

# A Structural Investigation of the Approximability of Polynomial-Time Problems

**Karl Bringmann**

Universität des Saarlandes, Saarland Informatics Campus, Saarbrücken, Germany  
Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

**Alejandro Cassis**

Universität des Saarlandes, Saarland Informatics Campus, Saarbrücken, Germany  
Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

**Nick Fischer**

Universität des Saarlandes, Saarland Informatics Campus, Saarbrücken, Germany  
Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

**Marvin Künnemann**

TU Kaiserslautern, Germany

---

## Abstract

---

An extensive research effort targets optimal (in)approximability results for various NP-hard optimization problems. Notably, the works of (Creignou'95) as well as (Khanna, Sudan, Trevisan, Williamson'00) establish a tight characterization of a large subclass of MAXSNP, namely Boolean MAXCSPs and further variants, in terms of their polynomial-time approximability. Can we obtain similarly encompassing characterizations for classes of *polynomial-time optimization problems*?

To this end, we initiate the systematic study of a recently introduced polynomial-time analogue of MAXSNP, which includes a large number of well-studied problems (including Nearest and Furthest Neighbor in the Hamming metric, Maximum Inner Product, optimization variants of  $k$ -XOR and Maximum  $k$ -Cover). Specifically, for each  $k$ ,  $\text{MAXSP}_k$  denotes the class of  $O(m^k)$ -time problems of the form  $\max_{x_1, \dots, x_k} \#\{y : \phi(x_1, \dots, x_k, y)\}$  where  $\phi$  is a quantifier-free first-order property and  $m$  denotes the size of the relational structure. Assuming central hypotheses about clique detection in hypergraphs and exact MAX-3-SAT, we show that for any  $\text{MAXSP}_k$  problem definable by a quantifier-free  $m$ -edge *graph* formula  $\phi$ , the *best possible* approximation guarantee in faster-than-exhaustive-search time  $O(m^{k-\delta})$  falls into one of four categories:

- optimizable to exactness in time  $O(m^{k-\delta})$ ,
- an (inefficient) approximation scheme, i.e., a  $(1 + \varepsilon)$ -approximation in time  $O(m^{k-f(\varepsilon)})$ ,
- a (fixed) constant-factor approximation in time  $O(m^{k-\delta})$ , or
- an  $m^\varepsilon$ -approximation in time  $O(m^{k-f(\varepsilon)})$ .

We obtain an almost complete characterization of these regimes, for  $\text{MAXSP}_k$  as well as for an analogously defined minimization class  $\text{MINSP}_k$ . As our main technical contribution, we show how to rule out the existence of *approximation schemes* for a large class of problems admitting constant-factor approximations, under a hypothesis for *exact* Sparse MAX-3-SAT algorithms posed by (Alman, Vassilevska Williams'20). As general trends for the problems we consider, we observe: (1) Exact optimizability has a simple algebraic characterization, (2) only few maximization problems do not admit a constant-factor approximation; these do not even have a subpolynomial-factor approximation, and (3) constant-factor approximation of minimization problems is *equivalent* to deciding whether the optimum is equal to 0.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Problems, reductions and completeness; Theory of computation  $\rightarrow$  Approximation algorithms analysis

**Keywords and phrases** Classification Theorems, Hardness of Approximation in P, Fine-grained Complexity Theory

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2022.30

**Category** Track A: Algorithms, Complexity and Games



© Karl Bringmann, Alejandro Cassis, Nick Fischer, and Marvin Künnemann;  
licensed under Creative Commons License CC-BY 4.0

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).

Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;

Article No. 30; pp. 30:1–30:20



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



**Related Version** *Full Version:* <https://arxiv.org/abs/2204.11681>

**Funding** *Karl Bringmann, Alejandro Cassis, Nick Fischer:* This work is part of the project TIPEA that has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 850979). *Marvin Künnemann:* Part of this research was performed at ETH Zürich, Institute for Theoretical Studies, supported by Dr. Max Rössler, by the Walter Haefner Foundation, and by the ETH Zürich Foundation.

## 1 Introduction

For many optimization problems, the best known exact algorithms essentially explore the complete search space, up to low-order improvements. While this holds true in particular for NP-hard optimization problems such as maximum satisfiability, also common polynomial-time optimization problems like Nearest Neighbor search are no exception. When facing such a problem, the perhaps most common approach is to relax the optimization goal and ask for *approximations* rather than optimal solutions. Can we obtain an approximate value significantly faster than exhaustive search, and if so, what is the best approximation guarantee we can achieve?

Even considering polynomial-time problems only, the study of such questions has led to significant algorithmic breakthroughs, such as locality-sensitive hashing (LSH) [49, 41, 11], subquadratic-time approximation algorithms for Edit Distance [56, 18, 17, 19, 14, 12, 26, 55, 21, 43, 13], scaling algorithms for graph problems [37, 70, 34], and fast approximation algorithms via the polynomial method [5, 6] (which lead to the currently fastest known exponential-time approximation schemes for MaxSAT).

These algorithmic breakthroughs have recently been complemented by exciting tools for proving hardness of approximation in P: Most notably, the distributed PCP framework [4] has lead to strong conditional lower bounds, including fundamental limits for Nearest Neighbor search [62], as well as tight approximability results for the Maximum Inner Product problem [28]. Other technical advances include evidence against deterministic approximation schemes for Longest Common Subsequence [1, 3, 29], strong (at times even tight) problem-specific hardness results such as [61, 22, 16, 25, 52], the first fine-grained equivalences of approximation in P results [30, 29], and related works on parameterized inapproximability [27, 51, 59], see [36] for a survey.

These strong advances from both sides shift the algorithmic frontier and the frontier of conditional hardness towards each other. Consequently, it becomes increasingly important to generalize isolated results – both algorithms and reductions – towards making these frontiers explicit: A more comprehensive description of current techniques might enable us to understand general trends underlying these works and to highlight the most pressing limitations of current methods. In fact, we view this as one of the fundamental tasks and uses of fine-grained algorithm design & complexity.

**Optimization Classes in P.** To study (hardness of) approximation in P in a general way, we study a class  $\text{MAXSP}_k$  recently introduced in [23] (see the full version of this paper for a comparison to the classic class MAXSNP, which motivated the definition of  $\text{MAXSP}_k$ ). This class consists of polynomial-time optimization problems of the form:

$$\max_{x_1, \dots, x_k} \#\{y : \phi(x_1, \dots, x_k, y)\},$$

where  $\phi$  is a quantifier-free first-order property.<sup>1</sup> One obtains an analogous minimization class  $\text{MINSP}_k$  by replacing maximization by minimization. A large number of natural and well-studied problems can be expressed this way: In particular, we may think of each  $x_i$  ranging over a set  $X_i$  of  $n$  vectors in  $\{0, 1\}^d$ , and the task is to maximize (or minimize) the number of coordinates  $y$  satisfying an arbitrary Boolean function over the coordinate values  $x_1[y], \dots, x_k[y]$  (here,  $\phi$  is defined using a binary relation  $R \subseteq (X_1 \cup \dots \cup X_k) \times Y$  that expresses whether the  $y$ -th coordinate of  $x_i$  is 0 or 1; see Section 2 for details). In particular, this class of problems includes:

