

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Provably Good Solutions to the Knapsack Problem via Neural Networks of Bounded Size*

Christoph Hertrich

London School of Economics and Political Science, c.hertrich@lse.ac.uk

Martin Skutella

Technische Universität Berlin, martin.skutella@tu-berlin.de

The development of a satisfying and rigorous mathematical understanding of the performance of neural networks is a major challenge in artificial intelligence. Against this background, we study the expressive power of neural networks through the example of the classical NP-hard Knapsack Problem. Our main contribution is a class of recurrent neural networks (RNNs) with rectified linear units that are iteratively applied to each item of a Knapsack instance and thereby compute optimal or provably good solution values. We show that an RNN of depth four and width depending quadratically on the profit of an optimum Knapsack solution is sufficient to find optimum Knapsack solutions. We also prove the following tradeoff between the size of an RNN and the quality of the computed Knapsack solution: for Knapsack instances consisting of n items, an RNN of depth five and width w computes a solution of value at least $1 - \mathcal{O}(n^2/\sqrt{w})$ times the optimum solution value. Our results build upon a classical dynamic programming formulation of the Knapsack Problem as well as a careful rounding of profit values that are also at the core of the well-known fully polynomial-time approximation scheme for the Knapsack Problem. A carefully conducted computational study qualitatively supports our theoretical size bounds. Finally, we point out that our results can be generalized to many other combinatorial optimization problems that admit dynamic programming solution methods, such as various Shortest Path Problems, the Longest Common Subsequence Problem, and the Traveling Salesperson Problem.

Key words: Neural Networks; Expressive Power; Knapsack Problem; Dynamic Programming

*An extended abstract of this article, including Figures 1–7, appeared in the Proceedings of the AAAI Conference on Artificial Intelligence, volume 35, 7685–7693, (Hertrich and Skutella 2021); see <https://ojs.aaai.org/index.php/AAAI/article/view/16939>; copyright © 2021, Association for the Advancement of Artificial Intelligence.

1. Introduction

Deep learning and neural networks (NNs) are at the heart of some of the greatest advances in modern computer science. They enable huge breakthroughs in applications like computer vision, translation, speech recognition, and autonomous driving, to name just a few; see, e.g., (LeCun et al. 2015). While numerous computational studies present impressive empirical proof of neural networks' computational power, we are still far away from a more rigorous theoretical explanation of these observations.

Apart from the popular applications named above, it has been shown that NNs have high potential for practically solving combinatorial optimization (CO) problems or empirically improving classical solution methods (Bengio et al. 2021). For example, Yang et al. (2018) and Xu et al. (2020b) utilize NNs in order to empirically enhance *dynamic programming*, a very classical CO method. While the methods used in these papers indeed provide fast and empirically near-optimal solutions, their use of NNs makes it virtually impossible to give theoretical optimality or worst-case approximation guarantees. Motivated by this imbalance, and focusing on the Knapsack Problem, which is a prime example of CO problems that can be solved via dynamic programming, we investigate the following question:

*Which neural network size is theoretically sufficient
to find solutions of provable quality for the Knapsack Problem?*

We give an answer to this question by presenting a class of carefully constructed NNs with provable quality guarantees and support our size bounds by a computational study. Finally, we argue that our approach is not at all specific for the Knapsack Problem, but can be generalized to many other CO problems, e.g., various Shortest Path Problems, the Longest Common Subsequence Problem, and the Traveling Salesperson Problem.

It is worth to note that it is by no means obvious that classical algorithms can be translated to neural networks with bounded size. In particular, as discussed by Hertrich and Sering (2021), even for efficiently (polynomial-time) solvable problems it is a priori unclear whether polynomially sized neural networks can compute provably good solutions.

The Knapsack Problem

The Knapsack Problem constitutes one of the oldest and most studied problems in Combinatorial Optimization (CO). Given a set of items with certain profit and size values, as well as a Knapsack capacity, the Knapsack Problem asks for a subset of items with maximum total profit such that the total size of the subset does not exceed the capacity.

The Knapsack Problem is one of Karp’s 21 original NP-complete problems (Karp 1972) and has numerous applications in a wide variety of fields, ranging from production and transportation, over finance and investment, to network security and cryptography. It often appears as a subproblem at the core of more complex problems; see, e.g., (Kellerer et al. 2004, Martello and Toth 1990). This fact substantiates the Knapsack Problem’s prominent importance as one of the key problems in CO. In particular, the Knapsack Problem is frequently used as a testbed for measuring the progress of various exact and heuristic solution approaches and computational methods such as, e.g., integer programming, constraint programming, or evolutionary algorithms. In integer programming, for example, the Knapsack Problem and so-called ‘Knapsack Inequalities’ play a central role, both with respect to theory as well as in the development of modern computational methods; see, e.g., (Bertsimas and Weismantel 2005, Fischetti and Lodi 2010). The Knapsack Problem is therefore a natural and important object of study in order to advance our theoretical understanding of neural networks and get closer to a rigorous explanation of their stunning success in so many applications, including miscellaneous optimization problems.

Related work

The idea of using neural networks (NNs) to practically solve CO problems became popular with the work of Hopfield and Tank (1985). Hopfield NNs are special versions of recurrent neural networks (RNNs) that find solutions to optimization problems by converging towards a minimum of an energy function. Smith (1999) reviews this early stream of research. While most authors mainly focus on the Traveling Salesperson Problem (TSP), Ohlsson et al. (1993) study a so-called mean field NN for (generalizations of) the Knapsack Problem and empirically assess the quality of its solutions.

While there has been less research at the intersection of CO and NNs in the 2000s, modern advances in the area of deep learning have boosted the interest in this direction again; see (Bengio et al. 2021) for a general review and (Cappart et al. 2021) for a focused review on graph neural networks. Common applications include speeding up solvers for mixed-integer linear programs, for instance, by automatically learning on which variables to branch in branch-and-bound algorithms; see (Lodi and Zarpellon 2017) for a survey. Machine learning has also been applied to modeling aspects of CO, as reviewed by Lombardi and Milano (2018), and to several specific CO problems, where the TSP is often one of them (Bello et al. 2016, Emami and Ranka 2018, Khalil et al. 2017, Kool et al. 2019,

Nowak et al. 2017, Vinyals et al. 2015). The different methods used by these authors include feedforward and recurrent neural networks, reinforcement learning, attention mechanisms, pointer networks, graph embeddings, and graph neural networks. Expressivity results similar to the paper at hand, but using completely different techniques, have been achieved for the Maximum Flow Problem by Hertrich and Sering (2021).

There have also been specific applications of neural networks to the Knapsack Problem. For example, Bello et al. (2016) utilize an RNN trained by reinforcement learning and Gu and Hao (2018) use a pointer network to find empirically good knapsack solutions. Li et al. (2021) derive heuristics inspired by game theory to solve a non-linear knapsack version whose objective function is given by a neural network.

Particularly related to our work are interactions between neural networks and dynamic programming algorithms. For example, Yang et al. (2018) and Xu et al. (2020b) use NNs to speed up dynamic programming algorithms for CO problems. The key difference to our work, however, is that NNs are used as heuristics in these papers, making it virtually impossible to give any meaningful worst-case performance guarantees. Another interesting research stream deals with the learnability of algorithms. In this context, Xu et al. (2020a) develop the concept of *algorithmic alignment* and show that dynamic programming algorithms align well with graph neural networks. Veličković et al. (2020) perform a more empirical study in this direction. These results concentrate on learnability in contrast to our focus on expressivity. Still, they agree with the message of our paper that dynamic programming is a good paradigm for bringing classical algorithms and neural networks closer together.

The recent success of deep neural networks has also triggered a lot of research on their general expressivity. As we do in this paper, many authors focus on the simple but practically powerful model of feedforward NNs with activations in the form of rectified linear units (ReLU). Since Glorot et al. (2011) corroborated their empirical success, such ReLU NNs have been established as a standard model in machine learning within the past decade. ReLU NNs can compute any continuous piecewise linear function (Arora et al. 2018, Goodfellow et al. 2013). This fact implies universal approximation properties. Recent progress concerning the classes of piecewise linear functions representable by ReLU neural networks with a certain architecture was made by Hertrich et al. (2021) and Haase et al. (2023). Results on the exact expressivity boosted success with respect to understanding the

computational complexity of the training problem for ReLU networks (Arora et al. 2018, Bertschinger et al. 2022, Froese et al. 2022, Goel et al. 2021, Khalife and Basu 2022). A variety of results has been achieved on depth vs. width tradeoffs (Arora et al. 2018, Eldan and Shamir 2016, Hanin 2019, Hanin and Sellke 2017, Liang and Srikant 2017, Nguyen et al. 2018, Safran and Shamir 2017, Telgarsky 2015, 2016, Yarotsky 2017). Closely related are investigations concerning the number and structure of linear regions that NNs with certain size and depth may have (Hanin and Rolnick 2019, Montufar et al. 2014, Pascanu et al. 2014, Raghu et al. 2017). Serra et al. (2018) use mixed-integer programming for precisely counting the number of such regions. Mukherjee and Basu (2017) prove size lower bounds to represent Boolean functions with NNs of limited depth.

Our contribution

We present a rigorous mathematical study on the expressivity of NNs through the example of the NP-hard Knapsack Problem. To this end, we show that there is a class of feedforward ReLU NNs of bounded size that compute *provably* good solutions to the NP-hard Knapsack Problem. In Section 3, we first present such an NN of depth $\mathcal{O}(n)$ and width $\mathcal{O}((p^*)^2)$ that always finds the exact value of an optimum Knapsack solution. Here, n is the number of items in the Knapsack instance, and p^* is an a priori known upper bound on the value of an optimum solution. More precisely, the optimum solution value is found by iteratively applying an RNN of depth four and width $\mathcal{O}((p^*)^2)$ to the n items of a Knapsack instance. As p^* can, e.g., be chosen as the total size of all items, the RNN's width is pseudo-polynomially bounded in the input size of the Knapsack instance. Due to the Knapsack Problem's NP-hardness, however, there is no polynomial-size NN with polynomial-time computable weights that always finds the optimum solution value, unless $P = NP$; compare the discussion in the next subsection.

In Section 4, we prove that the width of the NNs can be drastically decreased while still obtaining solution values of provable quality in the worst case. We construct an RNN of depth five and fixed width w which, when applied iteratively to the n items of a Knapsack instance, always produces a solution value of at least $1 - \mathcal{O}(n^2/\sqrt{w})$ times the optimum solution value. In particular, an ε -approximate solution value can be guaranteed by choosing width $w \in \mathcal{O}(n^4/\varepsilon^2)$. If we require the weights of the neural network to be polynomial-time computable, then the dependence of the width on ε is unavoidable, unless $P = NP$; compare the discussion in the next subsection. To the best of our knowledge, our results

establish the first rigorous tradeoff between the size of neural networks for CO problems and their worst-case solution quality.

Even though we cannot show theoretical lower bounds beyond what is directly implied by NP-hardness, we provide empirical evidence for the superlinear dependence on p^* (and $1/\varepsilon$) in Section 5.

The idea behind our construction of the NNs is to mimic the classical dynamic program for the Knapsack Problem. More precisely, the output neurons of the RNN can be seen as elements of the dynamic programming state space while the hidden neurons and the network itself implement the recursive dynamic programming formula. Here, the main technical difficulty is to always filter out the correct entries of the previous state space (input neurons) needed in the recursive formula. In addition, our NNs of fixed width rely on a subtle variant of the rounding procedure that turns the pseudo-polynomial dynamic program into a fully polynomial-time approximation scheme for the Knapsack Problem.

In this paper, the Knapsack Problem mainly serves as a prominent showcase for a novel approach to the rigorous analysis of neural networks' expressivity. This approach is by no means specific for the Knapsack Problem. In Section 6, we discuss how it can be applied to NNs for other combinatorial optimization problems that can be solved via dynamic programming. In particular, we establish similar results for the Longest Common Subsequence Problem, the Single-Source and All-Pairs Shortest Path Problems, as well as the NP-hard Traveling Salesperson Problem and the Constrained Shortest Path Problem. For the latter problem one can show similar results on the tradeoff between the size of NNs and their solution quality.

