

LORM: Learning to Optimize for Resource Management in Wireless Networks with Few Training Samples

Yifei Shen, *Student Member, IEEE*, Yuanming Shi, *Member, IEEE*, Jun Zhang, *Senior Member, IEEE*, and Khaled B. Letaief, *Fellow, IEEE*

Abstract—Effective resource management plays a pivotal role in wireless networks, which, unfortunately, typically results in challenging mixed-integer nonlinear programming (MINLP) problems. Machine learning-based methods have recently emerged as a disruptive way to obtain near-optimal performance for MINLPs with affordable computational complexity. There have been some attempts in applying such methods to resource management in wireless networks, but these attempts require huge amounts of training samples and lack the capability to handle constrained problems. Furthermore, they suffer from severe performance deterioration when the network parameters change, which commonly happens and is referred to as the *task mismatch* problem. In this paper, to reduce the sample complexity and address the feasibility issue, we propose a framework of Learning to Optimize for Resource Management (LORM). In contrast to the end-to-end learning approach adopted in previous studies, LORM learns the optimal pruning policy in the branch-and-bound algorithm for MINLPs via a sample-efficient method, namely, *imitation learning*. To further address the task mismatch problem, we develop a transfer learning method via self-imitation in LORM, named *LORM-TL*, which can quickly adapt a pre-trained machine learning model to the new task with only a few additional *unlabeled* training samples. Numerical simulations demonstrate that LORM outperforms specialized state-of-the-art algorithms and achieves near-optimal performance, while providing significant speedup compared with the branch-and-bound algorithm. Moreover, LORM-TL, by relying on a few unlabeled samples, achieves comparable performance with the model trained from scratch with sufficient labeled samples.

Index Terms—Resource allocation, mixed-integer nonlinear programming, wireless communications, few-shot learning, transfer learning.

I. INTRODUCTION

A. Motivations

In wireless networks, effective management of radio resources is vital for performance optimization [3]. Unfortunately, typical resource management problems, such as sub-carrier allocation in OFDMA [4], user association [5], access

The materials in this paper were presented in part at the IEEE Global Conference on Signal and Information Processing (GlobalSIP), 2018 [1], and IEEE International Conference on Communications (ICC), 2019 [2]. This work was supported by the General Research Funding (Project No. 16210719) from the Research Grants Council of Hong Kong.

Y. Shen and K. B. Letaief are with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong (E-mail: {yshenaw, eekhaled}@ust.hk). Y. Shi is with the School of Information Science and Technology, ShanghaiTech University, Shanghai 201210, China (E-mail: shiym@shanghaitech.edu.cn). J. Zhang is with the Department of Electronic and Information Engineering, The Hong Kong Polytechnic University, Hong Kong (E-mail: jun-eie.zhang@polyu.edu.hk). (The corresponding author is J. Zhang.)

point selection [6], and computation offloading [7], are mixed integer nonlinear programming (MINLP) problems, which are NP-hard. The complexity of global optimization algorithms, e.g., the branch-and-bound algorithm, is exponential. Thus, most of the existing studies focused on sub-optimal or heuristic algorithms, whose performance gaps to the optimal solution are difficult to quantify and control.

Machine learning has recently emerged as a disruptive technology to balance the computational complexity and the performance gap for solving NP-hard problems, and has attracted lots of attention from the mathematical optimization community [8]. This trend has also inspired researchers to apply machine learning-based methods to solve optimization problems in wireless networks [9]. For example, to solve the power control problem in the interference channel, it was proposed in [10] to use a neural network to approximate the classic weighted minimum mean square error (WMMSE) algorithm to speed up the computation. However, as the training samples are obtained via a sub-optimal algorithm, i.e., WMMSE, there is a performance gap compared with the optimal solution. To achieve near-optimal performance, unsupervised learning methods have been proposed in [11], [12], which do not rely on any existing resource management algorithm. Besides improving performance and computational efficiency, machine learning techniques have also been applied to other problems in resource management. In particular, deep reinforcement learning [13], spatial deep learning [14], and deep neural network parameterization [15] have been proposed to deal with the scenarios with delayed channel state information (CSI), without CSI but only geographical locations, and with unknown resource allocation functions, respectively.

While the attempts [10], [11], [12] have demonstrated the great potential of the “learning to optimize” approach for resource management in wireless networks, applying them into real systems faces additional difficulties. A prominent shortcoming of these methods is that they require large amounts of training problem instances, e.g., millions of samples are needed for a small-size system [11]. This incurs a significant cost for sample acquisition, and may not be feasible in practice. Secondly, resource management problems are constrained by nature, but the ability of existing machine learning-based methods in dealing with constraints is limited [8]. Finally, wireless networks are inherently dynamic, e.g., both the locations and number of users are changing dynamically. Thus,

a pre-trained machine learning model may be useless or suffer from severe performance deterioration as the network setting changes. This challenge can be characterized as *task mismatch*, i.e., the test setting is different from the trained one, and it has not been treated in previous studies. To address these challenges, in this paper we develop a novel machine learning framework for MINLP resource management problems in wireless networks, which only requires a small number of training problem instances, achieves near-optimal performance with feasibility guarantee, and is able to effectively handle task mismatch.

B. Literature Review

In the past several decades, optimization-based algorithms have been extensively studied for wireless resource management problems [16]. For instance, different algorithms have been developed for power control (e.g., water filling [17]), transmit beamforming design (e.g., WMMSE [18] and conic optimization-based algorithms [19]), user association (e.g., the primal dual optimization-based algorithm [5]), and dynamic offloading (e.g., the Lyapunov optimization-based algorithm [7]). However, as most of the wireless resource management problems are NP-hard, the optimization-based algorithms are either too time-consuming or have non-negligible performance gaps. To enable efficient and near-optimal resource allocation, researchers have recently shifted their attention to the “learning to optimize” paradigm.

The investigation of “learning to optimize” sits at the intersection of wireless communications, mathematical optimization, and machine learning. To set the stage, we first review available results on “learning to optimize” in these different areas. The goal of “learning to optimize” is to obtain near-optimal algorithms with affordable computational complexity for challenging optimization problems. This topic has been developed for decades in the machine learning community, and has also attracted significant attention from mathematical optimization and operation research [8].

There are different approaches to apply machine learning methods, and a straightforward one is “end-to-end learning”. The idea is to regard the optimization algorithm as a black-box, and directly learn the input-output mapping [8], as illustrated in Fig. 1(a). For example, the pointer networks [20], [21] have been proposed for approximating the output of algorithms for combinatorial optimization problems. In wireless networks, the study in [10] also applied this idea. The adopted neural networks in these studies have generic architectures and perform well in practical problems. Nevertheless, they require large amounts of samples for training, which impedes using optimal solutions as labels when solving MINLPs. In this approach, the machine learning model is used as an approximator for an existing optimization algorithm, and thus the performance highly depends on the baseline algorithm. Besides, the feasibility of constraints is guaranteed by simply projecting the solution to the feasible set, which will result in performance loss. Furthermore, such simple projection trick will have difficulties in handling general constraints, even convex ones.

One major objective of this study is to develop a “learning to optimize” framework with few training problem instances, which falls in the paradigm of *few-shot learning*, a method that mimics the way humans learn new concepts [22]. One effective way to reduce the sample complexity is to consider a specific optimization algorithm and exploit its algorithmic structure [8]. Such “optimization policy learning” approach learns the optimal policy in the given algorithm, and thus has little influence on the feasibility guarantee of that algorithm, as illustrated in Fig. 1(b). It has recently attracted lots of attention, especially for mixed-integer linear programs (MILPs) [23], [24] and problems on graphs [25]. Besides reducing the sample complexity and providing feasibility guarantee, there are other advantages by exploiting the algorithm structure. Firstly, by exploration, the performance of the learned policy is not bounded by the performance of heuristic algorithms that generate labels [8], [26], which helps to achieve close to the optimal performance. Therefore, we can develop an optimization policy with better performance than the learned one. Secondly, these methods enjoy good generalization abilities, e.g., they can scale up to larger problem sizes and generalize to different problems [23], [25]. Inspired by the unique advantages as mentioned above, our proposed approach is based on “optimization policy learning”, instead of “end-to-end learning”.

Optimization policy learning has been applied to solve combinatorial optimization problems. One such study, which is related to our investigation, is [23] for MILPs. While [23] achieves certain speedup compared with the branch-and-bound algorithm for MILPs, its computational complexity is still high, and more importantly, it cannot be directly applied to constrained MINLP problems. First, compared with MILPs, it is much more time-consuming to solve the relaxed problems in MINLPs for resource management in wireless networks. Second, in [23], to address the feasibility problem, the search space is kept large for each problem in order to guarantee feasibility, which brings additional complexity. In this paper, we have made great efforts to improve the computational efficiency over [23]. Specifically, to address the first issue, we use neural networks and problem dependent features to avoid solving relaxed problems. To address the second issue, we propose an adaptive strategy to control the search space and guarantee the feasibility, which leads to a small search space in average. These two innovations allow our proposed framework to achieve low computational complexity with near-optimal performance.

C. Contributions

In this paper, we focus on finding near-optimal solutions for MINLP resource management problems in wireless networks using machine learning. The main innovation is a Learning to Optimize framework for Resource Management, called LORM, which is based on optimization policy learning. The merits of this framework and major contributions of this paper are summarized as follows:

- 1) We develop the LORM framework to solve MINLPs for resource management in wireless networks. The proposed

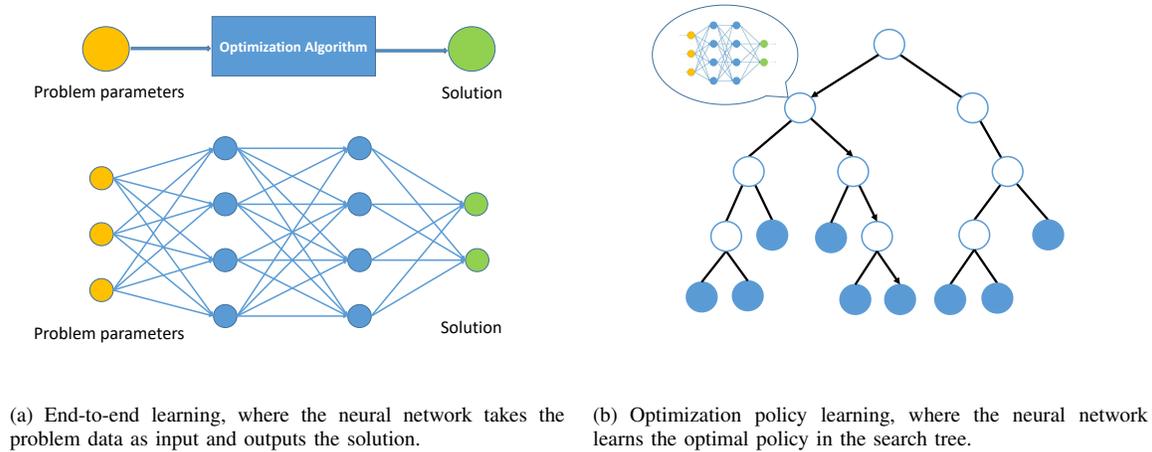


Fig. 1. A comparison of the end-to-end learning and optimization policy learning.

approach accelerates the branch-and-bound algorithm by casting the branch-and-bound algorithm as a sequential decision problem and learning the optimal pruning policy via imitation learning. This framework achieves near-optimal performance with few training problem instances and low computational complexity, by exploiting the structure of the branch-and-bound algorithm and problem data. Furthermore, the feasibility of constraints is guaranteed by iteratively increasing the search space.