- (Offline) Furthest/Nearest Neighbor search in the Hamming metric  
 $(\max_{x_1, x_2} / \min_{x_1, x_2} d_H(x_1, x_2)),$
- Maximum/Minimum Inner Product; the latter is the optimization formulation of the Most-Orthogonal Vectors problem [2]  
 $(\max_{x_1, x_2} / \min_{x_1, x_2} \langle x_1, x_2 \rangle),$
- A natural similarity search problem that we call Maximum  $k$ -Agreement  
 $(\max_{x_1, \dots, x_k} \#\{y : x_1[y] = \dots = x_k[y]\}),$
- Maximum  $k$ -Cover [35, 31, 59] and its variation Maximum Exact- $k$ -Cover [54]  
 $(\max_{x_1, \dots, x_k} \#\{y : x_i[y] = 1 \text{ for some } i\}, \max_{x_1, \dots, x_k} \#\{y : x_i[y] = 1 \text{ for exactly one } i\})^2,$
- The canonical optimization variants of the  $k$ -XOR problem [50, 33]  
 $(\max_{x_1, \dots, x_k} / \min_{x_1, \dots, x_k} \#\{y : x_1[y] \oplus \dots \oplus x_k[y] = 0\}).$

By the standard split-and-list technique [66], it is easy to see that any  $c$ -approximation for Maximum  $k$ -XOR or Minimum  $k$ -XOR in time  $O(n^{k-\delta} \text{poly}(d))$  gives a  $c$ -approximation for MAX-LIN (maximize the number of satisfied constraints of a linear system over  $\mathbb{F}_2$ , see [45, 66, 6]) or the Minimum Distance Problem<sup>3</sup> (finding the minimum weight of a non-zero code word of a linear code over  $\mathbb{F}_2$ , see [63]), respectively, in time  $O(2^{n(1-\delta')})$ .

By a simple baseline algorithm, all of these problems can be solved in time  $O(m^k)$ , where  $m$  denotes the input size (for the above setting of  $k$  sets of  $n$  vectors in  $\{0, 1\}^d$  we have  $m = O(nd)$ ). A large body of work addresses problems with  $k = 2$ , typically inventing or adapting strong techniques to each specific problem as needed:

- Abboud, Rubinfeld, and Williams introduced the distributed PCP in P framework and ruled out almost-polynomial approximations for the Maximum Inner Product problem assuming the Strong Exponential Time Hypothesis (SETH). Subsequently, Chen [28] strengthened the lower bound and gave an approximation algorithm resulting in tight bounds on the approximability in strongly subquadratic time, assuming SETH. Corresponding inapproximability results for its natural generalization to  $k$ -Maximum Inner Product have been obtained in [51].
- In contrast, strong approximation algorithms have a rich history for the Nearest Neighbor search problem: Using LSH, we can obtain a  $(1 + \varepsilon)$ -approximation in time  $O(n^{2-\Theta(\varepsilon)})$  [44, 9, 10, 15]. Using further techniques, the dependence on  $\varepsilon$  has been improved to  $O(n^{2-\Omega(\sqrt{\varepsilon})})$  [64] and  $O(n^{2-\tilde{\Omega}(\frac{1}{\sqrt{\varepsilon}})})$  [5, 6]. On the hardness side, Rubinfeld [62] shows that the dependence on  $\varepsilon$  cannot be improved indefinitely, by proving that for every  $\delta$ , there exists an  $\varepsilon$  such that  $(1 + \varepsilon)$ -approximate Nearest Neighbor search requires time  $\Omega(n^{2-\delta})$  assuming SETH. In particular, this rules out  $\text{poly}(1/\varepsilon)n^{2-\delta}$ -time algorithms with  $\delta > 0$ .

<sup>1</sup> Note that [23] more generally defines classes  $\text{MAXSP}_{k,\ell}$  for  $\ell \geq 1$ . We focus on  $\text{MAXSP}_k = \text{MAXSP}_{k,1}$ , as it was determined as computationally harder than  $\text{MAXSP}_{k,\ell}$  with  $\ell \geq 2$ , and contains many natural problems (see below).

<sup>2</sup> These problems are typically studied in the setting where  $k$  is part of the input, while we consider them for a fixed constant  $k$ .

<sup>3</sup> In fact, even for the Nearest Codeword Problem over  $\mathbb{F}_2$ .

- For the dual problem of Furthest Neighbor search (i.e., diameter in the Hamming metric), [20] gives a  $O(n^{2-\Theta(\varepsilon^2)})$ -time algorithm, which was improved to  $O(n^{2-\Theta(\varepsilon)})$  via reduction to Nearest Neighbor search in [47]. Following further improvements [42, 48], also here [5, 6] give an  $O(n^{2-\tilde{\Omega}(\sqrt[3]{\varepsilon})})$ -time algorithm. Analogous inapproximability to Rubinfeld's result are given in [30].
- For Minimum Inner Product, the well-known Orthogonal Vectors hypothesis [65] (which is implied by SETH [66]) is precisely the assumption that already distinguishing whether the optimal value is 0 or at least 1 cannot be done in strongly subquadratic time. Interestingly, Chen and Williams [30] show a converse: a refutation of the Orthogonal Vectors hypothesis would give a subquadratic-time constant-factor approximation for Minimum Inner Product.

In the above list, we focus on the difficult case of moderate dimension  $d = n^{o(1)}$  (when measuring the time complexity with respect to the input size). Lower-dimensional settings such as  $d = \Theta(\log n)$ ,  $d = \Theta(\log \log n)$  or even lower are addressed in other works [69, 30].

While the above collection of results gives a detailed understanding of isolated problems, we know little about general phenomena of (in)approximability in  $\text{MAXSP}_k$  and  $\text{MINSP}_k$  using faster-than-exhaustive-search algorithms: Are there problems for which constant-factor approximations are best possible? Is maximizing (or minimizing) Inner Product the only problem without a constant-factor approximation? Which problems can we optimize exactly?

There are precursors to our work that show fine-grained *equivalence classes* of approximation problems in P [29, 30]. However, establishing membership of a problem in these classes requires a problem-specific proof, while we are interested in *syntactically* defined classes, where class membership can be immediately read off from the definition of a problem. Our aim is to understand the approximability landscape of such classes fully. Finally, the previous works either focus on lower-dimensional settings [30]<sup>4</sup>, or target more powerful problems than we consider, such as Closest-LCS-Pair [29].

**MaxSP<sub>k</sub> as Polynomial-Time Analogue of MaxSNP.** Investigating NP-hard optimization problems, Papadimitriou and Yannakakis [60] introduced the class MAXSNP, which motivates the definition of  $\text{MAXSP}_k$  as a natural polynomial-time analogue (see [23] and the full version of this paper for details). As a general class containing prominent, constant-factor approximable optimization problems, MAXSNP was introduced to give the first evidence that MAX-3-SAT does not admit a PTAS, by proving that MAX-3-SAT belongs to the hardest-to-approximate problems in MAXSNP.

Ideally, one would like to understand the approximability landscape in MAXSNP fully and give tight approximability results for each such problem. Major advances towards this goal have been achieved by Creignou [32] and Khanna, Sudan, Trevisan, and Williamson [53] who gave a complete classification of a large subclass of MAXSNP, namely, maximum Boolean Constraint Satisfaction Problems (MAXCSP): Each Boolean MAXCSP either is polynomial-time optimizable or it does not admit a PTAS unless  $P = NP$ , rendering a polynomial-time constant-factor approximation best possible. For minimization analogues, including Boolean MINCSPs, the situation is more diverse with several equivalence classes needed to describe the result [53].

