

Finding One Local Optimum Is Easy—but What About Two?

Yasuaki Kobayashi¹, Kazuhiro Kurita², Yutaro Yamaguchi³

¹Hokkaido University

²Okayama University

³Osaka University

koba@ist.hokudai.ac.jp, k-kurita@okayama-u.ac.jp, yutaro.yamaguchi@ist.osaka-u.ac.jp

Abstract

The class PLS (Polynomial Local Search) captures the complexity of finding a solution that is locally optimal and has proven to be an important concept in the theory of local search. It has been shown that local search versions of various combinatorial optimization problems, such as MAXIMUM INDEPENDENT SET and MAX CUT, are complete for this class. Such computational intractability typically arises in local search problems allowing *arbitrary weights*; in contrast, for *unweighted* problems, locally optimal solutions can be found in polynomial time under standard settings. In this paper, we pursue the complexity of local search problems from a different angle: We show that computing two locally optimal solutions is NP-hard for various natural unweighted local search problems, including MAXIMUM INDEPENDENT SET, MINIMUM DOMINATING SET, MAX SAT, and MAX CUT. We also discuss several tractable cases for finding two (or more) local optimal solutions.

Introduction

Local search is one of the most popular heuristics for solving hard combinatorial optimization problems, which is frequently used in both theory and practice (Aarts and Lenstra 2003; Martí, Pardalos, and Resende 2018; Glover and Kochenberger 2003; Michiels, Aarts, and Korst 2007; Monien, Dumrauf, and Tscheuschner 2010; Williamson and Shmoys 2011), as well as in numerous studies in the field of Artificial Intelligence (Garvardt et al. 2023; Gaspers et al. 2012; Selman, Kautz, and Cohen 1994; Selman, Levesque, and Mitchell 1992; Sun et al. 2024; Zhang and Looks 2005). A primitive implementation of local search methods is the following hill-climbing algorithm: Starting from an arbitrary initial (feasible) solution X , the algorithm iteratively replaces the current solution X with a (strictly) improved solution that can be found in a local solution space $\mathcal{N}(X)$, called a *neighborhood*, defined around the current solution X , as long as such an improvement can be found. A solution that cannot be further improved with this procedure is called a *local optimal solution*. A plausible reason local search methods are frequently used is that finding a local optimal solution seems much easier than finding a *global optimal solution* since every global optimal solution is locally optimal by

definition. In fact, however, there are several known obstacles to finding a local optimal solution.

Even if a local improvement in a neighborhood can be found in polynomial time, it is not easy to find a local optimal solution in general. For example, the well-known k -opt heuristics require exponentially many improvement steps to find a local optimal solution in TRAVELING SALESPERSON PROBLEM for all $k \geq 2$ (Chandra, Karloff, and Tovey 1999; Englert, Röglin, and Vöcking 2014). A similar phenomenon can be observed in several combinatorial problems with imposed neighborhood structures, such as WEIGHTED MAX CUT with FLIP neighborhood (Schäffer and Yannakakis 1991). The complexity of finding a local optimal solution under this setting is captured by the class PLS, which is introduced by (Johnson, Papadimitriou, and Yannakakis 1988a), and various local search problems (with particular neighborhood structures) are shown to be complete in this class (Komusiewicz and Morawietz 2024; Krentel 1989; Schäffer and Yannakakis 1991). In particular, Schäffer and Yannakakis (Schäffer and Yannakakis 1991) showed that the problem of finding a stable cut in edge-weighted graphs is PLS-complete, where a cut $\{X, Y\}$ in an edge-weighted graph G is *stable* if $w(X, Y) \geq w(X \triangle \{v\}, Y \triangle \{v\})$ for all $v \in X \cup Y$, where $w(X, Y)$ denotes the total weight of edges between X and Y .¹ In other words, a cut is stable if it is locally optimal under FLIP neighborhood (Schäffer and Yannakakis 1991). Komusiewicz and Morawietz recently showed that the problem of finding a locally optimal weighted independent set under 3-swap neighborhood is PLS-complete (Komusiewicz and Morawietz 2024), where an independent set Y is a *3-swap neighbor* of an independent X if Y is obtained from X by exchanging at most three vertices, that is, $|X \triangle Y| \leq 3$. We would like to mention that both problems are polynomial-time solvable in *unweighted* or *polynomially weighted*² graphs as local optima can be obtained with a polynomial number of improvement steps from an arbitrary initial solution.

Another obstacle is the complexity of finding a local improvement in a large-scale neighborhood. This is frequently

¹Here, \triangle denotes the symmetric difference of two sets.

²A polynomially weighted graph is a vertex- or edge-weighted graph such that all the weights are positive integers bounded above by a polynomial in the size of the graph.

discussed through the lens of parameterized complexity theory (de Berg et al. 2021; Fellows et al. 2012; Garvardt et al. 2023; Gaspers et al. 2012; Guo et al. 2013; Komusiewicz and Morawietz 2022; Marx 2008; Marx and Schlotter 2011; Szeider 2011). In this context, we are given an instance and its feasible solution X of a combinatorial problem and asked to find a better feasible solution Y in a parameterized neighborhood $\mathcal{N}_k(X)$. For the parameterized local search version of INDEPENDENT SET, given a graph G , an independent set X of G , and an integer k , the neighborhood $\mathcal{N}_k(X)$ of X is defined as the collection of independent sets Y of G with $|X \Delta Y| \leq k$, and hence the goal is to find an independent set $Y \in \mathcal{N}_k(X)$ with $|Y| > |X|$. This problem is known to be W[1]-hard when parameterized by k (Fellows et al. 2012), meaning that it is unlikely to exist an $f(k)|V(G)|^{O(1)}$ -time algorithm. Similarly to this, the local search version of MAX CUT is also intractable (Garvardt et al. 2023).

These obstacles can be immediately circumvented as long as we seek to find a local optimal solution for *unweighted* or *polynomially weighted* combinatorial optimization problems under *polynomial-time computable neighborhood*, where a neighborhood function \mathcal{N} is polynomial-time computable if given a solution X of an instance I , its neighborhood $\mathcal{N}(X)$ can be computed in polynomial time in $|I|$. In this paper, we explore the complexity of local search from yet another perspective, particularly focusing on (UNWEIGHTED) INDEPENDENT SET, MAX CUT on unweighted multigraphs, and MAX SAT on unweighted CNF formulas. An independent set S of a graph G is said to be *2-maximal* if it is (inclusion-wise) maximal and, for every $v \in S$ and distinct $u, w \notin S$, $(S \setminus \{v\}) \cup \{u, w\}$ is not an independent set of G , that is, it is locally optimal under 3-swap neighborhood. For a CNF formula, a truth assignment α is said to be *unflippable* if the number of satisfied clauses under α is not smaller than the number of satisfied clauses under the truth assignment obtained from α by flipping the assignment of each single variable. As mentioned above, we can find a 2-maximal independent set of G , a stable cut of an unweighted graph, and an unflippable truth assignment for a CNF formula in polynomial time by straightforward hill-climbing algorithms.

