

Learning chordal extensions

D. Liu,
A. Lodi, M. Tanneau

G-2019-78

October 2019

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée : D. Liu, A. Lodi, M. Tanneau (Octobre 2019). Learning chordal extensions, Rapport technique, Les Cahiers du GERAD G-2019-78, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2019-78>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: D. Liu, A. Lodi, M. Tanneau (October 2019). Learning chordal extensions, Technical report, Les Cahiers du GERAD G-2019-78, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2019-78>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2019
– Bibliothèque et Archives Canada, 2019

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2019
– Library and Archives Canada, 2019

Learning chordal extensions

Defeng Liu ^{a, c}

Andrea Lodi ^{a, b, c}

Mathieu Tanneau ^{a, b, c}

^a Canada Excellence Research Chair in Data Science for Real-time Decision-making, Polytechnique Montréal, Montréal (Québec) Canada, H3C 3A7

^b GERAD, Montréal (Québec), Canada, H3T 2A7

^c Department of Mathematics and Industrial Engineering, Polytechnique Montréal, Montréal (Québec) Canada, H3C 3A7

defeng.liu@polymtl.ca

andrea.lodi@polymtl.ca

mathieu.tanneau@polymtl.ca

October 2019
Les Cahiers du GERAD
G–2019–78

Copyright © 2019 GERAD, Liu, Lodi, Tanneau

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract: A highly influential ingredient of many techniques designed to exploit sparsity in numerical optimization is the so-called chordal extension of a graph representation of the optimization problem. The definitive relation between chordal extension and the performance of the optimization algorithm that uses the extension is not a mathematically understood task.

For this reason, we follow the current research trend of looking at Combinatorial Optimization tasks by using a Machine Learning lens, and we devise a framework for learning elimination rules yielding high-quality chordal extensions. As a first building block of the learning framework, we propose an on-policy imitation learning scheme that mimics the elimination ordering provided by the (classical) minimum degree rule.

The results show that our on-policy imitation learning approach is effective in learning the minimum degree policy and, consequently, produces graphs with desirable fill-in characteristics.

1 Introduction

A simple undirected graph $G = (V, E)$ is *chordal* if, for every cycle c of length at least four, there exists an edge $e \in E$ that connects two non-consecutive vertices of c . A *chordal extension* of a graph G is a chordal graph H such that G is a sub-graph of H , i.e., one can obtain H by adding edges to G . A practical way of constructing chordal extensions is via *graph elimination* [29], which consists in sequentially eliminating the nodes of the graph. At each step, a node v is selected, new edges are inserted so as to make the neighbors of v into a clique, then v is removed (i.e., eliminated). This process is repeated until all nodes have been eliminated, and one obtains a chordal extension by adding to the original graph all edges that were inserted in the process. The order in which nodes were eliminated is thereby called an *elimination ordering*.

This work focuses on the role of chordal extensions and graph elimination within optimization frameworks. Indeed, there is a direct connection between chordal extensions, which are typically computed via graph elimination, and a number of classical sparsity-exploiting techniques [30, 29]. In particular, we seek to devise a framework for *learning* elimination rules that yield high-quality chordal extensions, as we illustrate below.

Our first motivating example is the computation of a fill-reducing ordering for sparse Cholesky factorization, a process that reduces to computing an elimination ordering [8, 29]. Crucially, Cholesky factorization underlies most implementations of interior-point algorithms for linear programming [30, 5], and the choice of ordering can have a major impact on the method’s performance [26]. Similarly, chordal graphs form the basis of *chordal decomposition* techniques to exploit sparsity in semi-definite optimization problems, see, e.g., [1, 20]. Specifically, a single, dense, semi-definite (SDP) constraint, can be decomposed into several smaller, yet coupled, semi-definite constraints. This reformulation also reduces to computing a chordal extension, and can dramatically improve the performance of both interior-point and first-order methods on large problems [32, 20]. More generally, a similar approach can be leveraged in linear conic optimization and convex optimization, see, e.g., [29].

Historically, efforts have focused on computing minimum chordal extensions, i.e., chordal extensions with a minimum number of additional edges [26, 4], which has been proven to be NP-complete [31]. This fostered the development of fast and efficient heuristics such as minimum degree [11] and nested dissection [10] orderings. State-of-the-art implementations of these methods are routinely used in most optimization and sparse linear algebra software, where they tackle problems with up to millions of variables. Nevertheless, which chordal extension is computed can significantly impact the subsequent performance of the optimization algorithm. Therefore, it is natural to seek a “best” chordal extension, i.e., one that maximizes performance.

Recently, the use of *Machine Learning* (ML) in *Combinatorial Optimization* (CO) became a popular research area with quite a number of contributions investigating many angles of such a connection. On the one hand, some research has been devoted to solve CO problems by ML, i.e., to devise new heuristic algorithms that perform the end-to-end learning of the solution of a CO problem. On the other hand, ML has been used to tackle some tasks within CO algorithms and software for which modern statistical learning has chances to improve the current performances, either because the known way of performing those tasks is computationally heavy or because they are poorly understood from the mathematical standpoint. The interested reader is referred to [3] for a methodological survey on this new research area.

Our work does not follow the first direction outlined in the previous paragraph. Indeed, the goal of this paper is *not* to compete with existing state-of-the-art heuristics for graph elimination. First, these heuristics leverage decades of development and clever engineering, and have been optimized for fast runtime and good solution quality. Second, and more importantly, the aim of our work is to eventually gain deeper understanding of the relation between the characteristics of a chordal extension and the behavior of optimization algorithms and software, a topic of interest in its own right and whose mathematical knowledge is currently insufficient. In order to achieve this ultimate goal, our present contributions are 1) to propose a mathematical framework for the problem of learning elimination

orderings, and 2) to provide methodological and practical insights on the learning process itself. As a byproduct, a better understanding of what “good” chordal extensions look like (for a specific task) can lead to an easier and more direct customization of elimination orderings to sets of similar graphs.

The rest of the paper is organized as follows. In Section 2, we introduce some relevant definitions and concepts. In Section 3, we present our methodology for learning chordal extensions. In Section 4, we report on numerical experiments. Section 5 gives further discussion and Section 6 concludes the paper.