We initiate the study of the same type of questions in the polynomial-time regime. Our aim is to achieve a detailed understanding of  $\text{MAXSP}_k$  and  $\text{MINSP}_k$  akin to the classification theorems achieved for MAXCSPs and MINCSPs [32, 53].

---

<sup>4</sup> Chen and Williams also give some results for the moderate-dimensional case; we discuss these in Section 3.

## 1.1 Our Results

We approach the classification of  $\text{MAXSP}_k$  and  $\text{MINSP}_k$  by considering the simplest, yet expressive case of a *single, binary relation* involved in the first-order formula; this route was also taken in earlier classification work for first-order properties [24]. We may thus view the relational structure as a graph, and call such a formula a *graph formula*. Note that despite its naming, this includes problems not usually viewed as graph problems, such as all examples given in the introduction, which are natural problems on sets of *vectors* in  $\{0, 1\}^d$ .

We obtain a classification of each graph formula into one of four regimes, assuming central fine-grained hardness hypotheses whose plausibility we detail in the full version of this paper. All of our hardness results are implied by the Sparse MAX-3-SAT hypothesis [7], which states that for all  $\delta > 0$  there is some  $c > 0$  such that MAX-3-SAT on  $n$  variables and  $cn$  clauses has no  $O(2^{n(1-\delta)})$ -time algorithm.<sup>5</sup> (Actually, most of our hardness results already follow from weaker assumptions.)

► **Theorem 1.** *Let  $\psi$  be a  $\text{MAXSP}_k$  or  $\text{MINSP}_k$  graph formula. Assuming the Sparse MAX-3-SAT hypothesis,  $\psi$  belongs to precisely one of the following regimes:*

**R1 Efficiently optimizable:**

*There is some  $\delta > 0$  such that  $\psi$  can be solved exactly in time  $O(m^{k-\delta})$ .*

**R2 Admits an approximation scheme, but not an efficient one:**

*For all  $\varepsilon > 0$ , there is some  $\delta > 0$  such that  $\psi$  can be  $(1 + \varepsilon)$ -approximated in time  $O(m^{k-\delta})$ . However, for all  $\delta > 0$ , there is some  $\varepsilon > 0$  such that  $\psi$  cannot be  $(1 + \varepsilon)$ -approximated in time  $O(m^{k-\delta})$ .*

**R3 Admits a constant-factor approximation, but no approximation scheme:**

*There are  $\varepsilon, \delta > 0$  such that  $\psi$  can be  $(1 + \varepsilon)$ -approximated in time  $O(m^{k-\delta})$ . However, there also exists an  $\varepsilon > 0$  such that for all  $\delta > 0$ ,  $\psi$  cannot be  $(1 + \varepsilon)$ -approximated in time  $O(m^{k-\delta})$ .*

**R4 Arbitrary polynomial-factor approximation is best possible (maximization):**

*For every  $\varepsilon > 0$ , there is some  $\delta > 0$  such that  $\psi$  can be  $O(m^\varepsilon)$ -approximated in time  $O(m^{k-\delta})$ . However, for every  $\delta > 0$ , there exists some  $\varepsilon > 0$  such that  $\psi$  cannot be  $O(m^\varepsilon)$ -approximated in time  $O(m^{k-\delta})$ .*

**No approximation at all (minimization):**

*For all  $\delta > 0$ , we cannot decide whether the optimum value of  $\psi$  is 0 or at least 1 in time  $O(m^{k-\delta})$ .*

Note that the characteristics of the fourth (i.e., hardest) regime differ between the maximization and the minimization case.

This theorem has immediate consequences for the approximability landscape in  $\text{MAXSP}_k$  and  $\text{MINSP}_k$ , based on fine-grained assumptions: In particular, while any  $\text{MAXSP}_k$  graph formula can be approximated within a subpolynomial factor  $O(m^\varepsilon)$ , we can rule out optimal approximation ratios that *grow with  $m$  but are strictly subpolynomial* (i.e., there are *no* graph formulas whose optimal approximability within  $O(m^{k-\Omega(1)})$  time is  $\Theta(\log \log m)$ ,  $\Theta(\log^2 m)$  or  $2^{\Theta(\sqrt{\log m})}$ ). Furthermore, there are no graph formulas with an  $f(1/\varepsilon)m^{k-\Omega(1)}$  approximation scheme that cannot already be optimized to exactness in time  $m^{k-\Omega(1)}$ .

In fact, beyond Theorem 1 we even give an almost complete characterization of each regime: Specifically, we introduce integer-valued hardness parameters  $0 \leq H_{\text{and}}(\psi) \leq H_{\text{deg}}(\psi) \leq k$  (defined in Section 3). As illustrated in Figure 1, we show how to place any graph formula  $\psi$

<sup>5</sup> This hypothesis is a stronger version of the MAX-3-SAT hypothesis [58].

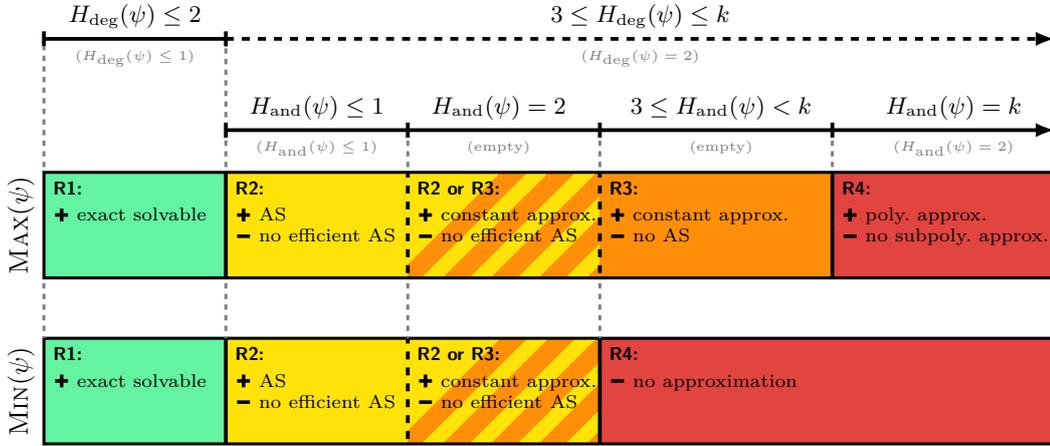


Figure 1 Visualizes our classification of first-order optimization problems  $\text{MAX}(\psi)$  and  $\text{MIN}(\psi)$  for all  $k \geq 3$ , in terms of the hardness parameters  $H_{\text{and}}$  and  $H_{\text{deg}}$  (as defined in Definitions 4 and 6). See Definition 10 for the precise definition of an (efficient) approximation scheme (AS). The pale labels indicate how to change the conditions to obtain the picture for  $k = 2$ .

Table 1 Some interesting examples classified according to Figure 1. For each problem  $\psi$ , an instance consists of  $k$  sets of  $n$  vectors  $X_1, \dots, X_k \subseteq \{0, 1\}^d$ . We write  $\psi = \max_{x_1, \dots, x_k} / \min_{x_1, \dots, x_k} \phi$  and only list the inner formulas  $\phi$  in the table.