This paper establishes several complexity results for finding *multiple* local optimal solutions in unweighted combinatorial optimization problems. We show that it is NP-complete to decide whether an input graph has at least two k -maximal independent sets for any fixed $k \geq 2$. It is worth mentioning that the problem of enumerating all maximal independent sets in a graph can be done in polynomial delay³ (Tsukiyama et al. 1977; Johnson, Papadimitriou, and Yannakakis 1988b), meaning that our result cannot be extended to the case $k = 1$. This hardness result immediately yields similar hardness results for k -maximal cliques and k -minimal vertex covers as well. We also show that it is NP-complete to decide whether an input graph has at

least two k -minimal dominating sets and k -minimal feedback vertex sets. For UNWEIGHTED MAX CUT and UNWEIGHTED MAX SAT, we show that the problems of finding two local optima (i.e., stable cuts and unflippable truth assignments, respectively) are NP-hard.

Our techniques used in this paper are not particularly new: We use standard polynomial-time reductions from known NP-complete problems. However, a more careful construction in the reduction is required. In a standard reduction for proving the NP-hardness of a problem P , given an instance I of an NP-hard problem, we construct an instance I' of P such that I has a solution if and only if I' has a solution. However, this argument does not work for the local search problems that we discuss in this paper: Every instance has at least one local optimum, which can be computed in polynomial time. To deal with this issue, we need to carefully construct an instance I' that has a unique “trivial solution” that is irrelevant to a solution of I and a “nontrivial solution” that is relevant to a solution of I . This would make the construction of I' and its proof more involved compared with standard NP-hardness reductions.

We would like to mention that our work is particularly relevant to practical aspects of local search. From a practical point of view, finding multiple local optimal solutions is crucial as some of them can be significantly bad compared to the global ones, while a simple hill-climbing algorithm may find such a bad solution. There are numerous approaches to avoid getting stuck in bad local optimal solutions, such as iterated local search (Lourenço, Martin, and Stützle 2003) and multi-start local search (Martí, Resende, and Ribeiro 2013), in which one aims to mitigate this issue by essentially finding multiple local optimal solutions. Our results would shed light on theoretical obstacles for such approaches.

On the positive side, we give a polynomial-time algorithm for deciding whether an input graph G has at least two 2-maximal matchings (i.e., 2-maximal independent sets in the line graph). Our algorithm also works for finding two k -maximal matchings for any fixed $k \geq 1$. Since we can solve the maximum matching problem in polynomial time, our result may not give an interesting consequence to this end. However, we believe that the result itself is nontrivial and still an interesting case that we can find multiple local optimal solutions in polynomial time, in contrast with our hardness results of k -maximal independent sets. We also give an efficient algorithm for finding multiple k -maximal independent sets in bounded-cliquewidth graphs.

Several proofs (marked with \star) are omitted due to space limitations. See the full version (Kobayashi, Kurita, and Yamaguchi 2025) for details.

Preliminaries

Let G be a (multi)graph. The vertex set and edge set of G are denoted by $V(G)$ and $E(G)$, respectively. For $v \in V(G)$, the set of neighbors of v in G is denoted by $N_G(v)$.

A vertex set $S \subseteq V(G)$ is an *independent set* of G if no pair of vertices in S are adjacent in G . An independent set S of G is said to be *maximal* if there is no vertex $v \in V(G) \setminus S$ such that $S \cup \{v\}$ is an independent set of G . The concept of maximality can be generalized as in the following way. For

³An enumeration algorithm runs in polynomial delay if the time elapsed between two consecutive events is bounded by a polynomial in the input size, where the events include the initiation and the termination of the algorithm, and the output of each solution.

$k \geq 1$, an independent set S is k -maximal if for every $X \subseteq S$ with $|X| \leq k-1$ and $Y \subseteq V(G) \setminus S$ with $|Y| \geq |X|+1$, $(S \cup Y) \setminus X$ is not an independent set of G . Clearly, every (inclusion-wise) maximal independent set is 1-maximal and vice versa. Moreover, every k -maximal independent set is k' -maximal for $k' \leq k$.

Observation 1. *Let S be a maximal independent set of G that is not 2-maximal. Then, there are a vertex $v \in S$ and distinct neighbors $u, w \in N_G(v) \setminus S$ such that $(S \cup \{u, w\}) \setminus \{v\}$ is an independent set of G .*

Let $D \subseteq V(G)$. A vertex $v \in V(G) \setminus D$ is *dominated* by D (or D *dominates* v) if $N_G(v) \cap D \neq \emptyset$. We say that D is a *dominating set* of G if every vertex in $V(G) \setminus D$ is dominated by D . Moreover, when D is simultaneously an independent set and a dominating set of G , we call it an *independent dominating set* of G . Let us note that a vertex set is a maximal independent set of G if and only if it is an independent dominating set of G .

A *cut* of G is an unordered pair of vertex sets $\{X, Y\}$ such that $X \cup Y = V(G)$ and $X \cap Y = \emptyset$. The (multi)set of edges between X and Y is denoted by $E(X, Y)$. A cut $\{X, Y\}$ is said to be *improvable* if there is a vertex $v \in V(G)$ such that $|E(X \triangle \{v\}, Y \triangle \{v\})| > |E(X, Y)|$. A cut that is not improvable is called a *stable cut* of G .

