



Title	Learning fine-grained search space pruning and heuristics for combinatorial optimization
Authors(s)	Lauri, Juho, Dutta, Sourav, Grassia, Marco, Ajwani, Deepak
Publication date	2023-05-08
Publication information	Lauri, Juho, Sourav Dutta, Marco Grassia, and Deepak Ajwani. "Learning Fine-Grained Search Space Pruning and Heuristics for Combinatorial Optimization." Springer, May 8, 2023. https://doi.org/10.1007/s10732-023-09512-z .
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/25093
Publisher's version (DOI)	10.1007/s10732-023-09512-z

Downloaded 2026-05-02 00:24:08

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information



Learning fine-grained search space pruning and heuristics for combinatorial optimization

Juho Lauri¹ · Sourav Dutta² · Marco Grassia³ · Deepak Ajwani⁴

Received: 21 July 2021 / Revised: 7 April 2023 / Accepted: 11 April 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

Abstract

Combinatorial optimization problems arise naturally in a wide range of applications from diverse domains. Many of these problems are NP-hard and designing efficient heuristics for them requires considerable time, effort and experimentation. On the other hand, the number of optimization problems in the industry continues to grow. In recent years, machine learning techniques have been explored to address this gap. In this paper, we propose a novel framework for leveraging machine learning techniques to scale-up *exact* combinatorial optimization algorithms. In contrast to the existing approaches based on deep-learning, reinforcement learning and restricted Boltzmann machines that attempt to directly learn the output of the optimization problem from its input (with limited success), our framework learns the relatively simpler task of pruning the elements in order to reduce the size of the problem instances. In addition, our framework uses only interpretable learning models based on intuitive local features and thus the learning process provides deeper insights into the optimization problem and the instance class, that can be used for designing better heuristics. For the classical maximum clique enumeration problem, we show that our framework can prune a large fraction of the input graph (around 99% of nodes in case of sparse graphs) and still detect almost all of the maximum cliques. Overall, this results in several fold speedups of state-of-the-art algorithms. Furthermore, the classification model used in our framework highlights that the chi-squared value of neighborhood degree

The present manuscript integrates and extends the works (Lauri and Dutta 2019; Grassia et al. 2019; Dutta and Lauri 2019), which appeared at AAAI'19, the DSO Workshop at IJCAI'19 and CIKM'19, respectively. Part of this work was done while all authors were at Nokia Bell Labs, Ireland.

✉ Juho Lauri
juho.lauri@gmail.com

¹ Helsinki, Finland

² Huawei Research, Dublin, Ireland

³ University of Catania, Catania, Italy

⁴ University College Dublin, Dublin, Ireland

has a statistically significant correlation with the presence of a node in a maximum clique, particularly in dense graphs which constitute a significant challenge for modern solvers. We leverage this insight to design a novel heuristic we call ALTHEA for the maximum clique detection problem, outperforming the state-of-the-art for dense graphs.

Keywords Combinatorial optimization · Machine learning · Maximum clique · Heuristics

1 Introduction

Combinatorial optimization problems typically involve finding groupings (subsets), orderings or assignments of a discrete, finite set of objects that satisfy certain conditions or constraints. Many optimization problems are NP-hard, implying that—unless $P = NP$ —no polynomial-time algorithms exist for these problems that can solve every instance of the problem to optimality. Over the last century, numerous approaches have been developed for these applications, including (i) exact algorithms with exponential time worst-case complexity (including integer linear programming and constraint programming solvers), (ii) approximation algorithms with formal guarantees on the solution quality, (iii) parameterized algorithms, (iv) carefully designed heuristics that leverage the structure often available in real-world instances and (v) meta-heuristic frameworks such as genetic algorithms or ant colony optimization. While exact algorithms have poor scalability, the design of approximation algorithms, parameterized algorithms and domain-specific heuristics require considerable development and design time. Similarly, meta-heuristics often require significant configuration time to select the best parameters and operators for a given optimization problem and can take considerable time to find a combination of elements close to the optimal solution.

In recent years, machine learning techniques have been used to design effective solutions for combinatorial optimization problems. Compared to the traditional meta-heuristic frameworks, machine learning techniques can enable us to explore the search space of solutions in a statistically principled way and this has the potential to explore the search space faster.

Early research on leveraging machine learning for combinatorial optimization focused on end-to-end solutions that use a combination of deep learning [e.g., pointer network (Vinyals et al. 2015)], reinforcement learning [e.g., neural combinatorial optimization (Bello et al. 2017) and greedy Q-learning for graph optimization problems (Khalil et al. 2017)] and unsupervised learning [e.g., restricted Boltzmann machines and Estimation of Distribution Algorithm (Probst et al. 2017)] techniques. These techniques learn the exact decision boundary to separate the elements in the optimal subset from the remaining elements. However, this is a challenging task requiring complex learning models with a large number of parameters. As a result, the learning models are not easy to interpret. Since the learned heuristic (mapping from input to output) is implicit in the complex learning model, this implies that the learned heuristic is itself not interpretable. This has the following consequences:

- With a large number of parameters, the learned algorithm is not easy for humans to understand and consequently there is little potential for mathematical analysis of the resultant algorithm.
- It is not clear if the learned model will still work if there is an additional constraint added to the problem. This is a major concern for applications in industry where the first modelling of a problem into the optimization objective and associated constraints is rarely enough and new constraints are incrementally discovered and added.
- It is not clear if the learned model will still work if a dataset from a slightly different distribution is used as an input. This has further implications for cross-domain generalizability of the learning model.
- It is not clear how to integrate the existing knowledge in the algorithm design and algorithm engineering community about the problem and the topological structure of the input instances, to improve upon the solution accuracy. In this context, we also note that in a recent survey of the area, Bengio et al. (2021) concluded that “end to end machine learning approaches to combinatorial optimization are not enough” and to be effective, learning frameworks should “benefit from the theoretical guarantees and state-of-the-art algorithms already available.”

In contrast, we propose a novel framework for solving combinatorial optimization problems that uses interpretable learning models based on intuitive local features. For the interpretable models to work well, we focus on the relatively simpler task of (non-exhaustive) pruning of the elements that are *not* in the optimal subset. This reduces the problem size, often significantly, enabling existing solvers to deal with considerably larger instances.

The complexity of exact algorithms generally depends only on the problem size. On the other hand, machine learning considers the ease and the confidence of identifying elements to either be or not to be in the optimal solution based on available information. Thus, our framework combines the best of both worlds—it uses machine learning to explore if an element of the problem instance can be confidently pruned based on its associated feature vector (thereby reducing the problem size with minimal effect on the optimal solution) and then it uses the exact algorithm to deal with the difficult core of the problem, where the learning model was uncertain.

The existing heuristics and algorithms for a problem and the characteristics of the input distribution can be encapsulated in the form of features in our learning framework. Thus, our proposed framework is able to leverage the domain knowledge and existing algorithmic literature on the problem.

To prune the elements further, we extend our framework to have multiple pruning stages. In each stage, the framework learns a new classification model for elements that were not pruned by earlier classification models (thereby increasingly focusing on harder elements to prune). Figure 1 illustrates the exact decision boundary and the multi-stage pruning framework.

Furthermore, the learning process in our framework provides deeper insights into the optimization problem and the instance class. In particular, it identifies the combination of simple features that are most indicative of an element belonging to an optimal

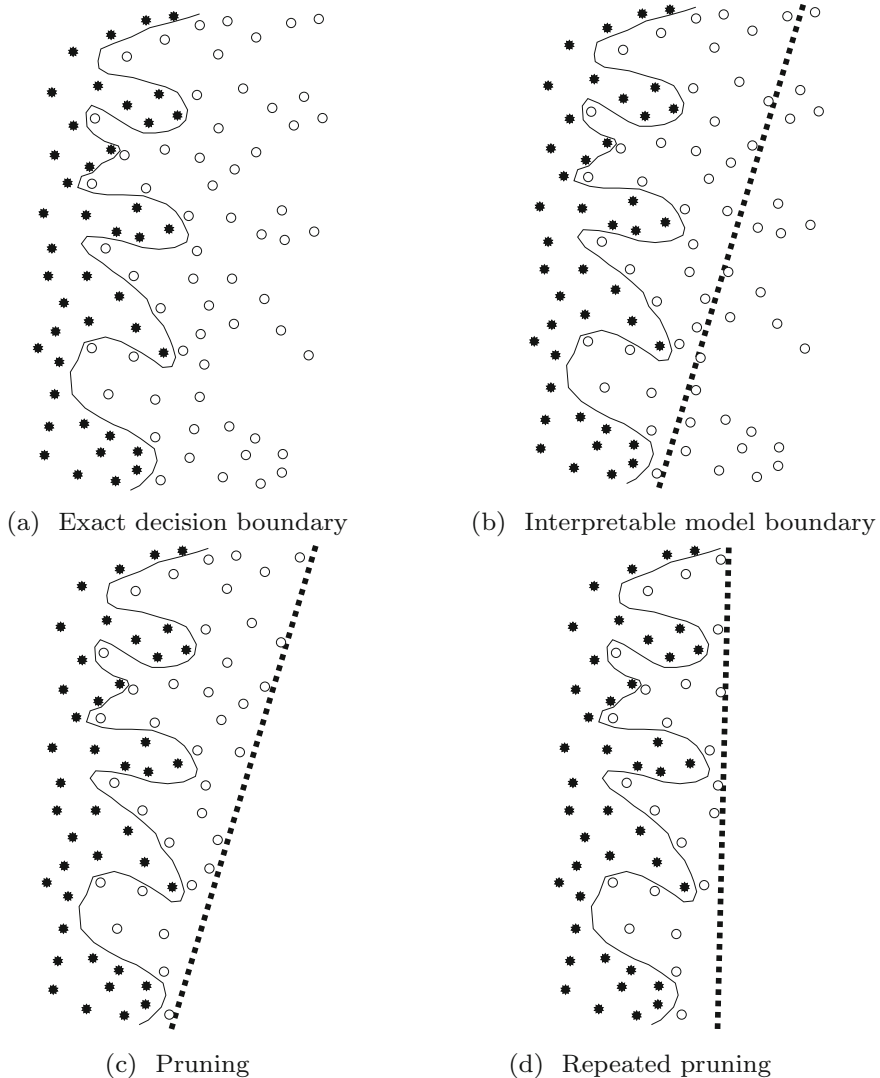


Fig. 1 Depiction of the overall framework. The black stars indicate the elements in the optimal subset while the white circles represent the elements not in the optimal subset. Earlier learning frameworks attempted to learn the exact decision boundary (drawn with a solid black line). Our framework use simple, interpretable classifier (shown by dashed lines) to repeatedly prune the white circles

subset. This insight can be leveraged to design better heuristics for the optimization problems.

For the classical maximum clique enumeration problem, we show that our framework can prune a large fraction of the graph (around 99% of nodes in case of sparse graphs) and still detect almost all of the maximum cliques. For dense graphs that have a large fraction of their edges in some maximum clique, it correctly predicts all of the

clique edges and for the dense graphs where a very small fraction of edges are in any maximum clique, our framework learns to prune a large fraction of edges while still retaining many of the maximum cliques in the graph. Overall, this results in several fold speedups of state-of-the-art algorithms on a large number of graph instances. Furthermore, the classification model used in our framework highlights that the chi-squared value of neighborhood degree has a statistically significant correlation with the presence of a node in a maximum clique, particularly in the case of dense graphs which constitute a major challenge for state-of-the-art solvers. We leverage this insight to design a novel heuristic we call ALTHEA for the maximum clique detection problem, outperforming the state-of-the-art for dense graphs. Here, dense graphs are of particular instance as many techniques effective in practice are known for sparse graphs (see Sect. 2.2). Furthermore, high-throughput data capture in applications in computational biology “often create graphs of extreme size and density” (Eblen et al. 2012). We emphasize that the above framework can be used on a wide range of combinatorial optimization problems. For instance, consider any binary integer linear program where the variables need to be fixed to 0 or 1. This can be viewed as a binary classification problem where we predict the variables to belong to either the zero class or the one class. Similarly, even though the solution of a problem such as travelling salesperson problem (TSP) involves all vertices, it can still be viewed as a classification problem where the edges of the graph need to be classified into whether or not they belong to an optimal TSP tour. Once we view our optimization problem as a classification problem, we can use our learning framework to prune a large number of elements (e.g., vertices or edges) or fix a large number of binary variables and solve the remaining problem using exact and/or heuristic approaches.