- 2) To tackle the task mismatch issue, we adopt a transfer learning method via self-imitation in LORM, which leads to the LORM-TL framework. Compared with the traditional transfer learning method, i.e., fine-tuning, self-imitation only requires a few additional *unlabeled* training problem instances. This is achieved by a proposed exploration policy to label the additional training problem instances in the new scenario, followed by fine-tuning the learned policy with these labels.
- 3) We test LORM and LORM-TL on a typical MINLP problem, i.e., network power minimization in Cloud-RANs [6]. Simulations will demonstrate the effectiveness of LORM and LORM-TL in the following aspects:
 - a) LORM outperforms the specialized state-of-the-art methods and achieves near-optimal performance under a variety of system configurations, with only tens of training problem instances. It also achieves significant speedup compared with the branch-and-bound and existing method [23].
 - b) It is demonstrated that, with LORM, a model trained on one setting can achieve state-of-the-art performance in moderately different settings, e.g., with different large-scale fading and different numbers of users or base stations. Such generalization capability is achieved via exploiting the algorithm structure and careful feature design.
 - c) When the variation of the system configuration becomes dramatic, directly generalizing the model trained via LORM to the new setting induces considerable performance degradation. In this case, LORM-TL is

able to achieve comparable performance with the model trained on sufficient samples for the test setting from scratch, and it only relies on tens of additional unlabeled training problem instances.

II. RESOURCE MANAGEMENT AND BRANCH-AND-BOUND

In this section, we first introduce a general formulation for mixed-integer resource management problems in wireless networks. A global optimization algorithm, i.e., branch-and-bound, is then introduced, followed by some observations.

A. Mixed-integer Resource Management

A wide range of resource management problems in wireless networks can be formulated as MINLP problems, which consist of a discrete optimization variable \mathbf{a} , e.g., indicating cell association or subcarrier management, and a continuous optimization variable \mathbf{w} , e.g., transmit power, subject to resource or performance constraints. Typical examples include subcarrier management in OFDMA [4], user association [5], network power minimization in Cloud-RANs [6], and joint task offloading scheduling and transmit power management [7]. A general formulation for such problems is given by

$$\begin{aligned}
 \mathcal{P} : & \underset{\mathbf{a}, \mathbf{w}}{\text{minimize}} && f(\mathbf{a}, \mathbf{w}) \\
 & \text{subject to} && Q(\mathbf{a}, \mathbf{w}) \leq 0 \\
 & && a[i] \in \mathbb{N}, w[i] \in \mathbb{C},
 \end{aligned} \tag{1}$$

where $f(\cdot, \cdot)$ is an objective function, e.g., the sum rate or network power consumption, $a[i]$ and $w[i]$ are the elements of \mathbf{a} and \mathbf{w} respectively, and $Q(\cdot, \cdot)$ represents constraints such as the QoS or power constraint. MINLP problems are NP-hard in general.

B. Branch-and-Bound

The branch-and-bound algorithm is one of the state-of-the-art algorithms to find an optimal solution to the MINLP problems [27]. The discrete variables and the continuous variables are first separated. The algorithm then determines

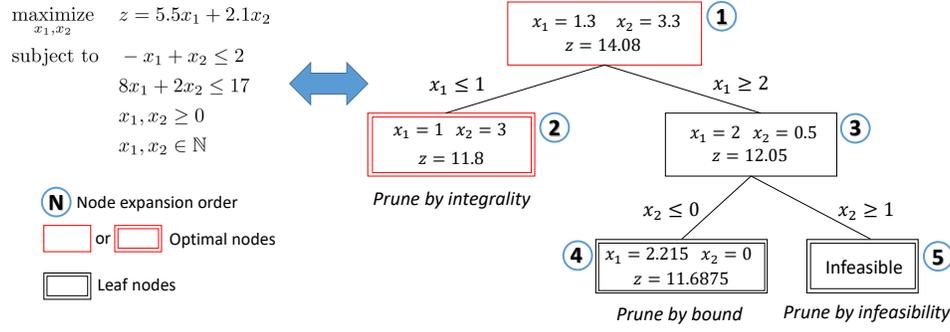


Fig. 2. An example of the branch-and-bound algorithm [27]. The solution of the relaxed linear programming problem is $x_1 = 1.3$, and $x_2 = 3.3$. It is used as the initial solution for the tree search. The objective value $z = 14.08$ provides an upper bound to the original problem. According to the branching rule, we branch out to the first variable x_1 and create two sub-problems. According to the node selection policy, we select the left child. The solution to the relaxed problem of the left child is $x_1 = 1$ $x_2 = 3$. Because they are integer values, we prune the node by integrality. The objective value $z = 11.8$ provides a lower bound for the original problem. We then turn to the right child of the root node. We solve it and branch out on x_2 , which generates two child nodes. The optimal objective value of the left child is $11.6875 < 11.8$, so we prune it by bound. The relaxed problem of the right child is infeasible and hence we prune it by infeasibility. Thus, the whole branch-and-bound tree is visited and we get the solution to the original problem as $x_1 = 1$ $x_2 = 3$.

Algorithm 1 The branch-and-bound algorithm.

```

1:  $\mathcal{L} \leftarrow \{N_1\}, c^* \leftarrow +\infty, t \leftarrow 0$ 
2: while  $\mathcal{L} \neq \emptyset$  do
3:    $t \leftarrow t + 1$ 
4:    $N_t \leftarrow$  selects a node from  $\mathcal{L}$ 
5:    $\mathcal{L} \leftarrow \mathcal{L} \setminus \{N_t\}$ 
6:    $(\mathbf{a}_t^*, \mathbf{w}_t^*), z_t^* \leftarrow$  solve the relaxed problem of  $\mathcal{P}_t$ 
7:   if  $\pi_p(N_t, \mathbf{a}_t^*, z_t^*) = \text{preserved}$  then
8:      $N_t^1, N_t^2 \leftarrow$  two children of  $N_t$ 
9:      $\mathcal{L} \leftarrow \mathcal{L} \cup \{N_t^1, N_t^2\}$ 
10:  end if
11:  if  $a_t^*[j], \forall j$  is integer then
12:     $c^* \leftarrow \min(c^*, z_t^*)$ 
13:  end if
14: end while

```

the values of the integer variables via a tree search process, followed by obtaining the values of the continuous variables via a convex solver. The tree search process consists of three main policies: Node selection policy, variable selection policy, and pruning policy, which together iteratively construct a binary search tree.

We first specify some notations. Each node in the binary search tree contains a MINLP problem and its relaxed problem. Let N_t denote the node selected at the t -th iteration, \mathcal{P}_t denote the MINLP problem at node N_t , and \mathcal{G}_t represent the feasible set of \mathcal{P}_t . The natural relaxation of \mathcal{G}_t , i.e., relax all the integer constraints into box constraints, is denoted as \mathcal{H}_t . For example, binary constraints $\mathcal{G}_t = \{(\mathbf{a}, \mathbf{w}) : a[i] \in \{0, 1\}, \forall i\}$ can be relaxed into box constraints $\mathcal{H}_t = \{(\mathbf{a}, \mathbf{w}) : a[i] \in [0, 1], \forall i\}$. Thus, the original problem and relaxed problems at node N_t can be represented respectively as $\min_{\mathbf{a}, \mathbf{w}} \{f(\mathbf{a}, \mathbf{w}) : (\mathbf{a}, \mathbf{w}) \in \mathcal{G}_t\}$ and $\min_{\mathbf{a}, \mathbf{w}} \{f(\mathbf{a}, \mathbf{w}) : (\mathbf{a}, \mathbf{w}) \in \mathcal{H}_t\}$. Let $(\mathbf{a}_t^*, \mathbf{w}_t^*)$ and z_t^* , respectively, denote the optimal solution and the optimal objective value of the relaxed problem at N_t , i.e.,

$$\begin{aligned}
 (\mathbf{a}_t^*, \mathbf{w}_t^*) &= \arg \min_{\mathbf{a}, \mathbf{w}} \{f(\mathbf{a}, \mathbf{w}) : (\mathbf{a}, \mathbf{w}) \in \mathcal{H}_t\} \\
 z_t^* &= \min_{\mathbf{a}, \mathbf{w}} \{f(\mathbf{a}, \mathbf{w}) : (\mathbf{a}, \mathbf{w}) \in \mathcal{H}_t\}.
 \end{aligned} \tag{2}$$

We use $a_t^*[j]$ to represent the j -th element of \mathbf{a}_t^* . The best integer solution to \mathcal{P} found before or in this iteration is denoted as c^* .

The branch-and-bound algorithm maintains an unexplored node list \mathcal{L} , and this list only contains the root node N_1 at the beginning. The MINLP problem at the root node is \mathcal{P} . At each iteration, the node selection policy first selects a node and pops it from the unexplored list. The relaxed problem at this node is solved to obtain \mathbf{a}_t^* and z_t^* . The pruning policy π_p then determines whether to preserve this node according to \mathbf{a}_t^* and z_t^* . The algorithm enters the next iteration if the pruning policy decides to preserve this node, i.e., no child node of this node will be considered. Otherwise, the variable selection policy chooses a variable index j such that $a_t^*[j]$ is fractional. Then, the feasible set \mathcal{G}_t is divided into two parts as

$$\begin{aligned}
 \mathcal{G}_t^1 &= \mathcal{G}_t \cap \{(\mathbf{a}, \mathbf{w}) : a[j] \leq \lfloor a_t^*[j] \rfloor\}, \\
 \mathcal{G}_t^2 &= \mathcal{G}_t \cap \{(\mathbf{a}, \mathbf{w}) : a[j] \geq \lceil a_t^*[j] \rceil\}.
 \end{aligned} \tag{3}$$

Two new MINLP problems are formed based on the partition

$$\begin{aligned}
 \mathcal{P}_t^1 &: \min_{\mathbf{a}, \mathbf{w}} \{f(\mathbf{a}, \mathbf{w}) : (\mathbf{a}, \mathbf{w}) \in \mathcal{G}_t^1\}, \\
 \mathcal{P}_t^2 &: \min_{\mathbf{a}, \mathbf{w}} \{f(\mathbf{a}, \mathbf{w}) : (\mathbf{a}, \mathbf{w}) \in \mathcal{G}_t^2\}.
 \end{aligned} \tag{4}$$

Two child nodes N_t^1 and N_t^2 , whose MINLP problems are \mathcal{P}_t^1 and \mathcal{P}_t^2 , respectively, are produced and put into the unexplored node list \mathcal{L} . These procedures repeat at each iteration until there are no nodes in \mathcal{L} . The pseudo-code is shown in Algorithm 1 and an illustration is shown in Fig. 2.

