine out of 12 tested TA-small instance classes and on all REC instance classes. The same holds for h_{itlb}.

Results for selected instance classes of the VFR-large instance set are shown in Figure 6.19, and Table 6.9. Due to the size of the instances and the resulting runtime, we



Figure 6.14: The figure shows the resulting approximation errors made by the machine learning models for n = 30 and m = 20. In total, 1000 random samples on different layers of the BS were created for each method independently by using the implementation of our replay buffer together with a fully trained NN. The beam width during the generation used for the NBS calls is $\beta' = 50$, samples are taken from BS runs with $\beta = 100$. For each configuration, the training samples and labels were generated once using the configuration itself for labeling and then evaluated once with every trained model, i.e., ten times. Each boxplot consists of a different number of samples due to the randomness of the generation. In total, there are $10 \cdot 10^3$ samples per configuration. On the vertical axis, the charts show the RPD to the labels. Thus, if values greater than zero are reported, the NN reported a higher objective value than it was, i.e., it overestimated the result. Values smaller than zero denote that the values reported by the NN are lower than the labels. The horizontal axes are fixed to a range from 100 to -100 and thus might not show all data.

performed tests only on instances up to n = 400 and $m \in \{20, 40, 60\}$. The MLP_{basic} was only tested on instances with n = 100. On these instances MLP_{ITLB} could on average outperform BS_{pour}. Further, MLP_{ITLB} also performs better than h_{itlb} with $\beta = 100$ on average. However, MLP_{basic} could not outperform BS_{pour} on average. This might be due to the difference in beam widths, that might have a bigger impact the bigger the instances get.

Overall we could outperform $BS_{LS,pour}$ on average on 10 out of 24 VFR-small instance classes, six out of 12 TA-small instance classes, four out of seven REC instance classes, and on zero out of 12 tested VFR-large instance classes. Further, we outperformed BS_{pour} on all instances of VFR-small, TA-small and REC, except for two TA-small instances with $n \in \{100, 500\}$ with $\beta = 100$.



zero denote that the values reported by the NN are lower than the labels. The horizontal axes are fixed to a range from 100 were generated once using the configuration itself for labeling and then evaluated once with every trained model, i.e., ten NBS m = 20. In total, 1000 random samples on different layers of the BS were created for each method independently by using the to -100 and thus might not show all data. zero are reported, the NN reported a higher objective value than it was, i.e., it overestimated the result. $10 \cdot 10^{3}$ times. Each boxplot consists of a different number of samples due to the randomness of the generation. In total, there are implementation of our replay buffer together with a fully trained NN. The beam width during the generation used for the Figure 6.15: The figure shows the resulting approximation errors made by the machine learning models for n =calls is samples per configuration. On the vertical axis, the charts show the RPD to the labels. Thus, if values greater than β = 50, samples are taken from BS runs with β = 100. For each configuration, the training samples and labels Values smaller than 40 and



Figure 6.16: The figure shows the resulting approximation errors made by the machine learning models for n = 50 and m = 20. In total, 1000 random samples on different layers of the BS were created for each method independently by using the implementation of our replay buffer together with a fully trained NN. The beam width during the generation used for the NBS calls is $\beta' = 50$, samples are taken from BS runs with $\beta = 100$. For each configuration, the training samples and labels times. Each boxplot consists of a different number of samples due to the randomness of the generation. In total, there are $10 \cdot 10^3$ samples per configuration. On the vertical axis, the charts show the RPD to the labels. Thus, if values greater than zero are reported, the NN reported a higher objective value than it was, i.e., it overestimated the result. Values smaller than zero denote that the values reported by the NN are lower than the labels. The horizontal axes are fixed to a range from 100 were generated once using the configuration itself for labeling and then evaluated once with every trained model, i.e., ten -100 and thus might not show all data. 0



Figure 6.17: The figure visualizes the results from Table 6.7. Each boxplot of the NNs represents results over ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots indicate the median and the triangles the mean results. Instances from the TA-small set are used. The dashed line shows the average RPD of BS_{pour} .

We did not test our approach on VFR-large and TA-large instances because the runtimes exceeded our time limits.

 $m \in \{10, 20\}$ are shown as they performed in general better than the NNs trained with n = 10. The results were produced with $\beta = 100$. The instances are from the VFR-small instance set. For easier comparison we added the results of h_{itb} as well as BS_{pour} and BS_{LS,pour}. The best results in each row are marked with ****** whereas we mark results that are the best except for BS_{LS,pour} with *. If an approach is the overall best approach, it is only marked by **. The last row of the table shows how often the individual approach was able to perform at least as good as BS_{LS,pour} and BS_{pour}. The Table continues in objective values, the standard deviations, and the runtimes. Only results for NNs trained on instance sizes n = 20 and Table 6.5: This Table shows the first part of the results shown in Figures 6.7, 6.8, 6.9, and 6.10. It contains the average Table 6.6.