Let φ be a CNF formula. In this paper, φ may contain multiple identical clauses. The set of variables in φ is denoted by $V(\varphi)$. Let $\alpha : V(\varphi) \rightarrow \{\mathfrak{t}, \mathfrak{f}\}$ be a truth assignment for φ . A clause is said to be *satisfied* under α if it is evaluated to \mathfrak{t} under α . Moreover, a clause is *NAE-satisfied* under α if it is satisfied under both α and its complement $\bar{\alpha}$, where $\bar{\alpha}(x) := \neg\alpha(x)$ for $x \in V(\varphi)$. The numbers of satisfied and NAE-satisfied clauses under α are denoted by $\#\text{Sat}_\varphi(\alpha)$ and $\#\text{NAESat}_\varphi(\alpha)$, respectively. We may omit the subscript φ when no confusion is possible. For $x \in V(\varphi)$, we define a truth assignment α_x as: for $y \in V(\varphi)$,

$$\alpha_x(y) = \begin{cases} \alpha(y) & y \neq x \\ \bar{\alpha}(y) & y = x \end{cases}$$

In other words, α_x is obtained from α by flipping the assignment of x . We say that α is (*NAE-*)*flippable* if the number of (NAE-)satisfied clauses under α is strictly smaller than that of satisfied clauses under α_x for some variable $x \in V(\varphi)$, that is, $\#\text{Sat}(\alpha_x) > \#\text{Sat}(\alpha)$ ($\#\text{NAESat}(\alpha_x) > \#\text{NAESat}(\alpha)$); otherwise it is (*NAE-*)*unflippable*.

Local Search Graph Problems

In this section, we investigate the intractability of finding multiple local optima for several graph problems. More precisely, we show that for any fixed $k \geq 2$, the problem of determining whether an input graph G has at least two k -maximal independent sets is NP-complete. We also prove that similar hardness holds for other local search graph problems. In contrast, we show several tractable cases for finding multiple k -maximal independent sets.

NP-Hardness

We start with the case $k = 2$ and then extend it to the general case $k \geq 2$.




Figure 1: The left figure illustrates the graph H . The edges between Z and $\{b, b'\}$ and some 4-cycles attached to vertices in Y are omitted and simplified due to visibility. The red circles indicate the vertices in 2-maximal independent sets of H .

Theorem 1. *It is NP-complete to determine whether an input graph has at least two 2-maximal independent sets.*

To prove Theorem 1, we perform a polynomial-time reduction from the following problem. Given a graph G and $X \subseteq V(G)$, **MAXIMAL INDEPENDENT SET EXTENSION** asks whether G has a maximal independent set D such that $D \cap X = \emptyset$. In other words, the vertex set D is an independent dominating set of G that has no vertices in X . This problem is known to be NP-complete (Casel et al. 2019; Conte and Tomita 2022). As observed in (Casel et al. 2019), we can assume that X is an independent set of G . Moreover, without loss of generality, we assume that G has no isolated vertices.

We construct a graph H as follows. Let $Y = V(G) \setminus X$, where $Y = \{y_1, \dots, y_t\}$. Starting from $H := G$, we add five vertices a, b, b', c, c' and add edges between a and b , between b and vertices in X , and $\{b, c\}, \{b, c'\}, \{b', c\}, \{b', c'\}$. The vertices b, b', c, c' form a 4-cycle. For each $y_i \in Y$, we add two vertices z_i and z'_i that are adjacent to y_i, b , and b' . We let $Z = \{z_i, z'_i : 1 \leq i \leq t\}$. Moreover, we add three vertices c_i, c'_i, y_i and edges that form a 4-cycle with y_i for each $y_i \in Y$. Figure 1 illustrates the constructed graph H .

The idea behind our construction is as follows. Suppose that there is a 2-maximal independent set S of H with $b \in S$. Since S is an independent set, we have $a \notin S$ and $X \cap S = \emptyset$. Due to the 2-maximality of S , $Y \cap S$ dominates all vertices in X : If $x \in X$ is not dominated by $Y \cap S$, $(S \setminus \{b\}) \cup \{a, x\}$ is an independent set of H . To prove Theorem 1, we need to ensure that H has at most one 2-maximal independent set S with $b \notin S$. The other parts are used for this.

We then show that G has an independent dominating set D with $D \cap X = \emptyset$ if and only if H has at least two 2-maximal independent sets. To this end, we start with showing several observations.

Observation 2. *Let S be a 2-maximal independent set of H with $b \in S$. Then, we have (1) $b' \in S$ and (2) $S \cap X = \emptyset$ and $S \cap Z = \emptyset$.*

Proof. Since $b \in S$, we have $c, c' \notin S$. When $b' \notin S$, S is not 2-maximal since $(S \cup \{c, c'\}) \setminus \{b\}$ is an independent

set of H . The second statement follows as $b \in N_G(v)$ for all $v \in X \cup Z$. \square

Observation 3. Let S be a 2-maximal independent set of H with $b \notin S$. Then, we have (1) $\{a, c, c'\} \subseteq S$ and (2) $S \cap Y = \emptyset$ and $X \cup Z \subseteq S$.

Proof. Since $b \notin S$, we have $a \in S$ due to the maximality of S . Suppose that $b' \in S$. Then, we have $c, c' \notin S$. This contradicts the 2-maximality of S : $(S \cup \{c, c'\}) \setminus \{b'\}$ is an independent set of H . Thus, we have $b' \notin S$. By the maximality of S , we have $\{c, c'\} \subseteq S$. For the second statement, suppose to the contrary that $y_i \in S$ for some $y_i \in Y$. Since $z_i, z'_i \notin S$ and $b, b' \notin S$, it holds that $(S \cup \{z_i, z'_i\}) \setminus \{y_i\}$ is an independent set of H , contradicting the 2-maximality of S . Hence $S \cap Y = \emptyset$. As $b, b' \notin S$, we have $X \cup Z \subseteq S$ due to the maximality of S . \square

The next observation follows from a similar discussion.

Observation 4. Let S be a 2-maximal independent set of H . For each $1 \leq i \leq t$, either $\{y_i, y'_i\} \subseteq S$ or $\{c_i, c'_i\} \subseteq S$.

Observations 3 and 4 yields the following observation.

Observation 5. Let $S_a = \{a, c, c'\} \cup X \cup Z \cup \{c_i, c'_i : 1 \leq i \leq t\}$. Then, S_a is a 2-maximal independent set of H . Moreover, H has no other 2-maximal independent set that contains a .

Proof. We show that S_a is a 2-maximal independent set of H . Observe that for $v \in S_a$, each neighbor of v is dominated by at least two vertices in S_a . Thus, by Observation 1, S_a is 2-maximal. The uniqueness of S_a immediately follows from Observations 3 and 4. \square

Now, we show that H has a 2-maximal independent set S with $S \neq S_a$ if and only if G has a maximal independent set (or, equivalently, an independent dominating set) D with $D \cap X = \emptyset$.

Lemma 1. Suppose that H has a 2-maximal independent set S with $S \neq S_a$. Then, G has an independent dominating set D with $D \cap X = \emptyset$.