Problem $\psi$	$\phi$	$H_{\text{deg}}(\psi)$	$H_{\text{and}}(\psi)$	Hardness regime
Maximum/Minimum 2-Agreement	$\#\{y : x_1[y] = x_2[y]\}$	2	1	<b>R2</b>
Maximum/Minimum 3-Agreement	$\#\{y : x_1[y] = x_2[y] = x_3[y]\}$	2	2	<b>R1</b>
Maximum $k$ -Agreement, $k \geq 4$	$\#\{y : x_1[y] = \dots = x_k[y]\}$	$2 \lfloor \frac{k}{2} \rfloor$	$k - 1$	<b>R3</b>
Minimum $k$ -Agreement, $k \geq 4$	$\#\{y : x_1[y] = \dots = x_k[y]\}$	$2 \lfloor \frac{k}{2} \rfloor$	$k - 1$	<b>R4</b>
Maximum/Minimum $k$ -XOR	$\#\{y : x_1[y] \oplus \dots \oplus x_k[y]\}$	$k$	1	<b>R2</b>
Maximum $k$ -Inner Product	$\#\{y : x_1[y] \wedge \dots \wedge x_k[y]\}$	$k$	$k$	<b>R4</b>
Minimum $k$ -Inner Product	$\#\{y : x_1[y] \wedge \dots \wedge x_k[y]\}$	$k$	$k$	<b>R4</b>

into its corresponding regime depending solely on  $H_{\text{and}}(\psi)$  and  $H_{\text{deg}}(\psi)$  – with the *single exception* of formulas  $\psi$  with  $H_{\text{and}}(\psi) = 2$  and  $H_{\text{deg}}(\psi) \geq 3$ . For these formulas (e.g., Maximum Exact-3-Cover), it remains open whether they belong to Regime 2 or 3, i.e., whether or not they admit an approximation scheme (see Open Problem 1).

In the full version we give natural problems for each regime (showing in particular that each regime is indeed non-empty). A particular highlight is that we can prove existence of *constant-factor approximable formulas that do not admit an approximation scheme*, such as Maximum-4-Agreement (see Section 3, Theorem 11 for a technical discussion of the lower bound). In the full version of this paper, we give the details on how to calculate the hardness parameters  $H_{\text{and}}(\psi)$  and  $H_{\text{deg}}(\psi)$  for each example in Table 1.

While the approximability of problems in the third and fourth regime seem rather unsatisfactory, we also consider the setting of additive approximation, and show that for every  $\text{MAXSP}_k$  and  $\text{MINSP}_k$  graph formula, there is an additive approximation scheme; however, assuming the Sparse MAX-3-SAT hypothesis it cannot be an efficient one unless we can optimize the problem exactly.

► **Theorem 2 (Additive Approximation).** *For every  $\psi$ , we give an additive approximation scheme, i.e., for every  $\varepsilon > 0$ , there is a  $\delta > 0$  such that we compute the optimum value up to an additive  $\varepsilon|Y|$  error in time  $O(m^{k-\delta})$ .*

*If  $\psi$  does not belong to Regime 1, we show that unless the Sparse MAX-3-SAT hypothesis fails, for every  $\delta > 0$ , there is some  $\varepsilon > 0$  such that we cannot compute the optimum value up to an additive  $\varepsilon|Y|$  error in time  $O(m^{k-\delta})$ .*

Finally, we remark that our classification identifies general trends in  $\text{MAXSP}_k$  and  $\text{MINSP}_k$ , including that exact optimizability is described by a simple algebraic criterion, and that constant-factor approximation for minimization problems is equivalent to testing whether the optimum is 0. We address these trends in Section 3.

**On Plausibility of the Hypotheses.** As tight unconditional lower bounds for polynomial-time problems are barely existent, we give *conditional* lower bounds, based on established assumptions in fine-grained complexity theory, such as SETH (see [65] for an excellent survey). The essentially only exception is the only recently introduced *Sparse MAX-3-SAT* hypothesis [7]; we use this hypothesis only for giving evidence against approximation schemes and for ruling out certain additive approximation schemes. While it is possible that this hypothesis could ultimately be refuted, our classification describes the frontier of the current state of the art. At the very least, our conditional lower bound for approximation schemes reveals a rather surprising connection: To obtain an *approximation scheme* for *polynomial-time problems* such as Maximum 4-Agreement, we need to give an *exponential-time* improvement for *exact* solutions for Sparse MAX-3-SAT!

## 1.2 Further Related Work

**Parameterized Inapproximability.** The problem settings we consider are related to a recent and strong line of research on parameterized inapproximability, including [27, 51, 31, 59, 57] (see [36] for a recent survey). In these contexts, one seeks to determine optimal approximation guarantees within some running time of the form  $f(k)n^{g(k)}$  (such as FPT running time  $f(k)\text{poly}(n)$  or running time  $n^{o(k)}$ ) for some parameter  $k$  (such as the solution size). Unfortunately, many of these results do not immediately establish hardness for specific values of  $k$ . An important exception is work by Karthik, Laekhanukit, and Manurangsi [51], which among other results establishes inapproximability of  $k$ -Dominating Set and the Maximum  $k$ -Inner Product within running time  $O(n^{k-\varepsilon})$ , assuming SETH. We give a detailed comparison of our setting to notions used in Karthik et al.’s work in the full version of this work.

**Fine-Grained Complexity of First-Order Properties.** Studying the fine-grained complexity of polynomial-time problem classes defined by first-order properties has been initiated in [68, 40], with different settings considered in [39, 38]. In particular, for model-checking first-order properties, [40] provides a completeness result (a fine-grained analogue of the Cook-Levin Theorem) and [24] provides a classification theorem (a fine-grained analogue of Schaefer’s Theorem, see also Section 3.1).

For *optimization* in P, [23] provides completeness theorems for  $\text{MAXSP}_k$  (a fine-grained analogue of MAXSNP-completeness of MAX-3-SAT [60]). In contrast, the present paper gives a classification theorem (building a fine-grained analogue of the approximability classifications of optimization variants of CSPs [32, 53]).

## 2 Preliminaries

For an integer  $k \geq 0$ , we set  $[k] = \{1, \dots, k\}$ . For a set  $A$  and integer  $k \geq 0$  we denote by  $\binom{A}{k}$  the set of all size- $k$  subsets of  $A$ . Let  $\phi(z_1, \dots, z_k)$  a Boolean function, and let  $S \subseteq [k]$ . Any function obtained by instantiating all variables  $z_i$  ( $i \notin S$ ) in  $\phi$  by constant values is called an  $S$ -restriction of  $\phi$ . Finally, we write  $\tilde{O}(\cdot)$  to hide poly-logarithmic factors and denote by  $\omega < 2.373$  the exponent of square matrix multiplication [8].

### 2.1 First-Order Model-Checking

A *relational structure* consists of  $n$  objects and relations  $R_1, \dots, R_\ell$  (of arbitrary arities) between these objects. A *first-order formula* is a quantified formula of the form

$$\psi = (Q_1 x_1) \dots (Q_{k+1} x_{k+1}) \phi(x_1, \dots, x_{k+1}),$$

where  $\phi$  is a Boolean formula over the predicates  $R(x_{i_1}, \dots, x_{i_\ell})$  and each  $Q_i$  is either a universal or existential quantifier. Given a  $(k+1)$ -partite structure on objects  $X_1 \cup \dots \cup X_{k+1}$ , the *model-checking problem* (or *query evaluation problem*) is to check whether  $\psi$  holds on the given structure, that is, for  $x_i$  ranging over  $X_i$  and by instantiating the predicates  $R(x_{i_1}, \dots, x_{i_\ell})$  in  $\phi$  according to the structure,  $\psi$  is valid.