Uniformity and Lower Bounds Implied by NP-Hardness

As mentioned earlier, NP-hardness makes it unlikely that we can get rid of the dependence on p^* (in the exact case) and ε (in the approximate case) for our constructions. However, strictly speaking, NP-hardness does not directly imply conditional lower bounds on the size of a neural network. Similar to the complexity of Boolean circuits (Arora and Barak 2009), one needs an additional *uniformity* assumption to deduce lower bounds conditioned on $P \neq NP$, which we will explain now.

We call a construction of neural networks *uniform*, if one can compute the weights of the neural network in polynomial time. Having a uniform neural network construction of polynomial size to solve a certain problem implies that the respective problem can be

solved in polynomial time: for a given instance, one first computes the neural network and then executes it to compute a solution to the instance. Hence, a uniform neural network construction of polynomial size for the Knapsack Problem cannot exist unless $P = NP$. This implies that the dependences on p^* (in the exact case) and ε (in the approximate case) are unavoidable if one aims for uniform constructions. Note that the constructions provided in this paper are trivially uniform because we explicitly provide the weights of the neural networks.

Without the uniformity assumption, however, it might be possible that $P \neq NP$ is true and there still exist polynomial-size neural networks for the Knapsack Problem. We might just be unable to compute these networks (in polynomial time or at all), which is why they cannot be used to obtain polynomial-time algorithms and thus obtain a contradiction to $P \neq NP$.

2. Preliminaries

Neural networks with rectified linear units

We use definitions and notations similar to (Shalev-Shwartz and Ben-David 2014, Chapter 20). A *feedforward neural network with rectified linear units*, abbreviated by ReLU NN, or simply NN, is a finite, directed, acyclic graph (V, E) , equipped with arc weights $w_{uv} \in \mathbb{R}$, for each $(u, v) \in E$, and node biases $b_v \in \mathbb{R}$, for each node $v \in V \setminus V_0$. Here, V_0 is the set of nodes with in-degree zero. The nodes in V are called *neurons*. The *depth* k is the length of a longest path in the graph. In the following we suppose that neurons are grouped into *layers* $V = V_0 \cup V_1 \cup \dots \cup V_k$ such that the layer index strictly increases along each arc.¹ Further, we assume that V_0 and V_k are precisely the sets of neurons with in-degree and out-degree zero, respectively. Consequently, they are called *input neurons* and *output neurons*, respectively. Neurons in $V \setminus (V_0 \cup V_k)$ are called *hidden neurons*. Let $n_\ell = |V_\ell|$ be the number of neurons in the ℓ -th layer. The *width* and *size* of the NN are defined to be $\max\{n_1, \dots, n_{k-1}\}$ and $\sum_{\ell=1}^{k-1} n_\ell$, respectively.

Every NN computes a function $\mathbb{R}^{n_0} \rightarrow \mathbb{R}^{n_k}$ as follows. Given an input vector $x \in \mathbb{R}^{n_0}$, we associate an activation $a(v)$ with every neuron $v \in V \setminus V_0$ and an output $o(v)$ with every neuron $v \in V \setminus V_k$. First, the output values $o(v)$, $v \in V_0$, of the n_0 input neurons equal the n_0

¹ Some authors only allow connections between successive layers. One can create such a structure by adding additional neurons propagating the values of neurons from former layers through the network. For our purposes, however, it is convenient to omit this restriction.




Figure 1 An NN with two input neurons, labeled x_1 and x_2 , one hidden neuron, labeled with the shape of the rectifier function, and one output neuron, labeled y . The arcs are labeled with their weights and all biases are zero. The network has depth 2, width 1, and size 1. It computes the function

$$x \mapsto y = x_2 - \max\{0, x_2 - x_1\} = \min\{x_1, x_2\}.$$

components of input vector x . Second, the activation of a neuron $v \in V \setminus V_0$ is the weighted sum of outputs of all predecessors plus its bias, that is, $a(v) = b_v + \sum_{u: (u,v) \in E} w_{uv} o(u)$. Third, for each hidden neuron $v \in V \setminus (V_0 \cup V_k)$, the output is determined by $o(v) = \sigma(a(v))$, where σ is the so-called *activation function*. In this paper, σ is always the *rectifier function* $\sigma(z) = \max\{0, z\}$ applied pointwise. Neurons having this activation function are called *rectified linear units* (ReLUs). Finally, the output vector $y \in \mathbb{R}^{n_k}$ consists of the n_k activation values $a(v)$ of the n_k output neurons $v \in V_k$. Figure 1 gives an example, which will also be used as a subnetwork in later sections.

Since feedforward NNs have a fixed input size, a common way of handling sequential inputs of arbitrary length is to use *recurrent neural networks* (RNNs). This type of NNs has become very popular, e.g., for tasks in language (Sutskever et al. 2014, Bahdanau et al. 2015) or speech processing (Graves et al. 2013). Essentially, an RNN is a feedforward NN that is used repeatedly for every piece of the input sequence and maintains a hidden state by passing (part of) its output in each step as an additional input to the next step. More precisely, in the i -th step, the input of the RNN consists of the i -th input vector x_i , as well as, the previous hidden state vector h_{i-1} . In the same manner as a feedforward NN described above, it then computes the i -th output vector y_i , as well as, the new hidden state vector h_i . The basic structure of an RNN is shown in Figure 2. Sometimes it holds that $y_i = h_i$, that is, the i -th output is actually equal to the i -th hidden state.

Notations and Algorithms for the Knapsack Problem

An instance of the Knapsack Problem consists of n items $1, 2, \dots, n$, where each item $i \in [n]$ comes with a given profit $p_i \in \mathbb{N}$ and size $s_i \in]0, 1]$, together with a Knapsack that can hold any subset $M \subseteq [n]$ of items of total size $\sum_{i \in M} s_i$ at most 1. The task is to find such a subset $M \subseteq [n]$ that maximizes the total profit $\sum_{i \in M} p_i$. Here and in the following, we




Figure 2 Basic structure of an (unfolded) RNN.

use $\mathbb{N} := \{1, 2, 3, \dots\}$ to denote the natural numbers (without zero), and for every $k \in \mathbb{N}$, we let $[k] := \{1, 2, \dots, k\}$.

We outline a classical dynamic programming formulation for the Knapsack Problem. Let p^* be an upper bound on the optimum solution value, e.g., $p^* = \sum_{i=1}^n p_i$. For $i \in [n]$ and $p \in [p^*]$, let

$$f(p, i) := \min \left\{ \sum_{j \in M} s_j \mid M \subseteq [i], \sum_{j \in M} p_j \geq p \right\}$$

be the minimum size of a subset of the first i items with total profit at least p . With $f(p, i) := 0$ for $p \leq 0$ and $f(p, 0) := +\infty$ for $p \in [p^*]$, the values of f can be computed recursively by

$$f(p, i) = \min \{ f(p, i - 1), f(p - p_i, i - 1) + s_i \} \quad (1)$$

for $i \in [n]$, $p \in [p^*]$, where the first option corresponds to not using the i -th item, while the second option corresponds to using it. The optimum solution value is $\max\{p \in [p^*] \mid f(p, n) \leq 1\}$, and the optimum subset can easily be found by backtracking. The runtime of the dynamic program is $\mathcal{O}(np^*)$, thus pseudo-polynomial in the input size.

Due to NP-hardness of the Knapsack Problem, one cannot expect to find an exact algorithm with polynomial running time. However, by carefully downscaling and rounding the profit values in the dynamic program, one can obtain a *fully polynomial-time approximation scheme* (FPTAS). That is, for each $\varepsilon > 0$, one can compute a feasible solution with guaranteed profit of at least $1 - \varepsilon$ times the optimal profit, with running time polynomial in the input size and $1/\varepsilon$. For more details, we refer to the books by Hochbaum (1997), Vazirani (2001), or Williamson and Shmoys (2011).

Usually, the Knapsack Problem is defined with integer size values $s_i \in \mathbb{N}$ and some Knapsack capacity $S \in \mathbb{N}$, bounding the total size of chosen items. Dividing all item sizes

by S transforms such an instance into an instance of the type considered here. For the case of integral item sizes, there is also a pseudo-polynomial dynamic programming formulation parameterized by the size instead of the profit values; see, e.g., (Kleinberg and Tardos 2006, Section 6.4). Our construction in Section 3 can analogously be applied to this formulation. This variant, however, does not easily extend to an FPTAS. We therefore stick to the variant parametrized by the profit values as introduced above.

3. An Exact RNN for the Knapsack Problem

In this section we introduce the *DP-NN*, an NN that exactly executes the dynamic program described in Section 2. In fact, the DP-NN is an RNN that receives the items one by one and computes the state space of the dynamic program for the items seen so far. In the first subsection we explain the high-level idea, before we provide further details in the second subsection.

High-level idea of the construction

Like the dynamic program described in Section 2, the DP-NN requires a fixed upper bound p^* on the optimal objective value of the Knapsack Problem. We relax this condition in Section 4, when we investigate how the FPTAS for the Knapsack Problem can be implemented as an NN.

In the i -th step, the DP-NN receives $p^* + 2$ inputs, namely $f(p, i - 1)$ for $p \in [p^*]$, as well as p_i and s_i . It computes p^* output values, namely $f(p, i)$ for $p \in [p^*]$. Hence, overall it has $p^* + 2$ input neurons and p^* output neurons. Figure 3 illustrates the recurrent structure of the NN, which computes the state space of the dynamic program.




Figure 3 Recurrent structure of the DP-NN to solve the Knapsack Problem.

In the following it is very important to distinguish fixed parameters of the NN from activation and output values of neurons that depend on the particular Knapsack instance. We denote the latter by bold symbols in order to make the difference visible. Moreover,




Figure 4 Desirable architecture for computing $f_{\text{out}}(p)$, $p \in [p^*]$, from the inputs. However, the existence of an edge (nonzero weight) depends critically on the input value p_{in} , which is not allowed.

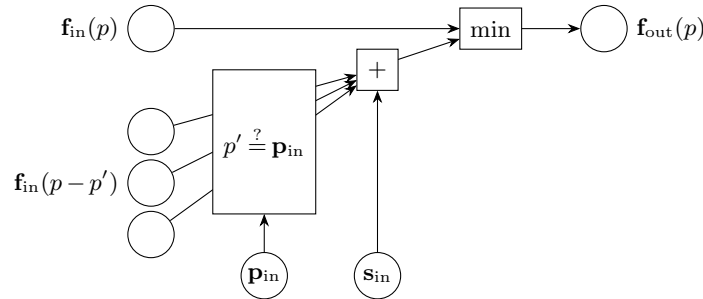


Figure 5 High-level idea how the DP-NN computes $f_{\text{out}}(p)$ for $p \in [p^*]$ from the inputs.

in order to make the recurrent structure of our NN obvious, we do not use the index i in the following description of the network. Instead, we denote the $n_0 = p^* + 2$ input values by $f_{\text{in}}(p)$ for $p \in [p^*]$, as well as p_{in} and s_{in} . The p^* output values are denoted by $f_{\text{out}}(p)$ for $p \in [p^*]$. The goal is to implement the recursion

$$f_{\text{out}}(p) = \min \{ f_{\text{in}}(p), f_{\text{in}}(p - p_{\text{in}}) + s_{\text{in}} \} \quad \text{for } p \in [p^*]$$

in an NN; cp. (1). It consists of an addition and taking a minimum, which are both simple operations for an NN. Hence, ideally, we would like to have an architecture as depicted in Figure 4 for computing $f_{\text{out}}(p)$ for every $p \in [p^*]$. The problem with this is, however, that the decision which component of f_{in} is accessed in order to compute the sum with s_{in} depends on the input value p_{in} . Since we aim for an architecture that is fixed and works for general input values p_{in} , we have to extend our construction as depicted in Figure 5. As we do not know the value of p_{in} in advance, we connect every input neuron $f_{\text{in}}(p - p')$, $p' \in [p - 1]$, to the unit that computes the sum $f_{\text{in}}(p - p_{\text{in}}) + s_{\text{in}}$. Since we only want to take the value $f_{\text{in}}(p - p_{\text{in}})$ into account, we need to add an additional unit that disables those connections if $p' \neq p_{\text{in}}$.

Due to the integrality of the profit values, this additional unit can be realized with two hidden layers and a constant number of neurons for every value of $p \in [p^*]$ and $p' \in [p - 1]$, as we show in a moment. Computing the minimum adds a third hidden layer. Hence,

the DP-NN has depth four while width and size are in $\mathcal{O}((p^*)^2)$. Unfolding the RNN and viewing it as a single feedforward NN executing the whole dynamic program results in depth $\mathcal{O}(n)$ and size $\mathcal{O}(n(p^*)^2)$. In the next subsection, we provide a detailed construction of the DP-NN and prove the following theorem.