,																											
	10	runt.	3.6	3.7	3.8	3.8	7.0	8.0	8.5	9.4	13.8	16.1	18.5	21.0	24.0	30.1	37.5	44.5	41.0	55.5	72.4	92.7	66.3	96.8	133.1	173.6	= 7
	$_{\rm B}, m =$	$\sigma_{\overline{x}}$	47.0	43.2	67.8	75.0	93.2	45.7	85.4	131.1	60.7	82.2	82.6	117.4	104.4	95.0	102.1	160.5	151.9	96.7	116.3	137.7	125.7	113.7	141.8	190.3	, BSrs n
	MLP _{ITI}	\overline{x}	804.1	1255.4	1645.6	1973.6	1490.8	2061.0	**2573.0	3038.9	2119.1	2816.7	3510.4	4093.7	2849.6	*3644.1	4449.3	5141.4	3479.0	*4461.7	5380.6	6249.1	4139.0	5362.0	6341.4	7248.8	$BS_{nour} = 24$
	0	runt.	3.1	3.1	3.1	3.1	4.5	4.4	4.4	4.3	7.0	7.2	7.2	7.1	10.1	10.2	10.4	10.6	15.7	15.3	15.3	14.9	22.3	22.8	22.7	22.4	0 = 0
	LP, $m = 2$	$\sigma_{\overline{x}}$	49.1	45.0	67.8	83.4	100.5	61.6	87.3	124.3	76.4	100.0	78.3	118.5	117.6	90.2	122.7	147.5	169.3	127.3	118.1	121.7	175.6	138.1	147.9	166.4	$12, BS_{LS}$
	Μ	\overline{x}	817.1	1277.6	1655.3	1990.4	1536.6	2117.5	2633.4	3104.5	2198.8	2898.9	3590.9	4156.8	2947.3	3738.0	4530.6	5240.1	3622.6	4579.2	5482.1	6333.5	4289.7	5519.0	6471.1	7419.3	BS _{non} =
.	0	runt.	3.1	3.1	3.1	3.1	4.4	4.4	4.4	4.4	6.7	6.6	6.8	6.7	10.2	10.1	10.1	9.9	15.0	15.0	14.8	15.3	22.2	22.2	21.7	21.5	= 0
	P, $m = 1$	$\sigma_{\overline{x}}$	46.0	34.6	69.8	86.4	92.6	60.9	91.3	130.9	70.5	94.9	92.1	144.1	117.2	80.7	131.6	172.2	166.1	108.4	122.9	181.7	149.2	124.8	133.3	231.2	8. BS _{LS}
	IW	\overline{x}	815.5	1274.5	1661.6	2005.7	1535.8	2120.3	2649.3	3118.6	2190.5	2898.7	3617.8	4220.4	2924.4	3734.2	4539.8	5290.2	3603.2	4564.5	5485.4	6409.2	4254.8	5481.4	6449.0	7466.7	BS _{nonr} =
-	20	runt.	2.6	2.7	2.8	2.8	6.7	7.8	8.4	9.2	17.0	19.6	22.0	24.5	37.6	44.4	51.5	59.0	74.3	88.3	106.2	126.5	136.3	165.7	203.0	245.3	ur = 3
	Id, $m =$	$\sigma_{\overline{x}}$	48.2	53.5	71.1	74.9	103.9	74.7	87.9	125.3	85.5	118.3	106.0	119.7	139.7	127.7	157.1	157.4	190.0	187.6	196.9	175.9	204.0	245.2	253.5	255.8	BS _{LS, nc}
	AJNN _{ac}	\underline{x}	825.2	1277.3	1648.3	1973.8	1569.4	2134.5	2614.7	**3027.8	2256.3	2952.2	3581.7	*4085.4	3063.1	3850.9	4607.0	5190.1	3791.2	4785.2	5662.1	6435.4	4504.7	5786.3	6715.4	7601.8	$BS_{nom} = 9$.
-	10	runt.	2.6	2.7	2.8	2.8	6.7	7.8	8.5	9.2	17.0	19.7	22.0	24.4	37.7	44.4	51.3	58.8	74.3	88.5	106.0	126.7	136.5	167.0	203.9	245.1	= 5
	Id, $m =$	$\sigma_{\overline{x}}$	45.5	42.3	69.4	76.8	102.7	43.5	109.2	121.2	75.7	81.9	105.9	137.0	131.9	107.2	173.9	229.7	173.3	158.2	196.2	246.9	195.6	193.2	259.3	305.4	BS _{LS nc}
	A JNN _{ac}	\overline{x}	815.9	1254.9	1647.2	1976.4	1538.4	**2053.5	2634.4	3129.4	2228.8	2824.0	3626.3	4319.1	3036.9	3693.8	4668.1	5548.8	3754.7	4586.2	5763.7	6858.0	4498.7	5568.8	6829.5	8105.6	$BS_{nour} = 9$.
ŀ	tances		ŋ	10	15	20	ŋ	10	15	20	IJ	10	15	20	IJ	10	15	20	r.	10	15	20	r.	10	15	20	prformed:
	Ins	u	10	10	10	10	20	20	20	20	30	30	30	30	40	40	40	40	50	50	50	50	60	60	00	60	Outpe