Following previous work in this line of research [40, 24], we usually assume that the input is represented sparsely – that is, we assume that the relational structure is written down as an exhaustive enumeration of all records in all relations; let  $m$  denote the total number of such entries<sup>6</sup>. The convention is reasonable as this data format is common in the context of database theory and also for the representation of graphs (where it is called the *adjacency list* representation), see Section 4 for a further discussion.

A first-order formula  $\psi$  is called a *graph formula* if it is defined over a single binary predicate  $E(x_i, x_j)$ . Many natural problems fall into this category; see [24] for a detailed discussion on the subject.

### 2.2 $\text{MaxSP}_k$ and $\text{MinSP}_k$

In analogy to first-order properties with quantifier structure  $\exists^k \forall$  (with maximization instead of  $\exists$  and counting instead of  $\forall$ ) and following the definition in [23], we now introduce the class of optimization problems. Let  $\text{MAXSP}_k$  be the class containing all formulas of the form

<sup>6</sup> By ignoring objects not occurring in any relation, we may always assume that  $n \leq O(m)$ .

$$\psi = \max_{x_1, \dots, x_k} \#_y \phi(x_1, \dots, x_k, y), \quad (1)$$

where, as before,  $\phi$  is a Boolean formula over some predicates  $R(x_{i_1}, \dots, x_{i_\ell})$ . We similarly define  $\text{MINSP}_k$  with “min” in place of “max”. Occasionally, we write  $\text{OPTSP}_k$  to refer to both of these classes simultaneously, and we write “opt” as a placeholder for either “max” or “min”. In analogy to the model-checking problem for first-order properties, we associate to each formula  $\psi \in \text{OPTSP}_k$  an algorithmic problem:

► **Definition 3** ( $\text{MAX}(\psi)$  and  $\text{MIN}(\psi)$ ). *Let  $\psi \in \text{MAXSP}_k$  be as in (1). Given a  $(k + 1)$ -partite structure on objects  $X_1 \cup \dots \cup X_k \cup Y$ , the  $\text{MAX}(\psi)$  problem is to compute*

$$\text{OPT} = \max_{x_1 \in X_1, \dots, x_k \in X_k} \#_{y \in Y} \phi(x_1, \dots, x_k, y).$$

We similarly define  $\text{MIN}(\psi)$  for  $\psi \in \text{MINSP}_k$ . Occasionally, for  $\psi \in \text{OPTSP}_k$ , we write  $\text{OPT}(\psi)$  to refer to both problems simultaneously.

For convenience, we introduce some further notation: For objects  $x_1 \in X_1, \dots, x_k \in X_k$ , we denote by  $\text{Val}(x_1, \dots, x_k) = \#_{y \in Y} \phi(x_1, \dots, x_k, y)$  the value of  $(x_1, \dots, x_k)$ .

The problem  $\text{OPT}(\psi)$  can be solved in time  $O(m^k)$  for all  $\text{OPTSP}_k$  formulas  $\psi$ , by a straightforward extension of the model-checking baseline algorithm; see [23] for details. As this is clearly optimal for  $k = 1$ , we will often implicitly assume that  $k \geq 2$  in the following.

For a clean analogy between model-checking and the optimization classes  $\text{MAXSP}_k$  and  $\text{MINSP}_k$ , we will from now view model-checking as “testing for zero”. More precisely, the model-checking problem of  $\exists x_1 \dots \exists x_k \forall y \neg \phi(x_1, \dots, x_k, y)$  is equivalent to testing whether  $\text{MIN}(\psi)$  has optimal solution  $\text{OPT} = 0$ , where  $\psi = \min_{x_1, \dots, x_k} \#_y \phi(x_1, \dots, x_k, y)$ . We refer to the latter problem as  $\text{ZERO}(\psi)$ .

Definition 3 introduces  $\text{MAX}(\psi)$  and  $\text{MIN}(\psi)$  as *exact* optimization problems (i.e.,  $\text{OPT}$  is required to be computed exactly). We say that an algorithm computes a *(multiplicative)  $c$ -approximation* for  $\text{MAX}(\psi)$  if it computes any value in the interval  $[c^{-1} \cdot \text{OPT}, \text{OPT}]$ . Similarly, a *(multiplicative)  $c$ -approximation* for  $\text{MIN}(\psi)$  computes a value in  $[\text{OPT}, c \cdot \text{OPT}]$ .

### 3 Technical Overview

This section serves the purpose of stating our results formally, to provide the main proof ideas and techniques, and to convey some intuition whenever possible. Due to space constraints, we omit precise definitions of the fine-grained hypotheses here and instead refer to the full version of this paper.

We first introduce our hardness parameters: the and-hardness  $H_{\text{and}}(\psi)$  borrowed from previous work [24] (reviewed below), as well as a novel algebraic parameter that we call degree-hardness  $H_{\text{deg}}(\psi)$ . In the subsequent sections, we give an overview over our various algorithmic and hardness results based on the values  $H_{\text{deg}}(\psi), H_{\text{and}}(\psi)$  (see Figure 1), with the formal proof of Theorem 1 given at the end of this section.

#### 3.1 Bringmann et al.’s Model-Checking Dichotomy

Bringmann, Fischer and Künnemann [24] established a fine-grained classification of all  $\exists^k \forall$ -quantified graph properties into computationally easy and hard model-checking problems. As our work extends that classification (and also since our results are of a similar flavor), we briefly summarize their results. The hardness parameter presented in Definition 4 forms the basis for the dichotomy.

Here, and for the remainder of this section, we write  $\psi_0$  to denote the Boolean function obtained from  $\psi$  in the following way: Let  $\psi = \text{opt}_{x_1, \dots, x_k} \#_y \phi(x_1, \dots, x_k, y)$  be any graph formula, where  $\phi$  is a propositional formula over the atoms  $E(x_i, x_j)$  ( $i, j \in [k]$ ) and  $E(x_i, y)$  ( $i \in [k]$ ). Consider the Boolean function obtained from  $\phi$  by replacing every atom  $E(x_i, x_j)$  ( $i, j \in [k]$ ) by *false*. What remains is a Boolean function over the  $k$  atoms  $E(x_i, y)$  ( $i \in [k]$ ) and we denote this formula by  $\psi_0$ . For example, if  $\psi = \max_{x_1, x_2} \#_y (\neg E(x_1, x_2) \wedge E(x_1, y) \wedge E(x_2, y))$  then we define  $\psi_0 : \{0, 1\}^2 \rightarrow \{0, 1\}$  by  $\psi_0(a_1, a_2) = a_1 \wedge a_2$ .

As apparent in Definition 4, the core difficulty of a formula  $\psi$  is captured by  $\psi_0$  and thus not affected by the predicates  $E(x_i, x_j)$ . In the full version we elaborate on this phenomenon.