THEOREM 1. *For a Knapsack instance with capacity $S = 1$, $s_i \in]0, 1]$, and $p_i \in \mathbb{N}$, for $i \in [n]$, with an upper bound p^* on the optimal solution value, the corresponding dynamic programming values $f(p, i)$, $i \in [n]$, $p \in p^*$, can be exactly computed by iteratively applying the DP-NN n times.*

As discussed in the introduction, due to the NP-hardness of the Knapsack Problem, the dependence on p^* cannot be avoided if one aims for exact, uniform neural network constructions, unless $P = NP$.

Details of the construction and correctness

Note that for size values larger than the Knapsack capacity, which is equal to 1 by our definition, we do not really care how large they actually are. Therefore, we define $\tilde{f}(p, i) = \min\{f(p, i), 2\}$ to be the values of the dynamic program truncated at 2. In other words, we replace all values in the interval $[2, +\infty]$ by 2. Note that the recursion

$$\tilde{f}(p, i) = \min\{\tilde{f}(p, i - 1), \tilde{f}(p - p_i, i - 1) + s_i\} \quad (2)$$

is still valid with starting values $\tilde{f}(p, i) = 0$ for $p \leq 0$ and $\tilde{f}(p, 0) = 2$ for $p \in [p^*]$. Instead of computing the actual values of f , the DP-NN computes the values of \tilde{f} .

The DP-NN has three hidden layers. After the $n_0 = p^* + 2$ input neurons $\mathbf{f}_{\text{in}}(p)$ for $p \in [p^*]$, \mathbf{p}_{in} , and \mathbf{s}_{in} , the first hidden layer consists of $n_1 = 2p^*$ neurons whose outputs are denoted by $\mathbf{o}_1^{(+)}(k)$ and $\mathbf{o}_1^{(-)}(k)$ for $k \in [p^*]$. Its role is to detect whether $k = \mathbf{p}_{\text{in}}$. If yes, then both $\mathbf{o}_1^{(+)}(k)$ and $\mathbf{o}_1^{(-)}(k)$ should be zero, otherwise at least one of them should be large (i.e., at least 2). In the second hidden layer, we have $n_2 = p^*(p^* - 1)/2$ neurons, denoted by $\mathbf{o}_2(p, k)$ for $p \in [p^*]$ and $k \in [p - 1]$. A neuron in this layer should output $\mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$ if $k = \mathbf{p}_{\text{in}}$ and zero otherwise. This way, the sum $\sum_{k=1}^{p-1} \mathbf{o}_2(p, k)$ equals $\mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$. The third hidden layer has $n_3 = p^*$ neurons, denoted by $\mathbf{o}_3(p)$ for $p \in [p^*]$. It is used for computing the minimum of $\mathbf{f}_{\text{in}}(p)$ and $\mathbf{s}_{\text{in}} + \mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$ using several copies of the construction shown in

Figure 1. Finally, the $n_4 = p^*$ output values are denoted by $\mathbf{f}_{\text{out}}(p)$ for $p \in [p^*]$. The following equations define the DP-NN.

$$\mathbf{o}_1^{(+)}(k) = \sigma(2(\mathbf{p}_{\text{in}} - k)), \quad k \in [p^*], \quad (3a)$$

$$\mathbf{o}_1^{(-)}(k) = \sigma(2(k - \mathbf{p}_{\text{in}})), \quad k \in [p^*], \quad (3b)$$

$$\mathbf{o}_2(p, k) = \sigma(\mathbf{f}_{\text{in}}(p - k) - \mathbf{o}_1^{(+)}(k) - \mathbf{o}_1^{(-)}(k)), \quad p \in [p^*], k \in [p - 1], \quad (3c)$$

$$\mathbf{o}_3(p) = \sigma\left(\mathbf{f}_{\text{in}}(p) - \left(\mathbf{s}_{\text{in}} + \sum_{k=1}^{p-1} \mathbf{o}_2(p, k)\right)\right), \quad p \in [p^*], \quad (3d)$$

$$\mathbf{f}_{\text{out}}(p) = \mathbf{f}_{\text{in}}(p) - \mathbf{o}_3(p), \quad p \in [p^*]. \quad (3e)$$

Our next goal is to prove Theorem 1, which states that the DP-NN indeed solves the Knapsack Problem exactly. We do so by going through the NN layer by layer and show what the individual layers do.

As mentioned, the role of the first hidden layer is to detect the input value of \mathbf{p}_{in} and to provide a large value for every k that is not equal to \mathbf{p}_{in} . The following lemma follows immediately from the construction and the properties of the rectifier function σ .

LEMMA 1. *Let $\mathbf{p}_{\text{in}} \in \mathbb{N}$. Then, for every $k \in [p^*]$, it holds that $\mathbf{o}_1^{(+)}(k) + \mathbf{o}_1^{(-)}(k) = 0$ if and only if $k = \mathbf{p}_{\text{in}}$, and $\mathbf{o}_1^{(+)}(k) + \mathbf{o}_1^{(-)}(k) \geq 2$ otherwise.*

The role of the second layer is to compute $\mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$, which is needed in the dynamic program. The main difficulty of this step is that it depends on the input \mathbf{p}_{in} which neuron to access. Therefore, this is computed for every possible value k of \mathbf{p}_{in} and set to zero if $k \neq \mathbf{p}_{\text{in}}$. The following lemma explains how this is established.

LEMMA 2. *Let $\mathbf{p}_{\text{in}} \in \mathbb{N}$ and $\mathbf{f}_{\text{in}}(p) \in]0, 2]$ for every $p \in [p^*]$. Then, for every $p \in [p^*]$ and every $k \in [p - 1]$, it holds that $\mathbf{o}_2(p, k) = \mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$ if and only if $k = \mathbf{p}_{\text{in}}$, and $\mathbf{o}_2(p, k) = 0$ otherwise.*

Proof. If $k = \mathbf{p}_{\text{in}}$, we obtain from Lemma 1 that $\mathbf{o}_1^{(+)}(k) + \mathbf{o}_1^{(-)}(k) = 0$. Thus, due to nonnegativity of $\mathbf{f}_{\text{in}}(p - k)$, we obtain $\mathbf{o}_2(p, k) = \sigma(\mathbf{f}_{\text{in}}(p - k)) = \mathbf{f}_{\text{in}}(p - k) = \mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$.

If $k \neq \mathbf{p}_{\text{in}}$, we obtain from Lemma 1 that $\mathbf{o}_1^{(+)}(k) + \mathbf{o}_1^{(-)}(k) \geq 2$. Thus, due to $\mathbf{f}_{\text{in}}(p - k) \leq 2$, we obtain $\mathbf{f}_{\text{in}}(p - k) - \mathbf{o}_1^{(+)}(k) - \mathbf{o}_1^{(-)}(k) \leq 0$, which implies $\mathbf{o}_2(p, k) = 0$. \square

The purpose of the third layer is to help calculating the final minimum. More precisely, it computes how much $\mathbf{f}_{\text{out}}(p)$ should be smaller than $\mathbf{f}_{\text{in}}(p)$ in the following way.

LEMMA 3. Let $\mathbf{p}_{\text{in}} \in \mathbb{N}$, $\mathbf{s}_{\text{in}} \in]0, 1]$, and $\mathbf{f}_{\text{in}}(p) \in]0, 2]$ for every $p \in [p^*]$. Then $\mathbf{o}_3(p) = \max\{0, \mathbf{f}_{\text{in}}(p) - \mathbf{s}_{\text{in}}\}$ for every $p \in [\mathbf{p}_{\text{in}}]$ and $\mathbf{o}_3(p) = \max\{0, \mathbf{f}_{\text{in}}(p) - (\mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}}) + \mathbf{s}_{\text{in}})\}$ for every $p \in \{\mathbf{p}_{\text{in}} + 1, \mathbf{p}_{\text{in}} + 2, \dots, p^*\}$.

Proof. If $p \leq \mathbf{p}_{\text{in}}$, we obtain from Lemma 2 that $\sum_{k=1}^{p-1} \mathbf{o}_2(p, k) = 0$. If $p > \mathbf{p}_{\text{in}}$, Lemma 2 implies $\sum_{k=1}^{p-1} \mathbf{o}_2(p, k) = \mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$. Thus, the claim follows by construction of the third layer and definition of σ . \square

Now, after we have investigated the functionality of each of the hidden layers, we are able to prove the main theorem of this section.

Proof of Theorem 1. Using Lemma 3, we obtain $\mathbf{f}_{\text{out}}(p) = \min\{\mathbf{f}_{\text{in}}(p), \mathbf{s}_{\text{in}}\}$ if $p \leq \mathbf{p}_{\text{in}}$ and $\mathbf{f}_{\text{out}}(p) = \min\{\mathbf{f}_{\text{in}}(p), \mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}}) + \mathbf{s}_{\text{in}}\}$ if $p > \mathbf{p}_{\text{in}}$. The claim follows due to the Recursion (2) with the respective starting values. \square

4. Smaller RNNs with Provable Approximation Guarantees

In order to overcome the drawback due to the dependence of the network width on p^* , we provide a construction, called *FPTAS-NN*, that uses less neurons, at the cost of losing optimality. Instead, we prove an approximation ratio (i.e., a worst-case bound) for the solution value computed by the FPTAS-NN. As in the standard Knapsack FPTAS (Hochbaum 1997, Vazirani 2001, Williamson and Shmoys 2011), the idea of this construction is to round the profit values if p^* becomes too large for an exact computation. Our approximation result can be interpreted as a tradeoff between the width of the NN and the quality of the Knapsack solution obtained. As in the previous section, we first give the high-level idea before providing details.

High-level idea of the construction

Let $P \in \mathbb{N}$ be a fixed number. The FPTAS-NN computes values $g(p, i)$ for every $p \in [P]$ and $i \in [n]$. These values are similar to the values $f(p, i)$ of the previous section, there is, however, one major difference. Let $p_i^* = \sum_{j=1}^i p_j$ be the total profit of the first i items. As soon as p_i^* exceeds P , we can no longer store a required size value for every possible profit value but have to round profits instead. The granularity we want to use for rounding is $d_i := \max\{1, p_i^*/P\}$. We construct the FPTAS-NN to compute values $g(p, i)$, $p \in [P]$, $i \in [n]$, such that we can guarantee the existence of a subset of $[i]$ that has size at most $g(p, i)$ and profit at least $p d_i$. Moreover, this is done in such a way that the optimal solution cannot

have a considerably higher profit value. That is, we prove a worst-case approximation guarantee for the solution found by the FPTAS-NN.

In addition to the values of g , the FPTAS-NN must also propagate the current total profit value p_i^* through the network in order to determine the rounding granularity in each step. Hence, in the i -th step, it receives $P + 3$ inputs, namely $g(p, i - 1)$ for $p \in [P]$, p_{i-1}^* , p_i , and s_i . It computes $P + 1$ outputs, namely $g(p, i)$ for $p \in [P]$ and p_i^* . Figure 6 illustrates the recurrent structure of this NN.

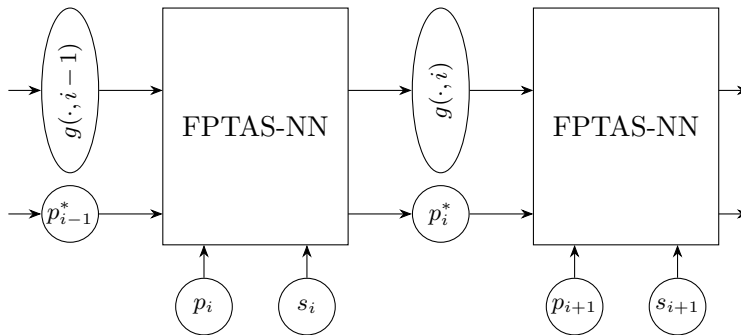


Figure 6 Recurrent structure of the FPTAS-NN for the Knapsack Problem.