		1	our = 5	4, $BS_{LS,p}$	$ BS_{pour} = 2$	$_{\rm S,pour} = 2$	$= 10, BS_{I}$	BSpour	pour = 4	0, $BS_{LS,I}$	$BS_{pour} = 1$	our = 5	4, $BS_{LS,F}$	$BS_{pour} = 2$	formed:	Outper
	135.3	**7152.9	143.5	205.1	7363.1	271.3	318.0	7502.8	268.4	343.3	7892.7	175.4	172.4	*7219.5	20	60
	136.2	**6247.3	110.2	126.6	6396.6	226.5	173.0	6627.0	224.8	210.6	6713.6	133.3	125.8	*6329.8	15	60
	81.6	**5304.5	85.1	138.6	5393.5	190.1	173.1	5732.6	189.4	154.0	5486.6	97.5	126.6	*5356.9	10	60
	120.6	**4072.9	68.2	131.4	*4123.0	161.0	196.8	4528.9	158.9	182.8	4445.7	66.0	138.5	4147.2	5	60
	96.2	**6157.8	84.6	168.6	6327.7	144.6	189.3	6354.4	143.7	270.0	6743.1	92.3	97.8	*6242.7	20	50
	101.7	**5311.7	66.9	97.6	5444.3	124.2	129.9	5563.7	123.7	185.0	5684.5	73.0	115.1	*5370.1	15	50
	113.7	**4436.4	59.1	106.6	4476.1	106.3	167.3	4719.6	107.1	125.0	4535.4	55.6	93.1	4470.4	10	50
	142.7	**3447.3	32.3	166.0	*3460.2	93.4	180.9	3807.7	92.8	172.4	3722.1	41.3	159.8	3498.9	5	50
0.77	151.7	**5097.4	40.1	181.6	5177.2	71.9	165.4	5225.7	71.0	247.2	5512.2	44.5	159.5	*5130.5	20	40
4	85.2	**4398.9	32.5	85.0	4468.6	63.7	121.6	4564.2	63.5	154.9	4625.5	37.4	98.5	*4446.3	15	40
ಲು	91.2	**3624.1	27.1	75.3	3678.7	56.7	111.0	3802.2	57.0	91.2	3658.5	30.3	77.1	3658.4	10	40
\sim	94.9	**2827.6	22.8	107.7	*2847.6	50.6	128.4	3077.8	50.5	134.4	3011.8	24.0	108.6	2866.0	5	40
4	96.3	**4076.0	24.8	106.9	4120.6	33.8	112.9	4088.0	33.8	186.9	4329.0	21.0	113.3	4088.8	20	30
ಲು	77.5	3503.0	23.5	77.9	3554.8	31.2	85.0	3587.9	31.1	120.1	3633.8	18.7	74.7	**3501.9	15	30
\sim	81.4	2825.2	21.9	81.0	2833.2	28.3	104.7	2931.7	28.5	82.3	**2815.1	16.0	90.5	2831.6	10	30
\mathbb{N}	57.8	2118.0	20.9	59.7	**2112.3	25.9	76.0	2261.3	25.9	83.8	2226.6	13.7	65.8	2127.0	5	30
cυ	130.1	3035.6	29.3	146.7	3055.1	15.9	126.1	3028.2	15.8	170.1	3170.8	9.3	132.8	3039.7	20	20
\sim	87.4	2573.2	29.3	92.1	2588.9	15.0	86.7	2642.0	15.0	120.6	2649.7	8.6	85.7	2579.2	15	20
\sim	51.8	2061.1	28.5	53.5	2071.7	14.2	61.1	2125.3	14.3	43.4	2055.0	7.9	50.7	2066.4	10	20
H	89.7	**1484.1	27.7	96.3	*1489.8	13.6	103.5	1559.6	13.6	104.1	1546.0	7.1	95.1	1495.6	5	20
÷	76.8	1977.9	25.5	74.9	**1973.4	8.6	74.6	1974.5	8.6	79.8	1982.4	3.8	74.9	1973.4	20	10
<u> </u>	73.3	1652.1	25.7	66.8	1645.7	8.6	74.5	1656.2	8.6	71.4	1650.2	3.7	67.4	**1645.6	15	10
⊢	45.5	1261.5	25.2	41.9	**1254.6	8.6	48.0	1280.2	8.6	42.9	1255.3	3.7	42.1	1254.8	10	10
	43.4	806.2	25.4	45.2	**802.6	8.5	44.6	823.2	8.6	52.0	821.0	3.6	47.0	804.8	5	10
	$\sigma_{\overline{x}}$	x	runt.	$\sigma_{\overline{x}}$	<i>x</i>	runt.	$\sigma_{\overline{x}}$	\overline{x}	runt.	$\sigma_{\overline{x}}$	$\frac{x}{x}$	runt.	$\sigma_{\overline{x}}$	\overline{x}	m	n
1	ur	BS _{LS,pot}	0	$\beta, \beta = 10$	ITLE	c = 20	$I_{nearest}, m$	NN	: 10	est, m =	NN _{near}	20	$_{\text{LB}}, m =$	MLPIT	nces	Insta

given in Table 6.5.
how often the individual approach was able to perform at least as good as $BS_{LS,pour}$ and BS_{pour} . The Tables first part is
for BS _{LS,pour} with *. If an approach is the overall best approach, it is only marked by ** . The last row of the table shows
as BS _{pour} and BS _{LS,pour} . The best results in each row are marked with ** whereas we mark results that are the best except
with $\beta = 100$. The instances are from the VFR-small instance set. For easier comparison we added the results of h _{ithb} as well
$m \in \{10, 20\}$ are shown as they performed in general better than the NNs trained with $n = 10$. The results were produced
objective values, the standard deviations, and the runtimes. Only results for NNs trained on instance sizes $n = 20$ and
Table 6.6: This Table shows the first part of the results shown in Figures 6.7, 6.8, 6.9, and 6.10. It contains the average

instances.
at least as good as BS _{LS,pour} and BS _{pour} . Due to the high runtimes during evaluation, not all NNs are evaluated over all
approach, it is only marked by ** . The last row of the table shows how often the individual approach was able to perform
are marked with ** whereas we mark results that are the best except for BS _{LS,pour} with * . If an approach is the overall best
instance set. For easier comparison we added the results of h _{itlb} as well as BS _{pour} and BS _{LS,pour} . The best results in each row
because of there good former performance. The results were produced with $\beta = 100$. The instances are from the TA-small
and the runtimes. Only results for MLP _{basic} and MLP _{TTLB} trained on instance sizes $n = 20$ and $m \in \{10, 20\}$ are shown
Table 6.7: This Table shows the results shown in Figure 6.17. It contains the average objective values, the standard deviations.