In the standard branch-and-bound algorithm, the pruning policy π_p prunes a node only if one of the following conditions is met:

- *Prune by bound:* $z_t^* \geq c^*$. z_t^* provides a lower bound for the optimal objective value of \mathcal{P}_t since $\mathcal{G}_t \subseteq \mathcal{H}_t$. Under such circumstance, the optimal solution to \mathcal{P}_t cannot be the best solution as the lower bound is worse than the best objective value found.
- *Prune by infeasibility:* The relaxed problem of \mathcal{P}_t is infeasible. This is a special case of *prune by bound* if the infeasibility is viewed as $z_t^* = +\infty$.

- *Prune by integrality*: $a_t^*[j], \forall j$, is an integer. It is unnecessary to search the children of N_t since we have found the optimal solution to \mathcal{P}_t .

C. Observations

Branch-and-bound has been widely employed in solving MINLPs in wireless networks, as it is capable to obtain globally optimal solutions. However, its computational complexity is exponential in the dimension of \mathbf{a} , which cannot be tolerated in practice.

The heavy computation of the branch-and-bound algorithm comes from two aspects.

- 1) The number of nodes in the branch-and-bound tree is exponential in the dimension of \mathbf{a} ;
- 2) Solving relaxed problems in MINLPs, e.g., second order cone programs (SOCPs) or semidefinite programs (SDPs), is computationally expensive.

These two observations form the basis for our proposed method. As the number of nodes in the search tree is controlled by the pruning policy, improving the policy helps to reduce the computation complexity. The pruning policy in the standard branch-and-bound algorithm is conservative because the optimality must be guaranteed, which is achieved by checking that all the other solutions are worse than the returned one. Therefore, most of the running time of the branch-and-bound algorithm is spent on checking the non-optimal nodes. Thus, we can dramatically boost the efficiency of the branch-and-bound algorithm by learning an aggressive pruning policy, without enforcing the optimality guarantee.

For the second issue, in the standard branch-and-bound algorithm, the relaxed problem must be solved because the features of the pruning policy merely consist of the solution of the relaxed problem. We can reduce the number of relaxed problems if we train the classifier based on the problem data and only the solution of the relaxed problem at the root node. While the tree size reduction via aggressive pruning has been discussed and attempted in previous studies, e.g., [23], our proposal to reduce the relaxed problems is new, and it can greatly speed up the computation.

III. LORM: LEARNING TO OPTIMIZE FOR RESOURCE MANAGEMENT

In this section, we first introduce the idea of learning the policies in the branch-and-bound algorithm via imitation learning. Then key components of the LORM framework, including classifier design, feature design, and algorithm design are then presented, followed by a complexity analysis and some discussions.

A. Learning Policies via Imitation Learning

The search procedure of the branch-and-bound algorithm can be regarded as a sequential decision problem. In theory, the policy in such finite-dimensional sequential decision problems can be learned via supervised learning if we have features and labels for all the nodes. However, this direct approach is inefficient because there is an exponential number of nodes,

which makes it tedious to obtain all the features. A too large dataset will also increase the complexity of training. To reduce the training effort and retain good performance, we instead employ DAgger (Data Aggregation) [28] to collect features and labels dynamically. The DAgger algorithm makes our approach deviate from supervised learning and fall into the category of *imitation learning*. In the following, key steps to apply imitation learning to learn the pruning policy will be introduced. The design of the classifier and features, as well as other implementation issues, will be discussed in the following subsections. One can refer to [29] for a more detailed background on imitation learning.

Modeling the tree search process as a sequential decision problem is a common method in machine learning [30]. A sequential decision problem consists of a state space \mathcal{X} , an action space \mathcal{Q} , and a policy space Π . A single trajectory τ consists of a sequence of states $x_1, \dots, x_T \in \mathcal{X}$, a sequence of actions $q_1, \dots, q_T \in \mathcal{Q}$, and a policy $\kappa \in \Pi$ and $\kappa(\cdot) : \mathcal{X} \rightarrow \mathcal{Q}, x_t \mapsto q_t$. The agent uses its own policy κ to make a sequence of decisions $q_t, t = 1, \dots, T$, based on the state it observed, i.e., $x_t, i = 1, \dots, T$. Specifically, at time step t , the agent is aware of the state x_t . It then uses its policy to take an action $\kappa(x_t) = q_t$. The action q_t might have some influences on the environment, i.e., the next state x_{t+1} is influenced by q_t .

The procedure of the branch-and-bound algorithm, essentially a tree search process, is a sequential decision problem because the pruning policy should make a decision at each iteration sequentially. We define ϕ as the feature map, i.e., $\phi(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^*$. When learning the pruning policy, i.e., to learn whether to prune a node N_t , the state x_t is the feature of N_t , denoted as $\phi(N_t)$, and the action is a binary value, i.e., $q_t \in \{\text{prune}, \text{preserve}\}$. Thus, the policy κ is a binary classifier that takes $\phi(N_t)$ as the input and outputs a binary value. We hope the learned policy can find the optimal solution with the minimal number of nodes expanded. Therefore, only the nodes whose feasible regions containing the optimal solutions should be preserved. These nodes are referred to as *optimal nodes*. We then label all the optimal nodes as *preserve* and other nodes as *prune*.

With the feature and label of every node, the policy can be learned via supervised learning. However, directly applying supervised learning is impractical even for medium-size wireless networks, as there would be a huge number of nodes. Meanwhile, including only part of the nodes will degrade the performance. Imitation learning differs from supervised learning in the way of collecting features and labels for the classifier. For supervised learning, all the features and labels are needed before the training process. In contrast, imitation learning iteratively collects data to correct the mistakes the learned policy makes, and the data collection algorithm is called DAgger [28]. With DAgger, at the i -th iteration, the standard policy in the branch-and-bound algorithm is replaced with the learned policy $\pi^{(i)} \in \Pi$. We collect the features of all the nodes explored by the learned policy and their corresponding labels. Then a new policy $\pi^{(i+1)}$ is learned via supervised learning on the dataset collected at previous iterations, and it corrects the mistakes made by $\pi^{(i)}$. In this way,

the learned policy approaches the optimal policy progressively. The detailed algorithm will be shown in Section III-D.

B. Neural Networks as Classifiers

In this paper, adopting imitation learning, we propose to use a neural network as the classifier for the pruning policy. One advantage of neural networks is that we can tune a parameter to control the computation-performance trade-off. This will not only enable us to use an iterative algorithm to guarantee the feasibility and outperform non-optimal labels, but also assist the development of the transfer learning approach in Section IV.

Recall that learning the pruning policy is a binary classification problem, where the input is the feature vector of the node and the output is a binary class in $\{prune, preserve\}$. We employ an L -layer multi-layer perceptron (MLP), a type of neural networks, as the binary classifier with L -layer. In each layer of MLP, the input is multiplied with a learned weight matrix, followed by an element-wise nonlinear transformation. Specifically, the k -th layer's output is calculated as:

$$\mathbf{g}^k = \text{Relu}(\mathbf{W}^k \mathbf{g}^{k-1} + \mathbf{b}^k), \quad (5)$$

where \mathbf{W}^k and \mathbf{b}^k are the learned parameters of the k -th layer. $\mathbf{g}^k, k = 1, \dots, L$, denotes the output of the k -th layer, and \mathbf{g}^0 is the input feature vector, which will be designed in the next subsection. $\text{Relu}(\cdot)$ is the rectified linear unit function, i.e., $\max(0, \cdot)$. The output indicates the probability of each class, which is a normalization of the last layer's output $\mathbf{g}^L \in \mathbb{R}^2$:

$$e[i] = \frac{\exp(\mathbf{g}^L[i])}{\sum_{j=1,2} \exp(\mathbf{g}^L[j])}, i = 1, 2, \quad (6)$$

where $e[i]$ indicates the i -th component of vector e .

In the training stage, the loss function is the weighted cross entropy, given by:

$$\ell = - \sum_{j=1,2} \mathbf{w}[j] \cdot \mathbf{y}[j] \log(e[j]), \quad (7)$$

where \mathbf{y} is the label, i.e., $\mathbf{y} = (1, 0)$ indicates that it belongs to the class *prune*, and $\mathbf{y} = (0, 1)$ otherwise. \mathbf{w} denotes the weight of each class, which is tuned by hand. Two parts contribute to the weight. First, if the number of non-optimal nodes is much larger than the number of optimal nodes, we should assign a higher weight to the class *preserve* in order to let the neural network not to ignore this class. We denote this part as \mathbf{o}_1 and it can be computed by:

$$\begin{aligned} \mathbf{o}_1[1] &= \frac{\# \text{ optimal nodes in the training set}}{\# \text{ nodes in the training set}}, \\ \mathbf{o}_1[2] &= 1 - \mathbf{o}_1[1]. \end{aligned} \quad (8)$$

Second, when the number of feasible solutions is small, we should assign a higher weight to the optimal nodes in the training dataset in order not to miss good solutions. This parameter, denoted as \mathbf{o}_2 , is tuned by hand to achieve good performance on the validation dataset. The total weight is calculated by $\mathbf{w} = \mathbf{o}_1 \odot \mathbf{o}_2$, where \odot is a hadamard product.

The classifier controls whether to prune a node. Pruning too aggressively leads to finding no feasible solutions. Pruning too

moderately will lead to a large number of preserved nodes and high computation cost. Thus, a desirable property of the classifier is an ability to dynamically control the search space during the test stage.

Control the search space: During the test phase, the MLP outputs e , which indicates the probability of each class, followed by setting a threshold indicating which class it belongs to. The threshold for the class *prune* is denoted as Λ . In the standard classification problem, $\Lambda = 0.5$. If $e[1] > 0.5$, the input should belong to the first class. Otherwise, it should belong to the second class. We can adjust the threshold to control the search space. Specifically, a pruning policy with a higher threshold will preserve more nodes than that with a lower threshold, which leads to a larger search space and better performance. In the simulations, we increase the threshold iteratively if we cannot get any feasible solution. Specifically, we increase the threshold in a moderately exponential way, i.e., $\Lambda_k = 1 - 0.5 \cdot 0.8^k$ at the k -th iteration. The specific method is shown in Algorithm 2. This method is shown to be very effective in the simulations, and the total number of iterations is smaller than 3 when the training task and test task are the same.

C. Feature Design

The input features of the neural network play a critical role in its accuracy, which determines the performance and computational complexity of the proposed framework. As a result, they should be carefully designed by leveraging domain knowledge. The available information consists of the problem data and the structure of the binary search tree. We propose to include two types of features, problem-independent features and problem-dependent features, which correspond to the structures of the tree and problem data respectively.