Since an early version of our work was released on a pre-print server (Lauri et al. 2020), our framework has also been successfully used for many other optimization problems such as TSP (Fitzpatrick et al. 2021), k-median (Tayebi et al. 2022), facility location (Tayebi et al. 2022), set cover (Tayebi et al. 2022), maximum coverage (Tayebi et al. 2022), minimum Steiner tree problem (Zhang and Ajwani 2022), and a variant of the electric vehicle routing problem (Fitzpatrick et al. 2023). For all of these problems, our framework is able to achieve a significant reduction in running time, while ensuring that the obtained solutions are close to optimal.

An important issue in this framework is to ensure that pruning does not render the remaining problem infeasible. This can be done by carefully selecting the threshold for pruning. Alternatively, we can use a heuristic or some fast approximation algorithm to obtain a feasible solution for the problem and not prune the variables/elements in that feasible solution (even if the classifier predicts so). For instance, in the case of TSP (Fitzpatrick et al. 2021), a solution obtained from the classical double-tree approximation was used as a set of edges that could not be pruned by the classifier. This ensured that the graph remained connected even after the pruning and thus, at least one feasible solution was guaranteed.

2 Related work

In this work, we propose a new framework to leverage machine learning for combinatorial optimization problems and apply this to the maximum clique enumeration problem to show its effectiveness. Therefore, as related work, we consider (i) the work related to leveraging machine learning for combinatorial optimization problems and (ii) the complexity of algorithms and preprocessing techniques to solve maximum clique enumeration and detection problem.

2.1 Machine learning for combinatorial optimization

For a survey of machine learning techniques used for combinatorial optimization, we refer the reader to the surveys by DiCarro (Caro 2019) and Bengio et al. (2021). The existing literature on learning techniques can be broadly categorized into the following classes.

Supervised learning to directly learn the solution of combinatorial optimization problems This set of techniques involve the use of deep-learning for learning the optimal solution. A seminal example in this category is the work by Vinyals et al. (2015), who viewed the task of learning combinatorial optimization solutions as a sequence-to-sequence learning problem. They aimed for directly learning the output solution from the input sequence for optimization problems such as convex hull and Delauney triangulation. For the sequence-to-sequence learning, they used a recurrent neural network (RNN). To deal with the issue of long-range correlations (elements far from each other in the input sequence affecting the same output element), they used an attention mechanism to augment the RNN model. To deal with the issue of a fixed vocabulary size required for the output of a recurrent neural network, they used pointers to elements in the input stream, resulting in the name pointer network.

Reinforcement learning for combinatorial optimization Deep learning requires a large amount of training data and to generate this, a large number of NP-hard problem instances need to be solved, limiting the applicability of these techniques. On the other hand, given a solution, it is relatively easy to evaluate the quality of the solution by computing the optimization objective. Thus, in recent years, reinforcement learning based techniques have been developed to solve optimization problems. In this framework, the goal is to learn a stochastic policy that samples solutions of high quality with high probability. In particular, Khalil et al. (2017) used the Q-learning technique to learn the solutions for graph optimization problems, specifically minimum vertex cover and maximum cut. They encode nodes using a graph embedding technique and then build a solution using a greedy construction meta-algorithm. The greedy decisions are based on an estimated Q-function parameterized by the embedding. The embedding parameters for the Q-function are updated step by step based on the partial solution computed. The GCOMB approach of Manchanda et al. (2020) follows the same framework, but claims to scale to very large graphs. Another example of this framework is the use of neural combinatorial optimization (Bello et al. 2017) for TSP. Very recent attempts in this direction include the development of general purpose

constraint programming solver guided by reinforcement learning (Chalumeau et al. 2021).

Unsupervised learning for combinatorial optimization Unsupervised approaches via restricted Boltzmann machines have also been used to deal with combinatorial optimization problems. An example of this framework is the Estimation of Distribution Algorithm by Probst et al. (2017). This approach iteratively builds and samples from a probabilistic model of candidate solutions. Intuitively, these approaches build information about the probability distribution of good candidate solutions. This model is built using contrastive divergence.

Limitations of the above frameworks The learning models used in these existing state-of-the-art frameworks are both hard to interpret and architecturally complex. For instance, the neural combinatorial optimization approach (Bello et al. 2017) is a combination of pointer networks (with two LSTM networks), a Monte Carlo policy gradient and an actor-critic architecture. The complexity of these approaches comes at a significant cost of interpretability. Since the learned algorithm (mapping from input to output) is implicit in the complex learning model, this implies that the learned heuristic is itself not easy for humans to understand. As a direct consequence, it is difficult to analyze the learned algorithm mathematically. Moreover, it is unclear what features of the input instances are being exploited by the learned heuristic and on which class of datasets will it perform well.

In contrast to using machine learning techniques as an end-to-end heuristic solver (see e.g., Khalil et al. 2017; Manchanda et al. 2020; Bello et al. 2017; Huang et al. 2019; Probst et al. 2017; Amizadeh et al. 2019; Vinyals et al. 2015; Li et al. 2018), these techniques have also been used as a component in an exact solver (see e.g., Lodi and Zarpellon 2017; Chmiela et al. 2021). We use a lightweight supervision as a preprocessing step to reduce the size of an instance and this can be followed by an exact approach on the pruned instance. From this perspective, our approach can also be viewed as a component in an exact solver, though it is not tied to any solver.

2.2 Maximum clique enumeration

We instantiate our framework for the maximum clique enumeration (MCE) problem. In this problem, the goal is to list all *maximum* (as opposed to maximal) cliques in a given graph. This listing variant of the problem, MCE, is encountered in computational biology in problems like searching for common cis-regulatory elements (Abu-Khzam et al. 2005), searching for putative molecular response networks (Eblen et al. 2012), and the detection of protein-protein interaction (Yeger-Lotem et al. 2004). For an overview of additional potential applications, see Bomze et al. (1999).

Note that for designing the heuristic in Sect. 7, we consider the maximum clique detection problem, where the goal is to just find the size of maximum clique (e.g., by identifying only one maximum clique). The maximum clique problem is one of the most heavily-studied combinatorial problems arising in various domains such as in the analysis of social networks (Faust and Wasserman 1995; Fortunato 2010; Palla et al. 2005; Papadopoulos et al. 2012), behavioral networks (Bernard et al. 1979), and

financial networks (Boginski et al. 2005). It is also relevant in clustering (Stix 2004; Yang et al. 2016) and cloud computing (Wang et al. 2014; Yao et al. 2013).

The computational aspects of this problem are well-studied. Indeed, it is NP-hard to even approximate the maximum clique problem within $n^{1-\epsilon}$ for any $\epsilon > 0$ (Zuckerman 2006). Furthermore, unless an unlikely collapse occurs in complexity theory, the problem of identifying whether a graph of n vertices has a clique of size k is not solvable in time $f(k)n^{o(k)}$ for any function f (Chen et al. 2006). As such, even small instances of this problem can be non-trivial to solve. Moreover, under reasonable complexity-theoretic assumptions, there is no polynomial-time algorithm that preprocesses an instance of k -clique to have only $f(k)$ vertices, where f is any computable function depending solely on k (see e.g., Cygan et al. 2015). These results indicate that it is unlikely that an efficient preprocessing method for MCE exists that can reduce the size of input instance drastically while guaranteeing to preserve all the maximum cliques. In particular, it is unlikely that polynomial-time sparsification methods (see e.g., Batson et al. 2013) would be applicable to MCE. This has led researchers to focus on heuristic pruning approaches.

On preprocessing for maximum clique For the discussion to follow, it will be useful to recall the concept of a k -core of a graph G . Here, the k -core of G is a maximal subgraph of G where every vertex in the subgraph has degree at least k in the subgraph. The *core number* of a vertex v is the largest k for which a k -core containing v exists. A typical preprocessing step in a state-of-the-art solver is the following: (i) quickly find a large clique (say of size k), (ii) compute the core number of each vertex of the input graph G , and (iii) delete every vertex of G with core number less than $k - 1$. This can be equivalently achieved by repeatedly removing all vertices with degree less than k . For example, the solver `pmc` (Rossi et al. 2015)—which is regarded as “the leading reference solver” (San Segundo et al. 2016)—use this as the only preprocessing method. However, there are two major downsides to this preprocessing step. First, it is crucially dependant on k , the size of a large clique found. Since the maximum clique size is NP-hard to approximate within a factor of $n^{1-\epsilon}$, maximum clique estimates with no formal guarantees are used. Second and more important, it is typical that even if the estimate k was equal to the size of a maximum clique in G , the core number of most vertices could be considerably higher than $k - 1$. This is particularly true in the case of dense graphs and it results in little or *no* pruning of the search space. Similarly, other preprocessing strategies [see e.g., Eblen (2010) for more discussion] depend on NP-hard estimates of specific graph properties and are not useful for pruning dense graphs.

These facts further motivate the quest for preprocessing methods that (i) are effective on dense graphs and (ii) work independently of any estimate k for the maximum clique size. Unfortunately, as described earlier, under widely-believed complexity-theoretic assumptions, no methods exist that can give strong guarantees for pruning arbitrary graphs. This raises the question if one can discover heuristic methods that can do significant pruning, in practice, on graphs from different domains, including dense graphs. Even more importantly, can we *learn* a heuristic to prune the input instance?

3 Our framework

In this section, we describe our framework for subset-based combinatorial optimization problems. For ease of exposition, we describe the framework in terms of the MCE problem. We stress that our approach is not restricted to MCE, but can be applied to other problems as well.

In our case, we assume the instance is represented as an undirected graph $G = (V, E)$. Moreover, in contrast to previous approaches, we view *individual vertices* of G as classification problems as opposed to G itself. That is, the problem is to induce a mapping $\gamma : V \rightarrow \{0, 1\}$ from a set of L training examples $T = \{f(v_i), y_i\}_{i=1}^L$, where $v_i \in V$ is a vertex, $y_i \in \{0, 1\}$ a class label, and $f : V \rightarrow \mathbb{R}^d$ a mapping from a vertex to a d -dimensional feature space. For reasons of scalability, we strive to keep d small and to ensure that f can be computed efficiently.

Single-stage sparsification To learn the mapping γ from T , we use a probabilistic classifier P which outputs a probability distribution over $\{0, 1\}$ for a given $f(u)$ for $u \in V$. A natural parameterized search strategy, which we call *probabilistic preprocessing* (or *single-stage sparsification*), for enhancing a search algorithm \mathcal{A} by P is as follows. Define a *confidence threshold* $q \in [0, 1]$. Delete from G each vertex predicted by P to *not* be in a solution with probability at least q , i.e., let $V(G') = V(G) \setminus V'$, where $V' = \{u \mid u \in V \wedge P(u = 0) \geq q\}$. Execute \mathcal{A} with G' as input instead of G . Here, the purpose of q is to control the error and pruning rate of preprocessing: (i) it is more acceptable to not remove a vertex that is *not* in a solution than to remove a vertex that *is* in a solution, and (ii) a lower value of q translates to a possibly higher pruning rate. Clearly, this strategy is a heuristic, i.e., it is possible that the cost of an optimal solution in G' differs from that in G .

Multi-stage sparsification A natural generalization of the probabilistic preprocessing strategy is the following approach that we call *multi-stage sparsification*. First, to gain some intuition, consider a probabilistic classifier P learning to prune vertices as a single-stage sparsifier. Plausibly, at first, P will learn to remove vertices of high degree. However, after this first stage, there are potentially still vertices of lower degree that could also be safely removed. To capture this, we can take a set of networks pruned by P and train a second classifier, say P' , to learn to remove additional vertices, this time perhaps of lower degree. Naturally, we can continue this process to train more classifiers (corresponding to subsequent stages) for more effective pruning.

To make our intuition precise, let \mathcal{G}_1 be the input set of networks. Consider a graph $G \in \mathcal{G}_1$. Let \mathcal{M} be the set of all maximum cliques of G , and denote by $V(\mathcal{M})$ the set of all vertices in \mathcal{M} . The positive examples in the training set T_1 consist of all vertices that are in some maximum clique ($V(\mathcal{M})$) and the negative examples are the ones in the set $V \setminus V(\mathcal{M})$. Since the training dataset can be highly skewed, we under-sample the larger class to achieve a balanced training data. A probabilistic classifier P_1 is trained on the balanced training data in stage 1. Then, in the next stage, we remove all vertices that were predicted by P_1 to be in the negative class with a probability above a predefined threshold q . At this point, the classifier P_1 is fixed.