Proof. Let S be a 2-maximal independent set of H with $S \neq S_a$ that maximizes $|S \cap Y|$. By Observation 5, we have $a \notin S$ and hence, by Observation 2, we have $b \in S$, $S \cap X = \emptyset$, and $S \cap Z = \emptyset$. Let $D = S \cap Y$. Clearly, D is an independent set of G with $D \cap X = \emptyset$. We show that D dominates all the vertices in $V(G) \setminus D$.

Observe first that all the vertices of $Y \setminus D$ are dominated by D . To see this, suppose that there is vertex $y_i \in Y$ that is not dominated by D . Then, y_i has no neighbor in $S \cap Y$, and hence $S' := (S \setminus \{c_i, c'_i\}) \cup \{y_i, y'_i\}$ is an independent set of H . This independent set S' is indeed 2-maximal as every neighbor of a vertex in S' is dominated by at least one other vertex in S' . This contradicts the assumption that S maximizes $|S \cap Y|$ as $|S' \cap Y| > |S \cap Y|$.

Suppose that $v \in X$ is not dominated by D . Then, $(S \cup \{a, v\}) \setminus \{b\}$ is an independent set of H , contradicting the 2-maximality of S . Therefore, D is a dominating set of G . \square




Figure 2: The graph H_4 . The red circles indicate maximal independent sets.

Lemma 2. Let $D \subseteq Y$ be an independent dominating set of G . Then, $\{b, b'\} \cup \{y_i, y'_i : y_i \in D\} \cup \{c_i, c'_i : y_i \notin D\}$ is a 2-maximal independent set of H .

Proof. Let $S = \{b, b'\} \cup \{y_i, y'_i : y_i \in D\} \cup \{c_i, c'_i : y_i \notin D\}$. We claim that S is a 2-maximal independent set of H . It is easy to verify that S is an independent set of H . Since all vertices in $X \cup Z \cup \{a, c, c'\}$ are dominated by b and $\{y_i, y'_i, c_i, c'_i\} \setminus S$ are dominated by $\{y_i, y'_i, c_i, c'_i\} \cap S$ for all i . Thus, S is a maximal independent set of H . To see 2-maximality, by Observation 1, it suffices to show that for each $v \in S$, there is at most one neighbor that is not dominated by $S \setminus \{v\}$. In the following, a vertex $v \in S$ is said to be *stable* in S if there is at most one neighbor that is not dominated by $S \setminus \{v\}$.

Since the neighbors of b' are also neighbors of b , it follows that b' is stable in S . Moreover, b is also stable as b' dominates the vertices in $Z \cup \{c, c'\}$ and D dominates the vertices in X . Note that a is not dominated by any other vertex in S . When one of c_i, c'_i, y'_i is included in S , it is stable as pairs $\{c_i, c'_i\}$ and $\{y_i, y'_i\}$ dominate each other. Finally, if $y_i \in S$, it is stable since c_i and c'_i are dominated by y'_i , the vertices in $X \cup Z$ are dominated by b , and each neighbor y_j in Y is dominated by c_j . As all vertices in S are stable, S is 2-maximal. \square

Hence, there are at least two 2-maximal independent sets in H if and only if G has an independent dominating set D with $D \cap X = \emptyset$, completing the proof of Theorem 1.

We can extend this proof to those for any fixed $k \geq 2$. From an instance (G, X) of MAXIMAL INDEPENDENT SET EXTENSION, we construct the graph H as above and convert it to a graph H_k as follows. We replace each vertex v of H with an independent set $M_k^v = \{v_1, \dots, v_{k-1}\}$ of $k-1$ vertices and add an edge between each vertex in M_k^v and each vertex in M_k^w if and only if v and w are adjacent in H . The graph obtained in this way is denoted by H_k . See Figure 2 for an illustration. Note that $H = H_2$.

Lemma 3 (\star). There is a bijection between the collection of 2-maximal independent sets in H and that of k -maximal independent sets in H_k . In particular, H has at least two 2-maximal independent sets if and only if H_k has at least two k -maximal independent sets.

This yields the following.

Theorem 2. For any fixed $k \geq 2$, it is NP-complete to determine whether an input graph has at least two k -maximal independent sets.

As an immediate corollary of Theorems 1 and 2, we have the NP-hardness of finding multiple local optima for the minimum vertex cover and the maximum clique problems. Let G be a graph. A *clique* is a pairwise adjacent vertex set in G . A *vertex cover* of G is a vertex set such that its complement is an independent set of G . A clique S is k -maximal if for $X \subseteq S$ with $|X| \leq k - 1$ and $Y \subseteq V(G) \setminus S$ with $|Y| > |X|$, $(S \cup Y) \setminus X$ is not a clique in G . Similarly, a vertex cover S is k -minimal if for $X \subseteq V(G) \setminus S$ with $|X| \leq k - 1$ and $Y \subseteq S$ with $|Y| > |X|$, $(S \setminus Y) \cup X$ is not a vertex cover of G . It is easy to see that a vertex set S is a k -maximal independent set of G if and only if it is a k -maximal clique in the complement graph of G . Moreover, S is a k -maximal independent set of G if and only if $V(G) \setminus S$ is a k -minimal vertex cover of G . Hence, the following corollary is immediate.

Corollary 1. For any fixed $k \geq 2$, it is NP-complete to determine whether an input graph has at least two k -maximal cliques. Moreover, it is NP-complete to determine whether an input graph has at least two k -minimal vertex covers.

Similarly, a dominating set D of G is k -minimal if for $X \subseteq V(G) \setminus D$ with $|X| \leq k - 1$ and $Y \subseteq D$ with $|Y| > |X|$, $(D \setminus Y) \cup X$ is not a dominating set of G . A slightly modified version of a well-known reduction from VERTEX COVER to DOMINATING SET proves the hardness of finding multiple k -minimal dominating sets.

Theorem 3 (\star). For any fixed $k \geq 2$, it is NP-complete to determine whether an input graph has at least two k -minimal dominating sets.

Finally, we show that finding k -minimal feedback vertex sets is hard. A *feedback vertex set* of a graph G is a vertex subset $X \subseteq V(G)$ whose removal makes the graph acyclic (i.e., a forest). The k -minimality for feedback vertex sets is defined analogously.

Theorem 4 (\star). For any fixed $k \geq 2$, it is NP-complete to determine whether an input graph has at least two k -minimal feedback vertex sets.

Positive Results for k -Maximal Independent Sets

To complement Theorem 1, we investigate the tractability of finding multiple local optima for a special case of the local search version of MAXIMUM INDEPENDENT SET.