1) *Problem-independent Features:* These kinds of features only contain information about the branch-and-bound tree and the solutions and objective values of the relaxed problems. They are universal to all MINLP problems and are invariant from problem to problem. Thus, the learning framework will have a good generalization capability if we only use problem-independent features. Such features can be categorized into three groups:

- 1) Node features, computed merely from the current node N_i . They contain the depth of N_i and the optimal objective value of the relaxed problem z_i .
- 2) Branching features, computed from the branching variable. They contain the branching variable's value at the current node, i.e., $a_i^*[j]$, that in the relaxed problem at its father node, and that in the relaxed problem of the root node, i.e., $a_1^*[j]$.
- 3) Tree features, computed from the branch-and-bound search tree. They contain the optimal objective value at the root node, i.e., z_1 , the number of solutions found ever, and the best objective value found ever, i.e., c^* .

Due to the significant variations among the objective value of \mathcal{P} under different network settings, all the objective values used as features in the branch-and-bound search tree should be normalized by the optimal objective value of the relaxed problem at the root node.

2) *Problem-dependent Features*: Problem-independent features are universal to all MINLP problems. Unfortunately, we cannot learn efficient policies with only problem-independent features, as proved in [31]. Thus, the problem-dependent features, are introduced. In a general MINLP, the problem-dependent features can be the coefficients of integer variables. As we consider problems in wireless communications, domain knowledge can be exploited and the problem dependent features can be channel state information (CSI) or some descriptions of the radio resources, e.g., power features. The design of problem dependent features for the network power minimization in Cloud-RANs is described in Section V. To maintain a good generalization capability, problem-dependent features are also normalized.

D. The LORM Framework

In this subsection, we present the algorithm of the LORM framework in details, which consists of a problem instance collection phase, a training phase, and a test phase.

a) *The problem instance collection phase*: We collect the problem parameters of each randomly generated wireless network, corresponding to one problem instance. For each problem instance, we employ the standard branch-and-bound algorithm to get the optimal solution. The nodes whose feasible set contains the optimal solution are marked as optimal nodes.

b) *The training phase*: The DAGger algorithm is used to collect data and train the classifier during the training phase. DAGger is an iterative algorithm, with a total of M iterations. we set $\pi^{(0)}$ as a random policy. At the i -th iteration, we have a learned policy $\pi^{(i-1)}$, i.e., a trained neural network classifier. We use a modified branch-and-bound algorithm to collect data and train the classifier. The modified branch-and-bound algorithm uses $\pi^{(i-1)}$ instead of the standard pruning policy introduced in Section II-B. Specifically, for each problem instance, the relaxed problem at the root node is firstly solved to get the problem-independent features $(\mathbf{a}_1^*, \mathbf{w}_1^*, z_1^*)$. A variable is selected, and two children nodes of the root node are expanded. For each child N_t , $\pi^{(i-1)}$ is used to decide whether to prune or preserve the node based on the problem-dependent feature and $(\mathbf{a}_1^*, \mathbf{w}_1^*, z_1^*)$. We then add the node feature and corresponding labels of the node $(\phi(N_t), \kappa^*(N_t))$ into a set \mathcal{D} . After all the nodes in this problem instance are explored or pruned, the optimal nodes are also added into \mathcal{D} . This is because the optimal nodes represent the desired behavior and we do not want to miss any of them. The new policy $\pi^{(i)}$ is trained on \mathcal{D} with supervised learning and stochastic gradient descent after all the problem instances are solved via the modified branch-and-bound algorithm. After M iterations, a sequence of policies $\pi^{(0)}, \dots, \pi^{(M)}$ are obtained and we select the policy κ that performs best on the validation dataset in the test.

c) *The test phase*: During the test phase, we replace the standard pruning policy in the branch-and-bound algorithm with the learned policy κ . For a non-leaf node N_t , the learned policy κ determines whether to prune the node or not. For the leaf node, the subproblem becomes a convex problem

Algorithm 2 LORM Branch-and-bound(κ)

```

1:  $\mathcal{L} \leftarrow \{N_1\}, \mathcal{D} \leftarrow \{\}, t \leftarrow 0, c^* = +\infty$ 
2:  $(\mathbf{a}_1^*, \mathbf{w}_1^*), z_1^* \leftarrow$  solve the relaxed problem of  $\mathcal{P}_1$ 
3: while  $c^* = \infty$  do
4:   Increase the threshold  $\Lambda$  of  $\kappa$ 
5:   while  $\mathcal{L} \neq \emptyset$  do
6:      $t \leftarrow t + 1$ 
7:      $N_t \leftarrow$  select a node from  $\mathcal{L}$ 
8:     if  $N_t$  is leaf node then
9:        $(\mathbf{a}_t^*, \mathbf{w}_t^*), z_t^* \leftarrow$  solve and save the relaxed problem of
        $\mathcal{P}_t$  or load from  $\mathcal{T}$ 
10:       $c^* \leftarrow \min(c^*, z_t^*)$ 
11:     else
12:        $f \leftarrow \phi(N_t)$ 
13:       if  $\kappa(f) = \text{preserve}$  then
14:          $N_t^1, N_t^2 \leftarrow$  two children of  $N_t$ 
15:          $\mathcal{L} \leftarrow \mathcal{L} \cup \{N_t^1, N_t^2\}$ 
16:       end if
17:     end if
18:   end while
19: end while
20: return  $c^*$ 

```

since all the integer variables are determined. It is solved via some convex solvers to check the feasibility and obtain the values for continuous variables. However, this may produce infeasible solutions if the pruning policy is too aggressive. Recall that in Section III-B, we propose to tune a parameter to control the search space. Specifically, we can iteratively increase the threshold Λ if the problem is infeasible. During this process, the convex problem at the leaf node needs to be solved several times, which is not time-efficient. We propose to build a lookup table to accelerate this process. The lookup table is indexed by the value of integer variables and is used to store the solutions of the relaxed problems. If we encounter a relaxed problem that has not been solved before, we solve it and save this problem instance and its solution into the lookup table. Otherwise, we directly extract the solution from the lookup table. The pseudo-code for our proposed algorithm is shown in Algorithm 2.

E. Complexity Analysis

In this subsection, we present some complexity analysis for the LORM framework. We follow the common assumptions and analysis for imitation learning [28], and study the expected number of nodes explored, which indicates the number of samples we have, and the expected number of convex problems solved, which indicates the computational complexity of the proposed algorithm.

Theorem 1. *Given the number of integer variables L , a node pruning policy which expands a non-optimal node with probability ϵ_1 and prunes an optimal node with probability ϵ_2 , the expected number of nodes explored and the number of relaxed problems solved are $O(L^2)$ and $O(L)$, respectively, when $\epsilon_1 \leq 0.5, \epsilon_2 \leq 1$. The expected number of nodes explored and the number of relaxed problems solved are $O(L)$ and $O(1)$, respectively, when $\epsilon_1 \leq \frac{1}{3}, \epsilon_2 \leq 1$.*

As shown in Theorem 1, the number of convex problems to solve is $O(L)$ if we have a proper choice of parameters and

features. Thus, LORM enjoys a much lower computational complexity compared with the standard branch-and-bound algorithm. When LORM is applied to the network power minimization problem in Cloud-RANs, which is mixed integer second order cone programming, the expected time complexity is $O(L^4 K^3 + L^3 K^2 \frac{1}{\epsilon})$ if the conic optimization-based convex optimization solver [32] is used. This is much more efficient than the standard branch-and-bound, and comparable to the state-of-the-art method for this problem [6], but with better performance, as to be shown in the simulations.

We further make a computational complexity comparison with the method in [23], referred to as ‘‘MILP-SVM’’, by comparing the number of relaxed problems to solve. The complexity analysis of their method is shown in Theorem I in [23]. In addition to ϵ_1 and ϵ_2 , they assume that the node selection policy ranks some non-optimal node higher than an optimal node with probability ϵ . The number of relaxed problems in MILP-SVM to solve then goes to ∞ compared to $1 + \epsilon_1 L$ in LORM when ϵ_1 goes to 0.5. The number of relaxed problems in MILP-SVM to solve is $O(L)$ compared to $O(1)$ in LORM when $\epsilon_1 \leq \frac{1}{3}$. This shows that LORM achieves significant speedup compared with MILP-SVM.

The number of training samples for the classifier is the sum of nodes explored in the branch-and-bound tree of each problem instance. Note that the number of problem instances is $|\mathcal{P}|$ and the number of nodes is $O(L^2)$. Therefore, the number of training samples for the classifier is $O(L^2 |\mathcal{P}|)$.

F. Discussions on LORM

The LORM framework enjoys a few unique advantages. Specifically, 1) it has the capability to generalize to different settings; 2) it is able to learn the optimization policy with few problem instances; and 3) can work with different baseline algorithms, as explained below.

Generalization to different settings: As the input and output dimensions of most learning algorithms, e.g., MLP, in the test stage must be the same as in the training stage, most of the end-to-end learning methods [10], [11], [12] cannot generalize to settings with different problem sizes than those in the training dataset. On the other hand, the input and output dimensions of the policies in the branch-and-bound algorithm are invariant when the problem size changes. Thus, generalization to different problem sizes is not a problem when learning the policies in the branch-and-bound algorithm. Such generalization capability also relies on the design of features (as shown in Section III-C), and will be tested via simulations.

Few-shot learning via learning policies: LORM works with few problem instances for two reasons. The first reason is that the available number of samples for learning the classifier is much larger than the problem instances. According to the analysis in Section III-E, the number of samples is $O(L^2 |\mathcal{P}|)$, where L is the number of integer variables and $|\mathcal{P}|$ is the number of problem instances. Thus, solving one problem instance produces many labeled samples for policy learning. The second reason is that to learn the pruning policy is much easier than directly learning the input-output mapping of an optimization algorithm, which is typically with high

dimensional inputs and outputs. These two reasons allow us to use a much lighter-weighted learning algorithm, with which we can achieve few-shot learning.

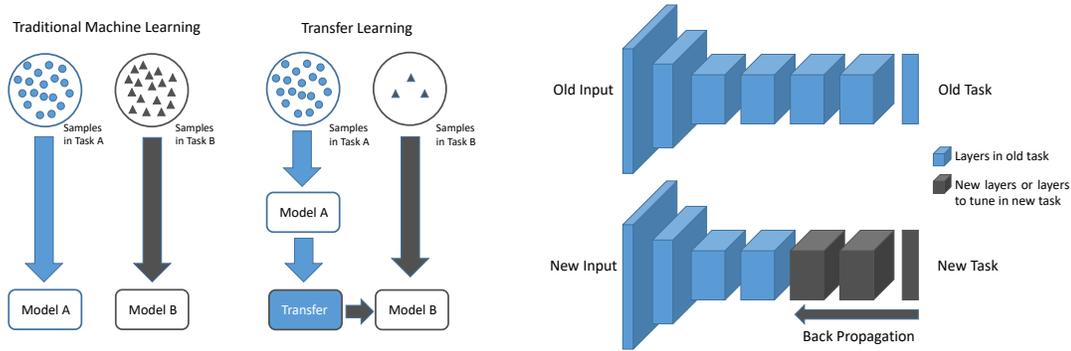
Sub-optimal labels: Since the branch-and-bound is time-consuming, we may also generate labels by sub-optimal algorithms. Suppose the solution obtained by the sub-optimal algorithm is $\mathbf{a}^*, \mathbf{w}^*$, and we label all the nodes whose feasible sets containing $\mathbf{a}^*, \mathbf{w}^*$ as *preserve* and other nodes as *prune*. As the tree search procedure is preserved, the LORM framework still applies, and thus it can work with different available algorithms for a given MINLP problem.