	$\sigma_{\overline{x}}$	61.3	102.0	66.6	111.6	114.3	111.2	136.7	179.9	170.3	162.4	229.0	324.2	
BS_{pow}	\overline{x}	1563.5	2072.9	3121.4	3486.5	4595.9	6321.3	6612.3	8587.9	11546.7	*16158.8	21281.1	*49405.2	
r	$\sigma_{\overline{x}}$	54.7	100.6	61.6	114.3	89.3	105.0	132.6	178.2	102.6	140.3	183.5	228.9	
$BS_{LS,pot}$	<u>x</u>	1531.9	2045.5	3052.7	3433.4	4489.5	**6151.3	6528.9	**8425.0	**11209.4	**15971.7	**20896.6	**48739.4	1
0	runt.	22.5	25.9	27.3	46.9	58.8	89.2	205.1	374.1	772.4	4993.4	12809.4	1	$o_{ur} = 5$
$\beta = 10$	$\sigma_{\overline{x}}$	42.7	107.2	72.6	98.5	129.1	98.4	153.9	155.4	177.3	159.9	158.6	1	, BS _{LS,pc}
ITLB.	\overline{x}	1522.5	2041.6	3044.2	**3426.7	4523.1	6302.7	**6494.7	8513.5	11421.4	16159.7	21434.3	ı	$BS_{pow} = 9$
= 20	runt.	7.6	8.5	10.1	43.2	61.3	107.3	370.0	686.1	1565.6	13266.4	47018.1	1	our = 4
LB, $m =$	$\sigma_{\overline{x}}$	43.6	99.7	68.1	116.4	111.0	139.8	197.3	179.5	169.4	195.9	286.2	1	, BS _{LS,p}
MLP _{ITI}	\overline{x}	1523.3	2026.4	**3033.5	3439.8	4487.4	*6209.0	6624.2	*8491.7	*11336.4	16321.5	*21272.0	1	$BS_{pour} = 9$
10	runt.	7.6	8.4	9.9	42.9	58.7	101.3	345.4	614.5	1360.9	'	I	1	ur = 4
B, $m =$	$\sigma_{\overline{x}}$	44.5	98.2	75.2	119.0	108.3	131.4	180.8	180.8	232.4	•	1	1	$BS_{LS,pc}$
MLP _{ITL}	\overline{x}	**1520.5	**2025.8	3034.4	3444.0	**4473.0	6228.9	6615.0	8506.2	11361.1	•	1	ľ	$BS_{pour} = 8$,
0 0	runt.	4.9	4.9	4.8	14.8	14.9	14.7	92.1	91.6	91.8	1083.0	1076.5	48100.6	our = 0
P, $m = 2$	$\sigma_{\overline{x}}$	52.6	116.5	62.1	138.1	114.4	126.3	247.7	219.9	286.1	328.0	544.7	1246.1	$4, BS_{LS,I}$
MI	\overline{x}	1557.6	2082.8	3121.0	3547.8	4583.0	6297.5	6886.3	8739.0	11566.1	16875.0	21849.6	52014.0	$BS_{pour} =$
10	runt.	5.0	4.9	4.8	15.4	15.7	15.6	91.7	91.8	91.4	1026.6	1030.4	1	pour = 0
P, $m =$	$\sigma_{\overline{x}}$	48.3	111.3	80.6	138.6	97.4	170.8	244.8	219.0	322.6	289.2	480.8	1	4, BS _{LS} ,
ML	\overline{x}	1546.3	2068.4	3115.0	3521.8	4531.9	6385.0	6832.9	8646.4	11635.8	16711.5	21637.8	•	$BS_{pour} =$
ces	m	r.	10	20	5	10	20	5	10	20	10	20	20	armed:
Instan	u	20	20	20	50	50	50	100	100	100	200	200	500	Outperfo



Figure 6.18: The figure visualizes the results from Table 6.8. Each boxplot of the NNs represents results over ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots indicate the median and the triangles the mean results. Instances from the TA-small set are used. The dashed line shows the average RPD of BS_{pour} .

Table 6.8: This Table shows the results shown in Figure 6.18. It contains the average objective values, the standard deviations, and the runtimes. Only results for MLP _{basic} and MLP _{ITLB} trained on instance sizes $n = 20$ and $m = 20$ are shown because of
there good former performance. The results were produced with $\beta = 100$. The instances are from the TA-small instance set. For easier comparison we added the results of h _{ith} as well as BS _{bour} and BS _{LS,pour} . The best results in each row are marked
with ** whereas we mark results that are the best except for BS _{LS,pour} with *. If an approach is the overall best approach, it
is only marked by ** . The last row of the table shows how often the individual approach was able to perform at least as good
as BS _{LS,pour} and BS _{pour} .

Instances	MLP,	$m=20,\ \beta$	= 100	MLP _{ITLB}	, $m = 20, \ \beta$	3 = 100	ΕI	LB, $\beta = 1$	100	ITLB	, $\beta = 60$	0	$\mathrm{BS}_{\mathrm{LS,pc}}$	our	$\mathrm{BS}_{\mathrm{pc}}$	ur
т п	<u>x</u>	$\sigma_{\overline{x}}$	runt.	\overline{x}	$\sigma_{\overline{x}}$	runt.	\overline{x}	$\sigma_{\overline{x}}$	runt.	\overline{x}	$\sigma_{\overline{x}}$	runt.	\overline{x}	$\sigma_{\overline{x}}$	\overline{x}	$\sigma_{\overline{x}}$
5	1503.3	64.1	4.5	1470.3	62.7	7.0	1478.3	52.8	20.9	*1469.7	53.3	27.6	**1460.7	66.7	1480.0	59.4
10	2041.3	77.5	4.4	2014.7	79.4	7.8	2024.0	68.2	21.7	**2014.0	60.1	28.6	2028.3	81.1	2048.7	74.1
15	2642.8	66.2	4.4	2585.7	52.7	8.5	2589.3	57.1	18.5	**2578.0	41.0	34.0	2606.0	36.4	2667.3	90.8
10	2984.6	90.3	6.6	2904.4	72.4	15.8	2881.7	72.9	18.6	**2854.0	62.5	55.7	2888.7	67.9	3007.7	90.8
15	3665.2	129.8	6.6	3557.4	170.8	18.0	3571.0	195.8	26.9	**3518.3	152.5	65.0	3537.3	142.6	3720.7	223.7
10	4701.1	94.0	13.5	4586.4	65.3	55.3	4581.0	50.6	49.0	*4560.7	19.7	165.0	**4538.3	32.8	4695.3	46.8
20	8982.5	326.6	35.5	*8752.8	280.5	404.3	8898.3	271.2	174.4	8788.0	254.1	1258.6	**8658.7	222.0	8895.3	239.1
utperformed:	BS_{pour}	$= 4, BS_{LS,E}$	oour = 0	$BS_{pour} =$	= 7, $BS_{LS,pot}$	$_{\rm Ir} = 2$	BS _{pour} =	$= 6, BS_{LS}$	$_{\rm pour} = 3$	$BS_{pour} = 7$, BS _{LS,p}	our = 4	1		1	



Figure 6.19: The figure visualizes the results from Table 6.9. Each boxplot of the NNs represents results over ten runs with individually trained networks over ten instances. The horizontal lines in the boxplots indicate the median and the triangles the mean results. Selected Instances from the VFR-large set are used. The dashed line shows the average RPD of BS_{pour}.