Next, we generate the set \mathcal{G}_2 of subgraphs (of graphs in \mathcal{G}_1) induced on the vertices not pruned by P_1 from the instances in \mathcal{G}_1 . This constitutes our training set T_2 . Here,

the positive examples in the training set T_2 consist of all vertices in some maximum clique ($V(\mathcal{M})$) and the negative examples are the ones in the set $V \setminus V(\mathcal{M})$. The training dataset is balanced by under-sampling and we use that balanced dataset to learn the probabilistic classifier P_2 . At this stage, the classifier P_2 is fixed. Again, we generate \mathcal{G}_3 and the training set T_3 using P_2 in a similar manner. We repeat the process for ℓ stages.

4 Computational features

In this section, we describe the computational features used in our framework.

Graph-theoretic features We use the following graph-theoretic features: **(F1)** number of vertices, **(F2)** number of edges, **(F3)** vertex degree, **(F4)** local clustering coefficient (LCC), and **(F5)** eigencentality.

The crude information captured by features (F1)–(F3) provides a reference for the classifier for generalizing to different distributions from which the graph might have been generated. Feature (F4), the LCC of a vertex is the fraction of its neighbors with which the vertex forms a triangle, encapsulating the well-known small world phenomenon. Feature (F5) eigencentality represents a high degree of connectivity of a vertex to other vertices, which in turn have high degrees as well. The *eigenvector centrality* \vec{v} is the eigenvector of the adjacency matrix A of G with the largest eigenvalue λ , i.e., it is the solution of $A\vec{v} = \lambda\vec{v}$. The i th entry of \vec{v} is the *eigencentality* of vertex v . In other words, this feature provides a measure of local “denseness”. A vertex in a dense region shows higher probability of being part of a large clique.

Statistical features In our *statistical features*, we focus on encoding the deviation of a vertex/neighborhood feature from the average characteristics as observed in the overall graph. Intuitively, any vertex that demonstrates a significant deviation from the expected behaviour (in terms of feature values), its statistical significance would increase, potentially indicating some observed phenomenon that cannot be attributed solely to randomness. Hence, we train our learning model to focus on such vertices. More details on how such a statistical measure would work can be found in Sect. 7.

To this end, additionally, we use the following statistical features: **(F6)** the χ^2 statistical significance value over vertex degree, **(F7)** average χ^2 value over neighbor degrees, **(F8)** χ^2 value over LCC, and **(F9)** average χ^2 value over neighbor LCCs.

The intuition behind (F6)–(F9) is that for a vertex v present in a large clique, its degree and LCC would deviate from the underlying expected distribution characterizing the graph. Further, the neighbors of v also present in the clique would demonstrate such behaviour. Indeed, statistical features have been shown to be robust in approximately capturing local structural patterns (Dutta et al. 2017).

Statistical significance is captured by the notion of p value (Read and Cressie 1988), and well-estimated (Read and Cressie 1989) by the *Pearson’s chi-square statistic*, χ^2 , computed as

$$\chi^2 = \sum_{v_i} \left[(O_i - E_i)^2 / E_i \right], \quad (1)$$




Fig. 2 The neighborhood of v contains a single color, so $\chi_d(v) = 1/3$. However, the coloring obtained by swapping the non-white colors in either triangle will only show that $\chi_d(v) \leq 2/3$. That is to say, finding an optimal proper coloring of a graph will not necessarily suffice to determine $\chi_d(v)$

where O_i and E_i are the observed and expected number of occurrences of the possible outcomes i .

Local chromatic density Let $G = (V, E)$ be a graph. A k -coloring of G is a function $c : V \rightarrow \{1, 2, \dots, k\}$. A coloring is a k -coloring for some $k \leq n$, where $n = |V|$. A coloring c is proper if $c(u) \neq c(v)$ for every edge $\{u, v\} \in E$. The chromatic number of G , denoted by $\chi(G)$, is the smallest k such that G has a proper k -coloring. We define the local chromatic density of a vertex $v \in V$, denoted by $\chi_d(v)$, as the ratio between the minimum number of distinct colors appearing in $N(v)$ among any optimal proper coloring of G and the chromatic number of G . Informally, the local chromatic density of v is the minimum possible number of colors in the immediate neighborhood of v in any optimal proper coloring of G (see Fig. 2).

We use the local chromatic density as the feature **(F10)**. A vertex v with high $\chi_d(v)$ means that the neighborhood of v is dense, as it captures the adjacency relations between the vertices in $N(v)$. Thus, a vertex in such a dense region has a higher chance of belonging to a large clique.

However, the problem of computing $\chi_d(v)$ is computationally difficult. In the decision variant of the problem, we are given a graph $G = (V, E)$, a vertex $v \in V$, and a ratio $q \in (0, 1)$. The task is to decide whether there is proper k -coloring of V such that $\chi_d(v) \geq q$. As shown in the following, the claimed hardness is straightforward to establish.

Theorem 1 Given a graph $G = (V, E)$, $v \in V$, and $q \in (0, 1)$, it is NP-hard to decide whether $\chi_d(v) \leq q$.

Proof The claim is established by giving a polynomial-time reduction from graph k -coloring, which is well-known to be NP-complete for every $k \geq 3$.

Let G be an instance of graph k -coloring, for any fixed $k \geq 3$. We construct $G' = G \cup K_k$, i.e., G' is the disjoint union of G and a complete graph on k vertices. Fix v to be an arbitrary vertex of the K_k . We claim that G has a proper k -coloring if and only if $\chi_d(v) \leq q$, where $q = \frac{k-1}{k}$.

If G admits a proper k -coloring, we map $\{1, 2, \dots, k\}$ bijectively to $V(K_k)$, implying that $\chi_d(v) = \frac{k-1}{k}$. On the other hand, a proper k -coloring of G' witnessing that $\chi_d(v) = \frac{k-1}{k}$ is clearly a proper k -coloring when restricted to G as well. \square

Despite its computational hardness, we can estimate $\chi_d(v)$ with the following simple heuristic. Compute a proper coloring for G using e.g., the well-known linear-time

greedy heuristic of Welsh and Powell (1967) and then estimate $\chi_d(v)$ as the ratio between the number of colors in $N(v)$ divided by the number of colors used by the greedy coloring algorithm. Note that we could use other graph coloring heuristics as well [see e.g., Lewis (2015) for an overview of the state-of-the-art].

Learning over edges Instead of individual vertices, we can view the framework also over *individual edges*. In this case, the goal is to find a mapping $\gamma' : E \rightarrow \{0, 1\}$, and the training set L' contains feature vectors corresponding to edges instead of vertices. We also briefly explore this direction in this work.

Edge features We use the following features (E1)–(E9) for an edge $\{u, v\}$. **(E1)** Jaccard similarity is the number of common neighbors of u and v divided by the number of vertices that are neighbors of at least one of u and v . **(E2)** Dice similarity is twice the number of common neighbors of u and v , divided by the sum of their degrees. **(E3)** Inverse log-weighted similarity is the number of common neighbors of u and v weighted by the inverse logarithm of their degrees. Formally, we compute $\sum_{x \in N(u) \cap N(v)} 1/\log(\deg(x))$. **(E4)** Cosine similarity is the number of common neighbors of u and v divided by the geometric mean of their degrees. The next three features are inspired by the vertex features: **(E5)** average LCC over u and v , **(E6)** average degree over u and v , and **(E7)** average eigencentrality over u and v . **(E8)** is the number of length-two paths between u and v . Finally, we use **(E9)** *local edge-chromatic density*, i.e., the number of distinct colors on the common neighbors of u and v divided by the total number of colors used in any optimal proper coloring.

The intuition behind (E1)–(E4) is well-established for community detection; see e.g., Adamic and Adar (2003) for more. For (E8), observe that the number of length-two paths is high when the edge is part of a large clique, and at most $n - 2$ when $\{u, v\}$ is an edge of a complete graph on n vertices. Notice that (E9) could be converted into a deterministic rule: the edge $\{u, v\}$ can be safely deleted if the common neighbors of u and v see less than $k - 2$ colors in any proper coloring of the input graph G , where k is an estimate for $\omega(G)$. To our best knowledge, such a rule has not been considered previously in the literature. Further, notice that there are situations in which this rule *can* be applied whereas the similar vertex rule uncovered from (F10) cannot. To see this, let G be a graph consisting of two triangles $\{a, b, c\}$ and $\{x, y, z\}$, connected by an edge $\{a, x\}$, and let $k = 3$. The vertex rule cannot delete a nor x , but the described edge rule removes $\{a, x\}$.

5 Experimental results

In this section, we describe how multi-stage sparsification is applied to the MCE problem and our computational results.

All experiments are executed on a machine equipped with an Intel Core i7-4770K CPU (3.5 GHz), 8 GB of RAM, running Ubuntu 16.04.

Training and test data All our datasets are obtained from Network Repository (Rossi and Ahmed 2015) (available at <http://networkrepository.com/>). We discard all vertex and edge weights and parallel edges (if any) and treat every directed edge as undirected.

For dense networks, we choose a total of 30 networks from various categories with the criteria that the edge density is at least 0.5 in each. We name this category “dense”. The test instances are in Table 1, chosen based on empirical hardness (i.e., they are solvable in reasonable amount of time).

For sparse networks, we choose our training data from four different categories: 31 biological networks (“bio”), 32 social networks (“soc”), 107 Facebook networks (“socfb”), and 13 web networks (“web”). In addition, we build a fifth category “all” that comprises all networks from the mentioned four categories. The test instances are in Table 2.

Feature computation We implement the feature computation in C++, relying on the `igraph` (Csardi and Nepusz 2006) C graph library. In particular, our feature computation is single-threaded with further optimization possible.

Domain oblivious training via local chromatic density To achieve a high classification accuracy, it is natural to assume that the classifier should be trained with networks coming from the same domain, and that testing should be performed on networks from that domain. Certainly, some similarity is needed between the two for training to be effective. For example, sparse networks (say trees) should not be representative of dense networks. However, we demonstrate in Table 3 that a classifier can be trained with networks from various domains, yet predictions remain accurate across different domains (see column “all”). The accuracy is boosted considerably by the introduction of the local chromatic density (F10) into the feature set (see Table 3). In particular, when generalizing across various domains, the impact on accuracy is almost 10%. For this reason, rather than focusing on network categories, we only consider networks by edge density (at least 0.5 or not).

Accuracy measures and setup For our experiments, the *vertex pruning ratio* is the ratio of the number of vertices removed from the instance to the number of vertices $|V|$ in the original instance. The *edge pruning ratio* is defined similarly, but for edges instead of vertices. We say *clique accuracy* is one precisely when the number of all maximum cliques of the instance G is equal to the number of all maximum cliques of the reduced instance G' and $\omega(G) = \omega(G')$.

State-of-the-art solvers for MCE To the best of our knowledge, the only publicly available maximum clique solvers able to list all maximum cliques¹ are `cliquer` (Östergård 2002), based on a branch-and-bound strategy; and `MoMC` (Li et al. 2017), introducing incremental maximum satisfiability reasoning to a branch-and-bound strategy. We use these solvers in our experiments.²

5.1 Dense networks

In this subsection, we show results for probabilistic preprocessing on dense networks (i.e., edge density at least 0.5).

¹ For instance, `pmc` (Rossi et al. 2015) does not have this feature.

² It is worth noticing that in principle, one could solve the problem by any algorithm that lists all *maximal cliques*. However, such algorithms solve a more general problem (i.e., every maximum clique is maximal but the opposite is not true in general) which usually comes with a significantly higher computational cost.