As in Section 3, we use bold symbols in order to distinguish input, activation, and output values that depend on the concrete Knapsack instance from fixed parameters of the network. We again drop the index i in order to make the recurrent structure obvious. We denote the $n_0 = P + 3$ input parameters by $\mathbf{g}_{\text{in}}(p)$, for $p \in [P]$, as well as \mathbf{p}_{in}^* , \mathbf{p}_{in} , and \mathbf{s}_{in} . The $P + 1$ output values are denoted by $\mathbf{g}_{\text{out}}(p)$, for $p \in [P]$, and $\mathbf{p}_{\text{out}}^*$. Similar to the DP-NN in Section 3, the basic idea is to implement a recursion of the type

$$\mathbf{g}_{\text{out}}(p) = \min\{\mathbf{g}_{\text{in}}(p^{(1)}), \mathbf{g}_{\text{in}}(p^{(2)}) + \mathbf{s}_{\text{in}}\} \quad \text{for } p \in [P],$$

where the first argument of the minimum represents the option of not using item i , while the second one corresponds to using it. Notice, however, that $p^{(1)}$ and $p^{(2)}$ cannot simply be calculated as p and $p - \mathbf{p}_{\text{in}}$, respectively, since we may have to round with different granularities in two successive steps. Therefore, the rough structure of the FPTAS-NN is as follows: first, \mathbf{p}_{in}^* and \mathbf{p}_{in} are used in order to calculate the old and new rounding granularities $\mathbf{d}_{\text{old}} = \max\{1, \mathbf{p}_{\text{in}}^*/P\}$, as well as, $\mathbf{d}_{\text{new}} = \max\{1, (\mathbf{p}_{\text{in}}^* + \mathbf{p}_{\text{in}})/P\}$. Since this computation consists of maxima and weighted sums only, it can easily be achieved by an NN with one hidden layer. Second, the granularities are used in order to select $\mathbf{g}_{\text{in}}(p^{(1)})$ and $\mathbf{g}_{\text{in}}(p^{(2)})$

from the inputs. Below we give some more details on how this is done. The value of $p^{(2)}$ also depends on \mathbf{p}_{in} . Third, the final recursion is established as in the DP-NN. In addition to $\mathbf{g}_{\text{out}}(p)$, for $p \in [P]$, we also output $\mathbf{p}_{\text{out}}^* = \mathbf{p}_{\text{in}}^* + \mathbf{p}_{\text{in}}$ in order to keep track of the rounding granularities in subsequent steps. An overview of the entire network structure is given in Figure 7.

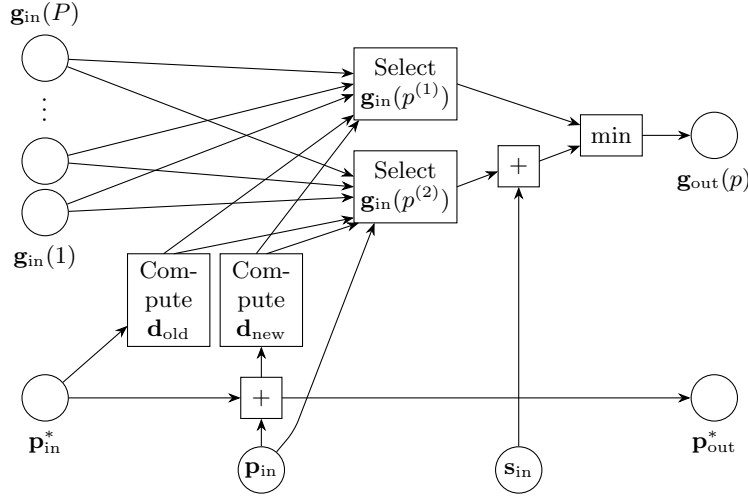


Figure 7 High-level idea how the FPTAS-NN computes $\mathbf{g}_{\text{out}}(p)$, $p \in [P]$, and $\mathbf{p}_{\text{out}}^*$ from the inputs.

Suppose we use the network for processing the i -th item. For each $p \in [P]$ we want to determine a (preferably small) value $\mathbf{g}_{\text{out}}(p)$ such that there is a subset of $[i]$ of total profit at least $p \mathbf{d}_{\text{new}}$ and total size at most $\mathbf{g}_{\text{out}}(p)$. For each $p' \in [P]$, we know that there is a subset of $[i - 1]$ of total profit at least $p' \mathbf{d}_{\text{old}}$ and total size at most $\mathbf{g}_{\text{in}}(p')$. We have two options: ignoring item i or using it. If we ignore it, then each $p^{(1)}$ with $p^{(1)} \mathbf{d}_{\text{old}} \geq p \mathbf{d}_{\text{new}}$ allows us to choose $\mathbf{g}_{\text{out}}(p) = \mathbf{g}_{\text{in}}(p^{(1)})$. If we do use the i -th item, however, then each $p^{(2)}$ with the property $p^{(2)} \mathbf{d}_{\text{old}} + \mathbf{p}_{\text{in}} \geq p \mathbf{d}_{\text{new}}$ allows us to choose $\mathbf{g}_{\text{out}}(p) = \mathbf{g}_{\text{in}}(p^{(2)}) + \mathbf{s}_{\text{in}}$. Hence, we want to choose $p^{(1)}$ and $p^{(2)}$ as small as possible such that these properties are fulfilled. Therefore, the units labeled ‘Select $\mathbf{g}_{\text{in}}(p^{(1)})$ ’ and ‘Select $\mathbf{g}_{\text{in}}(p^{(2)})$ ’ in Figure 7 are constructed by setting all other connections to zero except for those belonging to the smallest values of $p^{(1)}$ and $p^{(2)}$ satisfying the above properties. Similar to how we computed $\mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$ in the previous section, this requires two hidden layers and $\mathcal{O}(P^2)$ neurons in total.

In total, the FPTAS-NN has depth 5. The first hidden layer computes the rounding granularities, two hidden layers are required to select $\mathbf{g}_{\text{in}}(p^{(1)})$ and $\mathbf{g}_{\text{in}}(p^{(2)})$ and a final

hidden layer computes the minimum in the actual recursion. The width and size of the FPTAS-NN are in the order of $\mathcal{O}(P^2)$. Unfolding the RNN and viewing it as a single feedforward NN executing the whole FPTAS results in depth $\mathcal{O}(n)$ and size $\mathcal{O}(nP^2)$.

In the next subsection, we provide a formal description of the FPTAS-NN as well as proofs of the following two theorems. The first one ensures that the FPTAS-NN produces only feasible Knapsack solutions, while the second one shows that the FPTAS-NN indeed provides a fully polynomial-time approximation scheme to solve the Knapsack Problem.

THEOREM 2. *Suppose the FPTAS-NN is applied to a Knapsack instance with capacity $S = 1$, $s_i \in]0, 1]$, and $p_i \in \mathbb{N}$, for $i \in [n]$. For every $i \in [n]$ and every $p \in [P]$ with $g(p, i) \leq 1$, there exists a subset of $[i]$ with profit at least pd_i and size at most $g(p, i)$.*

THEOREM 3. *For a Knapsack instance with capacity $S = 1$, $s_i \in]0, 1]$, $p_i \in \mathbb{N}$, for $i \in [n]$, and for $\varepsilon \in]0, 1]$, set $P := \lceil n^2/\varepsilon \rceil$. Let p^{OPT} be the profit of the optimal solution and $p^{\text{NN}} = \max\{pd_n \mid g(p, n) \leq 1\}$ be the best possible profit found by the FPTAS-NN. Then $p^{\text{NN}} \geq (1 - \varepsilon)p^{\text{OPT}}$.*

Theorem 3 implies a tradeoff between the width of the NN and the precision of the Knapsack solution in the following sense. For achieving an approximation ratio of $1 - \varepsilon$, an NN of width $\mathcal{O}(P^2) = \mathcal{O}(n^4/\varepsilon^2)$ is required. In other words, the FPTAS-NN with fixed width w achieves a worst-case approximation ratio of $1 - \mathcal{O}(n^2/\sqrt{w})$.

As discussed in the introduction, assuming $P \neq \text{NP}$, it is clear that the size of a uniform NN must grow if ε tends to zero. Hence, complexity theory implies that a width-quality trade-off cannot be avoided. Still, it remains as an open question whether the growth rates implied by our construction are best possible.

Details of the construction and correctness

We now formally describe the FPTAS-NN and prove that it provides provable approximation guarantees.

As for the DP-NN, we truncate the values of g at 2, that is, instead of any value larger than 2 including $+\infty$, we just use the value 2. The FPTAS-NN is applied to a Knapsack instance in the following way. Using the initialization $g(p, 0) = 2$ for $p \in [P]$ and $p_0^* = 0$, for $i = 1, \dots, n$, we feed the inputs $\mathbf{g}_{\text{in}}(p) = g(p, i - 1)$ for $p \in [P]$, $\mathbf{p}_{\text{in}}^* = p_{i-1}^*$, $\mathbf{p}_{\text{in}} = p_i$, and $\mathbf{s}_{\text{in}} = s_i$ into the network and store the outputs as $g(p, i) = \mathbf{g}_{\text{out}}(p)$ for $p \in [P]$ and $p_i^* = \mathbf{p}_{\text{out}}^*$.

The FPTAS-NN has four hidden layers. After the $n_0 = P + 3$ input neurons $\mathbf{g}_{\text{in}}(p)$ for $p \in [P]$, \mathbf{p}_{in}^* , \mathbf{p}_{in} , and \mathbf{s}_{in} , the first hidden layer consists of $n_1 = 2$ neurons $\mathbf{o}_1^{\text{old}}$ and $\mathbf{o}_1^{\text{new}}$ which help to compute the rounding granularities \mathbf{d}_{old} and \mathbf{d}_{new} . They are defined as follows:

$$\mathbf{o}_1^{\text{old}} = \sigma \left(\frac{\mathbf{p}_{\text{in}}^*}{P} - 1 \right), \quad (4a)$$

$$\mathbf{o}_1^{\text{new}} = \sigma \left(\frac{\mathbf{p}_{\text{in}}^* + \mathbf{p}_{\text{in}}}{P} - 1 \right), \quad (4b)$$

$$\mathbf{d}_{\text{old}} = \mathbf{o}_1^{\text{old}} + 1, \quad (4c)$$

$$\mathbf{d}_{\text{new}} = \mathbf{o}_1^{\text{new}} + 1. \quad (4d)$$

The granularities \mathbf{d}_{old} and \mathbf{d}_{new} are just affine transformations of $\mathbf{o}_1^{\text{old}}$ and $\mathbf{o}_1^{\text{new}}$. Hence, they do not form an own hidden layer, because we do not apply the ReLU activation function there. The correct functionality of the first layer is ensured by the following lemma.

LEMMA 4. *The first layer of the FPTAS-NN correctly computes the rounding granularities $\mathbf{d}_{\text{old}} = \max\{1, \frac{\mathbf{p}_{\text{in}}^*}{P}\}$ and $\mathbf{d}_{\text{new}} = \max\{1, \frac{\mathbf{p}_{\text{in}}^* + \mathbf{p}_{\text{in}}}{P}\}$.*

Proof. This follows from the fact that $\sigma(x - 1) + 1 = \max\{0, x - 1\} + 1 = \max\{1, x\}$, where x equals either $\frac{\mathbf{p}_{\text{in}}^*}{P}$ or $\frac{\mathbf{p}_{\text{in}}^* + \mathbf{p}_{\text{in}}}{P}$. \square

Hence, in the i -th step, if we feed the inputs $\mathbf{p}_{\text{in}}^* = p_{i-1}^*$ and $\mathbf{p}_{\text{in}} = p_i$ into the network, \mathbf{d}_{old} and \mathbf{d}_{new} equal d_{i-1} and d_i , respectively.