TADIC 0.3. THIS TADIC SHOWS MIC ICOMES SHOWN IN TIGUE 0.13. IN COMMAND MIC AVELAGE OPICENTE SMALLES, MIC SCALLEN, UCVIDAND,

ITLB, $\beta = 100$ BSLS,pour BSpour	$\overline{x} \sigma_{\overline{x}} \overline{x}$	155.9 113.1 11462.2	176.7 15817.7	02.6 19331.2	7 * 21236.6	*28734.3	*34390.4	30771.9	11174.8	9366.3	0.171.9	346.8	718.4
ITLB, $\beta = 100$ BSLS,pour	$\overline{x} = \sigma_{\overline{x}}$	155.9 113.1	176.7	2.6				*	*	*	*4(*53	*63
ITLB, $\beta = 100$ BS _{LS,pou}	x	155.9		19	180.	161.9	190.9	222.7	172.7	443.6	180.2	161.6	361.4
ITLB, $\beta = 100$		**11 ⁻	**15408.3	**18769.2	**20807.3	**28009.3	**33601.8	**30261.0	**40334.4	**48323.4	**39544.8	**52380.8	**62679.5
ITLB, $\beta = 1$	runt.	769.2	1469.2	2320.1	11203.6	'	'	1	'	'	1	'	'
TEI	$\sigma_{\overline{x}}$	140.2	150.0	279.3	186.2	ı	1	•	1	1	•	ı	ı
	\overline{x}	11394.0	15776.6	19432.2	21277.0	I	'	•	'	'	•	1	1
20	runt.	1391.4	2917.4	4952.6	1	1	1	1	1	1	1	1	1
.B, <i>m</i> =	$\sigma_{\overline{x}}$	173.0	259.7	336.3	1	1	1	•	1	1	1	1	1
MLP _{ITI}	\overline{x}	*11292.4	*15714.5	*19223.8	1	I	ı	1	ı	I	1	1	1
20	runt.	92.6	92.5	92.9	1075.0	1070.7	1056.9	5376.0	5442.1	5376.6	17794.8	17544.7	17546.4
P, $m = 1$	$\sigma_{\overline{x}}$	289.6	513.0	532.1	522.0	844.2	923.7	825.9	1145.5	1340.4	1061.7	1486.2	1601.3
ML	le s	59.7	105.8)623.3	1840.2	3487.8	5485.0	2065.6	2586.3	1306.7	2045.4	5511.0	6500.7
Istances	<u>z</u>	115,	16	16	21	ಸ	ನ	က်	4	S	4	ŝ	9

CHAPTER

7

Conclusion

In this work, we considered a variant of the FSP which is the NWFSP-RT. It was shown how the NWFSP could be reduced to the ATSP, and we described which auxiliary data is precomputed to speed up later the state transition in the BS and the dominance check. Further, we explained BS and presented the LBS framework. We described how these general methods can be effectively tailored towards the NWFSP-RT. More specifically this included states for the NWFSP-RT, the branching in BS, how the state transitions work, and we gave a dominance definition applicable to two BS nodes on the same layer. Following that, several lower bounds for the NWFSP-RT including a novel lower bound called ITLB were described. The novel bound combines the idea of Taillard [Tai93] and a lower bound for the ATSP. To integrate the NWFSP-RT into the LBS framework, we described how the training data is generated, and proposed NN types, including the two novel NN types AJNN_{add} and NN_{nearest}. The latter two aim to aggregate the individual "contribution" of each job. They follow the idea of GNNs. The $AJNN_{add}$ uses a MLP and applies it to all job-specific and global features. The NN_{nearest} extends this idea further by passing internal states of the κ nearest neighbors to each job. As a third option, we suggested using a MLP. For all three NN types, we constructed feature sets and combined them to four observations. The observations for the $AJNN_{add}$ and the NN_{nearest} consists of global parts derived from the instance and the state as a whole and job-specific parts that provide further data about each unscheduled job.

We conducted tests on how well the NN types perform when used to guide a BS on the same instance size as they were trained. The tests were performed for every NN types four trained configurations, and the results were compared to using ITLB as a guidance function for a BS. The AJNN_{add} and the NN_{nearest} performed significantly better on several configurations tested with different beam widths than ITLB.

Further, we performed tests on how well the NNs generalize when run for a different number of machines and number of jobs than there were trained. Results show that MLP_{basic} and MLP_{ITLB} generalize well over the number of jobs and also over the number of jobs and machines at the same time. The $AJNN_{add}$ and the $NN_{nearest}$ do not perform well when the number of machines they are trained on is changed. The configurations trained on n = 20 generalize well over the number of jobs. Note, that no NN was able to significantly outperform using ITLB as guidance function.

To better understand why some NNs generalize well and others not, we evaluated the approximation error rates when generalizing over the number of jobs. We observed that the AJNN_{add}, the NN_{nearest}, and the MLP_{ITLB} have error rates within +/-25%RPD for layers resulting in instance sizes they were trained. The MLP_{ITLB} and the MLP_{basic} stay within the +/-25%RPD range. This even holds for BS layers yielding instance sizes MLP_{ITLB} and the MLP_{basic} are not trained. However, the MLP_{basic} has increasing error rates when the instance sizes yielded by the layers go below nine.