Table 1 Experiments for dense graphs

Instance	$ V $	$ E $	ω	n, ω	k -core	1-stage	Pruning	cliquer	MoMC
brock200-1	200	14.8 K	21 (20)	2 (16)	-	0.34	<0.01	0.39 (53.07)	0.04 (44.57)
keller4	171	9.4 K	11*	2304 (37)	-	0.30	<0.01	<0.01 (38.11)	0.02 (5.68)
keller5	776	226 K	27*	1000 (5)	-	0.28	0.19	t/o	1421.24 (>2.53)
p-hat300-3	300	33.4 K	36*	10*	-	0.38	0.02	87.1 (9.12)	0.05 (6.00)
p-hat500-3	500	93.8 K	50*	62 (40)	-	0.34	0.07	t/o	2.51 (5.98)
p-hat700-1	700	61 K	11*	2*	-	0.36	0.03	0.08 (1.22)	0.05 (1.30)
p-hat700-2	700	121.7 K	44*	138*	-	0.36	0.45	t/o	1.35 (-)
p-hat1000-1	1 K	122.3 K	10*	276 (165)	-	0.36	0.08	0.86 (2.22)	0.71 (1.67)
p-hat1500-1	1.5 K	284.9 K	12 (11)	1 (376)	-	0.33	0.43	13.18 (-)	3.2 (1.54)
fp	7.5 K	841 K	10*	1001*	-	0.06	0.29	0.65 (-)	5.19 (1.13)
nd3k	9 K	1.64 M	70*	720*	-	0.23	0.28	t/o	7.05 (1.09)
raefsky1	3.2 K	291 K	32*	613 (362)	-	0.33	0.38	2.80 (-)	0.31 (1.36)
HFE18_96_in	4 K	993.3 K	20*	2*	<1e-4	0.26	0.27	58.88 (1.05)	4.30 (1.18)
heart1	3.6 K	1.4 M	200*	45 (26)	<1e-4	0.19	0.25	t/o	19.37 (-)
cgeb2802	2.8 K	137.3 K	60*	101 (38)	0.09	0.39	0.46	0.05 (-)	0.15 (1.61)
movielens-1m	6 K	1 M	31*	147*	0.05	0.007	0.23	31.31 (-)	2.85 (1.14)
ex7	1.6 K	52.9 K	18*	199 (127)	0.02	0.01	0.28	0.01 (-)	0.1 (1.29)
Tre14	15.9 K	2.87 M	16*	99*	0.16	0.009	0.15	3.62 (-)	0.35 (-)

The column “ ω ” is the max. clique size and the column “ n, ω ” is the number of such cliques. In both, * means the quantity is preserved in the preprocessed instance; otherwise the new quantity is in parenthesis. The multicolumns “ k -core” and “1-stage” give the vertex pruning ratio followed by the edge pruning ratio when preprocessed by removing vertices of core number $< \omega - 1$ and our preprocessor, respectively. For the last three columns, all runtimes are in seconds averaged over three independent runs. The column “Pruning” is the time for feature computation and pruning. The two remaining columns give the runtime of a solver, containing the runtime on the pruned instance with the speedup obtained in parenthesis. We denote by t/o killed execution after an hour and - denotes no speedup. Here, note that $|E|$ is the number of directed edges even though we treat the edges as undirected

Table 2 Experiments for sparse graphs

Instance	$ V $	$ E $	ω	n, ω	k -core	5-stage	Pruning	c.li.quer	MoMC	
bio-WormNet-v3	16 K	763 K	121*	18*	0.868	0.602	0.975	0.36	0.37 (-)	0.40 (3.94)
ia-wiki-user-edits-page	2 M	9 M	15*	15*	0.958	0.641	0.946	1.12	1.16 (29.94)	s
rt-retweet-crawl	1 M	2 M	13*	26*	0.979	0.863	0.997	0.38	0.41 (5.66)	s
soc-digg	771 K	6 M	50*	192*	0.969	0.496	0.998	4.80	4.91 (1.78)	s
soc-flixster	3 M	8 M	31*	752*	0.986	0.834	0.999	1.32	1.41 (3.86)	s
soc-google-plus	211 K	2 M	66*	24*	0.986	0.785	0.998	0.35	0.35 (-)	0.41 (3.98)
soc-lastfm	1 M	5 M	14*	330 (324)	0.933	0.625	0.993	2.24	2.57 (10.56)	s
soc-pokec	2 M	22 M	29*	6*	0.824	0.595	0.940	17.59	24.40 (45.80)	s
soc-themarker	69 K	2 M	22*	40*	0.713	0.151	0.972	2.03	4.95 (-)	s
soc-twitter-higgs	457 K	15 M	71*	14*	0.852	0.540	0.986	9.52	9.85 (1.92)	s
soc-wiki-Talk-dir	2 M	5 M	26*	141*	0.993	0.830	0.999	1.09	3.47 (1.25)	s
socfb-A-anon	3 M	24 M	25*	35*	0.879	0.403	0.984	28.49	38.05 (55.95)	s
socfb-B-anon	3 M	21 M	24*	196*	0.884	0.378	0.986	28.33	35.49 (67.46)	s
socfb-Texas84	36 K	2 M	51*	34*	0.540	0.322	0.957	1.04	1.07 (1.32)	s
tech-as-skitter	2 M	11 M	67*	4*	0.997	0.971	1.000	0.28	0.28 (-)	0.36 (4.31)
web-baidu-baike	2 M	18 M	31*	4*	0.933	0.618	0.992	9.67	11.00 (7.48)	s
web-google-dir	876 K	5 M	44*	8*	1.000	0.999	1.000	< 0.00	< 0.00 (-)	< 0.00 (2.06)
web-hudong	2 M	15 M	267 (266)	59 (1)	1.000	0.996	1.000	0.997	0.10 (-)	0.1 (9.99)
web-wikipedia2009	2 M	5 M	31*	3*	0.999	0.988	1.000	0.03	0.03 (-)	0.03 (4.28)

The columns are precisely as in Table 1, with the exception that we show pruning ratios for 5 stages. All ratios are rounded to three decimal places. Ratios of 1.000 are between 1 and 0.999. An s marks a segmentation fault. Here, note that $|E|$ is the number of directed edges even though we treat the edges as undirected

Table 3 The effect of introducing the feature (F10) the local chromatic density into the feature set

bio		soc		socfb		web		all	
W/o	With	W/o	With	W/o	With	W/o	With	W/o	With
0.95	0.98	0.89	0.99	0.90	0.95	0.96	0.99	0.87	0.96

The column “W/o” is the vertex classification accuracy of the classifier of Sect. 5.2 without (F10) while column “With” is the same with (F10)

Classification framework for dense networks For training, we get 4762 feature vectors from our “dense” category. As a baseline, a 4-fold cross validation over this using logistic regression results in an accuracy of **0.73**. We improve on this by obtaining an accuracy of **0.81** with gradient boosted trees (further details omitted), found with the help of `auto-sklearn` (Feurer et al. 2015).

Search strategies Given the empirical hardness of dense instances, one should not expect a very high accuracy with polynomial-time computable features such as (F1)–(F10). For this reason, we set the confidence threshold $q = 0.98$ here.

The failure of k -core decomposition on dense graphs It is common that widely-adopted preprocessing methods like the k -core decomposition cannot prune any vertices on a dense network G , even if they had the computationally expensive knowledge of $\omega(G)$. This is so because the degree of each vertex is higher than than the maximum clique size $\omega(G)$.

We showcase precisely this poor behaviour in Table 1. For most of the instances, the k -core decomposition with the exact knowledge of $\omega(G)$ cannot prune any vertices. In contrast, the probabilistic preprocessor prunes typically around 30% of the vertices and around 40% of the edges.

Accuracy Given that around 30% of the vertices are removed, how many mistakes do we make? For almost all instances we retain the clique number, i.e., $\omega(G') = \omega(G)$, where G' is the instance obtained by preprocessing G (see column “ ω ” in Table 1). In fact, the only exceptions are brock200-1 and p-hat1500-1, for which $\omega(G') = \omega(G) - 1$ still holds. Importantly, for about half of the instances, we retain *all* optimal solutions.

For some graphs, the pruning rate might appear low. This is due to the fact that (i) even a single mispredicted edge can result in the loss of several maximum cliques and (ii) in some cases, a large majority of edges are part of one or the other maximum clique, thereby limiting the number of edges that can be pruned. For instance, we achieve an edge prune ratio of 0.25 for heart1 yet still miss around half of its maximum cliques. However, on this graph instance, the maximum possible prune ratio would be around 0.184 if we were not to remove any edge part of a maximum clique. Note that in such cases, even though the pruning rate is not impressive, the solver may still gain considerable speed-up as most of the “non-clique edges” get pruned.

Speedups We show speedups for the solvers after executing our pruning strategy in Table 1 (last two columns). We obtain speedups as large as 53x and 38x for brock200-1 and keller4, respectively. This might not be surprising, since in both cases we lose

some maximum cliques (but note that for `keller4`, the size of a maximum clique is still retained). For `p-hat300-3`, the preprocessor makes no mistakes, resulting in speedups of upto 9x. The speedup for `keller5` is *at least* 2.5x, since the original instance was not solved within 3600 s, but the preprocessed instances was solved in roughly 1421 s.

Most speedups are less than 2x, explained by the relative simplicity of instances. Indeed, it seems challenging to locate dense instances of MCE that are (i) structured and (ii) solvable within a reasonable time.

5.2 Sparse networks

In this subsection, we show results for probabilistic preprocessing on sparse networks (i.e., edge density below 0.5).

Classification framework for sparse networks We use logistic regression trained with stochastic gradient descent. We use a standard L2 regularizer, and use 0.0001 as the regularization term multiplier determined by a systematic grid search. The classifier is trained for 400 epochs. Unlike for dense graphs, we did not observe a significant improvement over logistic regression by using other models such as gradient boosted trees.

Implementing the k -core decomposition Recall the exact state-of-the-art preprocessor: (i) use a heuristic to find a large clique (say of size k) and (ii) delete every vertex of G of core number less than $k - 1$. For sparse graphs, a state-of-the-art solver `pmc` has been reported to find large cliques, i.e., typically k is at most a small additive constant away from $\omega(G)$.³ Further, given that some real-world sparse networks are scale-free (many vertices have low degree) the k -core decomposition can be effective in practice.

To ensure highest possible prune ratios for the k -core decomposition method, we supply it with the number $\omega(G)$ instead of an estimate provided by any real-world implementation. This ensures *ideal conditions*: (i) the method always prunes as aggressively as possible, and (ii) we further assume its execution has zero cost. We call this method the ω -oracle.

Test instance pruning Before applying our preprocessor on the sparse test instances, we prune them using the ω -oracle. This ensures that the pruning we report is highly non-trivial, while also speeding up feature computation.

Search strategies We experiment with the following two multi-stage search strategies:

- *Constant confidence (CC)*: at every stage, perform probabilistic preprocessing with confidence threshold q .
- *Increasing confidence (IC)*: at the first stage, perform probabilistic preprocessing with confidence threshold q , progressing q by d for every later stage.

Our goal is two-fold: to find (i) a number of stages ℓ and (ii) parameters q and d , such that the strategy never errs while pruning as aggressively as possible. We do a systematic search over parameters ℓ , q , and d . For the CC strategy, we let $\ell \in \{1, 2, \dots, 8\}$ and $q \in \{0.55, 0.6, \dots, 0.95\}$. For the IC strategy, we try $q \in \{0.55, 0.60, 0.65\}$, $d = 0.05$, and set ℓ so that in the last stage the confidence is 0.95.

³ A table of results seen at <http://ryanrossi.com/pmc/download.php>.

We find the CC strategy with $q = 0.95$ to prune the highest while still retaining all optimal solutions. Thus, for the remaining experiments, we use a CC strategy with $q = 0.95$.

Our 5-stage strategy outperforms, almost always safely, the ω -oracle (see Table 2). In particular, note that even if the difference between the vertex pruning ratios is small, the impact for the number of edges removed can be considerable (see e.g., all instances of the “soc” category). We note that the runtime is not sensitive to the number of stages ℓ . In fact, already the first step of pruning makes the graph so small that further stages add comparatively very small amounts to the overall runtime.

Speedups We show speedups for the solvers in Table 2. We use as a baseline the solver executed on an instance pruned by the ω -oracle, which renders many of the instances easy already. Most notably, this is *not* the case for soc-pokec, socfb-A-anon, and socfb-B-anon, all requiring at least 5 min of solver time. The largest speedup is for socfb-B-anon, where we go from requiring 40 min to only 7 s of solver time. For MoMC, most instances report a segmentation fault (marked with an s) for an unknown reason.

5.3 Edge-based classification

For edges, we do a similar training as that described for vertices. For the category “dense”, we obtain 79472 feature vectors. Further, for this category, the edge classification accuracy is **0.83**, which is 1% higher than the vertex classification accuracy using the same classifier as in Sect. 5.1. However, we note that the edge feature computation is noticeably slower than that for vertex features. This reason combined with the fact that the classification accuracy is almost the same, we omit further experiments with the edge features due to smaller speedups.