Let G be a graph. A set of edges $M \subseteq E(G)$ is called a *matching* of G if every pair of edges in M does not share their end vertices. For an integer $k \geq 1$, a matching M is said to be k -maximal if for $X \subseteq M$ with $|X| \leq k - 1$ and $Y \subseteq E(G) \setminus X$ with $|Y| > |X|$, $(M \cup Y) \setminus X$ is not a matching of G . In this section, we give a polynomial-time algorithm for deciding whether an input graph G has at least two k -maximal matchings for every fixed k .

Let M be a matching of G . In the following, we may not distinguish between an edge set of G and the subgraph induced by them. We say that a vertex v is *matched* in M if there is an edge incident to v in M . A path $P = (v_1, \dots, v_\ell)$

is said to be M -alternating if v_1 is not matched in M and $\{v_i, v_{i+1}\} \in M$ for even i . An M -alternating path $P = (v_1, \dots, v_\ell)$ is said to be M -augmenting if v_ℓ is not matched in M , that is, ℓ is even. It is well known that a matching M is maximum if and only if it has no M -augmenting paths (e.g., (Korte and Vygen 2018)). In particular, for a matching M' with $|M'| < |M|$, there is a path component in $M' \triangle M$ that is M' -augmenting, as $M' \triangle M$ consists of a disjoint union of paths and cycles.

Lemma 4 (\star). Let M be a matching in G and let k be a positive integer. Then, M is k -maximal if and only if G has no M -augmenting path of length at most $2k - 1$.

Given a graph G , it is easy to determine if G has at least two maximum matchings by using a polynomial-time algorithm for computing a maximum matching (see (Gabow, Kaplan, and Tarjan 2001) for a more sophisticated algorithm). Obviously, we can determine if G has at least two k -maximal matchings in polynomial time when there are at least two maximum matchings in G . Hence, we assume otherwise that G has a unique maximum matching M^* . A component in a graph is *nontrivial* if it has at least one edge.

Lemma 5. Let M^* be a unique maximum matching of G . Suppose that G has a k -maximal matching other than M^* . Then, G has a k -maximal matching M of size $|M^*| - 1$. Moreover, there is exactly one nontrivial component in $M \triangle M^*$, which is an M -augmenting path.

Proof. Let M be a k -maximal matching of G with $M \neq M^*$. We assume that $|M| < |M^*| - 1$ as otherwise we are done. Let P^* be an M -augmenting path component in $M \triangle M^*$. Since P^* is an M -augmenting path in G , by Lemma 4, P^* contains more than $2k - 1$ edges. Let $M' = M^* \triangle P^*$. Clearly, P^* is an M' -augmenting path, and M' is a matching with $|M'| = |M^*| - 1$. Moreover, there is no M' -augmenting path other than P^* , as otherwise such an M' -augmenting path P would imply a maximum matching $M' \triangle P$ distinct from M^* . Again, by Lemma 4, M' is k -maximal. \square

By Lemma 5, it suffices to find a k -maximal matching M with size $|M^*| - 1$. Suppose that the M -augmenting path P^* in the proof of Lemma 5 has $2k + 1$ edges. (Recall that each augmenting path has an odd number of edges.) In this case, we try to check all the possibilities of paths P with $|E(P)| = 2k + 1$ such that $M^* \triangle P$ is a k -maximal matching of G . This can be done in polynomial time, as k is fixed. Suppose otherwise that P^* has at least $2k + 3$ edges. In this case, we claim that the above algorithm also finds a k -maximal matching of G if it exists. To see this, let P be a subpath of P^* containing one of the end vertices of P^* such that $|E(P)| = 2k + 1$. Then, $M := M^* \triangle P$ is a matching of G with $|M| = |M^*| - 1$. Moreover, P is a unique M -augmenting path due to the uniqueness of M^* . Therefore, M is k -maximal.

Theorem 5. For each integer $k \geq 1$, there is an $n^{O(k)}$ -time algorithm for determining whether an input n -vertex graph has at least two k -maximal matchings.

For another tractable case, we would like to mention that a variant of Courcelle's theorem allows us to enumerate all k -maximal independent sets on bounded-cliquewidth graphs (and hence bounded-treewidth graphs) in polynomial delay for any fixed k .

Proposition 1 (*). *There is a polynomial-delay algorithm for generating all k -maximal independent sets in bounded-cliquewidth graphs.*

Local Search for MAX SAT and MAX CUT

In this section, we focus on the local search versions of MAX SAT and MAX CUT with FLIP neighborhood (Schäffer and Yannakakis 1991). For a CNF formula φ and truth assignments α, α' of φ , we say that α is *adjacent* to α' in the FLIP neighborhood if $\alpha' = \alpha_x$ for some $x \in V(\varphi)$. Similarly, for a graph G and two cuts $C = \{X, Y\}$ and $C' = \{X', Y'\}$, we say that C is *adjacent* to C' in the FLIP neighborhood if $X' = X \Delta \{v\}$ and $Y' = Y \Delta \{v\}$ for some $v \in V(G)$. Under these definitions, a truth assignment of φ is locally optimal if it is unflippable and a cut of G is locally optimal if it is stable. We show that the problems of finding multiple unflippable assignments and multiple stable cuts are NP-hard. To this end, we start with MAX NAESAT.

We observe that every 3-CNF formula has at least two NAE-unflippable assignments. To see this, let α be a truth assignment of φ maximizing the number of NAE-satisfied clauses in φ . This assignment is indeed NAE-unflippable, as it is a global optimum for MAX NAESAT, and its complement $\bar{\alpha}$ is also NAE-unflippable and satisfies the same set of clauses. The following theorem suggests, however, that finding a third NAE-unflippable assignment is hard.

Theorem 6. *It is NP-complete to determine whether an input 3-CNF formula has at least three NAE-unflippable assignments.*

The proof of Theorem 6 is also done by performing a polynomial-time reduction from MAXIMAL INDEPENDENT SET EXTENSION. Let us note that the isolated vertices in $V(G) \setminus X$ are contained in any maximal independent set D . By adding sufficiently many isolated vertices to G , we can assume that $|D| \geq |V(G) \setminus D|$ holds for every maximal independent set D with $D \cap X = \emptyset$.