► **Definition 4 (And-Hardness).** *Let  $\phi$  be a Boolean function on  $k$  inputs. The and-hardness  $H_{\text{and}}(\phi)$  of  $\phi$  is the largest integer  $0 \leq h \leq k$  such that, for any index set  $S \in \binom{[k]}{h}$ , there exists some  $S$ -restriction of  $\phi$  with exactly one satisfying assignment. (Set  $H_{\text{and}}(\phi) = 0$  for constant-valued  $\phi$ .) For an  $\text{OPTSP}_k$  graph formula  $\psi$ , we define  $H_{\text{and}}(\psi) = H_{\text{and}}(\psi_0)$ .*

This hardness parameter essentially specifies the computational hardness of the model-checking problems  $\text{ZERO}(\phi)$ ; here,  $H_{\text{and}}(\psi) \leq 2$  is the critical threshold:

- **Theorem 5 (Model-Checking, [24]).** *Let  $\psi$  be a  $\text{MINSP}_k$  graph formula.*
- *If  $H_{\text{and}}(\psi) \leq 2$  and  $H_{\text{and}}(\psi) < k$ , then  $\text{ZERO}(\psi)$  can be solved in time  $O(m^{k-\delta})$  for some  $\delta > 0$ .*
  - *If  $3 \leq H_{\text{and}}(\psi)$  or  $H_{\text{and}}(\psi) = k$ , then  $\text{ZERO}(\psi)$  cannot be solved in time  $O(m^{k-\delta})$  for any  $\delta > 0$  unless the  $\text{MAX-3-SAT}$  hypothesis fails.*

## 3.2 Exact Optimization

We are now ready to detail our results. Our main contribution is a dichotomy for the exact solvability and approximability of  $\text{MAX}(\psi)$  and  $\text{MIN}(\psi)$  for all  $\text{MAXSP}_k$  and  $\text{MINSP}_k$  graph formulas  $\psi$ . For the exact case, the decisive criterion for the hardness of some formula  $\psi$  can be read off the polynomial *extension* of the function  $\psi_0$ . Specifically, for any Boolean function  $\phi$  there exists a (unique) multilinear polynomial with real coefficients that computes  $\phi$  on binary inputs. By abuse of notation, we refer to that polynomial by writing  $\phi$  as well. The *degree*  $\text{deg}(\phi)$  of  $\phi$  is the degree of its polynomial extension.

As an example, consider the Exact-3-Cover property:  $\phi(z_1, z_2, z_3)$  is true if and only if exactly one of its inputs  $z_1, z_2$  or  $z_3$  is true. Then its unique multilinear polynomial extension is

$$\phi(z_1, z_2, z_3) = 3z_1z_2z_3 - 2(z_1z_2 + z_2z_3 + z_3z_1) + (z_1 + z_2 + z_3),$$

and therefore  $\text{deg}(\phi) = 3$ .

► **Definition 6 (Degree Hardness).** *Let  $\phi$  be a Boolean function on  $k$  inputs. The degree hardness  $H_{\text{deg}}(\phi)$  of  $\phi$  is the largest integer  $0 \leq h \leq k$  such that, for any index set  $S \in \binom{[k]}{h}$ , there exists some  $S$ -restriction of  $\phi$  of degree  $h$ . For an  $\text{OPTSP}_k$  graph formula  $\psi$ , we define  $H_{\text{deg}}(\psi) = H_{\text{deg}}(\psi_0)$ .*

It always holds that  $H_{\text{and}}(\psi) \leq H_{\text{deg}}(\psi)$ , but in general these parameters behave very differently. With  $H_{\text{deg}}$  in place of  $H_{\text{and}}$ , we are able to recover the same classification as Theorem 5 for both exact maximization and minimization:

► **Theorem 7** (Exact Optimization). *Let  $\psi$  be an  $\text{OPTSP}_k$  graph formula.*

- *If  $H_{\text{deg}}(\psi) \leq 2$  and  $H_{\text{deg}}(\psi) < k$ , then  $\text{OPT}(\psi)$  can be solved in time  $O(m^{k-\delta})$  for some  $\delta > 0$ .*
- *If  $3 \leq H_{\text{deg}}(\psi)$  or  $H_{\text{deg}}(\psi) = k$ , then  $\text{OPT}(\psi)$  cannot be solved in time  $O(m^{k-\delta})$  for any  $\delta > 0$  unless the MAX-3-SAT hypothesis fails.*

The algorithmic part is proven along the same lines as [24]: We first brute-force over all but  $k = 3$  quantifiers, and solve the remaining problem by a reduction to maximum-weight triangle detection with small edge weights. Intuitively, since the degree of the resulting 3-variable problem is at most 2, we can express the objective as a sum of three parts depending on two variables each (this approach is also used in a similar context in [67, Chapter 6.5]). This allows us to label the edges of a triangle instance with corresponding parts, which can be shown to yield average edge weight  $O(1)$ .

The conditional lower bound is more interesting. Our reduction is inspired by a standard argument proving quadratic-time hardness of the Maximum Inner Product problem (MAXIP). That lower bound is based on the OV hypothesis, so it consists of a reduction from an OV instance  $X_1, X_2 \subseteq \{0, 1\}^d$  to an instance of MAXIP. The idea is to use a *gadget* that maps every entry in  $x_i \in X_i$  to a constant number of new entries;<sup>7</sup> let  $x'_i \in \{0, 1\}^{O(d)}$  denote the vector after applying the gadget coordinate-wise. The crucial property is that  $\langle x'_1, x'_2 \rangle = d - \langle x_1, x_2 \rangle$ , and thus a pair of orthogonal vectors  $x_1, x_2$  corresponds to a pair of vectors  $x'_1, x'_2$  of maximum inner product.

To mimic the reduction for all problems which are hard in the sense of Theorem 7, we settle for the weaker but sufficient property that the value of  $\langle x'_1, x'_2 \rangle$  equals  $\beta_1 d - \beta_2 \langle x_1, x_2 \rangle$ , for some positive integers  $\beta_1, \beta_2$ . It follows from our algebraic characterization of hard functions that a gadget with such guarantees always exists. Ultimately, our hardness proof makes use of that gadget in a similar way as for MAXIP.

The hardness part of Theorem 7 can in fact be stated in a more fine-grained way: If some problem  $\text{OPT}(\psi)$  has degree hardness  $h = H_{\text{deg}}(\psi) \geq 3$ , then the hardness proof can be conditioned on the weaker  $h$ -UNIFORM HYPERCLIQUE assumption. That connection can be complemented by a partial converse, thus revealing a certain *equivalence* between exact optimization and hyperclique detection. An analogous equivalence for model-checking graph formulas could not be proved and was left as an open problem in [24].

► **Theorem 8** (Equivalence of  $\text{OPT}(\psi)$  and HYPERCLIQUE). *Let  $\psi$  be an  $\text{OPTSP}_k$  graph formula of degree hardness  $h = H_{\text{deg}}(\psi) \geq 2$ .*

- *If  $\text{OPT}(\psi)$  can be solved in time  $O(m^{k-\delta})$  for some  $\delta > 0$ , then, for some (large)  $k' = k'(k, h, \delta)$ ,  $h$ -UNIFORM  $k'$ -HYPERCLIQUE can be solved in time  $O(n^{k'-\delta'})$  for some  $\delta' > 0$ .*
- *If  $h$ -UNIFORM  $(h + 1)$ -HYPERCLIQUE can be solved in time  $O(n^{h+1-\delta})$  for some  $\delta > 0$ , then  $\text{OPT}(\psi)$  can be solved in time  $O(m^{k-\delta'})$  for some  $\delta' > 0$ .*

### 3.3 Approximation

Since Theorem 7 gives a complete classification for the exact solvability of  $\text{OPT}(\psi)$ , the next natural question is to study the approximability of properties which are hard to compute exactly. Unlike the exact case, we use different techniques and tools to give our classification for maximization and minimization problems.