IV. LORM-TL: TRANSFER LEARNING VIA SELF-IMITATION

While the LORM framework helps to achieve few-shot learning for MINLP resource management problems, its practical implementation still faces a few key challenges, which will be illustrated in this section. We will propose a transfer learning method to address these challenges.

A. Limitations of Machine Learning Methods

An essential assumption of most machine learning algorithms is that the training and future testing data are in the same feature space with the same distribution, i.e., the same task [33]. Narrowing down to problems in wireless networks, the ‘‘distribution’’ includes ‘‘structure’’, e.g., large-scale fading and signal-to-noise-ratio (SNR), and ‘‘size’’, e.g., the numbers of users, base stations, and antennas. The performance deteriorates when the machine learning task to be tested is different from the trained one. This is referred to as the *task mismatch* problem. How much will the performance deteriorate is determined by the similarity between the tasks, which is little known and an on-going research problem in machine learning. A straightforward way to resolve this issue is to collect enough additional labeled training samples to train a neural network for the new setting from scratch. This will achieve good performance, but it is impractical in real systems because training neural networks requires a large amount of samples and long computing time. To cope with the dynamics of wireless networks, it is highly desirable to reduce the training time for the new task, i.e., when the network setting changes. As discussed in Section III.F, LORM enjoys certain generalization capability, and can handle minor scenario changes. But its performance deteriorates when the setting changes dramatically. Note that although the tasks are different, they share something in common, i.e., the same structure of the underlying optimization problem. In other words, the knowledge learned in the old task can be helpful for the new task. Thus, it is possible to train a new model with only a few additional samples if we can effectively transfer such knowledge. Such an approach is called *transfer learning*, which can significantly reduce the training time and achieve good performance in the new task. The difference between traditional machine learning and transfer learning is shown in Fig. 3(a).



(a) The key difference between transfer learning and traditional machine learning is that transfer learning can tackle the task mismatch problem with few additional training samples. This is achieved by transferring the knowledge form the old task into the new task.

(b) An illustration of transfer learning via fine-tuning. The neural network for the new task keeps the first several layers from the neural network for the old task. The last several layers are trained from scratch or tuned from those for the old task with a small learning rate.

Fig. 3. Illustrations of transfer learning.

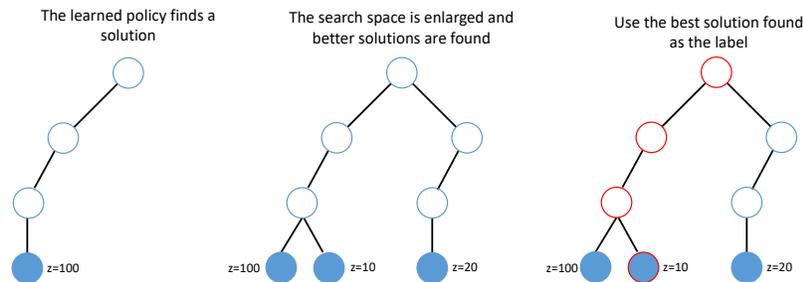


Fig. 4. An illustration of the self-exploration algorithm. The learned policy from the original task fails to find a good solution. Then, the search space is enlarged to find more solutions by tuning the threshold Λ . The best solution found is used as the label for LORM.

B. Self-imitation with Unlabeled Samples

In this subsection, we first present transfer learning via fine-tuning and discuss its limitations. Then, we introduce an approach called *self-imitation* to address the problems for fine-tuning.

Fine-tuning is the most frequently employed method for transfer learning in neural networks. Neural networks are usually trained via stochastic gradient descent (SGD). For different layers, we can have different learning rates, i.e., the step size of the SGD. Fine-tuning is to tune the learning rate of each layer to refine the pre-trained neural network on the additional training dataset. We can train the pre-trained neural network with a small learning rate. The knowledge learned on the original dataset can serve as a good initialization point. The learning rate for fine-tuning is small because we expect that the initial weights are relatively good, so distorting them too much and too quickly is not a smart choice. The illustration of fine-tuning is shown in Fig. 3(b). Fine-tuning reduces the training time, but it needs additional labeled samples. The time cost is still expensive as the computational complexity of branch-and-bound to generate the labels, i.e., the optimal solutions, is exponential. This implies that we even have difficulty in generating a small amount of training labels if the network size of the new setting is large. Hence, it will be desirable if we can refine the model with unlabeled data.

Algorithm 3 DAGger(κ)

```

1:  $\mathcal{D} \leftarrow \{\}, c^{*(j)} \leftarrow +\infty, \mathcal{N}_{opt}^{(j)} \leftarrow \{\}, \mathcal{T}^{(j)} \leftarrow \{\}, \pi^{(0)} \leftarrow \kappa$ 
2: for  $k \leftarrow 1, \dots, M$  do
3:   for  $j \leftarrow 1, \dots, |\mathcal{P}|$  do
4:     Set a high threshold  $\Lambda$  in  $\pi^{(k-1)}$ 
5:      $c^{*(j)}, \mathcal{N}_{opt}^{(j)}, \mathcal{D}^{(kj)}, \mathcal{T}^{(j)}, \leftarrow$ 
       COLLECT-SI( $\pi^{(k-1)}, \mathcal{P}_j, \mathcal{N}_{opt}^{(j)}, \mathcal{T}^{(j)}, c^{*(j)}$ )
6:      $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}^{(kj)}$ 
7:   end for
8:    $\pi^{(k)} \leftarrow$  fine-tune the classifier  $\pi^{(k)}$  using data  $\mathcal{D}$ 
9: end for
10: return best  $\pi^{(k)}$  on validation set

```

We propose to use transfer learning via self-imitation [26] to address the task mismatch problem with a few additional unlabeled samples. Self-imitation takes advantage of the exploration in reinforcement learning, which leads to an unsupervised paradigm. The key idea is to explore the search tree, and use the best solution found during the exploration as the label for imitation learning. In this way, the search space is enlarged and may contain better solutions than the learned policy. To transfer the knowledge in the learned policy, we harness the method proposed in Section III-B to enable a larger and better search space. This exploration policy is more effective than traditional exploration policies, e.g., ϵ -greedy

Algorithm 4 COLLECT ($\kappa, p, N_{old}, \mathcal{T}, c_{old}^*$)

```

1:  $\mathcal{L} \leftarrow \{N_1\}, \mathcal{D} \leftarrow \{\}, t \leftarrow 0, N_{opt} \leftarrow \{N_1\}$ 
2: Perform Algorithm 2 with  $\kappa$ 
    $c^* \leftarrow$  best solution found,  $N_{opt} \leftarrow$  the node containing optimal
   solution
   Save all visited nodes and features into  $\mathcal{D}$ 
   Save solutions of solved convex problems into  $\mathcal{T}$ 
3: if  $c^* < c_{old}^*$  then
4:    $N_{opt} \leftarrow$  the nodes from root to  $N_{opt}$ 
5: else
6:    $N_{opt} \leftarrow N_{old}, c^* \leftarrow c_{old}^*$ 
7: end if
8: for  $j \leftarrow 1, \dots, |\mathcal{D}|$  do
9:   if  $\text{node}(D_j) \in N_{opt}$  then
10:     $\mathcal{D}_j \leftarrow \{D_j.f, \text{preserve}\}$ 
11:   else
12:     $\mathcal{D}_j \leftarrow \{D_j.f, \text{prune}\}$ 
13:   end if
14: end for
15: return  $c^*, N_{opt}, \mathcal{D}, \mathcal{T}$ 

```

policy [34].

The old policy $\pi^{(0)}$ was trained on the original training dataset. As the task changes, we collect a few additional unlabeled samples. The performance of the learned policy may not be good on the new training dataset. We set the learned policy with a higher threshold Λ as the exploration policy. Algorithm 2 is performed on the new training dataset with the exploration policy. It will achieve better performance than the learned policy because the search space is larger. For each problem, we record the best solution found c^* during the exploration process. The nodes on the path from the root to the one containing the best solution are marked as the optimal nodes, i.e., labeled as *preserve*. Other nodes are labeled as *prune*. A new policy $\pi^{(1)}$ is trained on the new training dataset. We repeat such a process with DAgger [28]. The performance of the labels is improved iteratively by repeating such a process, and thus the performance of the learned policy will be improved iteratively. DAgger goes through the training dataset several times and we need to solve the convex problems at the leaf node each time. Thus, some convex problems are solved repeatedly in the training process. We propose to use the lookup table \mathcal{T} proposed in Section III-D to store the solutions. This will speed up the training process.

Remark: The idea of self-imitation has been proposed and applied to solve medium-size MILPs in [26]. There are a few challenges to apply it to the LORM framework. Firstly, an incremental strategy was adopted in [26] to transfer to a different size. For example, in order to transfer from a network with 5 users to a network with 15 users, we must obtain training data of the networks with 5, 6, \dots , 15 users. It is impractical to get these data in wireless networks. Indeed, transferring from the network with 5 users to the network with 15 users directly will not provide a good performance because of the ineffectiveness of traditional exploration policies. Thus, we propose to increase the threshold Λ of the pruning policy as the exploration policy. Such a policy increases the search space, which is the requirement of the exploration policy, and possesses the knowledge in the original task. In addition,

since some relaxed problems are solved repeatedly during the training, we further use the lookup table proposed in Section III-D to accelerate the training process.

V. SIMULATIONS

In this section, we test the performance of LORM and LORM-TL, with the network power minimization problem in Cloud-RANs as an example.

A. Network Power Minimization in Cloud-RANs

The network power minimization problem in Cloud-RANs [6] is a typical MINLP problem in wireless networks, as it consists of discrete variables (i.e., the selection of remote radio heads (RRHs) and fronthaul links), continuous variables (i.e., downlink beamforming coefficients), and QoS constraints and power constraints.