2 Basic notations and concepts

In this section, we review some notations and concepts used in the remainder of the paper. Section 2.1 introduces basic notations and definitions of graphs and Section 2.2 introduces a commonly used model for sequential decision problems. In Sections 2.3 and 2.4, we briefly go through some Machine Learning concepts, in order to help the reader be familiar with relevant ML methods and properly locate the methodology we propose in Section 3.

In all that follows, for an arbitrary set S , we use the notation $\mathfrak{P}(S)$ to denote the set of all probability distributions over S .

2.1 Graph-theoretic notations

In this paper, all considered graphs are simple, undirected graphs. A graph is denoted by $G = (V, E)$, where V (resp. E) denotes the set of its nodes (resp. its edges). For an edge $e = (v, w)$, we say that e is *incident* to v and w , that v, w are the *extremities* of e , and that v, w are *adjacent*. The neighborhood of v , denoted by $\mathcal{N}_G(v)$, is defined as the set of nodes that are adjacent to v .

The degree of node v in G is denoted by $\delta_G(v)$. We say that node v is of *minimum degree* if $\delta_G(v) \leq \delta_G(w)$ holds for all nodes $w \in V$. Note that several nodes of minimum degree may exist.

In what follows, we will drop the subscript G whenever context is sufficiently clear, and write for example $\delta(v)$ rather than $\delta_G(v)$.

2.2 Markov Decision Processes

Markov Decision Processes (MDPs) [16] are widely used to formulate sequential decision problems. An MDP is characterized by a set of possible *states* \mathcal{S} , a set of possible actions \mathcal{A} , the dynamics of the system, and a cost function.

Formally, an MDP is defined by a tuple $(\mathcal{S}, \mathcal{A}, P, c)$, where P encodes the system’s dynamics, and c encodes the cost function. For any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, $P(s, a)$ is a probability distribution over the state space \mathcal{S} . That is, if one takes action a in state s , the next observed state s' will be sampled from the distribution $P(s, a) \in \mathfrak{P}(\mathcal{S})$. Finally, the cost of executing action a in state s is denoted by $c(s, a)$.

Given an MDP $(\mathcal{S}, \mathcal{A}, P, c)$, a *policy* is a decision rule for selecting an action a , given a current state s . Specifically, a policy π is a function

$$\pi : \mathcal{S} \mapsto \mathfrak{P}(\mathcal{A}) \tag{1}$$

$$s \mapsto \pi(s), \tag{2}$$

and the next action a is sampled from $\pi(s) \in \mathfrak{P}(\mathcal{A})$. If a policy π maps each state to a *single* action, i.e., if $\pi(s)$ is a degenerate distribution for every $s \in \mathcal{S}$, then π is called a *deterministic policy*; otherwise it is called a *stochastic policy*. For the remainder of this paper, we will only consider stochastic policies, and we define the expected *immediate cost* for policy π in state $s \in \mathcal{S}$ as

$$C_\pi(s) = \mathbb{E}_{a \sim \pi(s)} [c(s, a)]. \tag{3}$$

A *trajectory* is a sequence of state-action pairs $((s_0, a_0), (s_1, a_1), \dots)$ where s_{t+1} is sampled from $P(s_t, a_t)$. In this paper, all trajectories will always be finite, although they may be of arbitrary length. We say that a trajectory is sampled from a policy π if each action a_t is sampled from $\pi(s_t)$. The total cost along a trajectory is then given by

$$\sum_{t \geq 0} c(s_t, a_t), \quad (4)$$

which is always finite since we only consider finite trajectories.

Finally, for a distribution of initial states $\mathcal{D}_0 \in \mathfrak{P}(\mathcal{S})$, we define the *expected total cost* of a policy π as

$$C_\pi^{tot} = \mathbb{E}_{s_0 \sim \mathcal{D}_0} \left(\mathbb{E}_\pi \left[\sum_{t \geq 0} c(s, a) \right] \right), \quad (5)$$

where \mathbb{E}_π denotes that trajectories are sampled from π . Furthermore, \mathcal{D}_0 and π induce a stationary distribution over states, which we denote by \mathcal{D}_π . Thus, we define the *expected average cost* of policy π as

$$\begin{aligned} C_\pi^{avg} &= \mathbb{E}_{s \sim \mathcal{D}_\pi} [\mathbb{E}_{a \sim \pi(s)} (c(s, a))] \\ &= \mathbb{E}_{s \sim \mathcal{D}_\pi} [C_\pi(s)]. \end{aligned} \quad (6)$$

In this work, we assume that the above two expectations are always finite; we make this mild technical assumption to ensure that the learning problems defined in Section 2.4 are well-defined. For ease of reading, we also drop the explicit dependency of C_π^{tot} and C_π^{avg} on \mathcal{D}_0 , since the latter will always be evident from the context.

2.3 Standard statistical learning

In statistical learning, the goal is to detect patterns from a set of observed data and make predictions about future data. We can formalize the standard statistical learning problem as follows. Given a variable space \mathcal{Z} and a set of examples $\mathcal{D}_\mathcal{Z} = \{z_1, z_2, \dots, z_n\}$ from the unknown distribution $\mathcal{P}(\mathcal{Z})$, the task is to find a function f over a family of functions \mathcal{F} , such that f “performs well” on $\mathcal{P}(\mathcal{Z})$. It is assumed that all the observed examples are drawn independent and identically distributed (i.i.d.) from the same distribution $\mathcal{P}(\mathcal{Z})$.

If a *loss* function $\mathcal{L} : \mathcal{F} \times \mathcal{Z} \mapsto \mathbb{R}$ is specified to measure the performance of f , then the goal can be described as finding $\hat{f} \in \mathcal{F}$ that minimizes the *expected loss* with respect to $\mathcal{P}(\mathcal{Z})$, i.e.,

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \mathbb{E}_{z \sim \mathcal{P}(\mathcal{Z})} [\mathcal{L}(f, z)]. \quad (8)$$

The learned function \hat{f} is then used to predict future data.