We construct a 3-CNF formula φ as follows. Let m be the number of edges in G . For each vertex $v \in V$, we associate a variable x_v , where the intention is that x_v being true indicates that vertex v is included in an independent set. We also use two additional variables x^* and y^* . For each edge $e = \{u, v\} \in E(G)$, we add clauses $(x^* \vee y^* \vee s_e)^2$, $(\neg s_e \vee \neg x_u \vee \neg x_v)^2$, and $(s_e \vee x^*)^3$ to φ . Here, for a clause c and a positive integer t , c^t indicates the conjunction of t copies of c . For each vertex $v \in V(G)$, we add a clause $(x_v \vee y^*)$ and for each $v \in X$, add clauses $(x_v \vee \neg x^* \vee y^*)^{4m}$ and $(x^* \vee y^*)^{4m|X|}$. The entire formula φ is defined as

$$\begin{aligned} \varphi &= \bigwedge_{e \in E(G)} \varphi_e \wedge \bigwedge_{v \in V(G)} \varphi_v \wedge \bigwedge_{v \in X} \varphi'_v \wedge (x^* \vee y^*)^{4m|X|}; \\ \varphi_e &= (x^* \vee y^* \vee s_e)^2 \wedge (\neg s_e \vee \neg x_u \vee \neg x_v)^2 \wedge (s_e \vee x^*)^3; \\ \varphi_v &= (x_v \vee y^*); \quad \varphi'_v = (x_v \vee \neg x^* \vee y^*)^{4m}, \end{aligned}$$

where $e = \{u, v\}$ in the second subformula φ_e .

There are two types of truth assignments α with (1) $\alpha(x^*) \neq \alpha(y^*)$ and (2) $\alpha(x^*) = \alpha(y^*)$. We first show that there are only two NAE-unflippable assignments of type (1).

Lemma 6. *There are exactly two NAE-unflippable assignments α of φ with $\alpha(x^*) \neq \alpha(y^*)$. Moreover, these two assignments are complements of each other.*

Proof. Suppose that $\alpha(x^*) = \mathbf{t}$ and $\alpha(y^*) = \mathbf{f}$. We first observe that $\alpha(s_e) = \mathbf{f}$ for all $e \in E(G)$. To see this, suppose that $\alpha(s_e) = \mathbf{t}$ for some $e \in E(G)$. Since clauses $(s_e \vee x^*)^3$ are not NAE-satisfied under α and clauses $(x^* \vee y^* \vee s_e)$ are still NAE-satisfied under α_{s_e} , we have $\#\text{NAESat}(\alpha_{s_e}) - \#\text{NAESat}(\alpha) \geq 1$, meaning that α is NAE-flippable. Thus, $\alpha(s_e) = \mathbf{f}$ for all $e \in E(G)$. Moreover, for $v \in V(G)$ with $\alpha(x_v) = \mathbf{f}$, flipping the assignment of x_v increases NAE-satisfied clauses by at least $4m - 2m + 1 = 2m + 1 \geq 1$, as φ'_v becomes NAE-satisfied under α_{x_v} . Hence, $\alpha(x_v) = \mathbf{t}$ for all v , and there is a unique NAE-unflippable assignment α with $\alpha(x^*) = \mathbf{t}$ and $\alpha(y^*) = \mathbf{f}$. Considering the symmetric case $\alpha(x^*) = \mathbf{f}$ and $\alpha(y^*) = \mathbf{t}$, the lemma holds. \square

Now, we show that G has a maximal independent set D avoiding X if and only if φ has an NAE-unflippable assignment α with $\alpha(x^*) = \alpha(y^*)$.

Lemma 7. *If there is a maximal independent set D of G with $D \cap X = \emptyset$, then φ has an NAE-unflippable assignment α with $\alpha(x^*) = \alpha(y^*)$.*

Proof. Suppose that G has a maximal independent set D with $D \cap X = \emptyset$. We define a truth assignment α for φ as:

$$\alpha(x) = \begin{cases} \mathbf{t} & x \in \{x_v : v \in D\} \cup \{s_e : e \in E(G)\} \\ \mathbf{f} & \text{otherwise} \end{cases}.$$

Note that all the clauses of the form $(x^* \vee y^* \vee s_e)$, $(\neg s_e \vee \neg x_u \vee \neg x_v)$, $(s_e \vee x^*)$, and $(x_v \vee \neg x^* \vee y^*)$ are NAE-satisfied under α , as D is an independent set of G . We show that α is NAE-unflippable at each variable. For $e \in E(G)$, when flipping the assignment of s_e to \mathbf{f} , it does not increase NAE-satisfied clauses at all but makes clauses $(x^* \vee y^* \vee s_e)$ and $(s_e \vee x^*)$ NAE-unsatisfied. For $u \in D$, α is NAE-unflippable at x_u since all the clauses appearing x_u are NAE-satisfied under α . For $u \in V(G) \setminus D$, there is a neighbor $v \in N_G(u)$ of u that belongs to D due to the maximality of D . If we flip the assignment of x_u to \mathbf{t} , clause $(x_u \vee y^*)$ becomes NAE-satisfied but clauses $(\neg s_e \vee \neg x_u \vee \neg x_v)^2$ become NAE-unsatisfied and clause $(x_u \vee \neg x^* \vee y^*)$ remains NAE-satisfied. Thus, flipping the assignment of x_u does not increase NAE-satisfied clauses, yielding α is NAE-unflippable at x_u for all $u \in V(G) \setminus D$. Finally, we conclude that α is NAE-unflippable at x^* and y^* : all $4m|X|$ clauses of $(x^* \vee y^*)$ become NAE-satisfied but all $4m$ clauses of $(x_v \vee \neg x^* \vee y^*)$ become NAE-unsatisfied for $v \in X$ under α_{x^*} ; the increment in the number of NAE-satisfied clauses is determined by the clauses in $\bigwedge_{v \in V(G)} (x_v \vee y^*)$, which is

$$\begin{aligned} & |\{v \in V(G) : \alpha(x_v) = \mathbf{f}\}| - |\{v \in V(G) : \alpha(x_v) = \mathbf{t}\}| \\ &= |V(G) \setminus D| - |D| \leq 0 \end{aligned}$$

as $|D| \geq |V(G) \setminus D|$. Therefore, α is NAE-unflippable. \square

Lemma 8 (\star). *If there is an NAE-unflippable assignment α with $\alpha(x^*) = \alpha(y^*)$ for φ , then G has an independent dominating set D with $D \cap X = \emptyset$.*

By Lemma 6, there are exactly two NAE-unflippable assignments α of φ with $\alpha(x^*) \neq \alpha(y^*)$ and, by Lemmas 7 and 8, there is an NAE-unflippable assignment α of φ with $\alpha(x^*) = \alpha(y^*)$ if and only if G has a maximal independent set D with $D \cap X = \emptyset$. Therefore, we can determine the existence of D by checking whether φ has at least three NAE-unflippable assignments. This proves Theorem 6.