Consider a Cloud-RAN with L RRHs and K single-antenna mobile users (MUs), where the l -th RRH is equipped with N_l antennas. The baseband unit (BBU) pool is connected to all the RRHs via a high-bandwidth, low-latency fronthaul network, and performs centralized signal processing. The corresponding SINR for the k -th MU is given by

$$\text{SINR}_k = \frac{|\sum_{l \in \mathcal{L}} \mathbf{h}_{kl}^H \mathbf{w}_{lk}|^2}{\sum_{i \neq k} |\sum_{l \in \mathcal{L}} \mathbf{h}_{kl}^H \mathbf{w}_{li}|^2 + \sigma_k^2}, \forall k \in \mathcal{S},$$

where $\mathbf{w}_{lk} \in \mathbb{C}^{N_l}$ denotes the transmit beamforming vector from RRH l to MU k , $\mathbf{h}_{kl} \in \mathbb{C}^{N_l}$ denotes the channel vector between the k -th MU and the l -th RRH, and σ_k^2 is the variance of additive noise.

Let a binary vector $\mathbf{a} = (a_1, \dots, a_L)$ with $a_i \in \{0, 1\}$ denote the mode of each RRH, i.e., $a_i = 1$ if the i -th RRH and the corresponding transport link are switched on. Each RRH has its own transmit power constraint

$$\sum_{k=1}^K \|\mathbf{w}_{lk}\|_2^2 \leq a_l \cdot P_l, l \in \{1, \dots, L\},$$

where $\|\cdot\|_2$ is the ℓ_2 -norm of a vector.

The network power consumption in Cloud-RAN consists of the relative fronthaul network power consumption and the total transmit power consumption [6]. Assuming that the target SINR for the k -th user is γ_k , the network power minimization problem is formulated as the following MINLP problem:

$$\begin{aligned} & \underset{\mathbf{w}, \mathbf{a}}{\text{minimize}} && \sum_{l=1}^L a_l \cdot P_l^c + \sum_{l=1}^L \sum_{k=1}^K \frac{1}{\eta_l} \|\mathbf{w}_{lk}\|_2^2 \\ & \text{subject to} && \sqrt{\sum_{i \neq k} |\mathbf{h}_{kl}^H \mathbf{w}_{li}|^2 + \sigma_k^2} \leq \frac{1}{\gamma_k} \Re(\mathbf{h}_{kl}^H \mathbf{w}_{lk}), \forall k \\ & && \sqrt{\sum_{k=1}^K \|\mathbf{A}_{lk} \mathbf{w}_{lk}\|_2^2} \leq a_l \cdot \sqrt{P_l}, \forall l \\ & && a_l \in \{0, 1\}, \forall l, \end{aligned} \quad (9)$$

where P_l^c is the relative fronthaul link power consumption, i.e., the power saved when both the RRH and the corresponding fronthaul link are switched off, and η_l is the drain efficiency of

the radio frequency power amplifier. Likewise, the aggregative beamforming vector is $\mathbf{w} = [\mathbf{w}_1^T, \dots, \mathbf{w}_K^T]^T \in \mathbb{C}^{NK}$ with $\mathbf{w}_k = [\mathbf{w}_{1k}^T, \dots, \mathbf{w}_{Lk}^T]^T \in \mathbb{C}^N$, $N = \sum_{l=1}^L N_l$, $\mathbf{h}_k = [\mathbf{h}_{1k}^T, \dots, \mathbf{h}_{Lk}^T]^T \in \mathbb{C}^N$, $\Re(\cdot)$ denotes the real part of a complex scalar, and $\mathbf{A}_{lk} \in \mathbb{C}^{N \times N}$ is a block diagonal matrix with an identity matrix \mathbf{I}_{N_l} as the l -th main diagonal block matrix and zeros elsewhere. This problem is reduced to a second-order cone programming (SOCP) problem if the integer variable \mathbf{a} is fixed. Thus, it belongs to the problem class \mathcal{P} and can be solved by LORM. In our simulations, we consider the channel model $\mathbf{h}_{kl} = 10^{-L(d_{kl})/20} \sqrt{\phi_{kl} s_{kl}} \mathbf{g}_{kl}$, where $L(d_{kl})$ is the path-loss at distance d_{kl} , ϕ_{kl} is the antenna gain, s_{kl} is the shadowing coefficient, and \mathbf{g}_{kl} is the small scale fading coefficient. One can refer to Table I in [6] for the detailed parameter settings. The RRHs and MUs are uniformly and independently distributed in the square region $[-1000, 1000] \times [-1000, 1000]$ meters, i.e., their locations are different in different samples. The fronthaul link power consumption is set to be a random permutation of $P_l^c = (5+l)W$, $l = 1, \dots, L$ in different samples.

B. Performance Comparison

We first compare the proposed methods with the state-of-the-art methods designed for the network power minimization problem in Cloud-RANs, in terms of the performance and computation time. For LORM and LORM-TL, we set the hidden units of the neural network as $\{32, 64, 16\}$. We use the branching variable's value, i.e., $a_i[j]$, and its value at the root problem, i.e., $a_0[j]$, as problem-independent features. Motivated by the observations in [6], we strive to open a small number of RRHs with a small fronthaul link power consumption if the problem is feasible. Thus, the problem-dependent features are selected as the fronthaul link power consumption, i.e., $\frac{P_j^c \cdot L}{\sum_{k=1}^L P_k^c}$. In addition, we use the iterative algorithm to guarantee the feasibility, as introduced in Section III-B, and set the threshold Λ as $1 - 0.5 \cdot 0.8^l$, where l is the iteration index. We set the number of DAgger as $M = 5$, and train 5 epochs for each DAgger iteration. All the methods run on Intel Core i7-6700 CPU @ 3.40Ghz and adopt SCS [32] as the SOCP solver.

The following three benchmarks are considered:

- 1) Relaxed Mixed-integer Nonlinear Programming based algorithm (RMINLP) [35]: This algorithm is a commonly used heuristic algorithm for MINLPs. It applies a deflation procedure to switch off RRHs one-by-one based on the solution of a relaxed problem. It needs to solve $3L$ SOCPs.
- 2) Iterative Group Sparse Beamforming (GSBF) [6]: This algorithm is designed for the network power minimization problem in Cloud-RANs, and is the state-of-the-art solutions for this problem. It needs to solve L SOCPs.
- 3) Branch-and-Bound algorithm (Optimal solution): This is the standard branch-and-bound algorithm with the best-first node selection policy.

1) *Performance of LORM*: We simulate the performance under the setting with the target SINR (TSINR) as 4dB, with 10 2-antenna RRHs and different numbers of users. For each setting, we generate 50 network realizations, i.e., samples,

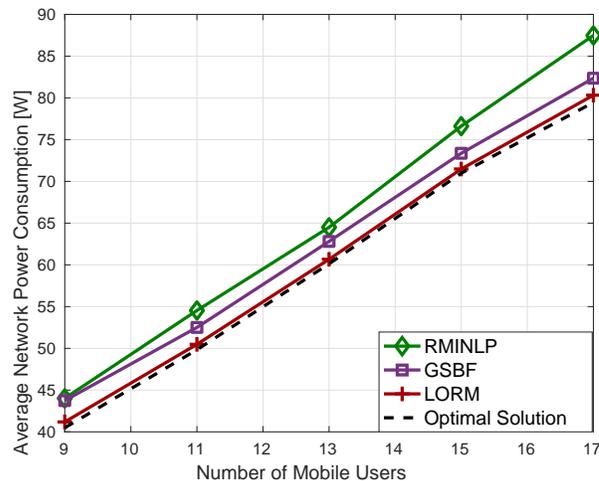


Fig. 5. The average network power consumption versus the number of mobile users.

for training and 50 network realizations for testing. The performance is shown in Fig. 5, and the computation time is shown in Table I. It is observed that LORM outperforms both GSBF and RMINLP, and achieves near-optimal performance. We also see that LORM has a much less running time than RMINLP and branch-and-bound, and achieves about 2x speedup compared with GSBF. The locations of RRHs and MUs in the training dataset and test dataset are different, and thus the results already demonstrate a certain generalization capability of LORM, i.e., generalization to different large-scale fading conditions.

TABLE I
AVERAGE RUNNING TIME FOR THE ALGORITHMS UNDER EACH SETTING (IN SECONDS).

Setting	Branch-and-Bound	LORM	GSBF	RMINLP
$L = 10, K = 13, \text{TSINR} = 4$	91.76s	0.892s	2.562s	5.264s
$L = 10, K = 15, \text{TSINR} = 4$	96.40s	1.157s	3.680s	8.136s
$L = 10, K = 17, \text{TSINR} = 4$	142.0s	2.920s	5.474s	12.64s

2) *Generalization of LORM*: We conduct the following experiment to test the generalization capability of LORM beyond different large-scale fading. We first train the model on 100 samples with $L = 6, K = 9, \text{TSINR} = 0$, and test it on $L = 10, \text{TSINR} = 4$, with different numbers of users. The results are shown in Table II, where “full training” means training with sufficient data for that setting. From the table, we see that even if trained on a different setting, LORM outperforms both GSBF and RMINLP, thus demonstrating its good generalization ability.

TABLE II
THE GAP FROM THE OPTIMAL OBJECTIVE VALUE.

Setting	LORM (full training)	LORM (train on $L = 6$)	RMINLP	GSBF
$L = 10, K = 7, \text{TSINR} = 4$	1.87%	4.39%	7.94%	12.9%
$L = 10, K = 9, \text{TSINR} = 4$	1.73%	2.97%	8.71%	8.04%
$L = 10, K = 11, \text{TSINR} = 4$	1.30%	1.72%	9.44%	5.36%
$L = 10, K = 13, \text{TSINR} = 4$	1.41%	1.41%	7.27%	4.45%
$L = 10, K = 15, \text{TSINR} = 4$	0.70%	1.48%	7.94%	3.36%

3) *Performance of LORM-TL*: Throughout experiments, we observed performance deterioration when the number of MUs change dramatically. In this part, we test the effectiveness of LORM-TL in this challenging situation, where LORM cannot achieve good performance. We conduct two experiments: When the number of MUs in the test dataset is either much less or much larger than the training dataset. In the first experiment, we train LORM-TL on 100 labeled samples with $L = 6$, $K = 6$, $TSINR = 0$. We then transfer on 10 unlabeled samples with $L = 10$, $K = 15$, and test on $L = 10$, $K = 15$. The results are shown in Table III. In the second experiment, we train LORM-TL on 50 labeled samples with $L = 10$, $K = 17$, $TSINR = 4$, and then transfer on 10 unlabeled samples with $L = 10$, $K = 7$, and test on $L = 10$, $K = 7$. For the self-imitation algorithm, the number of iterations for self-imitation is set as $M = 10$. The threshold Λ is set as 0.9 at the first iteration, and is halved if this iteration’s performance is the same as the last iteration. The results are shown in Table IV. As shown in the tables, equipped with transfer learning, LORM-TL outperforms all competing methods, and achieves comparable performance with LORM trained on sufficient samples of the same distribution with the test dataset.