However, the expected loss cannot be computed exactly due to the fact that $\mathcal{P}(\mathcal{Z})$ is unknown. In practice, if a subset of examples $\mathcal{D}_\mathcal{Z}$ sampled from $\mathcal{P}(\mathcal{Z})$ is available, a number of learning methods turn to minimize the *empirical loss* on $\mathcal{D}_\mathcal{Z}$. Then, \hat{f} is obtained by solving

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \mathcal{L}(f, z_i). \quad (9)$$

The differences in forms and contents of \mathcal{Z} , \mathcal{F} , \mathcal{L} result in different learning tasks. Here, we only introduce *supervised learning* that is the relevant task for the current state of our work.

Supervised learning. In supervised learning, the variable space \mathcal{Z} consists of $\mathcal{X} \times \mathcal{Y}$, where \mathcal{X} is the space of input variables and \mathcal{Y} is the space of output variables. The family of functions \mathcal{F} is a set of mappings $f : \mathcal{X} \mapsto \mathcal{Y}$. For any sample $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the loss function \mathcal{L} measures the discrepancy between $f(x)$ and y . Ideally, the output \mathcal{Y} can be in any form or intent. However, most tasks assume that \mathcal{Y} is categorical or nominal. The former characterizes the task as classification, whereas the latter induces regression.

2.4 Imitation learning for sequential decision problems

Imitation learning (IL) [2, 23, 25, 27] is an extension of supervised learning from problems satisfying i.i.d. assumption to sequential decision problems, see, e.g., [22, 24, 28].

In IL, the target policy learns its decision rule from an expert policy. More precisely, given the class of candidate policies Π , we seek to find a target policy $\pi \in \Pi$ that matches the expert policy π^* . The target and expert policy are often referred as the *learner* and the *expert*. The cost is defined by a loss function $\mathcal{L}(\pi, \pi^*)$, a measure of discrepancy between π and π^* . If a *behavior policy* π' is set to generate trajectories of states, the goal is to find a policy $\hat{\pi}$ that minimizes the expected loss with respect to the distribution of states induced by π' , namely

$$\hat{\pi} = \operatorname{argmin}_{\pi \in \Pi} \mathbb{E}_{s \sim \mathcal{D}_{\pi'}} [\mathcal{L}(\pi(s), \pi^*(s))]. \quad (10)$$

In the literature, depending on the choice of behavior policy, IL algorithms can be divided into two classes: *off-policy* and *on-policy* methods. The off-policy methods are the ones with behavior policy π' independent of the learner π . For instance, the supervised approach for imitation learning falls into this class. On the other hand, if the behavior policy depends on the learner, the algorithm becomes an on-policy method. Data Aggregation (DAGGER)[25] and its variants are examples of on-policy methods.

Supervised imitation learning. The supervised approach for IL fixes the behavior policy as the expert, i.e., $\pi' = \pi^*$. Then the learner is trained under the distribution induced by the expert, given by

$$\hat{\pi} = \operatorname{argmin}_{\pi \in \Pi} \mathbb{E}_{s \sim \mathcal{D}_{\pi^*}} [\mathcal{L}(\pi(s), \pi^*(s))]. \quad (11)$$

Data Aggregation. DAGGER is an iterative algorithm that improves the learner by executing a mixed behavior policy combined with the learner and the expert. At each iteration, the collected data will be aggregated into an accumulated dataset and the learner will be trained by all the data collected from previous iterations.

3 Methodology

In this section, we present our methodology for learning chordal extensions. In this work, we focus on how to learn elimination rules for graph elimination with imitation learning. More precisely, we propose an on-policy imitation learning scheme that mimics the elimination ordering provided by the ordering heuristic we choose.

3.1 MDP formulation

We begin by formulating graph elimination as a Markov decision process. First, the state space \mathcal{S} is the set of simple undirected graphs. Then, for a given graph $G = (V, E)$, the corresponding set of possible actions is identified by the nodes of the graph. Transitions are deterministic: if an action $a = v$ is performed in state G , i.e., if node v is eliminated from graph G , then the new state is uniquely defined as the graph obtained from the elimination of node v . Note that the number of nodes decreases

by one at each step. Hence, even though the initial graph may be of arbitrary size, trajectories are always finite.

Thus, a policy π maps a graph to a probability distribution over its set of nodes V . Therefore, if $V = \{1, \dots, n\}$, then $\pi(G)$ is a n -dimensional non-negative vector, whose i -th coordinate denotes the probability that node i be eliminated.

Finally, one may select a cost function according to the problem at hand, for example, the number of additional edges, i.e., fill-in. In that case, finding a policy that minimizes the expected total cost reduces to finding a policy that yields minimum chordal extensions. Rather than trying to minimize fill-in, which is an NP-hard problem for which efficient heuristics already exists, we adopt a more generic imitation learning scheme as discussed below.

3.2 On-policy imitation learning

Although the goal of imitation learning is to find a learner that best matches the expert, any parameterized stochastic policy will inevitably have chances to make occasional mistakes by choosing an action different from the expert. In the supervised imitation learning approach, where the learner is only trained under the distribution of states induced by the expert, the learner may not be able to correct its behavior from deviations induced by its bad choice of action.

Moreover, any possible state is observable in our problem setting and the expert is always accessible for querying any state. Therefore, we choose on-policy approach in order to be more robust to alleviate the deviation problem. Specifically, we set the learner as the behavior policy, i.e., $\pi' = \pi$. As a result, the goal is to find a policy $\hat{\pi}$ such that

$$\hat{\pi} = \operatorname{argmin}_{\pi \in \Pi} \mathbb{E}_{s \sim \mathcal{D}_\pi} [\mathcal{L}(\pi(s), \pi^*(s))]. \quad (12)$$

Since, in general, the expected loss cannot be computed analytically, we estimate it by sampling trajectories of states from finite dataset using π . Given a dataset \mathcal{G} of M graphs, we train the learner with our practical algorithm, namely *one-step on-policy imitation learning*, as described in Algorithm 1.