During our tests, we constantly outperformed the state-of-the-art results of Pourhejazy et al. [Pou+20] on single instances even though our tests were performed with way smaller beam widths than $\beta = 600$ and without a local search. Overall when running with $\beta = 100$ our approach MLP_{ITLB} trained on instances of size n = 20 and m = 20 outperformed BS_{pour} on 43 out of 46 tested instance classes on average and was able to outperform the former state-of-the-art method BS_{LS,pour} on 11 instance classes. Stable performance and guidance behavior were also shown by MLP_{basic} which is also faster than MLP_{ITLB}. Overall when running with $\beta = 100$, BS_{pour} was on average outperformed on 20 out of 55 tested instance classes by MLP_{basic}.

CHAPTER 8

Future Work and Open Questions

It took up to multiple days to train the NNs with LBS on instances of size n = 20, m = 20. Most of the time was spent creating and labeling new training samples. It might be possible to parallelize the training sample generation to reduce the time needed. Further options may be to perform the nested beam search calls only up to a certain depth and then use a greedy method to complete the solution or use the trained model to approximate the remaining solution cost.

The calculation of the features used as an input for the NNs has a high theoretical computational time complexity and takes most of the time the overall beam search needs. Finding additional promising features that can be computed faster or improving the computation of the current features is important to apply the approach to instances with hundreds of jobs and is left as an open question. Another suggestion is to use a reduced graph to compute the features, i.e., a graph where only a maximum number of neighbors or arcs for each job reside. Reducing the number of neighbors was already done by Joshi, Laurent, and Bresson [JLB19] for the euclidean TSP. Reducing the time needed for computing features would also enable the possibility to train the NNs on instances with more than n = 20 jobs and to use them to solve instances of hundreds of jobs or more.

To further improve the solution quality, one may also combine several NN types during search and profit from their different strengths. For example, the good generalization over the number of jobs of MLP_{ITLB} and the good guidance on instance sizes trained of $NN_{nearest}$ and $AJNN_{add}$ may be combined. This may happen based on the number of unscheduled nodes, or within another NN. Through tighter approximations it might be possible to improve the solution quality further.

A natural next step is the application of the presented approaches and features to other FSP variants. It would be interesting how the presented setup performs for them and if there is a feature set that can be used to learn guidance functions for many variants of the FSP. If successful this might on the long term eliminate the need to design and evaluate guidance functions individually for every variant of a problem.

During our tests no GPU support was used for the evaluation and training of the NNs. This could speed up the evaluation of the NNs possibly a lot, especially when evaluating all nodes of one layer in batches. Testing the impact of batch evaluations on the GPU should therefore be considered in future works.

List of Algorithms

3.1	General BS Algorithm	17
3.2	General LBS Algorithm [HR21]	22
3.3	Data generation for LBS [HR21]. This piece of code is included in Algorithm 3.1 at line 32, if BS is called with data generation. For simplicity, it is assumed that all arguments additionally needed are passed to the BS algorithm.	22

Glossary

- ${\bf MLP_{ITLB}}$ Denotes a MLP that uses observation $O_{\rm glob,itlb}..$ 44, 45, 53, 58–61, 67, 69, 71, 73–75

Acronyms

- AJNN_{add} aggregated jobs neural network. 2, 3, 34–36, 41, 42, 44, 45, 53, 57–59, 73–75
- **BS_{LS.pour}** BS with local search from [Pou+20]. 44–49, 58, 59, 61, 65–67, 69, 71, 74
- BS_{pour} BS from [Pou+20]. 45–49, 51, 52, 54–61, 64–71, 74
- h_{itlb} ITLB as heuristic function. 43–45, 50, 53–56, 59–61, 65–67, 69, 71
- ${\bf NN_{nearest}}$ nearest neighbors neural network. 2, 3, 35–38, 41, 42, 44, 45, 53, 57, 59, 73–75
- ACO Ant Colony Optimization. 14
- ATSP asymmetric Traveling Salesperson Problem. 2, 9, 13, 20, 73
- **B&B** Branch and Bound. 12
- **BS** beam search. 2, 3, 11, 14, 21, 40, 41, 43–45, 54, 59–63, 73, 74
- **FSP** Flow Shop Schedueling Problem. 1, 6, 11–14, 18, 19, 34, 73, 75, 76
- GA Genetic Algorithm. 12, 13
- GNN graph neural network. 2, 73
- **ITLB** improved Taillard lower bound. 2, 3, 41, 43–49, 51–57, 73, 74
- **LB** lower bound. 12, 18, 39, 40
- **LBS** learning beam search. 2, 3, 11, 21, 22, 33, 34, 44, 45, 73, 75
- LCS Longest Common Subsequence Problem. 21
- ML Machine Learning. 11
- MLP multilayer perceptron. 2, 34, 35, 37, 41, 44, 73

- **NBS** nested BS. 21, 44, 59–63
- ${\bf NEH}\,$ Nawaz-Enscore-Ham. 12
- NN neural network. 3, 18, 19, 21, 33–35, 37, 40, 43–68, 70, 71, 73–76
- **NWFSP** no-wait FSP. 6, 9, 13, 73
- **NWFSP-RT** no-wait FSP with release times. 1, 2, 6, 11, 14, 23, 34, 43, 73
- **PFSP** Permutation Flow Shop Schedueling Problem. 12–14, 26
- ${\bf RL}$ reinforcement learning. 21
- **RPD** relative percentage difference. 44–49, 51–57, 59–64, 68, 70
- ${\bf TS}\,$ Tabu Search. 12, 13
- **VRP** Vehicle Routing Problem. 20