We next show that the problem remains hard even for positive CNF formulas, that is, each clause contains only positive literals. The proof is almost analogous to the standard reduction from NAE3SAT to Positive NAE3SAT. Let φ be a CNF formula with m clauses. We replace each negative literal $\neg x$ with a fresh variable x' and add $(x \vee x')^k$ for sufficiently large k , say $k = m + 1$, for each variable $x \in V(\varphi)$. The obtained positive CNF formula is denoted by φ' . The following lemma ensures that the new variable x' serves as the negation of x for any NAE-unflippable assignment.

Lemma 9 (\star). *For every NAE-unflippable assignment α for φ' , it holds that $\alpha(x) \neq \alpha(x')$ for all $x \in V(\varphi)$.*

The above lemma implies that there is a bijection between the set of NAE-unflippable assignments for φ and those for φ' , proving the following theorem.

Theorem 7. *It is NP-complete to determine whether an input positive 3-CNF formula has at least three NAE-unflippable assignments.*

Now, we turn our attention to MAX CUT. Let φ be a positive 3-CNF formula. We can assume that φ has no unit clauses, as they are always NAE-unsatisfied. From φ , we construct a multigraph G as follows. The vertex set of G corresponds to $V(\varphi)$. For each clause with two literals, say $(x \vee y)$, we add two parallel edges between x and y , and for each clause with three literals, say $(x \vee y \vee z)$, we add edges $\{x, y\}, \{y, z\}, \{z, x\}$, forming a triangle. An *ordered* cut of G is an ordered pair (X, Y) such that $X \cup Y = V(G)$ and $X \cap Y = \emptyset$. An ordered cut (X, Y) is said to be stable if its unordered counterpart $\{X, Y\}$ is stable.

Lemma 10. *There is a bijection between the set of NAE-unflippable assignments for φ and the set of ordered stable cuts of G .*

Proof. From a truth assignment α of φ , we can naturally define an ordered cut (X, Y) as $X = \{x \in V(\varphi) : \alpha(x) = \mathbf{t}\}$ and $Y = \{x \in V(\varphi) : \alpha(x) = \mathbf{f}\}$. Conversely, we can define a truth assignment of φ from an ordered cut of G in the same way as above. Note that the complement of α defines the ordered cut (Y, X) and vice versa. Now, we show that α is NAE-unflippable if and only if (X, Y) is stable.

Suppose that α is NAE-unflippable. Observe that when a clause c is NAE-satisfied, the variables contained in c contribute exactly 2 to the cut (X, Y) . Thus, for $x \in X \cup Y$,

$$\begin{aligned} |E(X, Y)| - |E(X \triangle \{x\}, Y \triangle \{x\})| \\ = 2 \cdot \#\text{NAESat}(\alpha) - 2 \cdot \#\text{NAESat}(\alpha_x) \geq 0, \end{aligned}$$

implying that (X, Y) is stable. This implication is reversible, proving the claim of the lemma. \square

It is easy to see that G has at least three ordered stable cuts if and only if it has at least two (unordered) stable cuts. By Theorem 7 and Lemma 10, the following theorem holds.

Theorem 8. *It is NP-complete to determine whether an input multigraph G has at least two stable cuts.*

Finally, we consider MAX 2SAT. The proof is similar to a standard reduction from MAX CUT to MAX 2SAT.

Theorem 9 (\star). *It is NP-complete to determine whether a 2-CNF formula has at least two unflippable assignments.*

Concluding Remarks

In this paper, we initiate the study of the complexity of finding multiple local optima in unweighted combinatorial optimization problems. Our results suggest that there are several natural local search problems for which one of the local optima is easy to find, but two (or three) of them are hard to find, and give rise to several interesting future directions.

- In Theorem 8 and Theorem 9, the constructed graphs and 2-CNF formulas have (unweighted) multiedges and multiset of clauses, respectively. Since these multiple objects can be encoded by a polynomially-weighted single object, these results also hold for simple polynomially-weighted graphs and formulas. It would be interesting to investigate that the problems are still hard even for unweighted simple graphs and formulas.
- The primal goal of this paper is to establish that, even on very simple neighborhood structures, finding multiple local optima is computationally intractable in natural combinatorial optimization problems. To extend our results, there are various local search problems with other neighborhood structures studied in the context of PLS-completeness. We believe that these local search problems are also intractable in our setting, while it would be interesting to investigate natural local search problems that are PLS-complete but tractable in our setting.
- We show that finding two k -maximal independents can be solved in polynomial time when the input graph is restricted to the class of line graphs (Theorem 5 and Proposition 1) or to that of bounded-cliquewidth graphs. However, these results might be of no importance when it comes to finding a global optimum for INDEPENDENT SET, as we can find a global one in polynomial time on these classes. A natural question is to investigate a class of graphs on which INDEPENDENT SET is NP-hard but where finding multiple 2-maximal independent sets is easy, while it is also an interesting question whether the opposite situation can occur: finding a global optimum is easy, but finding multiple local optima is hard.

Acknowledgments

This work is partially supported by JSPS KAKENHI Grant Numbers JP23K28034, JP24H00686, JP24H00697, JP22H03549, JP25K21273, JP25K03080, JP25K00136, JP20K19743, JP20H00605, and JP25H01114, and by JST CRONOS Japan Grant Number JPMJCS24K2.