In the second hidden layer, we have a total of $n_2 = 2P^2 + 2P$ hidden neurons, denoted by $\mathbf{o}_2^{(1+)}(p, k)$ and $\mathbf{o}_2^{(1-)}(p, k)$ for $p, k \in [P]$ with $p \leq k$, as well as, $\mathbf{o}_2^{(2+)}(p, k)$ and $\mathbf{o}_2^{(2-)}(p, k)$ for $p, k \in [P]$ with $p \geq k$, defined as follows:

$$\mathbf{o}_2^{(1+)}(p, k) = \sigma(2P(p\mathbf{d}_{\text{new}} - k\mathbf{d}_{\text{old}})), \quad p, k \in [P], p \leq k, \quad (4e)$$

$$\mathbf{o}_2^{(1-)}(p, k) = \sigma(2P((k - 1)\mathbf{d}_{\text{old}} - p\mathbf{d}_{\text{new}}) + 2), \quad p, k \in [P], p \leq k, \quad (4f)$$

$$\mathbf{o}_2^{(2+)}(p, k) = \sigma(2P(p\mathbf{d}_{\text{new}} - k\mathbf{d}_{\text{old}} - \mathbf{p}_{\text{in}})), \quad p, k \in [P], p \geq k, \quad (4g)$$

$$\mathbf{o}_2^{(2-)}(p, k) = \sigma(2P((k - 1)\mathbf{d}_{\text{old}} + \mathbf{p}_{\text{in}} - p\mathbf{d}_{\text{new}}) + 2), \quad p, k \in [P], p \geq k. \quad (4h)$$

For a fixed $p \in [P]$, let $p^{(1)}$ and $p^{(2)}$ be the smallest possible integers with $p^{(1)}\mathbf{d}_{\text{old}} \geq p\mathbf{d}_{\text{new}}$ and $p^{(2)}\mathbf{d}_{\text{old}} + \mathbf{p}_{\text{in}} \geq p\mathbf{d}_{\text{new}}$, respectively. The task of the second layer is to detect the values $p^{(1)}$ and $p^{(2)}$, as formalized by the following two lemmas.

LEMMA 5. *For each $p, k \in [P]$ with $p \leq k$, we have $\mathbf{o}_2^{(1+)}(p, k) + \mathbf{o}_2^{(1-)}(p, k) = 0$ if and only if $k = p^{(1)}$. Otherwise, we have $\mathbf{o}_2^{(1+)}(p, k) + \mathbf{o}_2^{(1-)}(p, k) \geq 2$.*

Proof. Obviously, it holds that $\mathbf{o}_2^{(1+)}(p, k) = 0$ if and only if $k\mathbf{d}_{\text{old}} \geq p\mathbf{d}_{\text{new}}$. On the other hand, using that \mathbf{d}_{old} and \mathbf{d}_{new} are integer multiples of $\frac{1}{P}$, we obtain

$$\begin{aligned} \mathbf{o}_2^{(1-)}(p, k) &= 0 \\ \Leftrightarrow (k-1)\mathbf{d}_{\text{old}} &\leq p\mathbf{d}_{\text{new}} - \frac{1}{P} \\ \Leftrightarrow (k-1)\mathbf{d}_{\text{old}} &< p\mathbf{d}_{\text{new}} \\ \Leftrightarrow \text{no integer } k' < k &\text{ satisfies } k'\mathbf{d}_{\text{old}} \geq p\mathbf{d}_{\text{new}}. \end{aligned}$$

This proves the first part of the claim. The second part follows because, again, \mathbf{d}_{old} and \mathbf{d}_{new} are integer multiples of $\frac{1}{P}$ and, hence, $\mathbf{o}_2^{(1+)}(p, k) + \mathbf{o}_2^{(1-)}(p, k)$ is an integer multiple of 2. \square

LEMMA 6. *For each $p, k \in [P]$ with $p \geq k$, we have $\mathbf{o}_2^{(2+)}(p, k) + \mathbf{o}_2^{(2-)}(p, k) = 0$ if and only if $k = p^{(2)}$. Otherwise, we have $\mathbf{o}_2^{(2+)}(p, k) + \mathbf{o}_2^{(2-)}(p, k) \geq 2$.*

Proof. Analogous to Lemma 5. \square

The third hidden layer consists of $n_3 = P^2 + P$ neurons $\mathbf{o}_3^{(1)}(p, k)$ for $p, k \in [P]$ with $p \leq k$, as well as $\mathbf{o}_3^{(2)}(p, k)$ for $p, k \in [P]$ with $p \geq k$. Moreover, we have again helping variables that do not form an own hidden layer because they are only affine transformations of the previous layers, namely $\mathbf{h}^{(1)}(p)$ and $\mathbf{h}^{(2)}(p)$ for $p \in [P]$.

$$\mathbf{o}_3^{(1)}(p, k) = \sigma(2 - \mathbf{g}_{\text{in}}(k) - \mathbf{o}_2^{(1+)}(p, k) - \mathbf{o}_2^{(1-)}(p, k)), \quad p, k \in [P], p \leq k, \quad (4i)$$

$$\mathbf{o}_3^{(2)}(p, k) = \sigma(\mathbf{g}_{\text{in}}(k) - \mathbf{o}_2^{(2+)}(p, k) - \mathbf{o}_2^{(2-)}(p, k)), \quad p, k \in [P], p \geq k, \quad (4j)$$

$$\mathbf{h}^{(1)}(p) = 2 - \sum_{k=p}^P \mathbf{o}_3^{(1)}(p, k), \quad p \in [P], \quad (4k)$$

$$\mathbf{h}^{(2)}(p) = \sum_{k=1}^p \mathbf{o}_3^{(2)}(p, k), \quad p \in [P]. \quad (4l)$$

The idea of this layer is to compute $\mathbf{g}_{\text{in}}(p^{(1)})$ and $\mathbf{g}_{\text{in}}(p^{(2)})$, as the following two lemmas show.

LEMMA 7. *For each $p \in [P]$, if $p^{(1)} \leq P$, we have $\mathbf{h}^{(1)}(p) = \mathbf{g}_{\text{in}}(p^{(1)})$. If $p^{(1)} > P$, we have $\mathbf{h}^{(1)}(p) = 2$.*

Proof. Note that $p^{(1)}$ is never smaller than p . If $p \leq p^{(1)} \leq P$, then $\mathbf{o}_3^{(1)}(p, p^{(1)}) = 2 - \mathbf{g}_{\text{in}}(p^{(1)})$ and $\mathbf{o}_3^{(1)}(p, k) = 0$ for each $k \neq p^{(1)}$ by Lemma 5. If $p^{(1)} > P$, then $\mathbf{o}_3^{(1)}(p, k) = 0$ for each k . Thus, the claim follows by the definition of $\mathbf{h}^{(1)}$. \square

LEMMA 8. For each $p \in [P]$, if $p^{(2)} \geq 1$, we have $\mathbf{h}^{(2)}(p) = \mathbf{g}_{\text{in}}(p^{(2)})$. If $p^{(2)} \leq 0$, we have $\mathbf{h}^{(2)}(p) = 0$.

Proof. We first show that $p^{(2)}$ is never larger than p by proving that $p\mathbf{d}_{\text{old}} + \mathbf{p}_{\text{in}} \geq p\mathbf{d}_{\text{new}}$. If $\mathbf{d}_{\text{new}} = 1$, then also $\mathbf{d}_{\text{old}} = 1$ holds and this statement is true. Otherwise, we have $\mathbf{d}_{\text{new}} = \frac{\mathbf{p}_{\text{in}}^* + \mathbf{p}_{\text{in}}}{P}$ and $\mathbf{d}_{\text{old}} \geq \frac{\mathbf{p}_{\text{in}}^*}{P}$. Hence, we obtain $p(\mathbf{d}_{\text{new}} - \mathbf{d}_{\text{old}}) \leq \frac{p\mathbf{p}_{\text{in}}}{P} \leq \mathbf{p}_{\text{in}}$. Therefore, in any case, $p\mathbf{d}_{\text{old}} + \mathbf{p}_{\text{in}} \geq p\mathbf{d}_{\text{new}}$ follows, and thus also $p^{(2)} \leq p$.

If $1 \leq p^{(2)} \leq p$, then it follows that $\mathbf{o}_3^{(2)}(p, p^{(2)}) = \mathbf{g}_{\text{in}}(p^{(2)})$ and $\mathbf{o}_3^{(2)}(p, k) = 0$ for each $k \neq p^{(2)}$ by Lemma 6. If $p^{(2)} \leq 0$, then $\mathbf{o}_3^{(2)}(p, k) = 0$ holds for each k . Thus, the claim follows by the definition of $\mathbf{h}^{(2)}$. \square

The fourth hidden layer is used to compute the minimum in the recursion and consists of $n_4 = P$ neurons $\mathbf{o}_4(p)$ for $p \in [P]$. Finally, we output the P values $\mathbf{g}_{\text{out}}(p)$ for $p \in [P]$, as well as $\mathbf{p}_{\text{out}}^* = \mathbf{p}_{\text{in}}^* + \mathbf{p}_{\text{in}}$.

$$\mathbf{o}_4(p) = \sigma(\mathbf{h}^{(1)}(p) - (\mathbf{s}_{\text{in}} + \mathbf{h}^{(2)}(p))), \quad p \in [P], \quad (4m)$$

$$\mathbf{g}_{\text{out}}(p) = \mathbf{h}^{(1)}(p) - \mathbf{o}_4(p), \quad p \in [P], \quad (4n)$$

$$\mathbf{p}_{\text{out}}^* = \mathbf{p}_{\text{in}}^* + \mathbf{p}_{\text{in}}. \quad (4o)$$

The following lemma ensures that the output value $\mathbf{g}_{\text{out}}(p)$ is indeed computed by the desired recursion, provided that $\mathbf{h}^{(1)}$ and $\mathbf{h}^{(2)}$ are computed properly.

LEMMA 9. For each $p \in [P]$, we have $\mathbf{g}_{\text{out}}(p) = \min\{\mathbf{h}^{(1)}(p), \mathbf{s}_{\text{in}} + \mathbf{h}^{(2)}(p)\}$.

Proof. This final layer of the FPTAS-NN is constructed exactly in the same way as the NN in Figure 1. \square

Equations (4a) to (4o) fully define the FPTAS-NN. We have shown several lemmas concerning the functionality of the individual layers. Now we turn towards the proofs of Theorems 2 and 3.

Proof of Theorem 2. We show that the claim even holds for all values of p and i with $g(p, i) < 2$ and not only for those with $g(p, i) \leq 1$.

We use induction on i . For the induction start ($i = 0$), nothing is to show due to the initialization $g(p, 0) = 2$ for all $p \in [P]$. For the induction step, suppose the claim is valid for all steps up to $i - 1$.

Fix some $p \in [P]$. By Lemma 9, the output $g(p, i) = \mathbf{g}_{\text{out}}(p)$ in the i -th step equals $\min\{\mathbf{h}^{(1)}(p), \mathbf{s}_{\text{in}} + \mathbf{h}^{(2)}(p)\}$. In the following, we distinguish two cases. Recall that $p^{(1)}$ and

$p^{(2)}$ are the smallest possible integers with $p^{(1)}\mathbf{d}_{\text{old}} \geq p\mathbf{d}_{\text{new}}$ and $p^{(2)}\mathbf{d}_{\text{old}} + \mathbf{p}_{\text{in}} \geq p\mathbf{d}_{\text{new}}$, respectively.

Case 1: $\mathbf{h}^{(1)}(p) \leq \mathbf{s}_{\text{in}} + \mathbf{h}^{(2)}(p)$. This implies $g(p, i) = \mathbf{h}^{(1)}(p)$. If $\mathbf{h}^{(1)}(p) = 2$, nothing is to show. Otherwise, by Lemma 7, we have $p^{(1)} \leq P$ with $p^{(1)}\mathbf{d}_{\text{old}} \geq p\mathbf{d}_{\text{new}}$ and $g(p, i) = \mathbf{h}^{(1)}(p) = \mathbf{g}_{\text{in}}(p^{(1)}) = g(p^{(1)}, i - 1)$. By induction, we obtain that there exists a subset of $[i - 1]$ with size at most $g(p, i)$ and profit at least $p^{(1)}\mathbf{d}_{\text{old}}$. Hence, using the same items yields a subset of $[i]$ with size at most $g(p, i)$ and profit at least $p\mathbf{d}_{\text{new}}$. Thus, the claim is proven in this case.