Bibliography

[ACE06]	A. Agarwal, S. Colak, and E. Eryarsoy. "Improvement heuristic for the flow-shop scheduling problem: An adaptive-learning approach". In: <i>European Journal of Operational Research</i> 169.3 (2006), pp. 801–815. DOI: 10.1016/j.ejor.2004.06.039.
[All16]	A. Allahverdi. "A survey of scheduling problems with no-wait in process". In: <i>European Journal of Operational Research</i> 255.3 (2016), pp. 665–686. DOI: 10.1016/j.ejor.2016.05.036.
[And+97]	E. J. Anderson et al. "Machine scheduling". In: <i>Local search in combinatorial optimization</i> 11 (1997). Publisher: Wiley Chichester, UK, pp. 361–414.
[BDG99]	L. Bianco, P. Dell'Olmo, and S. Giordani. "Flow Shop No-Wait Scheduling With Sequence Dependent Setup Times And Release Dates". In: <i>INFOR:</i> <i>Information Systems and Operational Research</i> 37.1 (1999). Publisher: Taylor & Francis, pp. 3–19. DOI: 10.1080/03155986.1999.11732365.
[Blu05]	C. Blum. "Beam-ACO—hybridizing ant colony optimization with beam search: an application to open shop scheduling". In: <i>Computers & Operations Research</i> 32.6 (2005), pp. 1565–1591. DOI: 10.1016/j.cor.2003.11.018.
[CDS70]	H. G. Campbell, R. A. Dudek, and M. L. Smith. "A Heuristic Algorithm for the n-Job, m-Machine Sequencing Problem". In: <i>Management Science</i> 16.10 (1970), B-630-B-637. DOI: 10.1287/mnsc.16.10.B630.
[Chr72]	N. Christofides. "Technical Note—Bounds for the Travelling-Salesman Prob- lem". In: <i>Operations Research</i> 20.5 (1972). Publisher: INFORMS, pp. 1044– 1056. DOI: 10.1287/opre.20.5.1044.
[DJ64]	R. A. Dudek and O. F. T. Jr. "Development of M-Stage Decision Rule for Scheduling n Jobs through M Machines". In: <i>Operations Research</i> 12.3 (1964), pp. 471–497. DOI: 10.1287/opre.12.3.471.
[EM95]	H. Emmons and K. Mathur. "Lot sizing in a no-wait flow shop". In: <i>Operations Research Letters</i> 17.4 (1995), pp. 159–164. DOI: 10.1016/0167-6377 (95) 00008-8.

- [FVVF18] V. Fernandez-Viagas, J. M. S. Valente, and J. M. Framinan. "Iteratedgreedy-based algorithms with beam search initialization for the permutation flowshop to minimise total tardiness". In: *Expert Systems with Applications* 94 (2018), pp. 58–69. DOI: 10.1016/j.eswa.2017.10.050.
- [GG64] P. C. Gilmore and R. E. Gomory. "Sequencing a One State-Variable Machine: A Solvable Case of the Traveling Salesman Problem". In: *Operations Research* 12.5 (1964), pp. 655–679. DOI: 10.1287/opre.12.5.655.
- [GJS76] M. R. Garey, D. S. Johnson, and R. Sethi. "The Complexity of Flowshop and Jobshop Scheduling". In: *Mathematics of Operations Research* 1.2 (1976), pp. 117–129. DOI: 10.1287/moor.1.2.117.
- [GML20] J. N. D. Gupta, A. Majumder, and D. Laha. "Flowshop scheduling with artificial neural networks". In: *Journal of the Operational Research Society* 71.10 (2020). Publisher: Taylor & Francis, pp. 1619–1637. DOI: 10.1080/ 01605682.2019.1621220.
- [Gra66] R. L. Graham. "Bounds for certain multiprocessing anomalies". In: The Bell System Technical Journal 45.9 (1966). Conference Name: The Bell System Technical Journal, pp. 1563–1581. DOI: 10.1002/j.1538-7305.1966. tb01709.x.
- [Gra+79] R. L. Graham et al. "Optimization and Approximation in Deterministic Sequencing and Scheduling: a Survey". In: Annals of Discrete Mathematics. Ed. by P. L. Hammer, E. L. Johnson, and B. H. Korte. Vol. 5. Discrete Optimization II. Elsevier, 1979, pp. 287–326. DOI: 10.1016/S0167– 5060 (08) 70356-X.
- [HR21] M. Huber and G. Raidl. "Learning Beam Search: Utilizing Machine Learning to Guide Beam Search for Solving Combinatorial Optimization Problems".
 In: Machine Learning, Optimization, and Data Science 7th International Conference, LOD 2021. Vol. 11943. LNCS. to appear. Springer, 2021.
- [HS96] N. G. Hall and C. Sriskandarajah. "A Survey of Machine Scheduling Problems with Blocking and No-Wait in Process". In: *Operations Research* 44.3 (1996). Publisher: INFORMS, pp. 510–525. DOI: 10.1287/opre.44.3.510.
- [JLB19] C. K. Joshi, T. Laurent, and X. Bresson. "An Efficient Graph Convolutional Network Technique for the Travelling Salesman Problem". In: arXiv:1906.01227 [cs, stat] (2019). DOI: 10.48550/arXiv.1906.01227.
- [Joh54] S. M. Johnson. "Optimal two- and three-stage production schedules with setup times included". In: *Naval Research Logistics Quarterly* 1.1 (1954), pp. 61–68. DOI: 10.1002/nav.3800010110.
- [Jos+21] C. K. Joshi et al. "Learning TSP Requires Rethinking Generalization". In: arXiv:2006.07054 (2021). DOI: 10.4230/LIPIcs.CP.2021.33.