We measure the running time for the full training case and the transfer learning case shown in Table V. For transfer learning, LORM-TL is trained on 100 labeled samples with $L = 6$, $K = 6$, $TSINR = 0$, transferred on 20 unlabeled samples with $L = 10$, $K = 15$, $TSINR = 4$, which correspond to the fourth row in Table III. We see that the self-imitation method achieves over 10x speedup to the full training. For the full training case, the majority of time is spent on generating labels. Self-imitation resolves this problem by learning from unlabeled samples. It is also shown that the time for self-imitation is longer than training neural networks in the full training case. This is because self-imitation requires exploration, which performs the training and inference for M times.

TABLE III

THE PERFORMANCE GAP OF EACH METHOD TO THE OPTIMAL OBJECTIVE VALUE.

Setting	LORM (train on $L = 6$, $K = 6$, $TSINR = 0$)	LORM-TL	LORM (full training)	RMINLP	GSBF
$L = 10$, $K = 15$ $TSINR = 0$	- ¹	0.63%	0.61%	7.07%	6.76%
$L = 10$, $K = 15$ $TSINR = 2$	-	0.38%	0.3%	7.92%	5.06%
$L = 10$, $K = 15$ $TSINR = 4$	-	0.68%	0.39%	7.94%	3.36%

C. Comparison with Existing “Learning to Optimize” Approaches

In this part, we further compare with other learning-based methods to show the effectiveness and efficiency of LORM. The main method we compare with is a modified version of [23], which is referred to as “Modified MILP-SVM”. The original method in [23] is proposed for MILPs, and it can

¹To control the time during the test, we set the maximum number of iterations as 5. For the method in this column, it only finds feasible solutions for 82% of the problems during the test. Hence, we do not record its performance here.

TABLE IV

THE PERFORMANCE GAP OF EACH METHOD TO THE OPTIMAL OBJECTIVE VALUE.

Setting	LORM (train on $L = 10$, $K = 17$, $TSINR = 4$)	LORM (full training)	LORM-TL	RMINLP	GSBF
$L = 10$, $K = 7$ $TSINR = 0$	57.1%	0.72%	2.48%	7.43%	10.0%
$L = 10$, $K = 7$ $TSINR = 2$	43.8%	2.30%	2.79%	5.57%	11.9%
$L = 10$, $K = 7$ $TSINR = 4$	29.9%	1.87%	2.95%	6.89%	12.9%

TABLE V

THE TIME FOR DIFFERENT METHODS

Full Training		Self-imitation
Time for Generating Labels	Time for Training	Total Time
4820s	65.99s	377.9s

be applied to MINLPs with minor modifications. Specifically, we remove some features that are available only for MILPs and add problem-dependent features. Furthermore, since the method in [23] has difficulty to guarantee feasibility, the benchmark method we implemented is tuned on the test dataset instead of the validation dataset for obtaining a feasible solution. We also implemented the method in [10], trained on 20,000 samples generated by the state-of-the-art heuristic algorithm [6], but the output solutions were not feasible. Therefore, we do not include it in the comparison.

The performance gap to the optimal solution and the computation time are recorded under the settings shown in Table VI. For each setting, we generate 50 network realizations, i.e., samples, for training and 50 network realizations for test. From the table, we see that MILP-SVM also accelerates the branch-and-bound for MINLP by tuning on the test dataset to guarantee the feasibility. Nevertheless, the computation burden of MILP-SVM is still high. It is shown that LORM achieves comparable performance with 10x speedup to MILP-SVM. This demonstrates the efficiency of LORM in MINLP problems, and verifies the computational complexity analysis in Section III-E. The acceleration mainly comes from the reduced number of SOCPs.

We also see that the learning-based methods achieves more speedup in the situation $L = 10$ than the situation $L = 6$. This is because the branch-and-bound tree is larger and we have more space for pruning when $L = 10$. When the feasible set become smaller, e.g., from $TSINR = 0$ to $TSINR = 4$, the neural network should prune nodes more carefully, which leads to a slight increase of computation time for learning-based methods.

VI. CONCLUSIONS

In this paper, we proposed two machine learning-based frameworks, namely LORM and LORM-TL, to achieve near-optimal performance for MINLP resource management problems in wireless networks. We first proposed LORM, an imitation learning-based approach, to learn the pruning policy in the optimal branch-and-bound algorithm. By exploiting the algorithm structure, LORM not only achieves near-optimal performance with few training samples, but also guarantees

TABLE VI
THE PERFORMANCE GAP TO THE OPTIMAL OBJECTIVE VALUE AND THE AVERAGE COMPUTATION TIME OF EACH METHOD.

Settings	Performance Gap		Time		
	LORM	Modified MILP-SVM	LORM	Modified MILP-SVM	Branch-and-Bound
$L = 6, K = 8, \text{TSINR}=0$	0.02%	0.02%	0.102s	0.68s	3.24s
$L = 6, K = 8, \text{TSINR}=2$	1.54%	1.63%	0.131s	0.75s	3.1s
$L = 6, K = 8, \text{TSINR}=4$	0.80%	0.62%	0.118s	1.098s	2.9s
$L = 10, K = 15, \text{TSINR}=0$	0.35%	0.63%	0.974s	6.634s	107.2s
$L = 10, K = 15, \text{TSINR}=2$	0.59%	0.32%	1.09s	10.08s	114.2s
$L = 10, K = 15, \text{TSINR}=4$	1.25%	1.17%	1.1s	11.04s	96.4s

feasibility of the constraints, and can be generalized to problem instances with larger sizes than the training dataset. Furthermore, we proposed LORM-TL to address the task mismatch issue for LORM by relying on only a few additional unlabeled training samples. Extensive experiments have been conducted to verify the effectiveness of both LORM and LORM-TL.

Different from previous approaches for “learning to optimize” that directly approximate the input-output mapping, our proposed framework achieves few-shot learning by exploiting the algorithm structure. As training samples are in general expensive to obtain, we expect that the proposed framework will facilitate the adoption of machine learning-based methods in wireless communications. For the next stage, it will be interesting to test the proposed methods on other design problems in wireless networks, and develop approaches to further reduce the sample complexity and computational complexity.

APPENDIX A PROOF OF THEOREM 1

We first derive the expected number of nodes explored. Consider an n -layer full binary tree. Denote the number of nodes whose root is an optimal node as $a(n)$, and the number of nodes whose root is not an optimal node as $b(n)$. Further assume that the node pruning policy expands a non-optimal node with probability ϵ_1 and prunes an optimal node with probability ϵ_2 . Define E_1 as the expected number of nodes explored and E_2 as the number of relaxed problems to solve.

Then, we have the following recurrence formula

$$\begin{cases} a(n) = (1 - \epsilon_2) \cdot a(n - 1) + \epsilon_1 \cdot b(n - 1) + 1 \\ b(n) = 2\epsilon_1 \cdot b(n - 1) + 1, \end{cases} \quad (10)$$

with initial states as $a(1) = b(1) = 1$.

The number $a(n)$ is needed for computing the number of nodes explored. Solving the recurrence formula, we obtain

$$a(n) = \begin{cases} (1 - \epsilon_2)^{n-1}(1 - \sigma_2 - \sigma_1) + \sigma_2 \cdot (2\epsilon_1)^{n-1} + \sigma_1, & \epsilon_2 \neq 0, \\ n + \frac{\epsilon_1}{1 - 2\epsilon_1} \left(n - 1 - \frac{2\epsilon_1}{1 - 2\epsilon_1} \right) + \frac{\epsilon_1(2\epsilon_1)^n}{(1 - 2\epsilon_1)^2}, & \epsilon_2 = 0, \end{cases} \quad (11)$$

where $\sigma_1 = \frac{1 - \epsilon_1}{\epsilon_2(1 - 2\epsilon_1)}$, $\sigma_2 = \frac{2\epsilon_1^2}{(1 - 2\epsilon_1 - \epsilon_2)(1 - 2\epsilon_1)}$.

Taking $E_1 = a(L + 1)$ completes the proof for the expected number of nodes. It is shown in (10) that $a(n)$ is monotonically increasing when ϵ_1 is increasing or ϵ_2 is decreasing. Thus, in order to prove the computational complexity bound, it is sufficient to show that $E_2 = O(L^2)$, when $\epsilon_1 = 0.5$ and $\epsilon_2 = 0$.

By letting $\epsilon_1 = 0.5$ in (11), we get $a(n) = 0.5(n - 1)^2 + n$. Therefore, $E_1 = a(L + 1) = O(L^2)$, when $\epsilon_1 \leq 0.5$. Similarly, taking $\epsilon_1 = 0.3$ and $\epsilon_2 = 1$, we get $a(n) = 2n - 3 + 3\left(\frac{2}{3}\right)^n = O(n)$. Therefore, $E_1 = a(L + 1) = O(L)$, when $\epsilon_1 \leq \frac{1}{3}$. This completes the proof for the expected number of nodes.

Next, we consider the number of relaxed problems to be solved. Consider an n -layer full binary tree. Also denote the number of relaxed problems whose root is an optimal node as $c(n)$, and the number of relaxed problems whose root is not an optimal node as $d(n)$. Since we only solve the convex problems at the leaf nodes, the recurrence formula can be represented as

$$\begin{cases} c(n) = (1 - \epsilon_2)c(n - 1) + \epsilon_1 d(n - 1) \\ d(n) = 2\epsilon_1 d(n - 1) \end{cases} \quad (12)$$

with initial states as $c(1) = d(1) = 1$.

Solving the recurrence formula, we get $c(n) \leq 1 + \epsilon_1 \left(\frac{1 - (2\epsilon_1)^{n-1}}{1 - 2\epsilon_1} \right)$ when $\epsilon_2 \leq 1, n \geq 2$. This implies $E_2 = c(L + 1) \leq 1 + \epsilon_1 L = O(L)$ when $\epsilon_1 \leq 0.5$, and $E_2 = c(L + 1) \leq 2 = O(1)$ when $\epsilon_1 \leq \frac{1}{3}$.

ACKNOWLEDGMENTS

The authors would like to thank anonymous reviewers and the associate editor for their constructive comments.

REFERENCES

- [1] Y. Shen, Y. Shi, J. Zhang, and K. B. Letaief, “Scalable network adaptation for Cloud-RANs: An imitation learning approach,” in *Proc. IEEE Global Conf. Signal Info. Process.*, Anaheim, CA, Dec. 2018.
- [2] Y. Shen, Y. Shi, J. Zhang, and K. B. Letaief, “Transfer learning for mixed-integer resource allocation problems in wireless networks,” in *Proc. IEEE Int. Conf. Comm.*, Shanghai, China, May 2019.
- [3] Z. Han and K. R. Liu, *Resource allocation for wireless networks: basics, techniques, and applications*. Cambridge, U.K.: Cambridge Univ. Press, 2008.
- [4] C. Y. Wong, R. S. Cheng, K. B. Letaief, and R. D. Murch, “Multiuser OFDM with adaptive subcarrier, bit, and power allocation,” *IEEE J. Sel. Areas Commun.*, vol. 17, pp. 1747–1758, Oct. 1999.
- [5] Q. Ye, B. Rong, Y. Chen, M. Al-Shalash, C. Caramanis, and J. G. Andrews, “User association for load balancing in heterogeneous cellular networks,” *IEEE Trans. Wireless Commun.*, vol. 12, pp. 2706–2716, Jun. 2013.
- [6] Y. Shi, J. Zhang, and K. B. Letaief, “Group sparse beamforming for green Cloud-RAN,” *IEEE Trans. Wireless Commun.*, vol. 13, pp. 2809–2823, May 2014.
- [7] Y. Mao, J. Zhang, and K. B. Letaief, “Dynamic computation offloading for mobile-edge computing with energy harvesting devices,” *IEEE J. Sel. Areas Commun.*, vol. 34, pp. 3590–3605, Dec. 2016.
- [8] Y. Bengio, A. Lodi, and A. Prouvost, “Machine learning for combinatorial optimization: a methodological tour d’horizon,” *arXiv preprint arXiv:1811.06128*, 2018.