Case 2: $\mathbf{h}^{(1)}(p) > \mathbf{s}_{\text{in}} + \mathbf{h}^{(2)}(p)$. This implies $g(p, i) = \mathbf{s}_{\text{in}} + \mathbf{h}^{(2)}(p)$. Note that this can only happen if $\mathbf{h}^{(2)}(p) < 2$ because $\mathbf{h}^{(1)}(p)$ has at most value 2. First, suppose $p^{(2)} \leq 0$. This implies $p_i = \mathbf{p}_{\text{in}} \geq p\mathbf{d}_{\text{new}}$. Hence, by using just item i , we obtain a subset of profit at least $p\mathbf{d}_{\text{new}}$ and size at most $s_i = \mathbf{s}_{\text{in}} \leq \mathbf{s}_{\text{in}} + \mathbf{h}^{(2)}(p) = g(p, i)$, and we are done. Second, if $p^{(2)} \geq 1$, then Lemma 8 implies that $g(p, i) = \mathbf{s}_{\text{in}} + \mathbf{h}^{(2)}(p) = \mathbf{s}_{\text{in}} + \mathbf{g}_{\text{in}}(p^{(2)}) = s_i + g(p^{(2)}, i - 1)$. By induction, we obtain that there exists a subset of $[i - 1]$ with size at most $g(p, i) - s_i$ and profit at least $p^{(2)}\mathbf{d}_{\text{old}}$. Hence, adding item i to this subset yields a subset of $[i]$ with size at most $g(p, i)$ and profit at least $p^{(2)}\mathbf{d}_{\text{old}} + p_i \geq p\mathbf{d}_{\text{new}}$. Thus, the claim is also proven in this case. \square

Proof of Theorem 3. Let M^{OPT} be an optimal solution to the Knapsack instance and $M_i^{\text{OPT}} = M^{\text{OPT}} \cap [i]$ be the subset of $[i]$ chosen by the optimal solution. Let $s_i^{\text{OPT}} = \sum_{j \in M_i^{\text{OPT}}} s_j$ be the size of M_i^{OPT} and $p_i^{\text{OPT}} = \sum_{j \in M_i^{\text{OPT}}} p_j$ be the profit of M_i^{OPT} . The idea of the proof is that in each step, we lose at most a profit of d_i compared to the optimal solution. Formally, we prove the following claim by induction on i : for every $i \in [n]$, and every $p \leq \left\lceil \frac{p_i^{\text{OPT}}}{d_i} \right\rceil - i$ we have $g(p, i) \leq s_i^{\text{OPT}}$.

The induction start is settled by extending the claim to $i = 0$, for which it is trivial. For the induction step, suppose the claim is valid for all steps up to $i - 1$. Fix a value $p \leq \left\lceil \frac{p_i^{\text{OPT}}}{d_i} \right\rceil - i$. Let again $p^{(1)}$ and $p^{(2)}$ be the smallest possible integers with $p^{(1)}d_{i-1} \geq pd_i$ and $p^{(2)}d_{i-1} + p_i \geq pd_i$, respectively. We distinguish two cases.

Case 1: $i \notin M^{\text{OPT}}$, i.e., the optimal solution does not use item i . Observe that

$$\begin{aligned} pd_i &\leq \left(\left\lceil \frac{p_i^{\text{OPT}}}{d_i} \right\rceil - i \right) d_i \\ &\leq p_i^{\text{OPT}} - (i - 1)d_i \\ &= p_{i-1}^{\text{OPT}} - (i - 1)d_i \end{aligned}$$

$$\begin{aligned}
&\leq p_{i-1}^{\text{OPT}} - (i-1)d_{i-1} \\
&\leq \left(\left\lceil \frac{p_{i-1}^{\text{OPT}}}{d_{i-1}} \right\rceil - (i-1) \right) d_{i-1}.
\end{aligned}$$

Hence, we obtain

$$p^{(1)} \leq \left\lceil \frac{p_{i-1}^{\text{OPT}}}{d_{i-1}} \right\rceil - (i-1) \quad (5)$$

by the definition of $p^{(1)}$. In particular, $p^{(1)} \leq \frac{p_{i-1}^*}{d_{i-1}} \leq P$ by the definition of d_{i-1} . Therefore, Lemmas 7 and 9 imply $g(p, i) \leq g(p^{(1)}, i-1)$. Due to Inequality (5), it further follows by induction that $g(p, i) \leq g(p^{(1)}, i-1) \leq s_{i-1}^{\text{OPT}} = s_i^{\text{OPT}}$, which settles the induction step in this case.

Case 2: $i \in M^{\text{OPT}}$, i.e., the optimal solution uses item i . Observe that

$$\begin{aligned}
pd_i &\leq \left(\left\lceil \frac{p_i^{\text{OPT}}}{d_i} \right\rceil - i \right) d_i \\
&\leq p_i^{\text{OPT}} - (i-1)d_i \\
&= p_{i-1}^{\text{OPT}} + p_i - (i-1)d_i \\
&\leq p_{i-1}^{\text{OPT}} + p_i - (i-1)d_{i-1} \\
&\leq \left(\left\lceil \frac{p_{i-1}^{\text{OPT}}}{d_{i-1}} \right\rceil - (i-1) \right) d_{i-1} + p_i.
\end{aligned}$$

Hence, we obtain

$$p^{(2)} \leq \left\lceil \frac{p_{i-1}^{\text{OPT}}}{d_{i-1}} \right\rceil - (i-1) \quad (6)$$

by the definition of $p^{(2)}$. If $p^{(2)} \leq 0$, Lemmas 8 and 9 imply $g(p, i) \leq s_i \leq s_i^{\text{OPT}}$. If $p^{(2)} \geq 0$, Lemmas 8 and 9 imply $g(p, i) \leq g(p^{(2)}, i-1) + s_i$. Due to Inequality (6), we can further conclude by induction that $g(p, i) \leq g(p^{(2)}, i-1) + s_i \leq s_{i-1}^{\text{OPT}} + s_i = s_i^{\text{OPT}}$, which finalizes the induction step.

Now, making use of the claim we have just proven by induction, we obtain that $g\left(\left\lceil \frac{p^{\text{OPT}}}{d_n} \right\rceil - n, n\right) \leq s_n^{\text{OPT}} \leq 1$. Therefore, it holds that

$$p^{\text{NN}} \geq \left(\left\lceil \frac{p^{\text{OPT}}}{d_n} \right\rceil - n \right) d_n \geq p^{\text{OPT}} - nd_n. \quad (7)$$

If $d_n = 1$, that is, $p^* \leq P$, then we have for all $i \in [n]$ that $d_i = 1$. Hence, in each step and for each $p \in [P]$, we have $p^{(1)} = p$ and $p^{(2)} = p - p_i$. Therefore, by Lemmas 7–9, the FPTAS-NN behaves like the DP-NN from Section 3 that executes the exact dynamic program and the theorem follows.

Otherwise, if $d_n > 1$, we have $d_n = \frac{p^*}{P}$. Since there must exist one item with profit at least $\frac{p^*}{n}$, we obtain $p^{\text{OPT}} \geq \frac{p^*}{n}$ and, hence, $nd_n = \frac{np^*}{P} \leq \frac{n^2 p^{\text{OPT}}}{P}$. Together with (7), this implies $\frac{p^{\text{NN}}}{p^{\text{OPT}}} \geq 1 - \frac{n^2}{P} \geq 1 - \varepsilon$. \square

5. Empirical Evidence for Superlinear Width

While the running time of the classical Knapsack dynamic program depends only linearly on p^* , the width of the DP-NN is of the order of $(p^*)^2$. In our construction, the quadratic factor arises from dynamically finding $\mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}})$ in a hard-coded network, as explained in Section 3. For similar reasons, the width of the FPTAS-NN grows with $1/\varepsilon^2$ instead of $1/\varepsilon$.

The natural question to ask is whether one can achieve a linear dependence with a different construction. While this question remains open from a purely theoretical point of view, in this section, we provide some empirical evidence that the required width might indeed grow in a superlinear manner.

We emphasize that the purpose of our empirical study is to validate our theoretical findings on the expressivity of neural networks. We do not propose a method to practically solve knapsack problems.

For details about the experimental setup, including used soft- and hardware, random data generation and systematic of seeds, training and validation setup, hyperparameters, as well as the source code, please refer to Appendix A. Here we only include the necessary information to understand the key findings.

Similar to the DP-NN of Section 3, we train an NN with three hidden layers and variable width to execute one step of the Knapsack dynamic program, that is, to map \mathbf{f}_{in} , \mathbf{p}_{in} , and \mathbf{s}_{in} to \mathbf{f}_{out} , for random Knapsack instances. For the 30 different values $\{3, 6, 9, \dots, 90\}$ of p^* , we increase the width neuron by neuron until a mean squared error (MSE) loss of at most 0.005 is reached on a validation set. The threshold 0.005 is carefully chosen such that NNs with reasonable width are empirically able to achieve it. In order to show that our results are robust with respect to the chosen threshold, we repeat the experiment with two different thresholds 0.00375 and 0.0025. Since smaller thresholds naturally make the training process more difficult, we use smaller values of p^* in these cases. The empirically required width to achieve an MSE of at most the respective thresholds is plotted in Figures 8, 9, and 10, respectively.

As one can see in these figures, the results of the experiments suggest that the dependence of the required width to reach the MSE threshold seems to grow superlinearly with p^* .

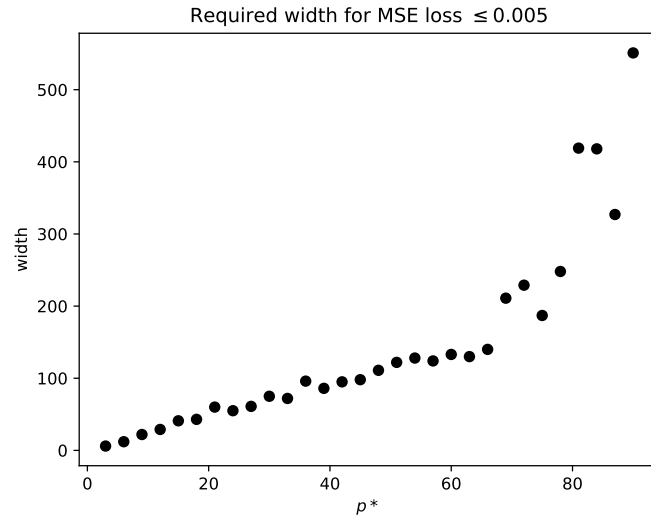


Figure 8 Required width to achieve a mean squared error of at most 0.005 as a function of p^* .

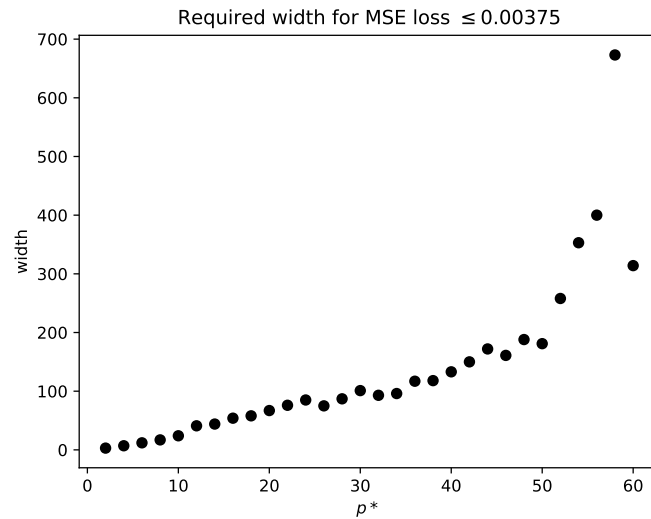


Figure 9 Required width to achieve a mean squared error of at most 0.00375 as a function of p^* .

We interpret this as a mild experimental validation of our theoretical construction which requires quadratic instead of linear width.

Of course, one should take these result with a grain of salt since the superlinear relation might have multiple reasons. For instance, it is unclear whether the difficulty to train larger networks has a stronger effect than the expressivity of ReLU NNs. Still, we find that this computational study supports our theoretical size bounds.

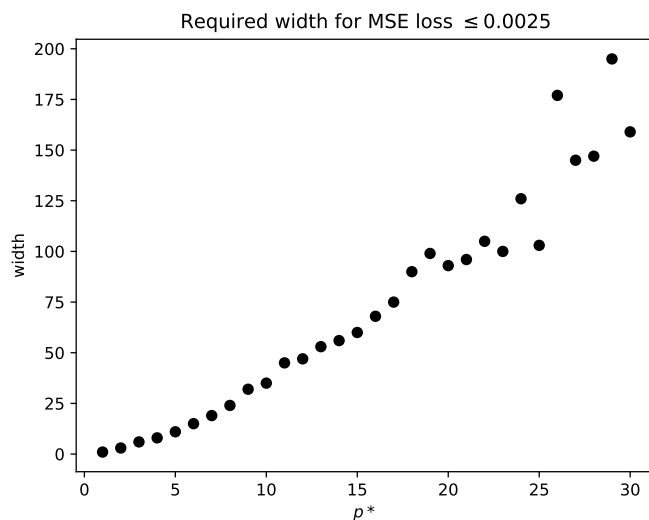


Figure 10 Required width to achieve a mean squared error of at most 0.0025 as a function of ρ^* .