Algorithm 1: One-step on-policy imitation learning

```

Input: Instance Dataset  $\mathcal{G} = \{\mathbf{g}_i\}_{i=1}^M$ 
Initialize  $\pi_\theta$  to any policy in  $\Pi$ ;
for  $i = 1$  to  $N$  do
  for each  $g \in \mathcal{G}$  do
    Initialize  $s$ ;
    for  $j = 1$  to  $T$  do
       $\theta \leftarrow \theta - \alpha \nabla_\theta \mathcal{L}(\pi_\theta(s), \pi^*(s))$ ;
       $a \leftarrow \pi_\theta(s)$ ;
      Take action  $a$ , observe next state  $s$ ;
    end
  end
end
return  $\pi_\theta$ 

```

The training proceeds as follows. At every epoch i (from 1 to N), each instance $G = (V, E) \in \mathcal{G}$ is used to generate one complete trajectory (with length $T = |V|$), and each state (i.e., graph) in this trajectory represents a training data point. The one-step updating is as follows. At every transition, i.e., at every elimination step, the learner π_θ is updated by taking a gradient step with respect to the loss $\mathcal{L}(\pi_\theta(s), \pi^*(s))$, where s is the current state. An action a is then sampled from the updated learner π_θ , which yields the next state s' . Note that, in this setting, the loss is computed by comparing the two distributions $\pi_\theta(s)$ and $\pi^*(s)$ directly. We do so because 1) we know analytically both the learner and the expert, and 2) it allows to exploit information from the entire distributions rather than sampling and comparing individual actions.

Expert. We consider the minimum degree heuristic [21, 11] as the expert. At each step, the minimum degree selects a node of minimum degree to be eliminated. Ties are broken arbitrarily, i.e., if several nodes have minimum degree, then one is selected uniformly at random among them. Let us note that today’s implementations include several additional features, such as smarter tie breaking or the simultaneous elimination of multiple nodes.

In all that follows, we denote π_{MD} the minimum-degree policy, i.e., for a given graph $G = (V, E)$, we have

$$\pi_{MD}(G)[v] = \begin{cases} \frac{1}{k} & \text{if } v \text{ is of minimum degree} \\ 0 & \text{otherwise} \end{cases},$$

where k is the number of nodes that have minimum degree.

Learner parameterization. Given that states are represented as graphs, with arbitrary size and topology, we propose to use [graph neural networks \(GNN\)](#) [13, 15] to parameterize the learner. Indeed, [GNN](#) is an expressive type of model to process graph-structured data and have been applied to a variety of representation learning tasks on graphs [7, 17, 14, 19, 9]. In [GNN](#) models, a graph is embedded according to the features of vertices and its topological structure. The encoding of a node is generated by propagating information from its neighborhood. One of the most appealing properties of [GNN](#) is that it is size-and-order invariant to input data, i.e., it can process graphs of arbitrary size, and the ordering of the input elements is irrelevant.

In our [GNN](#) architecture, the embedding of the graph in the $(l + 1)$ -th layer is computed by aggregating, for each node, the features of its neighborhood from the l -th layer, i.e.,

$$\mathbf{H}^{l+1} = f\left(\mathbf{A}\mathbf{H}^l\mathbf{W}^l + \mathbf{I}_{n \times 1}\mathbf{B}^l\right), \quad (13)$$

where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the adjacency matrix of the graph, $\mathbf{I}_{n \times 1}$ is the matrix of ones with size of $n \times 1$, d^l and d^{l+1} are the dimension of the features in layer l and $(l + 1)$, $\mathbf{H}^{l+1} \in \mathbb{R}^{n \times d^{l+1}}$ and $\mathbf{H}^l \in \mathbb{R}^{n \times d^l}$ are the embeddings of layer $(l + 1)$ and l , $\mathbf{W}^l \in \mathbb{R}^{d^l \times d^{l+1}}$ and $\mathbf{B}^l \in \mathbb{R}^{1 \times d^{l+1}}$ are the parameters in layer l , and $f(\cdot)$ specifies the activation function. For $f(\cdot)$, we apply *Softmax* in the output layer and *Relu* in the rest.

Loss function. Since the expert is the minimum degree heuristic, we can compute the exact $\pi_{MD}(s)$ given a state s of a graph as shown before. To measure the distance between two distributions given by the learner and the expert respectively, we compute the Kullback-Liebler (KL) divergence [18] between $\pi_{MD}(s)$ and $\pi(s)$ as the loss.

4 Numerical experiments

In this section, we report the details of our computational investigation. More precisely, Section 4.1 specifies the data generation and collection. In Section 4.2, we discuss the experimental setting and, finally, Section 4.3 reports the computational results.

4.1 Data collection

We evaluate our approach on four different datasets, which comprise graphs that vary in size and structural characteristics.

4.1.1 Erdos-Renyi graphs

We first build two datasets of Erdos-Renyi graphs, a simple and well-known class of random graphs. We use the notation $G(n, p)$ to denote a (random) Erdos-Renyi graph with n nodes, and such that edges are selected with probability $p \in [0, 1]$ independently of each other. Note that, for given n and p , $G(n, p)$ is a random variable whose realizations are graphs of size n . While n controls the size of the graph, p controls its sparsity.

We form two datasets of Erdos-Renyi graphs: one of smaller graphs, denoted ER_S , and the other one of larger graphs, denoted ER_L .

Each graph in ER_S , is sampled from $G(n, p)$, where n is drawn uniformly between 100 and 300, and p is sampled between 0.1 and 0.3 with uniform probability. This is done to introduce some variability in size and density in the dataset. Overall, ER_S contains 600 graphs. We follow the same methodology for ER_L , except that n is drawn uniformly between 300 and 500. Overall, ER_L contains 200 graphs.

4.1.2 SuiteSparse matrix collection

The SuiteSparse matrix collection¹ [6] is a dataset of (sparse) matrices collected from a number of real-life applications, and is routinely used as benchmark for numerical linear algebra software. Given a matrix M , we construct a non-oriented graph whose adjacency matrix corresponds exactly to the sparsity structure of M . We only consider square matrices, and any non-symmetric matrix is transferred into symmetric by adding its transpose to it.

First, we select square matrices of size between 50 and 500. This yields a dataset of 278 graphs, which we denote by SS_S . Similarly, we select square matrices of size between 1000 and 2000, and obtain a second dataset, denoted by SS_L , which contains 295 graphs.