5.4 Model analysis

Gradient boosted trees (used with dense networks in Sect. 5.1) naturally output feature importances. We apply the same classifier for the sparse case to allow for a comparison of feature importance. In both cases, the importance values are distributed among the ten features and sum up to one (See Fig. 3).

Unsurprisingly, for sparse networks, the local chromatic density (F10) dominates (importance 0.22). In contrast, (F10) is ineffective for dense networks (importance 0.08), since the chromatic number tends to be much higher than the maximum clique size. In both cases, (F5) eigencentrality has relatively high importance, justifying its expensive computation.

For dense networks, (F7) average χ^2 over neighbor degrees has the highest importance (importance 0.23), whereas in the sparse case it is least important feature (importance 0.03). This is so because all degrees in a dense graph are high and the degree distribution tends to be tightly bound or coupled. Hence, even slight deviations from the expected (e.g., vertices in large cliques) depict high statistical significance scores. We will capitalize on this observation later on in Sect. 7.




Fig. 3 The feature importance for **a** dense nets (category “dense”) and **b** sparse nets (category “all”)

6 On supervised learning for finding cliques

The goal of this section is two-fold: (i) to explain the high accuracy of our proposed framework, even when it was trained with small instances, and (ii) as a consequence, argue that supervised learning is a viable approach for solving structured instances of certain hard problems.

To ensure that the input instances are, at some point, “structure-free” we turn to the following heavily-studied variant of the maximum clique problem. This serves as a representative of the *worst-case input* for our preprocessing strategy. Also, observe that in case the input graph has a unique maximum clique, MCE is equivalent to finding the (single) maximum clique. For simplicity, we restrict ourselves to single stage sparsification in these experiments.

6.1 Planted clique

In the *planted clique problem* (Jerrum 1992; Kučera 1995), we are given an Erdős-Rényi random graph $H := G(n, p)$, i.e., an n -vertex graph where the presence of each edge is determined independently with probability p (see Erdős and Rényi 1959). In addition, the problem is parameterized by an integer k such that a random subset of k vertices has been chosen from H and a clique added on it. On this input, the task is to identify (with the knowledge of the value of k) the k vertices containing the planted clique.

The problem is easy for $k \leq \log_2(n)$. In particular, as shown in Bollobás (2013), the clique number of $G(n, p)$ as $n \rightarrow \infty$ is almost surely w or $w + 1$ where w is the greatest natural number such that

$$\binom{n}{w} p^{\binom{w}{2}} \geq \log(n), \quad (2)$$

where w is roughly $2 \log_2(n)$. Even when a clique of such size is known to exist (whp), we only know how to find a clique of size $\log_2(n)$ efficiently,⁴ and also solve the problem in polynomial-time when k is large enough. Specifically, it is known that several algorithmic techniques such as spectral methods [see e.g., Feldman et al. (2017) for more] produce efficient algorithms for the problem when $k = \Omega(\sqrt{n})$.

However, settling the complexity of the problem is a notorious open problem when k is between $2 \log_2(n)$ and \sqrt{n} . Next, we will focus precisely on this difficult region.

6.2 Pushing the limits of preprocessing

In this subsection, we explore the limits of scalability and robustness of our framework on the planted clique problem. All experiments are done on an Intel Core i5-6300U CPU (2.4 GHz), 8 GB of RAM, running Ubuntu 16.04, differing only slightly from the earlier hardware configuration. For all experiments here, we use only the `igraph` algorithm.

Generation of synthetic data We use the `genrang` utility program (McKay and Piperno 2014) to sample a random graph $H := G(n, p)$. To plant a clique of size k , we sample uniformly at random k vertices, denoted by K , from H and insert all corresponding at most $\binom{k}{2}$ missing edges into H .

For each H , we compute the features described in Sect. 4 with the following differences: we replace (F10) the local chromatic density with the order-four LCC and modify (F8) and (F9) to consider order-four LCC instead of the LCC. This brings more predictive power while still remaining computationally feasible for small graphs. The values E_i in Eq. 1 for (F6) and (F7) are the expected degree $n \cdot p$, while for (F8) and (F9) they are the expected order- k LCC given as $\binom{n-1}{k-1} p^{\binom{k}{2}} \cdot 1 / \binom{np}{k-1}$. To ensure a balanced dataset, we sample (i) k label-0 examples from K and (ii) k label-1 examples from $G \setminus K$, both uniformly at random.

For training, we consider $n \in \{64, 128, 256, 512\}$ because the clique number grows roughly logarithmically with n (see Eq. 2). We fix $p = 1/2$. For every n , we compute w from Eq. 2, and sample graphs $G(n, p)$ with a planted clique of size $k \in \{w + 2, \dots, w + 6\}$ such that each pair (n, k) gives a dataset of size at least 100 K feature vectors. When planting a clique of size at least $w + 2$, we try to guarantee the existence of a unique maximum clique in the graph. However, this procedure does not always succeed due to randomness, but we do not discard such rare outcomes.

Vertex classification accuracy We study the accuracy of our classifiers for distinguishing vertices that are and are not in a maximum clique. Specifically, we train a classifier for each pair $(n, k) \in \{(64, 10), (128, 12), (256, 12)\}$, and test for unseen graphs with the same n but growing planted clique size $k' = k + 1, \dots, k + 4$. The results are shown in Fig. 4a. As expected, the classification task becomes easier once k' increases. This

⁴ It is conjectured (Karp 1976; Feldman et al. 2017) that there is no polynomial-time algorithm for finding a clique of size $(1 + \varepsilon) \log_2 n$ for any $\varepsilon > 0$ in $G(n, 1/2)$.




Fig. 4 The vertex accuracy, pruning ratio, and clique accuracy of our framework when trained with $G(n, 1/2)$ with three different parameter pairs (64, 10), (128, 12), and (256, 13). The predictions are for independent, distinct samples with the same n , but growing planted clique size k

is also supported the fact that multiple algorithms solve the planted clique problem in polynomial-time for large enough k' (see Sect. 6.1). In addition, as n grows larger, we see accuracy deterioration caused by the converge of the local properties towards their expected values. For example, since an edge is present with probability p in $G(n, p)$, the expected degree of a vertex is $n \cdot p$. As n grows large, it will become increasingly difficult for a classifier to learn to remove vertices based on degree, as degree deviations tend to zero providing little signal for classification. Especially for small values of k' , the injection of the planted clique is not substantial enough to cause significant deviations from the expected values.

Pruning ratio and clique accuracy We study the effectiveness of our framework as a probabilistic preprocessor for the planted clique instances. We fix the confidence threshold $q = 0.55$ and use the same set of classifiers and test data. The average pruning ratios over all instances are shown in Fig. 4b. We see pruning ratios as high as at most 0.6, while always discarding more than 40% of the vertices.

Now, it is possible that P makes an erroneous prediction causing the deletion of a vertex, which in turn lowers the size of a maximum clique in the instance. The average clique accuracies over all instances are shown in Fig. 4c. Here, we see that for $n = 256$, the vertex accuracy (Fig. 4a) is still above 0.7, but the clique accuracy drops to above 0.4. As the vertex accuracy decreases, the probability of deleting a vertex present in a maximum clique increases, translating to a higher chance of error in extracting a maximum clique. However, while not completely error-free, we observe that even in the case of (256, 13) we always delete at most two members of a maximum clique, whereas in the case of (512, 16), 95% of the time, we extract a maximum clique of size at least 13 (see Fig. 5).

Robustness and speedups The robustness and speedups obtained using the `igraph` algorithm are given in Table 4. Here, the clique accuracy and runtime are obtained as the average over 200 samples for each (n, k) except for $(n, k) = (512, 15)$ for which there are 20 independent samples. We see the drop in clique accuracy when a classifier P is trained with $(n, k) \in \{(256, 12), (512, 15)\}$ and is predicting for the same n but increasing k . The clique accuracy is a strict measure, so to quantify the severeness of the erroneous predictions made by P we show the distributions of the extracted maximum clique sizes in Fig. 5 for some pairs (n, k) . Again, we observe the effects of growing n causing the convergence of local properties, consequently decreasing the




Fig. 5 Distribution of extracted maximum clique size, with black bars denoting the size of the planted clique. Both **(a, b)** are over 200 samples, while **c** is over 20 samples. In each, the predicting classifier has been trained with 64-vertex random graphs with a planted clique of size 10

predictive power of P . For $(n, k) = (128, 13)$, 73% of the runs still produce an optimal solution (here, one can also observe the rare event of having a maximum clique of size 14 when the planted clique was of size 13).

We observed that in such scenarios, the performance of the other baseline heuristic solvers were lower in terms of accuracy, as well as suffered from higher runtime. In fact, in certain cases (both for real and synthetic datasets), the heuristic solvers takes longer and also extracted a maximum clique of a much smaller size, as compared to our framework. Such performance characteristics and comparison can be observed in Fig. 6b–e.

The case for supervised learning on intractable problems As n grows, the instances get increasingly time-consuming to solve even for state-of-the-art solvers for suitable k , as there is no exploitable structure. Consequently, obtaining optimally labeled data becomes practically impossible for large enough n . However, in our experiments, we find that random graphs with $n = 64$ and $k = 10$ are representative of the input for moderately larger graphs as well, up to a point. Further, obtaining the optimal label for such small graphs is fast.

We show the deviation in vertex classification accuracy in Table 5. The column “Trained acc.” corresponds to the accuracy of the classifier trained with the values n and k mentioned in the two first columns, while the column “Rob. acc.” is the accuracy of a classifier trained with smaller $(n, k) = (64, 10)$ instances, and predictions are made for the specified (n, k) with planted clique size $k + 1$. A key observation is that the difference between the two accuracies in a single row in Table 5 is small enough not to warrant training on larger instances. This offers an explanation for the perfect clique accuracy with limited training, observed earlier for sparse real-world networks. This observation reduces the need of labeling costly data points for training.

Discussion. We showcase how a relatively lightweight machine learning approach (e.g., gradient boosted trees) for vertex/edge pruning can provide significant improvements in the performance of existing solvers. In fact, deep learning techniques like *graph neural networks* (GNN) (Scarselli et al. 2009) might also be an alternative in this regard, and could potentially remove the necessity of computing hand-crafted features [for an overview, see Zhou et al. (2020)]. While such graph learning techniques in the context of combinatorial optimization have been recently explored (Mirhoseini et al. 2021), we steer away from such GNN architectures owing to the following reasons:

Table 4 Robustness and speedups with fixed n and increasing k

n	k	$k + 1$			$k + 2$			$k + 3$					
		Pr.	Acc.	Time	Speedup	Pr.	Acc.	Time	Speedup	Pr.	Acc.	Time	Speedup
64	10	0.530	0.905	0.068	0.132	0.548	0.965	0.068	0.135	0.564	0.995	0.068	0.135
128	12	0.506	0.710	0.301	0.759	0.517	0.875	0.296	0.774	0.525	0.935	0.297	0.784
256	13	0.489	0.170	3.261	3.264	0.493	0.190	3.233	3.304	0.493	0.310	3.260	3.315
512	15	0.492	0.05	70.587	12.994	0.492	0.05	70.086	12.816	0.491	0.100	70.562	12.722

The leftmost two columns show the data (n, k) used to train a classifier P . For each planted clique size $k + 1, k + 2,$ and $k + 3$, we show the average pruning ratio (column “Pr.”), the average clique accuracy (column “Acc.”), the average runtime of `igraph` on the reduced instance obtained from our framework using P (column “Time (s)”), and the average speedup over executing the same algorithm on the original instance