<sup>7</sup> Every entry  $a$  in  $x_1$  is replaced by three new coordinates  $(a, 1 - a, 1 - a)$  and every entry  $b$  in  $x_2$  is replaced by  $(1 - b, b, 1 - b)$ . The contribution to the inner product of the new vectors is equal to  $a(1 - b) + (1 - a)b + (1 - a)(1 - b) = 1 - ab$ .

### 3.3.1 Maximization

We obtain a simple classification of all constant-factor approximable maximization problems, conditioned on the Strong Exponential Time Hypothesis (SETH).

- **Theorem 9** (Constant Approximation – Maximization). *Let  $\psi$  be a  $\text{MAXSP}_k$  graph formula.*
- *If  $H_{\text{and}}(\psi) < k$  (or equivalently, if  $\psi_0$  does not have exactly one satisfying assignment), then there exists a constant-factor approximation for  $\text{MAX}(\psi)$  in time  $O(m^{k-\delta})$  for some  $\delta > 0$ .*
- *Otherwise, if  $H_{\text{and}}(\psi) = k$  (or equivalently,  $\psi_0$  has exactly one satisfying assignment), then there exists no constant-factor approximation for  $\text{MAX}(\psi)$  in time  $O(m^{k-\delta})$  for any  $\delta > 0$ , unless SETH fails.*

We give a high-level explanation of our proof by focusing on two representative problems: On the one hand, strong conditional hardness results have been shown for the MAXIP problem [4, 28]. When  $\psi_0$  has a single satisfying assignment,  $\text{MAX}(\psi)$  is equivalent (up to complementation) to MAXIP, so we adapt the hardness results for our setting.

On the other hand, consider the FURTHEST NEIGHBOR problem: Given two sets of bit-vectors  $X_1, X_2 \subseteq \{0, 1\}^d$ , compute the maximum Hamming distance between vectors  $x_1 \in X_1$  and  $x_2 \in X_2$ . There exists a simple linear-time 3-approximation for this problem: Fix some  $x_1 \in X_1$  and compute its furthest neighbor  $x_2 \in X_2$ . Then compute the furthest neighbor  $x'_1 \in X_1$  of  $x_2$  and return the distance between  $x'_1$  and  $x_2$  as the answer. By applying the triangle inequality twice, it is easy to see that this indeed yields a 3-approximation. That argument generalizes for approximating  $\text{MAX}(\phi)$  whenever  $\psi_0$  satisfies the following property: If  $\alpha$  is a satisfying assignment of  $\psi_0$ , then the component-wise negation of  $\alpha$  is also satisfying. Finally, if  $\psi_0$  has at least two satisfying assignments, then  $\text{MAX}(\psi)$  can be reduced to this special case via a reduction which worsens the approximation ratio by at most a constant factor. The essential insight for that last step is that we can always “cover” all satisfying assignments by only two satisfying assignments, as there always exists one satisfying assignment which contributes a constant fraction (depending on  $k$ ) to the optimal value.

We give a finer-grained view of the classification in Theorem 9 in two ways. First, we want to isolate properties which admit *arbitrarily good* constant-factor approximations. We make that notion precise in the following definition:

- **Definition 10** (Approximation Scheme). *Let  $\psi$  be an  $\text{OPTSP}_k$  formula. We say that  $\text{OPT}(\psi)$  admits an approximation scheme if for any  $\varepsilon > 0$  there exists some  $\delta > 0$  and an algorithm computing a  $(1 + \varepsilon)$ -approximation of  $\text{OPT}(\psi)$  in time  $O(m^{k-\delta})$ .<sup>8</sup>*

In the following theorem, we identify some formulas which admit such an approximation scheme, and some formulas for which this is unlikely:

- **Theorem 11** (Approximation Scheme – Maximization). *Let  $\psi$  be a  $\text{MAXSP}_k$  graph formula.*
- *If  $H_{\text{and}}(\psi) \leq 1$ , then there exists a randomized approximation scheme for  $\text{MAX}(\psi)$ .*
- *If  $3 \leq H_{\text{and}}(\psi)$  or  $H_{\text{and}}(\psi) = k$ , then there exists no approximation scheme for  $\text{MAX}(\psi)$  unless the Sparse MAX-3-SAT hypothesis fails.*

Unfortunately, we were not able to close the gap in Theorem 11, and it remains open whether problems with  $H_{\text{and}}(\phi) = 2$  admit approximation schemes. In Section 4, we give a specific example falling into this category.

---

<sup>8</sup> We stress that in our work,  $\varepsilon$  is always a constant and cannot depend on  $m$ .

For the first item, we give approximation schemes for any  $\text{MAX}(\psi)$  with  $H_{\text{and}}(\psi) = 1$  via a reduction to FURTHEST NEIGHBOR, which, as mentioned in the introduction, admits an approximation scheme.

The lower bound is more interesting: By Theorem 9 we know that if  $\psi$  has and-hardness  $H_{\text{and}}(\psi) < k$ , then it admits a constant-factor approximation. Nevertheless, the lower bound in Theorem 9 only addresses formulas of and-hardness  $H_{\text{and}}(\psi) = k$ . In particular, Theorem 11 identifies a class of problems for which some (fixed) constant-factor approximation is *best-possible in terms of approximation*.

At a technical level, we make use of interesting machinery to obtain the lower bound. The starting point is the *distributed PCP* framework, which was introduced in [4] and further strengthened in [62, 28] to give hardness of approximation for Maximum  $k$ -Inner Product ( $k$ -MAXIP). In this case, we cannot use this tool directly, since it crucially relies on the fact that the *target* problem  $\text{MAX}(\psi)$  has full hardness  $H_{\text{and}}(\psi) = k$ .<sup>9</sup> Instead, we first show MAX-3-SAT-hardness of the following intermediate problem:

► **Problem 12** ( $(k, 3)$ -OV). *Given sets of  $n$  vectors  $X_1, \dots, X_k \subseteq \{0, 1\}^d$ , where each coordinate  $y \in [d]$  is associated to three active indices  $a, b, c \in [k]$ , detect if there are vectors  $x_1 \in X_1, \dots, x_k \in X_k$  such that for all  $y \in [d]$ , it holds that  $x_a[y] \cdot x_b[y] \cdot x_c[y] = 0$  where  $a, b, c$  are the active indices at  $y$ .*

We then provide a gap introducing reduction from  $(k, 3)$ -OV to  $\text{MAX}(\psi)$ , in the same spirit as the PCP reduction gives such a reduction from  $k$ -OV to  $k$ -MAXIP. This involves several technical steps as outlined in Figure 2. At a high level, we decompose a  $(k, 3)$ -OV instance into a combination of multiple 3-OV instances and use the PCP reduction as a black box on each of these. After combining the outputs of the reduction, we obtain instances of  $\text{MAX}(\psi)$  with the desired gap. The issue with this approach is that the PCP reduction blows up the dimension of the input vectors exponentially,<sup>10</sup> which makes the reduction inapplicable to our case if we start from a moderate-dimensional  $(k, 3)$ -OV instance. To show that  $(k, 3)$ -OV does not even have a  $O(m^{k-\delta})$ -time algorithm when the dimension is  $d = O(\log n)$ , we use the stronger *Sparse* MAX-3-SAT hypothesis.<sup>11</sup> See the full version of this paper for further discussion of this hypothesis.