4.2 Experimental settings

Our experiments were conducted on a dual Intel Xeon Gold 6126@2.60GHz, 768BG RAM machine running Linux and equipped with Nvidia Tesla V100 GPUs. Our code² is written in Python 3.6, and we use Pytorch 0.4 for modeling and training GNNs.

Datasets. We split the ER_S dataset into $\{training, validation, test\}$, each containing 200 graphs. Our GNN policy is trained and validated only with the *training* and *validation* set of ER_S , respectively. Then, we test the generalization performance of the trained model with the *test* set of ER_S , ER_L , SS_S and SS_L .

GNN setting. We apply the GNN architecture described in Section 3.2 with 2 layers. As initially the vertices of the graphs have no attribute, we initialize the feature of each vertex with the same value. Specifically, we take $h_v^0 = 1, \forall v \in V$. As a result, the encoding of each vertex only depends on the topological structure of its neighborhoods. The dimension of features in all layers is the same. For each layer, the weights are initialized from Xavier normal distribution [12] and we initialize the bias with zero.

Performance metrics. For a finite dataset \mathcal{G} , the first metric computes the average KL loss, given by

$$\hat{\mathcal{L}}_{KL} = \frac{1}{\sum_{g \in \mathcal{G}} n_g} \sum_{g \in \mathcal{G}} \sum_{i=1}^{n_g} \mathcal{L}_{KL}(\pi_{\theta}(s_i), \pi_{MD}(s_i)), \quad (14)$$

where n_g is the size of each graph $g \in \mathcal{G}$ and $\mathcal{L}_{KL}(\cdot)$ specifies the KL divergence between $\pi_{\theta}(s_i)$ and $\pi_{MD}(s_i)$ in state s_i of g .

To measure the fill-in of a policy, the second metric computes the average number of fill-in per graph. For each graph $g \in \mathcal{G}$, we denote the total number of fill-in by c_{fillin}^g . Then, the average fill-in per graph is given by

$$\hat{C}_{fillin} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} c_{fillin}^g. \quad (15)$$

¹SuiteSparse matrix collection was formerly known as the University of Florida sparse matrix collection.

²The code repository as well as the instances are in the process of being made publicly available.

Training and validation. We train our **GNN** policy with Algorithm 1. At each epoch, we randomly shift the training set and sample single-graph mini batches. For learning rate tuning, we experiment different learning rates from 10^{-5} to 10^{-3} . The validation result is shown by plotting the average KL loss and the average fill-in per graph in Figure 1. Observing that 10^{-4} yields fast and smooth convergence, we train the model with the learning rate of 10^{-4} for 20 epochs. Moreover, we also observe a plateau effect for larger step size in Figure 1, notably, sudden decrease with larger step size. This effect will be discussed in Section 5.

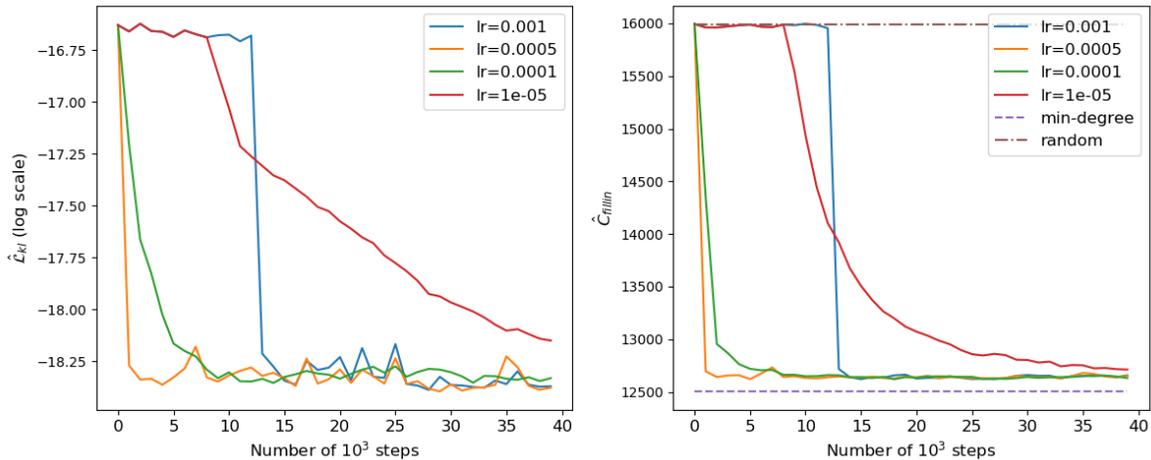


Figure 1: Validation results of imitation learning. We plot the average KL loss in *log* scale (left) and the average fill-in per graph (right) on the validation set of ER_S . For fill-in, we compare **GNN** with minimum degree and random policy.

Test. We test the generalization performance of trained **GNN** with four test sets as specified before. Additionally, to evaluate the performance of **GNN** models at different stages of training, we first save the trained model at the end of each epoch. Then, we test the performance of each saved model on four test sets.

4.3 Results

In this section, we compare the predictive performance of **GNN** with the two metrics introduced in the previous section. The results on the training set and four test sets are shown in Figure 2. Specifically, we plot the curves of two metrics over the entire training period (20 epochs), in order to compare the performance of **GNN** models at different stages of training.

From the results of the training set (shown in the first row of Figure 2), we observe that the loss significantly decreases and stabilizes after about 15 epochs of training. Moreover, for fill-in, the **GNN** also matches the minimum degree heuristic.

Comparing the results of different test sets, we observe that our **GNN** generalizes well, both to larger size graphs and to different distributions. First of all, loss curves of all datasets show same decreasing tendency over the entire training period, although the magnitude of values can be different across datasets. In terms of fill-in, we have similar and consistent results. Moreover, by comparing the curves in each row, we also observe a strong correlation between KL loss (i.e., how good we replicate the minimum degree expert), and the actual fill-in (which is only observed, we never learn anything from it).

It is worth noting that, although the initial graphs can be i.i.d., the other states in the trajectory always depend on previous states and actions, which indicates the induced distribution of states depends on the behavior policy itself.