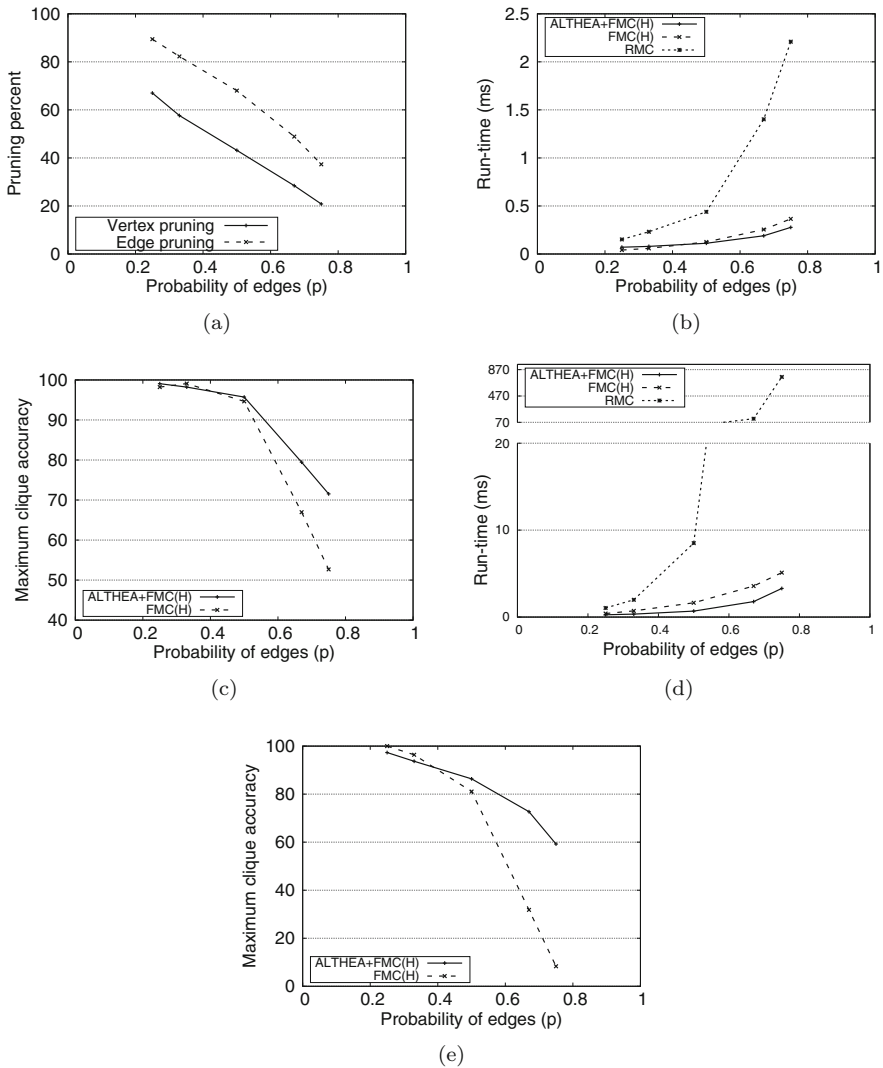


Fig. 6 Performance comparison of competing approaches on ER random graphs $G(64, p)$ with varying density p based on **a** vertex and edge pruning rates, **b** run-time, and **c** maximum clique accuracy; and on $G(256, p)$ with varying density p , based on **d** run-time and **e** maximum clique accuracy

- Training data limitation—Deep learning techniques rely on a large amount of training data for the model training to converge, as well as for the model to capture fine-grained patterns present in the data. Such large amounts of training data might not be readily available in many application domains, particularly involving computationally difficult problems, limiting such approaches;
- Explainability—Deep learning architectures are traditionally *black boxes*, and do not provide any explainability as to why the corresponding output was generated. This poses a bottleneck in many scenarios (e.g., active learning), and hence

Table 5 Deviation in vertex classification accuracy

n	k	Trained acc.	Rob. acc.
128	12	0.858	0.844
256	13	0.747	0.728
512	15	0.678	0.665

- provides limited opportunity to analyze the results and make subsequent modifications to the learning model. Further, in cases of evolving constraints in the applications, model explainability would help better understanding of feature space addition/modification to suit the needs. In this regard, we use gradient boosted trees, which provide the importance values of each feature (for the decision of vertex/edge pruning), enabling explainability along with performance improvements;
- Generalizability—GNNs have been shown to be difficult to generalize beyond the training data (Knyazev et al. 2019; Yehudai et al. 2021). Although transfer learning provides an avenue to mitigate the problem, the performance of such techniques are low in general. This, compounded with the need for large training data, renders such approaches infeasible for diverse application domains; and
 - Computation time—GNNs typically suffer from comparatively large inference time (i.e., time required to predict to prune vertices/edges individually), which might be infeasible for web-scale graphs (e.g., biological and social network graphs). Thus, the use of GNNs (as a preprocessing step) along with an existing solver might increase the overall runtime (than directly applying an existing heuristic based solver).

Thus, based on the above, we argue that a lightweight machine learning approach suits better the problem at hand, as depicted in this work. However, for future work, we aim to investigate GNN approaches to the problem of finding maximum cliques. For a comprehensive review on GNN applications in this domain, see Cappart et al. (2021).

7 ALTHEA: maximum clique-finding heuristic for dense graphs

In this section, we capitalize on the observation we made in Sect. 5.4. In particular, we describe a heuristic we call *ALTHEA* (*mAximum cLique exTraction using cHi-square stAtistics*) for extracting an approximate maximum clique from a *simple* input graph $G = (V, E)$.

7.1 Description of ALTHEA

ALTHEA hinges on categorizing the *degree* of each vertex in G based on its deviation from the average degree of G . Each vertex is subsequently represented by a sequence of category symbols encoding its neighbourhood, which are then used for computing its statistical significance score. Any vertex depicting the maximum χ^2 value (along with its neighbourhood) forms a candidate region for containing a maximum clique in G . *ALTHEA* comprises the following five steps.

1. *Initialization* We compute the following three degree characteristics of G .

- (i) $\Delta(G)$: the maximum degree of any vertex in G ,
- (ii) $a(G)$: the average degree of the vertices in G ; and
- (iii) $\sigma(G)$: the standard deviation of the vertex degrees of G .

Formally, we define

$$\Delta(G) = \max\{\deg(v) : \forall v \in V\}, \tag{3}$$

$$a(G) = \frac{\sum_{v \in V} \deg(v)}{|V|}, \tag{4}$$

and

$$\sigma(G) = \sqrt{\frac{\sum_{v \in V} (\deg(v) - a(G))^2}{|V| - 1}}, \tag{5}$$

where $\deg(v)$ is degree of vertex v .

2. *Symbol categorization* *ALTHEA* captures the nature of vertex degree deviation (in the number of standard deviations) from the underlying degree distribution of G . The number of *category symbols* τ^G is $\lceil (\Delta(G) - a(G))/\sigma(G) \rceil + 1$. The obtained set of category symbols is $\Gamma^G = \{\gamma_1, \gamma_2, \dots, \gamma_{\tau^G}\}$, where γ_i is the multiple of $\sigma(G)$ by which the degree of v_i deviates from $a(G)$. Next, we compute the expected probability of occurrence for the symbols in Γ^G .

To this end, we use *Chebyshev's inequality* (Tchebichef 1867), which for a random variable X and a real number $k > 0$ states that $\Pr(|X - \mu| \geq k\delta) \leq 1/k^2$. Here, μ and δ are the mean and standard deviation, respectively, of the distribution from which the random variable X is drawn. Thus, the occurrence probability of the symbol γ_i is given by $\Pr(\gamma_i) = 1/i^2 - 1/(i + 1)^2$.

Other tail distribution bounds or domain-dependent probability distributions capturing the underlying characteristics of G might also be used depending on the application. This makes *ALTHEA* robust to diverse domains, applicable to different input distributions.

3. *Vertex symbol sequence* For each vertex $v \in V$, we extract its *closed neighbourhood* $N[v] = \{v\} \cup \{u : \{u, v\} \in E\}$. The vertex v is then represented by a *sequence of category symbols* $\text{Seq}(v)$ of length $|N[v]|$ based on the symbol categorization of the degree of the vertices in its neighbourhood $N[v]$. Formally, this is given as

$$\text{Seq}(v) = \{\gamma(u) : \forall u \in N[v]\}, \tag{6}$$

where $\gamma(u)$ is the unique $\gamma_i \in \Gamma^G$ and $i \in \{1, 2, \dots, \tau^G\}$, for which the inequality $i \leq (\deg(u) - a(G))/\sigma(G) + 1 < i + 1$ holds.

4. *Statistical significance computation* For each vertex v , *ALTHEA* computes the χ^2 statistical significance score using $\text{Seq}(v)$ and the associated symbol probabilities. For each category symbol $\gamma_i \in \Gamma^G$, its expected occurrence count for vertex v is computed as $E_{\gamma_i}^v = \Pr(\gamma_i) \times |N[v]|$. Similarly, the corresponding *observed occurrence count*

$O_{\gamma_i}^v$ of the category symbol γ_i for v can be obtained from $\text{Seq}(v)$. Combining the above steps, the statistical significance of v is

$$\chi^2(v) = \sum_{\forall \gamma_i \in \Gamma^G} \frac{(O_{\gamma_i}^v - E_{\gamma_i}^v)^2}{E_{\gamma_i}^v}. \quad (7)$$

5. Approximate maximum clique extraction After computing the statistical significance of the vertices, *ALTHEA* selects the vertex v' demonstrating the maximum statistical significance (chosen arbitrarily in case of ties), as the best candidate whose neighbourhood contains an (approximate) maximum clique for G . Intuitively, a vertex and its neighbours that are a part of a maximum clique in G would exhibit the largest variation in the degree distribution characteristic compared to the average (or expected) characteristic of G , which is captured by the notion of statistical significance. Finally, the subgraph induced by the neighbourhood $N[v']$ is fed to a maximum clique solver for extracting a large clique of G .

Discussion In a dense graph, the degree of any vertex is high, and the degree distribution tends to be tightly bound (or coupled). Hence, even slight deviations from the expected behaviour (in cases of vertices that are a part of large cliques) depict high statistical significance scores. This enables *ALTHEA* to effectively identify large maximum cliques, as we will experimentally show next.

7.2 Experimental evaluation

Baselines We benchmark the performance of *ALTHEA* against the following existing state-of-the-art approaches:

- (i) *igraph* (Csardi and Nepusz 2006) C library's implementation of the exact modified Bron-Kerbosch algorithm (Bron and Kerbosch 1973),
- (ii) *MoMC* (Li et al. 2017) employing a branch-and-bound pruning strategy,⁵
- (iii) *FMC (E)* (Pattabiraman et al. 2015) using exact hierarchical pruning strategy,⁶
- (iv) *FMC (H)*—the fast heuristic variant of *FMC(E)*,
- (v) *RMC* (Lu et al. 2017)—randomized heuristic based on “binary search” with optimum-bounding (obtained from the authors), and
- (vi) *PMC* (Rossi and Gleich 2015; Rossi et al. 2013)—parallelized heuristic based on vertex colouring and k-core computation.⁷

By definition, the exact algorithms give the maximum clique size ω .

Note that the final pruned subgraph obtained by *ALTHEA* is presented to a maximum clique solver. We couple *ALTHEA* with either the exact *MoMC* solver, or the fast *FMC(H)* heuristic. Such a coupling is denoted as e.g., *ALTHEA+MoMC* or *ALTHEA+FMC(H)*, respectively. The approaches are evaluated on run-time efficiency

⁵ Available from URL: www.home.mis.u-picardie.fr/~cli/EnglishPage.html.

⁶ Available from URL: www.cucis.ece.northwestern.edu/projects/MAXCLIQUE/.

⁷ Available from URL: www.github.com/ryanrossi/pmc.

and accuracy of extracting a maximum clique. Our implementation of *ALTHEA* is in C, and all experiments are run on an Intel Xeon E5-2680 CPU (2.80 GHz) with 8 cores and 32 GB of RAM.

7.3 Real datasets

We experiment on structured datasets from diverse domains such as biological networks, financial graphs, social interaction and blog conversations. Again, our instances are obtained from Network Repository (Rossi and Ahmed 2015).

Easy instances We selected 17 dense graphs (see Table 6) with varying sizes of upto 30 K vertices and 1 M edges. Table 6 reports the vertex and edge pruning achieved by *ALTHEA+FMC(H)* in addition to run-time and the maximum clique size extracted. We see that *ALTHEA* is highly accurate in identifying regions that contain a maximum clique. In fact, it is successful in extracting an optimal maximum clique in 13 of the instances, while in the remaining 4 instances, it extracts larger cliques than standalone *FMC(H)*.