6. Neural Networks for Other CO Problems

In this section we demonstrate that our approach is by no means bound to the Knapsack Problem. In fact, for many other CO problems it is possible to convert a dynamic programming solution method into a provably correct NN. For certain NP-hard CO problems, a dynamic programming solution even implies the existence of a fully polynomial-time approximation scheme (Woeginger 2000). This, in turn, might shed light on the tradeoff between size of corresponding NNs and their solution quality, as for the Knapsack Problem in Section 4. In the following we provide several examples in order to support these claims.

Longest Common Subsequence

First, consider the problem of finding the length of the longest common subsequence of two finite integer sequences x_1, \dots, x_m and y_1, \dots, y_n . A standard dynamic programming procedure, see, e.g., (Cormen et al. 2001, Section 15.4), computes values $f(i, j)$ equal to the length of the longest common subsequence of the partial sequences x_1, x_2, \dots, x_i and y_1, y_2, \dots, y_j by applying the recursion

$$f(i, j) = \begin{cases} f(i-1, j-1) + 1 & \text{if } x_i = y_j, \\ \max\{f(i-1, j), f(i, j-1)\} & \text{if } x_i \neq y_j. \end{cases}$$

Since the sequence consists of integers, the check whether x_i equals y_j can be performed similarly to checking whether $p' = \mathbf{p}_{\text{in}}$ in Section 3. The remainder of the recursion only consists of maxima and sums. Hence, computing $f(i, j)$ from $f(i-1, j-1)$, $f(i-1, j)$,

$f(i, j - 1)$, x_i , and y_j can be realized via an NN of constant size. These basic units can be plugged together in a two-dimensional way for computing all values $f(i, j)$, $i \in [m]$, $j \in [n]$. The resulting NN can be seen as a two-dimensional RNN of constant size that is applied in an m by n grid structure, an architecture introduced by Graves et al. (2007). Unfolding the RNN results in a feedforward NN of depth $\mathcal{O}(m + n)$ and size $\mathcal{O}(mn)$ for computing the length of the longest common subsequence.

Single-Source Shortest Path Problem

As a second example, we consider the Bellman-Ford algorithm for the Single-Source Shortest Path Problem, see, e.g., (Kleinberg and Tardos 2006, Section 6.8). If $(c_{uv})_{u,v \in V}$ is the length matrix of a graph with vertex set V and $s \in V$ is the source vertex, this algorithm recursively computes values $f(i, v)$ determining the shortest possible length of a path from s to v using at most i arcs by

$$f(i, v) = \min_{u \in V} \{f(i - 1, u) + c_{uv}\}.$$

Since this recursion consists only of sums and minima, it can be easily implemented in an NN. The sequential time complexity of the Bellman-Ford algorithm on complete digraphs with $n = |V|$ is $\mathcal{O}(n^3)$, which can naturally be parallelized into $\mathcal{O}(n)$ rounds. Since the best known NNs for computing the minimum of n numbers require $\mathcal{O}(\log n)$ depth (Arora et al. 2018), there exists an NN executing the Bellman-Ford algorithm with depth $\mathcal{O}(n \log n)$ and size $\mathcal{O}(n^3 \log n)$. Observe that in each round $i \in [n]$, the computation mapping the values $f(i - 1, v)$, $v \in V$, to $f(i, v)$, $v \in V$, is the same. Therefore, this NN can also be seen as an RNN of depth $\mathcal{O}(\log n)$ and size $\mathcal{O}(n^2 \log n)$ that is applied n times.

All-Pairs Shortest Path Problem

Third, recall that the All-Pairs Shortest Path Problem can be solved by computing the $(n - 1)$ -th min-plus matrix power of the length matrix $(c_{uv})_{u,v \in V}$, see, e.g., (Leighton 1991, Section 2.5.4). By repeated squaring, this can be achieved with only $\mathcal{O}(\log n)$ min-plus matrix multiplications. For a single multiplication it is required to compute $\mathcal{O}(n^2)$ times in parallel the minimum of n numbers. One of these minimum computations requires depth $\mathcal{O}(\log n)$ and size $\mathcal{O}(n \log n)$. Putting them in parallel to execute one min-plus matrix product results in depth $\mathcal{O}(\log n)$ and size $\mathcal{O}(n^3 \log n)$. Note that the whole execution consists of $\mathcal{O}(\log n)$ repetitions of the same procedure, namely squaring a matrix in the

min-plus sense. Hence, this can again be viewed as an RNN with depth $\mathcal{O}(\log n)$ and size $\mathcal{O}(n^3 \log n)$ that is applied $\mathcal{O}(\log n)$ times. Unfolding results in a single feedforward NN with depth $\mathcal{O}(\log^2 n)$ and size $\mathcal{O}(n^3 \log^2 n)$ for solving the All-Pairs Shortest Path Problem.

Constrained Shortest Path Problem

Next, consider a common generalization of the Shortest Path Problem and the Knapsack Problem, namely the NP-hard Constrained Shortest Path Problem. Here, in addition to a (nonnegative) length matrix $(c_{uv})_{u,v \in V}$, the input graph is also equipped with a (non-negative) resource matrix $(r_{uv})_{u,v \in V}$. The task is to find a minimum length path P from a source vertex s to any other vertex, but this time subject to a resource constraint $\sum_{(u,v) \in P} r_{uv} \leq R$ for a given resource limit R . An extensive overview of solution approaches to this problem can be found, e.g., in the dissertation by Ziegelmann (2001). Similar to the Knapsack Problem, there exist two natural pseudo-polynomial dynamic programs, one of them parametrized by length values and the other one by resource values. Both can be implemented on an NN by combining the ideas from Section 3 with the NN for the Bellmann-Ford algorithm above. We showcase this for the variant parametrized by the length values. This dynamic program recursively calculates values $f(c, v)$ representing the minimum amount of resource needed for a path from s to v with length at most c by

$$f(c, v) = \min\{f(c-1, v), \min_{u \in V \setminus \{v\}} \{f(c - c_{uv}, u) + r_{uv}\}\}.$$

For fixed c , u , and v , the term $f(c - c_{uv}, u) + r_{uv}$ can be calculated by a similar construction as we computed $\mathbf{f}_{\text{in}}(p - \mathbf{p}_{\text{in}}) + \mathbf{s}_{\text{in}}$ in Section 3. Assuming an upper bound c^* on the optimal objective value, this can be achieved with constant depth and $\mathcal{O}(c^*)$ width. Hence, it remains to compute a minimum of at most n numbers in order to compute $f(c, v)$. Thus, each single value $f(c, v)$ can be computed with depth $\mathcal{O}(\log n)$ and size $\mathcal{O}(nc^* \log n)$. We have to compute $\mathcal{O}(nc^*)$ of these values, but for fixed c , all these values can be computed in parallel. Therefore, the whole dynamic program can be executed on an NN with depth $\mathcal{O}(c^* \log n)$ and a total size of $\mathcal{O}(n^2(c^*)^2 \log n)$. This is pseudo-polynomial, which is the best we can hope for due to the NP-hardness of the problem, if we aim for uniform NN constructions. Moreover, similar to the Knapsack Problem, this dynamic program can be turned into an FPTAS by downscaling and rounding the length values. This observation can be used to obtain a width-quality tradeoff for the Constrained Shortest Path Problem similar to what we have shown in Section 4.

Traveling Salesperson Problem

Finally, consider the Bellman-Held-Karp algorithm for solving the (asymmetric) Traveling Salesperson Problem (TSP); see (Bellman 1962, Held and Karp 1962). Given a (complete, directed) graph with vertex set V and distances c_{uv} from vertex $u \in V$ to vertex $v \in V$, the TSP asks for the shortest round-trip visiting each vertex exactly once. Choosing an arbitrary starting vertex $s \in V$, the Bellman-Held-Karp algorithm recursively computes values $f(T, v)$ for each $T \subseteq V \setminus \{s\}$, $v \in T$, corresponding to the length of the shortest s - v -path visiting exactly the nodes in $T \cup \{s\}$ by the formula

$$f(T, v) = \min_{u \in T \setminus \{v\}} \{f(T \setminus \{v\}, u) + c_{uv}\}.$$

The length of the shortest TSP tour is then given by $\min_{u \in V \setminus \{s\}} \{f(V \setminus \{s\}, u) + c_{us}\}$. While the sequential time complexity of this algorithm on digraphs with $n = |V|$ is $\mathcal{O}(n^2 2^n)$, in an NN we can compute the values of f for all T with equal cardinality in parallel. As before, computing the minimum introduces an additional factor of $\log n$ in the depth and size of the network. Hence, in total, the TSP can be solved with an NN of depth $\mathcal{O}(n \log n)$ and size $\mathcal{O}(n^2 2^n \log n)$. In particular, a polynomially deep NN suffices to solve the NP-hard (asymmetric) TSP.

7. Conclusions and Future Work

An obvious open problem is to improve the obtained bounds on the required width of our neural network constructions. In particular, an interesting question is whether meaningful lower bounds beyond those immediately implied by the NP-hardness of the Knapsack Problem can be obtained, as suggested by our experimental results.

Another interesting direction would be to design and conduct a more sophisticated empirical study to analyze the expressivity of neural networks for the Knapsack Problem and related problems. For example, one could aim to find out whether (larger) real-world instances have a certain structure allowing that smaller neural networks capture already the dynamic programming steps for these instances. Note that an exhaustive search for the width as performed in this paper is computationally infeasible for larger instances, so such a study would require a different but still robust experimental setup.

Notice that our networks only output the solution value but not the corresponding solution, i.e., subset of items. It is easy to see that, as for the dynamic program solving

the Knapsack Problem, the subset of items can be obtained in a straightforward way via backtracking. On the other hand, notice that it is impossible for a ReLU NN (without threshold gates) to output (the characteristic vector of) the optimum subset of items: while the function computed by a ReLU NN is piecewise linear and continuous (Arora et al. 2018, Goodfellow et al. 2013), already infinitesimal changes of the input (i.e., the profit values of items) might change the optimum subset of items.

Finally, an exciting direction for future research is to generalize our results of Section 6 by describing a general procedure to convert dynamic programs into ReLU NNs. Ideally, one could exactly classify the type of dynamic programs that guarantee the existence of a corresponding ReLU NN. Similar in spirit, Woeginger (2000) classifies dynamic programs that guarantee the existence of a fully polynomial-time approximation scheme.

Acknowledgments

The authors would like to thank the referees for their insightful comments which helped to improve the presentation of this paper.

A large portion of this research was conducted while the first author was affiliated with TU Berlin and supported by DFG-GRK 2434 “Facets of Complexity”. Furthermore, this work is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement ScaleOpt–757481) and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy — The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

Appendix

A. Detailed Experimental Setup

In this section we describe in detail how we conducted the experiments of Section 5 in the main paper.

Hard- and software

All our experiments were conducted on a Windows 10 Enterprise (version 1909) machine with an Intel Core i5-8500 6-Core 64-bit CPU and 16 GB RAM. We use Python 3.8.13 with Numpy 1.22.3, and Tensorflow 2.3.0 in CPU-only mode.

Generation of random Knapsack instances

For a given value of p^* we sample a set of items of total profit $\sum p_i = p^*$ in the following way: the profit of the i -th item is always chosen uniformly at random among all integer values between 1 and $p^* - \sum_{i'=1}^{i-1} p_{i'}$. This is repeated until all profits sum up to p^* . We chose this procedure in order to make it likely to have both, very profitable and less profitable items within one instance. Finally, we shuffle the order of the items.

For each item, we then pick a size value uniformly in the interval $[0, 1]$ and normalize these values such that their sum equals a uniformly at random chosen value $\sum s_i \in]1, 2[$. We certainly want $\sum s_i > 1$, because otherwise all items would fit into the Knapsack. On the other hand, $\sum s_i < 2$ makes sense, because in our DP-NN (compare Section 3), we use 2 as a placeholder for $+\infty$.

Preparation of the training set

Since we can generate arbitrarily many random Knapsack instances, we use an infinite training set and never train on the same data point twice. A Knapsack instance with n items yields n training points, namely one for each step of the dynamic program. In order to avoid having the n training points belonging to one instance in successive order, we generate training points belonging to several instances and shuffle them.