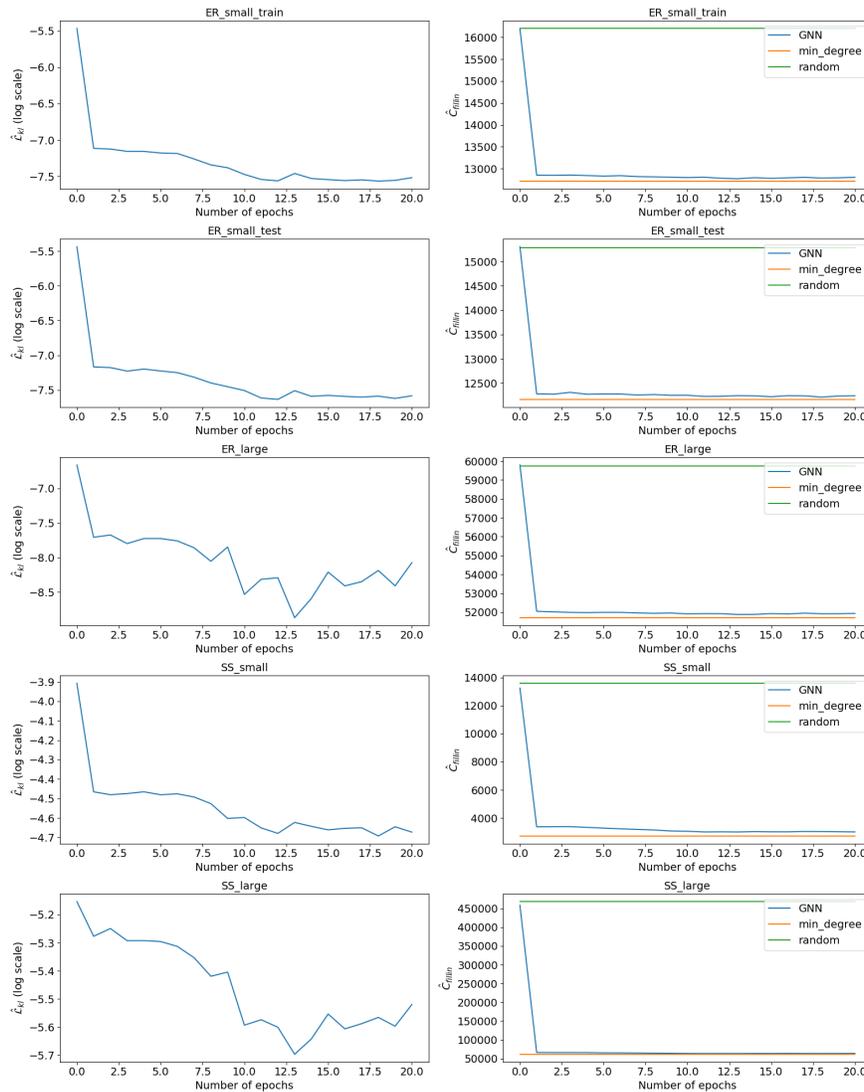


Figure 2: Test results of imitation learning. We plot the average KL loss in *log scale* (left) and the average fill-in per graph (right) on the training set and four test sets. For fill-in, we compare our **GNN** with minimum degree and random policy.

Since we use the on-policy imitation learning approach, the loss is measured under the distribution of states induced by our **GNN** policy. As **GNN** model changes during training, this metric is actually measured under different distributions. As a result, the decrease of loss over the training period (shown on the left side of Figure 2) only shows a tendency that the **GNN** replicates the minimum degree expert better on an evolutionary distribution induced by itself. The predictive performance of the **GNN** still needs to be validated by the actual fill-in, which is precisely done on the right side of Figure 2.

5 Further discussion

We now seek to further explain the sharp drops in loss that were observed during training, e.g., in Figure 1, and the stark correlation between imitation loss and fill-in.

To do so, we consider the following **GNN** with two layers:

$$x_i^0 = 1 \quad \forall i \in V, \tag{16}$$

$$h_i^1 = \sum_{j \in \mathcal{N}(i)} w_1 x_j^0 \quad \forall i \in V, \quad (17)$$

$$x_i^1 = \text{Relu}(1 + h_i^1) \quad \forall i \in V, \quad (18)$$

$$h_i^2 = \sum_{j \in \mathcal{N}(i)} w_2 x_j^1 \quad \forall i \in V, \quad (19)$$

$$x^2 = \text{Softmax}(h^2), \quad (20)$$

where $w_1, w_2 \in \mathbb{R}$ are the only two scalar parameters of the GNN, and x^0, h^1, x^1, h^2, x^2 are vectors of size $|V|$. The input vector is x^0 with all coordinates equal to one and, by definition of *Softmax*, the coordinates of the output vector x^2 are all non-negative and sum to one. Also note that, for every node i , since $x_i^0 = 1$, we have $h_i^1 = w_1 \delta(i)$. It follows that, by setting $w_1 = 0$, we get $x_i^1 = \text{Relu}(1 + 0) = 1$, and then $h_i^2 = w_2 \times \delta(i)$. Therefore, as w_2 approaches $-\infty$, x^2 becomes arbitrarily close to a minimum degree distribution.

5.1 Landscape of the loss function

We begin by plotting the landscape of the expected average KL loss, evaluated on the training set. This landscape is represented in Figure 3. Although this corresponds to a simpler model than the one that yielded the results in Section 4.3, it gives several insights into the behavior during training.