We observe that *ALTHEA* aggressively prunes the search space (with high accuracy), achieving vertex and edge prunings as high as 99%—with more than 80% vertex/edge pruning on 11 instances. This enables our framework to be very efficient in practice, showcasing consistent speedups of around 3× compared to the best performing heuristic and upto 10× with respect to the exact algorithms. On the other hand, *RMC* is able to extract the maximum clique size in nearly all the instances, but suffers from large run-time in general (compared to other heuristics), owing to its dependency on vertex coloring and independent set computation. We also observe the performance of our approach to be consistently comparable with that of *PMC* in terms of finding the maximum clique in the graph. However, in terms of run-time *PMC* based on a single execution thread performs worse in most of the instances. However, considering the inherent parallel nature, *PMC* provides the best run-time among the approaches when executed with at least 3 or 4 threads. In fact, for very high density graphs with large clique sizes, the pruning factor offered by *ALTHEA* outweighs the parallel nature of *PMC*, when compared in terms of run-time improvement. For example, in Table 6, the *moreno-blogs* dataset has an average density of around 1K with a maximum clique of size 1490. For this instance, *ALTHEA* (even with a small pruning ratio) offers a run-time improvement of 10× compared to *PMC* with 8 execution threads.

Hard instances We select 8 additional hard instances, on which exact algorithms were unable to run to completion with a timeout of 5 min. Table 7 tabulates these instances and the performance of the competing approaches. Here, we also evaluate the performance of *ALTHEA* when coupled with the exact *MoMC* solver.

Similar to our previous observations from Table 6, *ALTHEA+FMC(H)* again performs better than the standalone *FMC(H)* heuristic, and extracts **better solutions**. Further, vertex and edge pruning (of around 40% on average) gives *ALTHEA* faster run-times with upto 5× speedups over *FMC(H)*. Again, *RMC* requires high computation time but extracts larger cliques.

Table 6 Performance comparison of *ALTHEA* on real-world datasets

Instance	Dataset characteristics			Heuristic approaches										Exact approaches					
	V	E	ω	ALTHEA+FMCH					FMCH		RMC		PMC		FMC(E)		MoMC		igraph
				Vert. Pr.	Edge Pr.	Time (s)	ω'	Time (s)	ω'	Time (s)	ω'	Time (s)	ω'	Time (s)	ω'	Time (s)	ω'	Time (s)	
bio-WormNet-v3-b	2 K	79 K	126	94.48	89.83	0.004	126	0.015	126	0.001	126	0.002	126	0.006	0.496	0.239			
bn-macaque	93	2 K	30	68.82	82.05	0.000	30	0.000	30	0.001	30	0.001	30	0.000	0.004	0.001			
bn-mouse_retina_I	1 K	91 K	51	77.74	80.30	0.013	39 [#]	0.044	35	0.184	51	0.165	48	-	0.26	-			
cari	1 K	77 K	200	68.28	11.19	0.783	200	0.883	200	0.176	200	0.097	200	-	0.656	4.933			
cavity26	5 K	71 K	19	98.66	99.00	0.002	19	0.009	19	0.038	19	0.04	19	0.017	1.148	0.132			
econ-psmigr1	3 K	411 K	144	90.16	92.20	0.138	116	0.728	114	-	-	0.455	140	-	1.38	-			
frb30-15-1	451	83 K	30	17.29	32.21	0.073	25	0.133	25	1.008	25	0.015	25	-	0.324	-			
hamming10-2	1 K	519 K	512	1.171	2.14	39.70	512	42.00	512	11.158	512	1.01	512	-	29.188	-			
light_in_tissue	29 K	188 K	6	99.94	99.97	0.008	6	0.012	6	10.894	6	0.053	6	0.031	-	0.63			
nasa2910	3 K	86 K	36	96.6	97.48	0.003	36	0.036	36	0.027	36	0.014	36	0.479	0.64	0.31			
robot24c1_mat5_J	405	14 K	24	88.89	95.49	0.001	21	0.003	20	0.003	24	0.004	24	-	0.016	0.53			
sec_enron-only	152	10 K	120	7.89	4.92	0.04	120	0.037	120	0.002	120	0.001	120	-	0.336	0.022			
sec_infect-dublin	11 K	176 K	84	99.12	97.57	0.012	84	0.032	84	0.021	84	0.002	84	2.715	-	0.864			
sec_twitter-copen	9 K	474 K	581	87.74	16.80	26.223	581	30.967	581	10.267	581	0.85	581	-	38.004	-			
Trec12	3 K	151 K	11	91.71	98.46	0.004	8	0.038	8	-	-	0.02	9	21.334	0.756	8.926			
polblogs	2 K	17 K	20	95.44	93.74	0.001	19 [#]	0.004	16	0.005	20	0.004	20	0.186	0.204	0.503			
moreno-blogs	1 K	1109 K	1490	0.134	0.134	0.405	1490	0.369	1490	16.75	1490	16.5	1490	1.645	-	0.24			

Here, (i) ω denotes the maximum clique size and ω' is the approximate maximum clique size found; (ii) results for approaches that were “killed” after 5 min of run-time (without output) are marked with -; (iii) for results marked with #, refer to Table 7 for additional results; (iv) averaged run-times over 5 runs are shown in sec; (v) vertex and edge pruning (Pr.) are in percentage of |V| and |E| resp.; and (vi) PMC results obtained with 1 thread. For reasons of space, the instance names bio-WormNet-v3-benchmark has been shortened to bio-WormNet-v3-b, and bn-macaque-rhesus_inter-cort-netw_2 to bn-macaque

Table 7 Performance comparison of *ALTHEA* on difficult real-world data

Instance	V			E			ALTHEA+PMC(H)			ALTHEA+MoMC			FMC(H)			RMC			PMC		
							Vert. Pr.	Edge Pr.	Time (s)	ω'	Time (s)	ω'	Time (s)	ω'	Time (s)	ω'	Time (s)	ω'	Time (s)	ω'	
bio-WormNet-v3	16K	763K	96.23	91.85	0.091	94	0.166	121	0.465	90	0.156	121	0.12	115							
brock800-1	801	208K	35.21	58.08	0.088	17	-	19*	0.289	17	-	17*	0.041	17							
C1000-9	1001	450K	7.69	14.64	2.41	51	-	59*	3.065	51	-	53*	0.144	53							
econ-psmigr1	3141	411K	90.16	92.20	0.138	116	0.615	122	0.728	114	-	-	0.455	140							
frb30-15-1	451	83K	17.29	32.21	0.073	25	0.866	29	0.133	25	-	25*	0.015	25							
frb50-23-5	1151	581K	16.33	29.77	2.1	42	-	48*	3.235	41	-	40*	0.162	41							
frb53-24-5	1273	714K	5.81	11.35	3.863	42	-	49*	4.65	42	-	42*	0.181	42							
p-hat1500-3	1501	847K	17.19	30.40	2.725	60	-	91*	4.589	60	-	60*	0.355	81							
bn-mouse_retina_1	1123	91K	77.74	80.30	0.012	39	0.026	51	0.044	35	0.184	51	0.163	48							
polblogs	1491	17K	95.44	93.74	0.001	19	0.197	20	0.004	16	0.005	20	0.004	20							

Here, (i) results marked with * denote the maximum clique size found by the heuristic before the cut-off time of 5 min.; (ii) RMC reported segmentation fault for the *econ-psmigr1* network, and is marked with -; (iii) average run-times over 5 runs are shown in sec.; (iv) vertex and edge pruning (Pr.) are given in % of |V| and |E|, respectively; and (v) PMC results obtained with 1 thread

From Table 7, we see that the pruning strategy of *ALTHEA* with *MoMC* provides an interesting trade-off between solution quality and run-time. This approach is able to identify significantly better solutions compared to others, in all instances. In fact, for the last two instances in Table 7 (also in Table 6), we are now able to extract the optimal solution. Although *ALTHEA+MoMC* consumes slightly more run-time (than *FMC(H)*), it is still faster than *RMC*. It is interesting to note, that *ALTHEA+MoMC* depicts a far better performance quality (in terms of maximum clique extraction) compared to *PMC*. However, *PMC* is much faster than the other heuristics, even with only a single execution, and provides a fast approximation, albeit with lower accuracy.

To summarize, we see that *ALTHEA* provides an efficient and robust pruning strategy for finding an approximate maximum clique with high accuracy in dense real-life graphs from diverse domains, providing an effective trade-off between quality with slight increase in run-time. Further, as is well-known, such dense instances constitute a major challenge for state-of-the-art solvers.

7.4 Synthetic datasets

We turn to study the robustness of *ALTHEA* on Erdős-Rényi (ER) random graphs, denoted as $G(n, p)$, which is an n -vertex graph where every edge is present with independent probability p . We observe the pruning ratio, run-time and accuracy of the approaches, by varying the two parameters n and p . Particularly for $p \geq 0.5$, random graphs present a challenging benchmark for pruning. Hence, we relax the accuracy measure by considering a heuristic *accurate* if the size of the clique returned is at most 1 less than the optimum.

Graph density The effect of density on the performance of the approaches is shown in Fig. 6a–c obtained on ER-graphs with 64 vertices with varying density of $p \in \{0.25, 0.33, 0.5, 0.67, 0.75\}$. In terms of pruning rate, we observe in Fig. 6a that *ALTHEA* effectively prunes nearly 50% of the edges (and vertices) even in dense random graphs (i.e., $p = 0.5$). However, the pruning rate decreases linearly with increase in density (to around 20% for $p = 0.75$). The high pruning rate enables *ALTHEA* (coupled with *FMC(H)* heuristic) to be superior than the other approaches in terms of run-time demonstrating upto $1.5\times$ speedups compared to the standalone *FMC(H)*. Similar to the real datasets, *RMC* suffers from high run-time (upto $10\times$ slower). Interestingly, we observe that *ALTHEA* exhibits higher accuracy compared to *FMC(H)* (Fig. 6c). For $G(64, 0.75)$, we report an accuracy of more than 70% compared to around 50% for *FMC(H)*. The accuracy of *FMC(H)* is seen to degrade significantly as density increases. For low density graphs (i.e., $p < 0.5$), both heuristics perform similarly. *RMC* has perfect accuracy, but infeasible running times for larger and denser graphs.

Graph size We assess the effect of varying n on the performance of *ALTHEA*. Figure 6d, e present the results for $n = 256$.

The approaches are seen to exhibit similar behaviour as above, with high pruning rates for *ALTHEA*, along with a large speedup in extracting large cliques compared to *FMC(H)*. From Fig. 6e, we observe that our approach depicts significantly superior

accuracy (compared to $FMC(H)$)—being nearly $6\times$ more accurate in identifying a maximum clique in dense input graphs. Finally, we remark that similar results were observed on ER-graphs for other parameter values of n and p , but omit further details.

8 Conclusions

We have proposed a novel framework for learning to scale-up combinatorial optimization algorithms. In contrast to the existing learning frameworks that use difficult-to-interpret learning models to learn the exact decision boundary, our proposed framework relies on interpretable learning models with local features to prune the elements that are not in any optimal solution(s). The deeper insights learned by our multi-stage pruning framework result in the identification of feature combinations relevant to the optimization problem and the instance class. This can result in better heuristics for the problem, as evidenced by our heuristic for the maximum clique detection problem.

Our framework has been designed primarily for combinatorial optimization problems that involve finding an optimal subset of elements. A crucial direction for future research is to explore if this framework can be extended to deal with combinatorial optimizations involving ordering and assignment problems. Other avenues for future research include the design of approaches to improve the accuracy of the pruning by incorporating problem constraints in the learning process.

Funding There are no sources of funding to be listed.

Data availability All data is freely available (see the corresponding references and links in the manuscript proper).