[9] A. Zappone, M. Di Renzo, M. Debbah, T. T. Lam, and X. Qian, "Model-aided wireless artificial intelligence: Embedding expert knowledge in deep neural networks towards wireless systems optimization," *arXiv preprint arXiv:1808.01672*, 2018.

[10] H. Sun, X. Chen, Q. Shi, M. Hong, X. Fu, and N. D. Sidiropoulos, "Learning to optimize: Training deep neural networks for interference management," *IEEE Trans. Signal Process.*, vol. 66, pp. 5438 – 5453, Oct. 2018.

[11] F. Liang, C. Shen, W. Yu, and F. Wu, "Towards optimal power control via ensembling deep neural networks," *arXiv preprint arXiv:1807.10025*, 2018.

[12] W. Lee, M. Kim, and D.-H. Cho, "Deep power control: Transmit power control scheme based on convolutional neural network," *IEEE Commun. Lett.*, vol. 22, pp. 1276–1279, Apr. 2018.

[13] Y. S. Nasir and D. Guo, "Deep reinforcement learning for distributed dynamic power allocation in wireless networks," *arXiv preprint arXiv:1808.00490*, 2018.

[14] W. Cui, K. Shen, and W. Yu, "Spatial deep learning for wireless scheduling," *arXiv preprint arXiv:1808.01486*, 2018.

[15] M. Eisen, C. Zhang, L. F. Chamon, D. D. Lee, and A. Ribeiro, "Learning optimal resource allocations in wireless systems," *arXiv preprint arXiv:1807.08088*, 2018.

[16] M. Hong and Z.-Q. Luo, "Signal processing and optimal resource allocation for the interference channel," in *Academic Press Library Signal Process.*, vol. 2, pp. 409–469, Elsevier, 2014.

[17] J. G. Proakis and M. Salehi, *Digital communications*, vol. 4. McGraw-hill New York, 2001.

[18] Q. Shi, M. Razaviyayn, Z. Luo, and C. He, "An iteratively weighted MMSE approach to distributed sum-utility maximization for a MIMO interfering broadcast channel," *IEEE Trans. Signal Process.*, vol. 59, pp. 4331–4340, Sept. 2011.

[19] Y. Shi, J. Zhang, B. O'Donoghue, and K. B. Letaief, "Large-scale convex optimization for dense wireless cooperative networks," *IEEE Trans. Signal Process.*, vol. 63, pp. 4729–4743, Sept. 2015.

[20] O. Vinyals, M. Fortunato, and N. Jaitly, "Pointer networks," in *Proc. Adv. Neural Inform. Process. Syst.*, pp. 2692–2700, Dec. 2015.

[21] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio, "Neural combinatorial optimization with reinforcement learning," in *Proc. Int. Conf. Learning Representation*, Apr. 2017.

[22] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*. MIT press Cambridge, 2016.

[23] H. He, H. Daume III, and J. M. Eisner, "Learning to search in branch and bound algorithms," in *Proc. Adv. Neural Inform. Process. Syst.*, pp. 3293–3301, Dec. 2014.

[24] E. B. Khalil, P. Le Bodic, L. Song, G. Nemhauser, and B. Dilkina, "Learning to branch in mixed integer programming," in *AAAI Conf. Artificial Intell.*, 2016.

[25] E. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song, "Learning combinatorial optimization algorithms over graphs," in *Proc. Adv. Neural Inform. Process. Syst.*, pp. 6348–6358, Dec. 2017.

[26] J. Song, R. Lanka, A. Zhao, Y. Yue, and M. Ono, "Learning to search via self-imitation with application to risk-aware planning," in *Proc. Adv. Neural Inform. Process. Syst. Workshop*, Dec. 2017.

[27] M. Conforti, G. Cornuéjols, and G. Zambelli, *Integer Programming (Graduate texts in mathematics)*, vol. 271. Springer Heidelberg, 2014.

[28] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proc. Int. Conf. Artificial Intell. Stat.*, pp. 627–635, Apr. 2011.

[29] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters, et al., "An algorithmic perspective on imitation learning," *Found. Trends Robotics*, vol. 7, pp. 1–179, Mar. 2018.

[30] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al., "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, p. 484, Jan. 2016.

[31] M.-F. Balcan, T. Dick, T. Sandholm, and E. Vitercik, "Learning to branch," in *Proc. Int. Conf. Mach. Learning*, vol. 80, pp. 344–353, Jul. 2018.

[32] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd, "Conic optimization via operator splitting and homogeneous self-dual embedding," *J. Optimization Theory Appl.*, vol. 169, pp. 1042–1068, Feb. 2016.

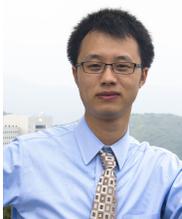
[33] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.

[34] V. Francois-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau, "An introduction to deep reinforcement learning," *Found. Trends Mach. Learning*, vol. 11, no. 3–4, 2018.

[35] Y. Cheng, M. Pesavento, and A. Philipp, "Joint network optimization and downlink beamforming for CoMP transmissions using mixed integer conic programming," *IEEE Trans. Signal Process.*, vol. 61, pp. 3972–3987, May 2013.



Yifei Shen (S'18) is currently working towards the Ph.D. degree in the Department of Electronic and Computer Engineering at the Hong Kong University of Science and Technology (HKUST).



Yuanming Shi (S'13-M'15) received the B.S. degree in electronic engineering from Tsinghua University, Beijing, China, in 2011. He received the Ph.D. degree in electronic and computer engineering from The Hong Kong University of Science and Technology (HKUST), in 2015. Since September 2015, he has been with the School of Information Science and Technology in ShanghaiTech University, where he is currently a tenured Associate Professor. He visited University of California, Berkeley, CA, USA, from October 2016 to February 2017. Dr.

Shi is a recipient of the 2016 IEEE Marconi Prize Paper Award in Wireless Communications, and the 2016 Young Author Best Paper Award by the IEEE Signal Processing Society. His research areas include optimization, statistics, machine learning, signal processing, and their applications to 6G, IoT and AI.



Jun Zhang (S'06-M'10-SM'15) received the B.Eng. degree in Electronic Engineering from the University of Science and Technology of China in 2004, the M.Phil. degree in Information Engineering from the Chinese University of Hong Kong in 2006, and the Ph.D. degree in Electrical and Computer Engineering from the University of Texas at Austin in 2009. He is an Assistant Professor in the Department of Electronic and Information Engineering at the Hong Kong Polytechnic University (PolyU). His research interests include wireless communications and network-

working, mobile edge computing and edge learning, distributed learning and optimization, and big data analytics.

Dr. Zhang co-authored the books *Fundamentals of LTE* (Prentice-Hall, 2010), and *Stochastic Geometry Analysis of Multi-Antenna Wireless Networks* (Springer, 2019). He is a co-recipient of the 2019 IEEE Communications Society & Information Theory Society Joint Paper Award, the 2016 Marconi Prize Paper Award in Wireless Communications (the best paper award of IEEE Transactions on Wireless Communications), and the 2014 Best Paper Award for the EURASIP Journal on Advances in Signal Processing. Two papers he co-authored received the Young Author Best Paper Award of the IEEE Signal Processing Society in 2016 and 2018, respectively. He also received the 2016 IEEE ComSoc Asia-Pacific Best Young Researcher Award. He is an Editor of IEEE Transactions on Wireless Communications, IEEE Transactions on Communications, and Journal of Communications and Information Networks.



Khaled B. Letaief (S'85–M'86–SM'97–F'03) received the BS degree with distinction in Electrical Engineering from Purdue University at West Lafayette, Indiana, USA, in December 1984. He received the MS and Ph.D. Degrees in Electrical Engineering from Purdue University, in Aug. 1986, and May 1990, respectively.

He is an internationally recognized leader in wireless communications and networks with research interest in wireless communications, artificial intelligence, big data analytics systems, mobile edge computing, 5G systems and beyond. In these areas, he has over 620 journal and conference papers and given keynote talks as well as courses all over the world. He also has 15 patents, including 11 US patents.

He is well recognized for his dedicated service to professional societies and in particular IEEE where he has served in many leadership positions, including President of IEEE Communications Society, the world's leading organization for communications professionals with headquarters in New York and members in 162 countries. He is also the founding Editor-in-Chief of the prestigious IEEE Transactions on Wireless Communications and served on the editorial board of other premier journals including the IEEE Journal on Selected Areas in Communications – Wireless Series (as Editor-in-Chief).

He has been a dedicated teacher committed to excellence in teaching and scholarship with many teaching awards, including the Michael G. Gale Medal for Distinguished Teaching.

He is the recipient of many other distinguished awards including the 2019 Distinguished Research Excellence Award by HKUST School of Engineering (Highest research award and only one recipient/3 years is honored for his/her contributions); 2019 IEEE Communications Society and Information Theory Society Joint Paper Award; 2018 IEEE Signal Processing Society Young Author Best Paper Award; 2017 IEEE Cognitive Networks Technical Committee Publication Award; 2016 IEEE Signal Processing Society Young Author Best Paper Award; 2016 IEEE Marconi Prize Paper Award in Wireless Communications; 2011 IEEE Wireless Communications Technical Committee Recognition Award; 2011 IEEE Communications Society Harold Sobol Award; 2010 Purdue University Outstanding Electrical and Computer Engineer Award; 2009 IEEE Marconi Prize Award in Wireless Communications; 2007 IEEE Communications Society Joseph LoCicero Publications Exemplary Award; and over 15 IEEE Best Paper Awards. Since 1993, he has been with HKUST where he has held many administrative positions, including the Head of the Electronic and Computer Engineering department. He also served as Chair Professor and HKUST Dean of Engineering. Under his leadership, the School of Engineering dazzled in international rankings (rising from # 26 in 2009 to # 14 in the world in 2015 according to QS World University Rankings.)

He is a Fellow of IEEE and a Fellow of HKIE. He is also recognized by Thomson Reuters as an ISI Highly Cited Researcher.