- [JSW06] F. Jin, S.-j. Song, and C. Wu. "A New Beam Search Algorithm for the Large-Scale Permutation FSP". In: 2006 International Conference on Machine Learning and Cybernetics. 2006, pp. 1–6. DOI: 10.1109/ICMLC.2006. 258806.
- [KHW19] W. Kool, H. van Hoof, and M. Welling. "Attention, Learn to Solve Routing Problems!" In: arXiv:1803.08475 (2019). DOI: 10.48550/arXiv.1803. 08475.
- [KK07] P. J. Kalczynski and J. Kamburowski. "On no-wait and no-idle flow shops with makespan criterion". In: European Journal of Operational Research 178.3 (2007), pp. 677–685. DOI: 10.1016/j.ejor.2006.01.036.
- [KV18] B. Korte and J. Vygen. Combinatorial Optimization: Theory and Algorithms.
 6th ed. Algorithms and Combinatorics. Berlin Heidelberg: Springer-Verlag, 2018. DOI: 10.1007/978-3-662-56039-6.
- [LH05] T. Ladhari and M. Haouari. "A computational study of the permutation flow shop problem based on a tight lower bound". In: Computers & Operations Research 32.7 (2005), pp. 1831–1847. DOI: 10.1016/j.cor.2003.12. 001.
- [Low76] B. T. Lowerre. "The Harpy speech recognition system". PhD thesis. Carnegie-Mellon Univ., Pittsburgh, PA., 1976.
- [LS00] I. Lee and M. J. Shaw. "A neural-net approach to real time flow-shop sequencing". In: Computers & Industrial Engineering 38.1 (2000), pp. 125– 147. DOI: 10.1016/S0360-8352 (00) 00034-6.
- [NEH83] M. Nawaz, E. E. Enscore, and I. Ham. "A heuristic algorithm for the m-machine, n-job flow-shop sequencing problem". In: Omega 11.1 (1983), pp. 91–95. DOI: 10.1016/0305-0483(83)90088-9.
- [Ni+21] F. Ni et al. "A Multi-Graph Attributed Reinforcement Learning based Optimization Algorithm for Large-scale Hybrid Flow Shop Scheduling Problem". In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining. KDD '21. New York, NY, USA: Association for Computing Machinery, 2021, pp. 3441–3451. DOI: 10.1145/3447548.3467135.
- [Pas00] C. A. S. Passos. "A Beam Search Based Algorithm to Solve Flowshop Scheduling Problems with Constraints on Shared Resources". In: *IFAC Proceedings Volumes* 33.17 (2000), pp. 675–679. DOI: 10.1016/S1474– 6670 (17) 39484–3.
- [Pie60] J. Piehler. "Ein Beitrag zum Reihenfolgeproblem". In: Unternehmensforschung
 4.3 (1960), pp. 138–142. DOI: 10.1007/BF01963410.
- [PKL00] Y. Park, S. Kim, and Y.-H. Lee. "Scheduling jobs on parallel machines applying neural network and heuristic rules". In: *Computers & Industrial Engineering* 38.1 (2000), pp. 189–202. DOI: 10.1016/S0360-8352(00) 00038-3.

- [Pou+20] P. Pourhejazy et al. "Improved Beam Search for Optimizing No-Wait Flowshops With Release Times". In: *IEEE Access* 8 (2020), pp. 148100–148124.
 DOI: 10.1109/ACCESS.2020.3015737.
- [PSW91] C. N. Potts, D. B. Shmoys, and D. P. Williamson. "Permutation vs. nonpermutation flow shop schedules". In: Operations Research Letters 10.5 (1991), pp. 281–284. DOI: 10.1016/0167-6377 (91) 90014-G.
- [Ram+11] T. R. Ramanan et al. "An artificial neural network based heuristic for flow shop scheduling problems". In: *Journal of Intelligent Manufacturing* 22.2 (2011), pp. 279–288. DOI: 10.1007/s10845-009-0287-5.
- [RR72] S. S. Reddi and C. V. Ramamoorthy. "On the Flow-Shop Sequencing Problem with No Wait in Process†". In: Journal of the Operational Research Society 23.3 (1972), pp. 323–331. DOI: 10.1057/jors.1972.52.
- [RYY21] J. Ren, C. Ye, and F. Yang. "Solving flow-shop scheduling problem with a reinforcement learning algorithm that generalizes the value function with neural network". In: *Alexandria Engineering Journal* 60.3 (2021), pp. 2787– 2800. DOI: 10.1016/j.aej.2021.01.030.
- [SB99] I Sabuncuoglu and M Bayiz. "Job shop scheduling with beam search". In: European Journal of Operational Research 118.2 (1999), pp. 390–412. DOI: 10.1016/S0377-2217 (98) 00319-1.
- [Sil+17] D. Silver et al. "Mastering the game of Go without human knowledge". In: Nature 550.7676 (2017). Publisher: Nature Publishing Group, pp. 354–359.
 DOI: 10.1038/nature24270.
- [Sul00] S. M. A. Suliman. "A two-phase heuristic approach to the permutation flowshop scheduling problem". In: *International Journal of Production Economics* 64.1 (2000), pp. 143–152. DOI: 10.1016/S0925-5273 (99) 00053-5.
- [Tai90] E. Taillard. "Some efficient heuristic methods for the flow shop sequencing problem". In: European Journal of Operational Research 47.1 (1990), pp. 65–74. DOI: 10.1016/0377-2217 (90) 90090-X.
- [Tai93] E. Taillard. "Benchmarks for basic scheduling problems". In: European Journal of Operational Research 64.2 (1993), pp. 278–285. DOI: 10.1016/ 0377-2217 (93) 90182-M.
- [Vai03] G. L. Vairaktarakis. "Simple Algorithms for Gilmore–Gomory's Traveling Salesman and Related Problems". In: *Journal of Scheduling* 6.6 (2003), pp. 499–520. DOI: 10.1023/A:1026200209386.
- [VRF15] E. Vallada, R. Ruiz, and J. Framinan. "New hard benchmark for flowshop scheduling problems minimising makespan". In: European Journal of Operational Research 240 (2015), pp. 666–677. DOI: 10.1016/j.ejor.2014. 07.033.

- [Wis72] D. A. Wismer. "Solution of the Flowshop-Scheduling Problem with No Intermediate Queues". In: Operations Research (1972). Publisher: INFORMS. DOI: 10.1287/opre.20.3.689.
- [ZLC15] S. Zhou, X. Li, and H. Chen. "An Estimation of Distribution Algorithm for Minimizing Makespan on a No-Wait Flow-Shop". In: *IIE Annual Conference*. *Proceedings* (2015), pp. 1070–1077. DOI: 10.1080/00207543.2016. 1140920.