Code availability All code will be made publicly available upon acceptance, except for the code related to Sect. 5.2 which is Nokia Bell Labs proprietary and cannot be shared.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

- Abu-Khzam, F.N., Baldwin, N.E., Langston, M.A., Samatova, N.F.: On the relative efficiency of maximal clique enumeration algorithms, with applications to high-throughput computational biology. In: International Conference on Research Trends in Science and Technology (2005)
- Adamic, L.A., Adar, E.: Friends and neighbors on the Web. *Soc. Netw.* **25**(3), 211–230 (2003)
- Amizadeh, S., Matuselych, S., Weimer, M.: PDP: A general neural framework for learning constraint satisfaction solvers. *CoRR* [arXiv:1903.01969](https://arxiv.org/abs/1903.01969) (2019)
- Batson, J., Spielman, D.A., Srivastava, N., Teng, S.-H.: Spectral sparsification of graphs: theory and algorithms. *Commun. ACM* **56**(8), 87–94 (2013)
- Bello, I., Pham, H., Le, Q.V., Norouzi, M., Bengio, S.: Neural combinatorial optimization with reinforcement learning. In: International Conference on Learning Representations (ICLR), Workshop Track (2017)
- Bengio, Y., Lodi, A., Prouvost, A.: Machine learning for combinatorial optimization: a methodological tour d’horizon. *Eur. J. Oper. Res.* **290**(2), 405–421 (2021)

- Bernard, H.R., Killworth, P.D., Sailer, L.: Informant accuracy in social network data IV: a comparison of clique-level structure in behavioral and cognitive network data. *Soc. Netw.* **2**(3), 191–218 (1979)
- Boginski, V., Butenko, S., Pardalos, P.M.: Statistical analysis of financial networks. *Comput. Stat. Data Anal.* **48**(2), 431–443 (2005)
- Bollobás, B.: *Modern Graph Theory*, vol. 184. Springer, Berlin (2013)
- Bomze, I., Budinich, M., Pardalos, P., Pelillo, M.: The maximum clique problem. In: *Handbook of Combinatorial Optimization*, vol. 4, pp. 1–74. Kluwer Academic Publishers (1999)
- Bron, C., Kerbosch, J.: Algorithm 457: finding all cliques of an undirected graph. *Commun. ACM* **16**(9), 575–577 (1973)
- Cappart, Q., Chételat, D., Khalil, E.B., Lodi, A., Morris, C., Veličković, P.: Combinatorial optimization and reasoning with graph neural networks. In: *IJCAI*, pp. 4348–4355 (2021)
- Caro, G.A.D.: A survey of machine learning for combinatorial optimization. In: *30th European Conference on Operational Research* (2019)
- Chalumeau, F., Coulon, I., Cappart, Q., Rousseau, L.-M.: SeaPearl: a constraint programming solver guided by reinforcement learning. In: *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pp. 392–409. Springer, Berlin (2021)
- Chen, J., Huang, X., Kanj, I.A., Xia, G.: Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.* **72**(8), 1346–1367 (2006)
- Chmiela, A., Khalil, E.B., Gleixner, A.M., Lodi, A., Pokutta, S.: Learning to schedule heuristics in branch-and-bound. *CoRR* [arXiv:2103.10294](https://arxiv.org/abs/2103.10294) (2021)
- Csardi, G., Nepusz, T.: The igraph software package for complex network research. *InterJournal Complex Syst.* **1695**, 1–9 (2006)
- Cygan, M., Fomin, F.V., Kowalik, Ł., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: *Parameterized Algorithms*. Springer, Berlin (2015)
- Dutta, S., Lauri, J.: Finding a maximum clique in dense graphs via χ^2 statistics. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, pp. 2421–2424 (2019). ISBN 978-1-4503-6976-3
- Dutta, S., Nayek, P., Bhattacharya, A.: Neighbor-aware search for approximate labeled graph matching using the chi-square statistics. In: *International Conference on World Wide Web (WWW)*, pp. 1281–1290 (2017)
- Eblen, J.D.: *The Maximum Clique Problem: Algorithms, Applications, and Implementations*. Ph.D. thesis, The University of Tennessee, Knoxville (2010)
- Eblen, J.D., Phillips, C.A., Rogers, G.L., Langston, M.A.: The maximum clique enumeration problem: algorithms, applications, and implementations. In: *BMC Bioinformatics*, vol. 13, p. S5. BioMed Central (2012)
- Erdős, P., Rényi, A.: On random graphs. I. *Publicationes Mathematicae (Debrecen)* **6**, 290–297 (1959)
- Faust, K., Wasserman, S.: *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge (1995)
- Feldman, V., Grigorescu, E., Reyzin, L., Vempala, S.S., Xiao, Y.: Statistical algorithms and a lower bound for detecting planted cliques. *J. ACM* **64**(2), 8:1–8:37 (2017)
- Feurer, M., Klein, A., Eggenberger, K., Springenberg, J., Blum, M., Hutter, F.: Efficient and robust automated machine learning. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 28*, pp. 2962–2970. Curran Associates, Inc. (2015)
- Fitzpatrick, J., Ajwani, D., Carroll, P.: Learning to sparsify travelling salesman problem instances. In: *18th International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)* (2021)
- Fitzpatrick, J., Ajwani, D., Carroll, P.: Learning to prune electric vehicle routing problems. In: *Proceedings of the 17th Learning and Intelligent Optimization Conference* (2023)
- Fortunato, S.: Community detection in graphs. *Phys. Rep.* **486**(3–5), 75–174 (2010)
- Grassia, M., Lauri, J., Dutta, S., Ajwani, D.: Learning multi-stage sparsification for maximum clique enumeration. *arXiv preprint* [arXiv:1910.00517](https://arxiv.org/abs/1910.00517) (2019)
- Huang, J., Patwary, M.M.A., Diamos, G.F.: Coloring Big Graphs with AlphaGoZero. *CoRR* [arXiv:1902.10162](https://arxiv.org/abs/1902.10162) (2019)
- Jerrum, M.: Large cliques elude the Metropolis process. *Random Struct. Algorithms* **3**(4), 347–359 (1992)
- Karp, R.M.: The probabilistic analysis of some combinatorial search algorithms. *Algorithms Complex. New Direct. Recent Results* **1**, 19 (1976)

- Khalil, E., Dai, H., Zhang, Y., Dilkina, B., Song, L.: Learning combinatorial optimization algorithms over graphs. In: *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6351–6361 (2017)
- Knyazev, B., Taylor, G.W., Amer, M.: Understanding attention and generalization in graph neural networks. In: *Advances in Neural Information Processing Systems*, vol. 32 (2019)
- Kučera, L.: Expected complexity of graph partitioning problems. *Discrete Appl. Math.* **57**(2), 193–212 (1995)
- Lauri, J., Dutta, S.: Fine-grained search space classification for hard enumeration variants of subset problems. In: *Thirty-Third AAAI Conference on Artificial Intelligence (AAAI)*, pp. 2314–2321 (2019)
- Lauri, J., Dutta, S., Grassia, M., Ajwani, D.: Learning fine-grained search space pruning and heuristics for combinatorial optimization. *CoRR arXiv:2001.01230* (2020)
- Lewis, R.: *A Guide to Graph Colouring*, vol. 7. Springer, Berlin (2015)
- Li, C.-M., Jiang, H., Manyà, F.: On minimization of the number of branches in branch-and-bound algorithms for the maximum clique problem. *Comput. Oper. Res.* **84**, 1–15 (2017)
- Li, Z., Chen, Q., Koltun, V.: Combinatorial optimization with graph convolutional networks and guided tree search. In: *NeurIPS*, pp. 539–548 (2018)
- Lodi, A., Zarpellon, G.: On learning and branching: a survey. *TOP* **25**, 207–236 (2017)
- Lu, C., Yu, J., Wei, H., Zhang, Y.: Finding the maximum clique in massive graphs. *VLDB* **10**(11), 1538–1549 (2017)
- Manchanda, S., Mittal, A., Dhawan, A., Medya, S., Ranu, S., Singh, A.: GCOMB: learning budget-constrained combinatorial algorithms over billion-sized graphs. In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual* (2020)
- McKay, B.D., Piperno, A.: Practical graph isomorphism, II. *J. Symb. Comput.* **60**, 94–112 (2014)
- Mirhoseini, A., Goldie, A., Y, M., et al.: A graph placement methodology for fast chip design. *Nature* **594**, 207–212 (2021)
- Östergård, P.R.: A fast algorithm for the maximum clique problem. *Discrete Appl. Math.* **120**(1), 197–207 (2002). Special Issue devoted to the 6th Twente Workshop on Graphs and Combinatorial Optimization
- Palla, G., Derényi, I., Farkas, I., Vicsek, T.: Uncovering the overlapping community structure of complex networks in nature and society. *Nature* **435**(7043), 814 (2005)
- Papadopoulos, S., Kompatsiaris, Y., Vakali, A., Spyridonos, P.: Community detection in social media. *Data Min. Knowl. Disc.* **24**(3), 515–554 (2012)
- Pattabiraman, B., Patwary, M., Gebremedhin, A., Liao, W., Choudhary, A.: Fast algorithms for the maximum clique problem on massive graphs with applications to overlapping community detection. *Internet Math.* **11**(4–5), 421–448 (2015)
- Probst, M., Rothlauf, F., Grahl, J.: Scalability of using restricted Boltzmann machines for combinatorial optimization. *Eur. J. Oper. Res.* **256**(2), 368–83 (2017)
- Read, T.R.C., Cressie, N.A.C.: *Goodness-of-fit statistics for discrete multivariate data*. Springer Series in Statistics (1988)
- Read, T., Cressie, N.: Pearson's χ^2 and the likelihood ratio statistic G^2 : a comparative review. *Int. Stat. Rev.* **57**(1), 19–43 (1989)
- Rossi, R.A., Ahmed, N.K.: The network data repository with interactive graph analytics and visualization. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence* (2015)
- Rossi, R., Gleich, A., F, D., Gebremedhin: Parallel maximum clique algorithms with applications to network analysis. *SIAM J. Sci. Comput.* **37**(5), C589–C616 (2015)
- Rossi, R.A., Gleich, D.F., Gebremedhin, A.H., Patwary, M.M.: A Fast Parallel Maximum Clique Algorithm for Large Sparse Graphs and Temporal Strong Components (2013). [arXiv:1302.6256](https://arxiv.org/abs/1302.6256)
- Rossi, R.A., Gleich, D.F., Gebremedhin, A.H.: Parallel maximum clique algorithms with applications to network analysis. *SIAM J. Sci. Comput.* **37**(5), C589–C616 (2015)
- San Segundo, P., Lopez, A., Pardalos, P.M.: A new exact maximum clique algorithm for large and massive sparse graphs. *Comput. Oper. Res.* **66**, 81–94 (2016)
- Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M., Monfardini, G.: The graph neural network model. *IEEE Trans. Neural Netw.* **20**(1), 61–80 (2009)
- Stix, V.: Finding all maximal cliques in dynamic graphs. *Comput. Optim. Appl.* **27**, 173–186 (2004)
- Tayebi, D., Ray, S., Ajwani, D.: Learning to prune instances of k-median and related problems. In: *Symposium on Algorithm Engineering and Experiments (ALENEX)*, pp. 184–194 (2022)
- Tchebichef, P.: Des valeurs moyennes. *Journal de Mathématiques Pures et Appliquées* **12**, 177–184 (1867)

- Vinyals, O., Fortunato, M., Jaitly, N.: Pointer networks. In: *Advances in Neural Information Processing Systems*, pp. 2692–2700 (2015)
- Wang, C., Schwan, K., Laub, B., Kesavan, M., Gavrilovska, A.: Exploring graph analytics for cloud troubleshooting. In: *ICAC*, pp. 65–71 (2014)
- Welsh, D.J., Powell, M.B.: An upper bound for the chromatic number of a graph and its application to timetabling problems. *Comput. J.* **10**(1), 85–86 (1967)
- Yang, L., Cao, J., Tang, S., Han, D., Suri, N.: Run time application repartitioning in dynamic mobile cloud environments. *IEEE Trans. Cloud Comput.* **4**(3), 336–348 (2016)
- Yao, Y., Cao, J., Li, M.: A network-aware virtual machine allocation in cloud datacenter. In: *IFIP International Conference on Network and Parallel Computing*, pp. 71–82. Springer, Berlin (2013)
- Yeger-Lotem, E., Sattath, S., Kashtan, N., Itzkovitz, S., Milo, R., Pinter, R.Y., Alon, U., Margalit, H.: Network motifs in integrated cellular networks of transcription-regulation and protein-protein interaction. *Proc. Natl. Acad. Sci. USA* **101**(16), 5934–5939 (2004)
- Yehudai, G., Fetaya, E., Meir, E., Chechik, G., Maron, H.: From local structures to size generalization in graph neural networks. In: Meila, M., Zhang, T. (eds.) *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11975–11986. PMLR, 18–24 (2021)
- Zhang, J., Ajwani, D.: Learning to Prune Instances of Steiner Tree Problem in Graphs. *arXiv preprint arXiv:2208.11985* (2022)
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., Sun, M.: Graph neural networks: a review of methods and applications. *AI Open* **57–81**, 2020 (2020)
- Zuckerman, D.: Linear degree extractors and the inapproximability of max clique and chromatic number. In: *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 681–690. ACM (2006)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.