References

- Aarts, E. H. L.; and Lenstra, J. K. 2003. *Local Search in Combinatorial Optimization*. Princeton, NJ: Princeton University Press, 2nd edition. ISBN 9780691115221.
- Casel, K.; Fernau, H.; Ghadikolaei, M. K.; Monnot, J.; and Sikora, F. 2019. Extension of Vertex Cover and Independent Set in Some Classes of Graphs. In Heggernes, P., ed., *Proceedings of CIAC 2019*, volume 11485 of *Lecture Notes in Computer Science*, 124–136. Springer.
- Chandra, B.; Karloff, H. J.; and Tovey, C. A. 1999. New Results on the Old k -opt Algorithm for the Traveling Salesman Problem. *SIAM J. Comput.*, 28(6): 1998–2029.
- Conte, A.; and Tomita, E. 2022. On the overall and delay complexity of the CLIQUES and Bron-Kerbosch algorithms. *Theor. Comput. Sci.*, 899: 1–24.
- de Berg, M.; Buchin, K.; Jansen, B. M. P.; and Woeginger, G. J. 2021. Fine-grained Complexity Analysis of Two Classic TSP Variants. *ACM Trans. Algorithms*, 17(1): 5:1–5:29.
- Englert, M.; Röglin, H.; and Vöcking, B. 2014. Worst Case and Probabilistic Analysis of the 2-Opt Algorithm for the TSP. *Algorithmica*, 68(1): 190–264.
- Fellows, M. R.; Fomin, F. V.; Lokshtanov, D.; Rosamond, F. A.; Saurabh, S.; and Villanger, Y. 2012. Local search: Is brute-force avoidable? *J. Comput. Syst. Sci.*, 78(3): 707–719.
- Gabow, H. N.; Kaplan, H.; and Tarjan, R. E. 2001. Unique Maximum Matching Algorithms. *J. Algorithms*, 40(2): 159–183.
- Garvardt, J.; Grüttemeier, N.; Komusiewicz, C.; and Morawietz, N. 2023. Parameterized Local Search for Max c -Cut. In *Proceedings of IJCAI 2023*, 5586–5594. ijcai.org.
- Gaspers, S.; Kim, E. J.; Ordyniak, S.; Saurabh, S.; and Szeider, S. 2012. Don't Be Strict in Local Search! In *Proceedings of AAAI 2012*, 486–492. AAAI Press.
- Glover, F. W.; and Kochenberger, G. A., eds. 2003. *Handbook of Metaheuristics*, volume 57 of *International Series in Operations Research & Management Science*. Kluwer / Springer. ISBN 978-1-4020-7263-5.
- Guo, J.; Hartung, S.; Niedermeier, R.; and Suchý, O. 2013. The Parameterized Complexity of Local Search for TSP, More Refined. *Algorithmica*, 67(1): 89–110.
- Johnson, D. S.; Papadimitriou, C. H.; and Yannakakis, M. 1988a. How Easy is Local Search? *J. Comput. Syst. Sci.*, 37(1): 79–100.
- Johnson, D. S.; Papadimitriou, C. H.; and Yannakakis, M. 1988b. On Generating All Maximal Independent Sets. *Inf. Process. Lett.*, 27(3): 119–123.
- Kobayashi, Y.; Kurita, K.; and Yamaguchi, Y. 2025. Finding One Local Optimum Is Easy - But What about Two? *CoRR*, abs/2507.07524.
- Komusiewicz, C.; and Morawietz, N. 2022. Parameterized Local Search for Vertex Cover: When Only the Search Radius Is Crucial. In Dell, H.; and Nederlof, J., eds., *Proceedings of IPEC 2022*, volume 249 of *LIPICs*, 20:1–20:18.
- Komusiewicz, C.; and Morawietz, N. 2024. Finding 3-Swap-Optimal Independent Sets and Dominating Sets is Hard. *ACM Transactions on Computation Theory*.
- Korte, B.; and Vygen, J. 2018. *Combinatorial Optimization: Theory and Algorithms*, volume 21 of *Algorithms and Combinatorics*. Springer, 6 edition. ISBN 978-3662560396.
- Krentel, M. W. 1989. Structure in Locally Optimal Solutions (Extended Abstract). In *30th Annual Symposium on Foundations of Computer Science, FOCS 1989*, 216–221. IEEE Computer Society.
- Lourenço, H. R.; Martin, O. C.; and Stützle, T. 2003. Iterated Local Search. In Glover, F. W.; and Kochenberger, G. A., eds., *Handbook of Metaheuristics*, volume 57 of *International Series in Operations Research & Management Science*, 320–353. Kluwer / Springer.
- Martí, R.; Pardalos, P. M.; and Resende, M. G. C., eds. 2018. *Handbook of Heuristics*. Springer. ISBN 978-3-319-07123-7.
- Martí, R.; Resende, M. G. C.; and Ribeiro, C. C. 2013. Multi-start methods for combinatorial optimization. *Eur. J. Oper. Res.*, 226(1): 1–8.
- Marx, D. 2008. Searching the k -change neighborhood for TSP is $W[1]$ -hard. *Oper. Res. Lett.*, 36(1): 31–36.
- Marx, D.; and Schlotter, I. 2011. Stable assignment with couples: Parameterized complexity and local search. *Discret. Optim.*, 8(1): 25–40.
- Michiels, W.; Aarts, E. H. L.; and Korst, J. H. M. 2007. *Theoretical aspects of local search*. Monographs in Theoretical Computer Science. An EATCS Series. Springer. ISBN 978-3-540-35853-4.
- Monien, B.; Dumrauf, D.; and Tscheuschner, T. 2010. Local Search: Simple, Successful, But Sometimes Sluggish. In Abramsky, S.; Gavaille, C.; Kirchner, C.; auf der Heide, F. M.; and Spirakis, P. G., eds., *Proceedings of the 37th International Colloquium on Automata, Languages and Programming, ICALP 2010, Part I*, volume 6198 of *Lecture Notes in Computer Science*, 1–17. Springer.
- Schäffer, A. A.; and Yannakakis, M. 1991. Simple Local Search Problems That are Hard to Solve. *SIAM J. Comput.*, 20(1): 56–87.
- Selman, B.; Kautz, H. A.; and Cohen, B. 1994. Noise Strategies for Improving Local Search. In Hayes-Roth, B.; and Korf, R. E., eds., *Proceedings of the 12th National Conference on Artificial Intelligence, Seattle, WA, USA, July 31 - August 4, 1994, Volume 1*, 337–343. AAAI Press / The MIT Press.
- Selman, B.; Levesque, H. J.; and Mitchell, D. G. 1992. A New Method for Solving Hard Satisfiability Problems. In Swartout, W. R., ed., *Proceedings of the 10th National Conference on Artificial Intelligence, San Jose, CA, USA, July 12-16, 1992*, 440–446. AAAI Press / The MIT Press.
- Sun, R.; Wang, Y.; Wang, S.; Li, H.; Li, X.; and Yin, M. 2024. Nukplex: An Efficient Local Search Algorithm for Maximum K -Plex Problem. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*, 7029–7037. ijcai.org.

Szeider, S. 2011. The parameterized complexity of k -flip local search for SAT and MAX SAT. *Discret. Optim.*, 8(1): 139–145.

Tsukiyama, S.; Ide, M.; Ariyoshi, H.; and Shirakawa, I. 1977. A New Algorithm for Generating All the Maximal Independent Sets. *SIAM J. Comput.*, 6(3): 505–517.

Williamson, D. P.; and Shmoys, D. B. 2011. *The Design of Approximation Algorithms*. Cambridge University Press. ISBN 978-0-521-19527-0.

Zhang, W.; and Looks, M. 2005. A Novel Local Search Algorithm for the Traveling Salesman Problem that Exploits Backbones. In Kaelbling, L. P.; and Saffiotti, A., eds., *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30 - August 5, 2005*, 343–350. Professional Book Center.