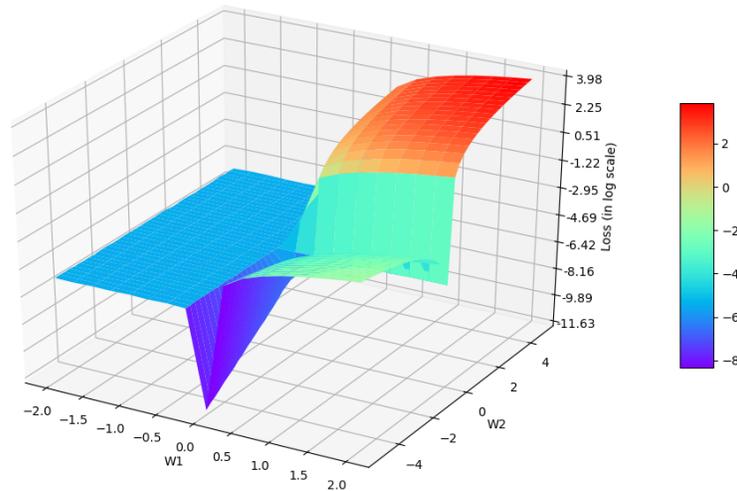


Figure 3: Landscape of the expected average KL loss (in log scale). For each (w_1, w_2) , we plot the expected average KL loss, estimated over the training set.

First, as expected, the average loss is minimized when $w_1 = 0$ and w_2 goes to $-\infty$.

Second, we observe that in the $w_1 \leq -1$ region, the average loss is flat. This region actually corresponds to the *Relu* of the first layer being inactive. Indeed, we have $h_i^1 = w_1 \delta(i)$, therefore, when $w_1 \leq -1$, we automatically get $1 + h_i^1 \leq 0$, which yields $x_i^1 = 0$. Consequently, the output of the GNN is a uniform distribution on the nodes of the graph, i.e., we obtain $x_i^2 = \frac{1}{n}$ for each node $i \in \{1, \dots, n\}$.

Third, observe that the landscape of the loss function displays fairly flat regions, which tend to be separated by sharp drops in the objective, e.g., around the $w_1 = 0$ region. This landscape most likely explains the shapes of the training curves in Figure 1, which displayed flat progression followed by sharp drops in the loss. Whether such behavior would carry out in larger dimensions remains an open question.

5.2 Landscape of the fill-in

We then plot the landscape of the expected total fill-in in Figure 4, also evaluated on the training set. While this gives us an insight into how fill-in correlates with the KL loss, let us formally restate that fill-in is never used during the training process. In particular, no gradient information is ever inferred from fill-in.

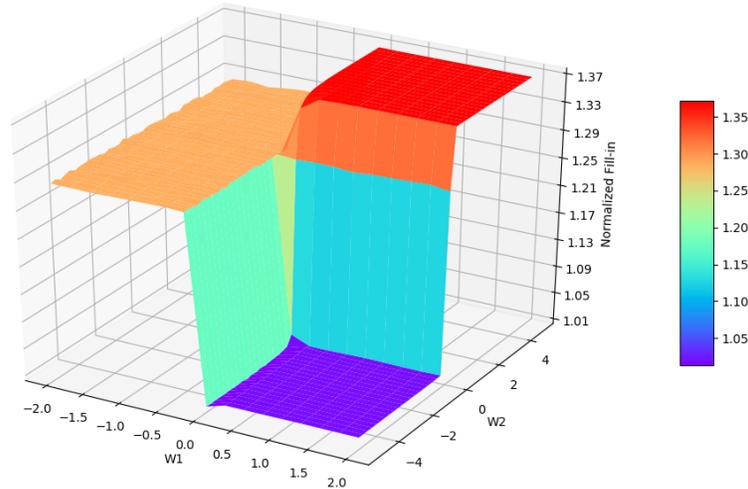


Figure 4: Landscape of the normalized total fill-in. For each (w_1, w_2) , we plot the average total fill-in of the corresponding GNN policy, divided by the expected total fill-in of the minimum degree heuristic. Both expectations are estimated over the training set.

First, unsurprisingly, similar to Figure 3, here we observe a flat landscape in the $w_1 \leq -1$ region. Recall indeed that setting $w_1 \leq -1$ means the GNN’s output reduces to a uniform policy. Second, the region $w_1 \geq 0, w_2 \geq 0$ displays high fill-in. This is not surprising either since this region essentially yields policies that select nodes with high degree, which is naturally detrimental to fill-in.

A third and more remarkable observation is the flat valley in the region $w_1 \geq 0, w_2 \leq 0$. While we know that the GNN policy converges to minimum degree when $w_1 = 0$ and w_2 takes large negative values, the plots in Figure 4 show that, when it comes to fill-in, the magnitude of w_2 does not matter as much.

Fourth and last, the minimum degree policy appears to be a minimizer of the expected total fill-in, among the set of policies that are representable by the class of GNN at hand. Although we cannot extrapolate to larger classes of models, nor to other datasets of graphs, this last observation has consequences if one were to train a GNN to minimize fill-in. Specifically, one would *need* models with higher representation power to achieve better fill-in than the minimum degree algorithm.

6 Conclusion

In this work, we have considered chordal extensions and graph elimination as major factors for devising sparsity-exploiting techniques for optimization algorithms. We have argued that, although effective heuristics to perform graph elimination (an NP-complete task) exist, there is no definitive understanding of the effect of the obtained chordal extension on the optimization algorithm using the final graph representation.

For this reason, we have followed the current research trend of looking at Combinatorial Optimization tasks by using a Machine Learning lens and we have devised a framework for learning elimination rules yielding high-quality chordal extensions. As a first building block of the learning framework, we

have proposed an on-policy imitation learning scheme that mimics the elimination ordering provided by the (classical) minimum degree rule.

The results have shown that our on-policy imitation learning approach is effective in learning the minimum degree policy and, consequently, produces graphs with desirable fill-in characteristics. In addition, the learned policy displays remarkable generalization performance, a desirable behavior since it allows to speed-up the learning process by training on smaller problems.

Finally, we identify two main research avenues for subsequent developments. On one hand, while GNNs are a good model prior for combinatorial problems over graphs, enlarging their representation power, for instance to represent hypernodes or to model multiple eliminations, will likely be key to handling practical tasks. On the other hand, the next logical step will be to learn elimination rules that explicitly address the performance of practical optimization algorithms, in conjunction with reinforcement learning-based approaches. In that regard, our future work will investigate chordal decomposition specially tailored to SDP optimization problems.