Neural network architecture

For a given value p^* and width w , the corresponding neural network used consists of an input layer with $p^* + 2$ neurons (corresponding to the p^* values of the previous dynamic programming state, as well as, the scalar profit and size values), three hidden layers with w neurons each and ReLU activations, as well as an output layer with p^* neurons (corresponding to the new state of the dynamic programming) without further activation function. That way, we choose the same depth as in the DP-NN (Section 3), but do not employ the specific knowledge about the widths of the three hidden layers. As in the DP-NN, each layer is not only connected to the previous layer, but also receives direct connections from the input layer. In total, by our results of Section 3, this architecture is theoretically able to exactly represent the dynamic programming transition function if $w \geq p^*(p^* - 1)/2$.

Training and validating a specific network

For a given value p^* and width w , we train the neural network architecture described above as follows. We train in epochs of 1000 batches with batch size 32 using mean squared error (MSE) loss and the Adam optimizer, which is a robust standard choice. It makes sense to use MSE loss as it punishes errors in both directions equally hard and large errors harder than small errors. All other (hyper-)parameters are left to the default settings of Tensorflow, which empirically works quite well for our problem type and size. It takes between 8 and 30 seconds to train one epoch with our machine setup. We train until there are two successive epochs without improvement in training loss, which empirically happens after 10 to 80 epochs. Using a validation set that is randomly generated in the same way as the training set, we evaluate the trained network on 1000 batches of size 32 each. The resulting mean squared error is our final result for the given values of p^* and w .

Finding the required width

For each of our three MSE thresholds and each value p^* , we always train networks with increasing widths $1, 2, 3, \dots$ as described above until a network achieves a validation MSE less or equal to the threshold.

Seed generation

In order to ensure reproducibility of our experiments, each time before we train and validate an NN with given value p^* and width w , we reset the random seeds of both Numpy and Tensorflow to $257 \cdot p^* + w$, where 257 is just an arbitrary prime number. Note that these packages only guarantee the same result of random experiments if the same package versions are used.

Source Code

The source code is publicly available at <https://github.com/ChristophHertrich/neural-knapsack>. There, the file `README.md` explains how the code can be used to reproduce the experiments of this paper.

References

- Arora R, Basu A, Mianjy P, Mukherjee A (2018) Understanding deep neural networks with rectified linear units. *International Conference on Learning Representations*.
- Arora S, Barak B (2009) *Computational Complexity: A Modern Approach* (Cambridge University Press).
- Bahdanau D, Cho KH, Bengio Y (2015) Neural machine translation by jointly learning to align and translate. *3rd International Conference on Learning Representations, ICLR 2015*.
- Bellman R (1962) Dynamic programming treatment of the travelling salesman problem. *Journal of the ACM* 9(1):61–63.
- Bello I, Pham H, Le QV, Norouzi M, Bengio S (2016) Neural combinatorial optimization with reinforcement learning. *arXiv:1611.09940* .
- Bengio Y, Lodi A, Prouvost A (2021) Machine learning for combinatorial optimization: A methodological tour d’horizon. *European Journal of Operational Research* 290(2):405–421.
- Bertschinger D, Hertrich C, Jungeblut P, Miltzow T, Weber S (2022) Training fully connected neural networks is $\exists\mathbb{R}$ -complete. *arXiv:2204.01368* .
- Bertsimas D, Weismantel R (2005) *Optimization over Integers* (Dynamic Ideas).
- Cappart Q, Chételat D, Khalil E, Lodi A, Morris C, Veličković P (2021) Combinatorial optimization and reasoning with graph neural networks. *arXiv:2102.09544* .
- Cormen TH, Leiserson CE, Rivest RL, Stein C (2001) *Introduction to algorithms* (MIT press), second edition.
- Eldan R, Shamir O (2016) The power of depth for feedforward neural networks. *Conference on Learning Theory*, 907–940.
- Emami P, Ranka S (2018) Learning permutations with sinkhorn policy gradient. *arXiv:1805.07010* .
- Fischetti M, Lodi A (2010) On the knapsack closure of 0-1 integer linear programs. *Electronic Notes in Discrete Mathematics* 36:799–804.
- Froese V, Hertrich C, Niedermeier R (2022) The computational complexity of ReLU network training parameterized by data dimensionality. *Journal of Artificial Intelligence Research* 74:1775–1790.
- Glorot X, Bordes A, Bengio Y (2011) Deep sparse rectifier neural networks. *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 315–323.
- Goel S, Klivans AR, Manurangsi P, Reichman D (2021) Tight hardness results for training depth-2 ReLU networks. *12th Innovations in Theoretical Computer Science Conference (ITCS)*.

- Goodfellow I, Warde-Farley D, Mirza M, Courville A, Bengio Y (2013) Maxout networks. *International Conference on Machine Learning*, 1319–1327.
- Graves A, Fernández S, Schmidhuber J (2007) Multi-dimensional recurrent neural networks. *International Conference on Artificial Neural Networks*, 549–558.
- Graves A, Mohamed Ar, Hinton G (2013) Speech recognition with deep recurrent neural networks. *2013 IEEE international conference on acoustics, speech and signal processing*, 6645–6649 (Ieee).
- Gu S, Hao T (2018) A pointer network based deep learning algorithm for 0–1 knapsack problem. *2018 Tenth International Conference on Advanced Computational Intelligence (ICACI)*, 473–477.
- Haase CA, Hertrich C, Loho G (2023) Lower bounds on the depth of integral ReLU neural networks via lattice polytopes. *The Eleventh International Conference on Learning Representations*.
- Hanin B (2019) Universal function approximation by deep neural nets with bounded width and ReLU activations. *Mathematics* 7(10):992.
- Hanin B, Rolnick D (2019) Complexity of linear regions in deep networks. *International Conference on Machine Learning*, 2596–2604.
- Hanin B, Sellke M (2017) Approximating continuous functions by ReLU nets of minimal width. *arXiv:1710.11278* .
- Held M, Karp RM (1962) A dynamic programming approach to sequencing problems. *Journal of the Society for Industrial and Applied Mathematics* 10(1):196–210.
- Hertrich C, Basu A, Di Summa M, Skutella M (2021) Towards lower bounds on the depth of ReLU neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, full version: *arXiv:2105.14835*.
- Hertrich C, Sering L (2021) ReLU neural networks of polynomial size for exact maximum flow computation. *arXiv:2102.06635* .
- Hertrich C, Skutella M (2021) Provably good solutions to the knapsack problem via neural networks of bounded size. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 7685–7693, URL <https://ojs.aaai.org/index.php/AAAI/article/view/16939>.
- Hochbaum DS (1997) Various notions of approximations: Good, better, best, and more. Hochbaum DS, ed., *Approximation Algorithms for NP-hard Problems*, 346–446 (PWS Publishing Co.).
- Hopfield JJ, Tank DW (1985) “Neural” computation of decisions in optimization problems. *Biological Cybernetics* 52(3):141–152.
- Karp RM (1972) Reducibility among combinatorial problems. Miller RE, Thatcher JW, eds., *Complexity of Computer Computations*, 85–103 (Springer).
- Kellerer H, Pferschy U, Pisinger D (2004) *Knapsack Problems* (Springer).
- Khalife S, Basu A (2022) Neural networks with linear threshold activations: structure and algorithms. *International Conference on Integer Programming and Combinatorial Optimization*, 347–360 (Springer).

- Khalil E, Dai H, Zhang Y, Dilkina B, Song L (2017) Learning combinatorial optimization algorithms over graphs. Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, eds., *Advances in Neural Information Processing Systems 30*, 6348–6358.
- Kleinberg J, Tardos E (2006) *Algorithm Design* (Pearson Education).
- Kool W, van Hoof H, Welling M (2019) Attention, learn to solve routing problems! *International Conference on Learning Representations*.
- LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521:436–444.
- Leighton FT (1991) *Introduction to parallel algorithms and architectures: Arrays · trees · hypercubes* (Morgan-Kaufmann).
- Li D, Liu J, Lee D, Seyedmazloom A, Kaushik G, Lee K, Park N (2021) A novel method to solve neural knapsack problems. *International Conference on Machine Learning*, 6414–6424.
- Liang S, Srikant R (2017) Why deep neural networks for function approximation? *International Conference on Learning Representations*.
- Lodi A, Zarpellon G (2017) On learning and branching: a survey. *TOP* 25(2):207–236.
- Lombardi M, Milano M (2018) Boosting combinatorial problem modeling with machine learning. *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 5472–5478.
- Martello S, Toth P (1990) *Knapsack Problems: Algorithms and Computer Implementations* (John Wiley & Sons).
- Montufar GF, Pascanu R, Cho K, Bengio Y (2014) On the number of linear regions of deep neural networks. Ghahramani Z, Welling M, Cortes C, Lawrence ND, Weinberger KQ, eds., *Advances in Neural Information Processing Systems 27*, 2924–2932.
- Mukherjee A, Basu A (2017) Lower bounds over boolean inputs for deep neural networks with ReLU gates. *arXiv:1711.03073* .
- Nguyen Q, Mukkamala MC, Hein M (2018) Neural networks should be wide enough to learn disconnected decision regions. *International Conference on Machine Learning*, 3737–3746.
- Nowak A, Villar S, Bandeira AS, Bruna J (2017) Revised Note on Learning Algorithms for Quadratic Assignment with Graph Neural Networks. *arXiv:1706.07450* .
- Ohlsson M, Peterson C, Söderberg B (1993) Neural networks for optimization problems with inequality constraints: The knapsack problem. *Neural Computation* 5(2):331–339.
- Pascanu R, Montufar G, Bengio Y (2014) On the number of inference regions of deep feed forward networks with piece-wise linear activations. *International Conference on Learning Representations*.
- Raghu M, Poole B, Kleinberg J, Ganguli S, Dickstein JS (2017) On the expressive power of deep neural networks. *International Conference on Machine Learning*, 2847–2854.

- Safran I, Shamir O (2017) Depth-width tradeoffs in approximating natural functions with neural networks. *International Conference on Machine Learning*, 2979–2987.
- Serra T, Tjandraatmadja C, Ramalingam S (2018) Bounding and counting linear regions of deep neural networks. *International Conference on Machine Learning*, 4565–4573.
- Shalev-Shwartz S, Ben-David S (2014) *Understanding machine learning: From theory to algorithms* (Cambridge University Press).
- Smith KA (1999) Neural networks for combinatorial optimization: A review of more than a decade of research. *INFORMS Journal on Computing* 11(1):15–34.
- Sutskever I, Vinyals O, Le QV (2014) Sequence to sequence learning with neural networks. *Advances in neural information processing systems* 27.
- Telgarsky M (2015) Representation benefits of deep feedforward networks. *arXiv:1509.08101* .
- Telgarsky M (2016) Benefits of depth in neural networks. *Conference on Learning Theory*, 1517–1539.
- Vazirani VV (2001) *Approximation Algorithms* (Springer).
- Veličković P, Ying R, Padovano M, Hadsell R, Blundell C (2020) Neural execution of graph algorithms. *International Conference on Learning Representations*.
- Vinyals O, Fortunato M, Jaitly N (2015) Pointer networks. Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R, eds., *Advances in Neural Information Processing Systems 28*, 2692–2700.
- Williamson DP, Shmoys DB (2011) *The design of approximation algorithms* (Cambridge University Press).
- Woeginger GJ (2000) When does a dynamic programming formulation guarantee the existence of a fully polynomial time approximation scheme (FPTAS)? *INFORMS Journal on Computing* 12(1):57–74.
- Xu K, Li J, Zhang M, Du SS, Kawarabayashi Ki, Jegelka S (2020a) What can neural networks reason about? *International Conference on Learning Representations*.
- Xu S, Panwar SS, Kodialam MS, Lakshman TV (2020b) Deep neural network approximated dynamic programming for combinatorial optimization. *AAAI Conference on Artificial Intelligence*, 1684–1691.
- Yang F, Jin T, Liu TY, Sun X, Zhang J (2018) Boosting dynamic programming with neural networks for solving np-hard problems. *Asian Conference on Machine Learning*, 726–739.
- Yarotsky D (2017) Error bounds for approximations with deep relu networks. *Neural Networks* 94:103–114.
- Ziegelmann M (2001) *Constrained Shortest Paths and Related Problems*. Ph.D. thesis, Universität des Saarlandes Saarbrücken.