References

- [1] Agler, J., Helton, W., McCullough, S., Rodman, L.: Positive semidefinite matrices with a given sparsity pattern. *Linear Algebra and its Applications* 107, 101–149 (1988). DOI [https://doi.org/10.1016/0024-3795\(88\)90240-6](https://doi.org/10.1016/0024-3795(88)90240-6). URL <http://www.sciencedirect.com/science/article/pii/0024379588902406>
- [2] Bagnell, J., Chestnutt, J., Bradley, D.M., Ratliff, N.D.: Boosting structured prediction for imitation learning. In: *Advances in Neural Information Processing Systems*, pp. 1153–1160 (2007)
- [3] Bengio, Y., Lodi, A., Prouvost, A.: Machine learning for combinatorial optimization: a methodological tour d’horizon. *arXiv preprint arXiv:1811.06128* (2018)
- [4] Bergman, D., Cardonha, C.H., Cire, A.A., Raghunathan, A.U.: On the minimum chordal completion polytope. *Operations Research* 67(2), 532–547 (2019). DOI 10.1287/opre.2018.1783
- [5] Bixby, R.E.: Solving real-world linear programs: A decade and more of progress. *Operations Research* 50(1), 3–15 (2002). DOI 10.1287/opre.50.1.3.17780
- [6] Davis, T.A., Hu, Y.: The university of florida sparse matrix collection. *ACM Trans. Math. Softw.* 38(1), 1:1–1:25 (2011). DOI 10.1145/2049662.2049663. URL <http://doi.acm.org/10.1145/2049662.2049663>
- [7] Duvenaud, D.K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., Adams, R.P.: Convolutional networks on graphs for learning molecular fingerprints. In: *Advances in neural information processing systems*, pp. 2224–2232 (2015)
- [8] Fomin, F.V., Philip, G., Villanger, Y.: Minimum fill-in of sparse graphs: Kernelization and approximation. *Algorithmica* 71(1), 1–20 (2015). DOI 10.1007/s00453-013-9776-1. URL <https://doi.org/10.1007/s00453-013-9776-1>
- [9] Gasse, M., Chételat, D., Ferroni, N., Charlin, L., Lodi, A.: Exact combinatorial optimization with graph convolutional neural networks. *arXiv preprint arXiv:1906.01629* (2019)
- [10] George, A.: Nested dissection of a regular finite element mesh. *SIAM Journal on Numerical Analysis* 10(2), 345–363 (1973). DOI 10.1137/0710032
- [11] George, A., Liu, J.W.: The evolution of the minimum degree ordering algorithm. *SIAM Review* 31(1), 1–19 (1989). DOI 10.1137/1031001. URL <https://doi.org/10.1137/1031001>
- [12] Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256 (2010)
- [13] Gori, M., Monfardini, G., Scarselli, F.: A new model for learning in graph domains. In: *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, vol. 2, pp. 729–734. IEEE (2005)
- [14] Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: *Advances in Neural Information Processing Systems*, pp. 1024–1034 (2017)
- [15] Hamilton, W.L., Ying, R., Leskovec, J.: Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584* (2017)
- [16] Howard, R.A.: *Dynamic programming and markov processes*. (1960)
- [17] Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016)

- [18] Kullback, S.: Information theory and statistics. Courier Corporation (1997)
- [19] Li, Z., Chen, Q., Koltun, V.: Combinatorial optimization with graph convolutional networks and guided tree search. In: Advances in Neural Information Processing Systems, pp. 539–548 (2018)
- [20] Majumdar, A., Hall, G., Ahmadi, A.A.: A survey of recent scalability improvements for semidefinite programming with applications in machine learning, control, and robotics. arXiv preprint arXiv:1908.05209 (2019)
- [21] Markowitz, H.M.: The elimination form of the inverse and its application to linear programming. *Management Science* 3(3), 255–269 (1957). DOI 10.1287/mnsc.3.3.255
- [22] Pan, Y., Cheng, C.A., Saigol, K., Lee, K., Yan, X., Theodorou, E., Boots, B.: Agile autonomous driving using end-to-end deep imitation learning. In: Robotics: science and systems (2018)
- [23] Ratliff, N.D., Silver, D., Bagnell, J.A.: Learning to search: Functional gradient techniques for imitation learning. *Autonomous Robots* 27(1), 25–53 (2009)
- [24] Ross, S., Bagnell, D.: Efficient reductions for imitation learning. In: Proceedings of the thirteenth international conference on artificial intelligence and statistics, pp. 661–668 (2010)
- [25] Ross, S., Gordon, G., Bagnell, D.: A reduction of imitation learning and structured prediction to no-regret online learning. In: Proceedings of the fourteenth international conference on artificial intelligence and statistics, pp. 627–635 (2011)
- [26] Rothberg, E., Hendrickson, B.: Sparse matrix ordering methods for interior point linear programming. *INFORMS Journal on Computing* 10(1), 107–113 (1998). DOI 10.1287/ijoc.10.1.107
- [27] Schaal, S.: Is imitation learning the route to humanoid robots? *Trends in cognitive sciences* 3(6), 233–242 (1999)
- [28] Silver, D., Bagnell, J., Stentz, A.: High performance outdoor navigation from overhead data using imitation learning. *Robotics: Science and Systems IV*, Zurich, Switzerland (2008)
- [29] Vandenberghe, L., Andersen, M.S., et al.: Chordal graphs and semidefinite optimization. *Foundations and Trends® in Optimization* 1(4), 241–433 (2015)
- [30] Wright, S.: Primal-Dual Interior-Point Methods. Society for Industrial and Applied Mathematics (1997). DOI 10.1137/1.9781611971453. URL <http://epubs.siam.org/doi/abs/10.1137/1.9781611971453>
- [31] Yannakakis, M.: Computing the minimum fill-in is np-complete. *SIAM Journal on Algebraic Discrete Methods* 2(1), 77–79 (1981). DOI 10.1137/0602010. URL <https://doi.org/10.1137/0602010>
- [32] Zheng, Y., Fantuzzi, G., Papachristodoulou, A., Goulart, P., Wynn, A.: Chordal decomposition in operator-splitting methods for sparse semidefinite programs. *Mathematical Programming* (2019). DOI 10.1007/s10107-019-01366-3. URL <https://doi.org/10.1007/s10107-019-01366-3>