The second way in which we get a closer look at Theorem 9 is by inspecting the hardest regime, i.e., when  $H_{\text{and}}(\psi) = k$ :

► **Theorem 13** (Polynomial-Factor Approximation). *Let  $\psi$  be a  $\text{MAXSP}_k$  graph formula of full and-hardness  $H_{\text{and}}(\psi) = k$ .*

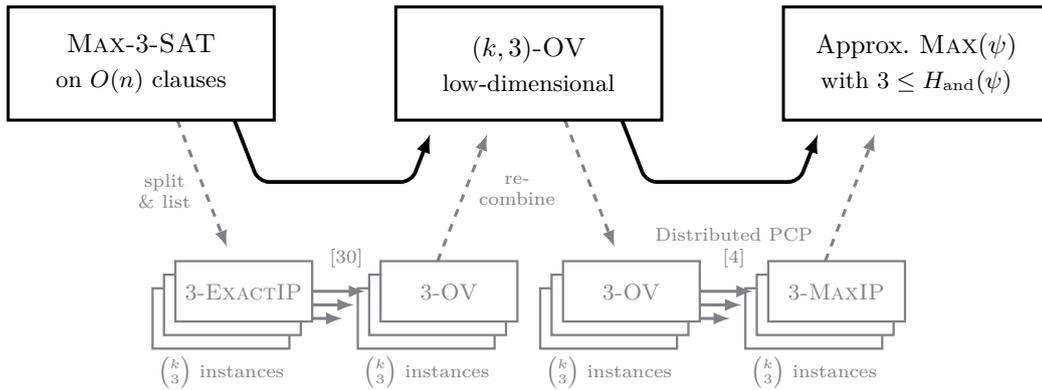
- For every  $\varepsilon > 0$ , there exists some  $\delta > 0$  such that an  $m^\varepsilon$ -approximation for  $\text{MAX}(\psi)$  can be computed in time  $O(m^{k-\delta})$ .
- For every  $\delta > 0$ , there exists an  $\varepsilon > 0$  such that there exists no  $m^\varepsilon$ -approximation for  $\text{MAX}(\psi)$  in time  $O(m^{k-\delta})$  unless SETH fails.

The lower bound is obtained by applying the subsequent improvements on the distributed PCP framework by [62, 28, 51], which improve the parameters of the reduction via algebraic geometry codes and expander graphs. For the upper bound, we give a simple algorithm which exploits the sparsity of the instances.

<sup>9</sup> More precisely, the reduction exploits the fact that  $\psi_0$  has a unique satisfying assignment to encode the communication protocol used in the reduction.

<sup>10</sup> An instance of  $k$ -OV on dimension  $d = c \log n$  is reduced to multiple instances of  $k$ -MAXIP on dimension  $\exp(c) \log n$ .

<sup>11</sup> Morally, just as SETH implies the hardness of low-dimensional OV, the *Sparse* MAX-3-SAT Hypothesis implies the hardness of *low-dimensional*  $(k, 3)$ -OV.



■ **Figure 2** The chain of reductions from Sparse MAX-3-SAT to  $V\text{MAX}(\psi)$  involving two main steps. Both proofs involve several intermediate steps as illustrated in the gray parts.

### 3.3.2 Minimization

There is an easier criterion for the hardness of approximating minimization problems. Namely, observe that giving a multiplicative approximation to a minimization problem  $\text{MIN}(\psi)$  is at least as hard as testing if the optimal value is zero (recall that we refer to this problem as  $\text{ZERO}(\psi)$ ). More precisely, suppose that we are given an instance with  $\text{OPT} = 0$ . Then any multiplicative approximation must return an optimal solution.

It turns out that this is the only source of hardness for  $\text{MIN}(\psi)$  problems. We show a fine-grained equivalence of deciding  $\text{ZERO}(\psi)$  and approximating  $\text{MIN}(\psi)$  within a constant factor:

► **Theorem 14** (Constant Approximation is Equivalent to Testing Zero). *Let  $\psi$  be a  $\text{MINSP}_k$  graph formula. Via a randomized reduction, there exists a constant-factor approximation algorithm for  $\text{MIN}(\psi)$  in time  $O(m^{k-\delta})$  for some  $\delta > 0$  if and only if  $\text{ZERO}(\psi)$  can be solved in time  $O(m^{k-\delta'})$  for some  $\delta' > 0$ .*

To reduce approximating  $\text{MIN}(\psi)$  to  $\text{ZERO}(\psi)$ , we make use of *locality-sensitive hashing* (LSH). This technique was for instance used to solve the Approximate Nearest Neighbors problem in Hamming spaces [44], and was recently adapted by Chen and Williams to show that Minimum Inner Product can be reduced to OV [30]. The latter result constitutes a singular known case for the general trend in  $\text{MINSP}_k$  revealed by Theorem 14. In comparison, our reduction is simpler and more general, but gives weaker guarantees on the constant factor.

We specifically use LSH for a reduction from approximating  $\text{MIN}(\psi)$  to the intermediate problem of *listing* solutions to  $\text{ZERO}(\psi)$  – the better the listing algorithm performs, the better the approximation guarantee. For instance, an algorithm listing  $L$  solutions in time  $O(m^{k-\delta} \cdot L^{\delta/k})$  for some  $\delta > 0$  results in a constant-factor approximation, while a listing algorithm in time  $\tilde{O}(m^{k-\delta} + L)$  leads to an approximation scheme. To finish the proof of Theorem 14, we show that any  $\text{ZERO}(\psi)$  problem with an  $O(m^{k-\delta})$ -time decider, for some  $\delta > 0$ , also admits a listing algorithm in time  $O(m^{k-\delta} \cdot L^{\delta/k})$ .

Since the hardness of  $\text{ZERO}(\psi)$  is completely classified [24], we obtain the following dichotomy as a consequence of Theorem 14:

- **Corollary 15** (Constant Approximation – Minimization). *Let  $\psi$  be a  $\text{MINSP}_k$  graph formula.*
- *If  $H_{\text{and}}(\psi) \leq 2$  and  $H_{\text{and}}(\psi) < k$ , then there exists a randomized constant-factor approximation algorithm for  $\text{MIN}(\psi)$  in time  $O(m^{k-\delta})$  for some  $\delta > 0$ .*
  - *If  $3 \leq H_{\text{and}}(\psi)$  or  $H_{\text{and}}(\psi) = k$ , then computing any approximation for  $\text{MIN}(\psi)$  in time  $O(m^{k-\delta})$  for any  $\delta > 0$  is not possible unless the  $\text{MAX-3-SAT}$  hypothesis fails.*

Similar to the maximization case, let us next consider approximation schemes for minimization problems. We can reuse the general framework outlined in the previous paragraphs to obtain an approximation scheme by giving an output-linear listing algorithm in time  $\tilde{O}(m^{k-\delta} + L)$ , for all formulas with and-hardness at most  $H_{\text{and}}(\psi) \leq 1$ .

- **Theorem 16** (Approximation Scheme – Minimization). *Let  $\psi$  be a  $\text{MINSP}_k$  graph formula. If  $H_{\text{and}}(\psi) \leq 1$ , then there exists a randomized approximation scheme for  $\text{MIN}(\psi)$ .*

### 3.4 Efficient (Multiplicative) Approximation Schemes

Theorems 11 and 16 show that if  $H_{\text{and}}(\psi) \leq 1$ , then we can give approximation schemes for  $\text{OPT}(\psi)$ . In the full version, we complement this result by ruling out the existence of *efficient* approximation schemes for most regimes. We say that  $\text{OPT}(\psi)$  admits an efficient (multiplicative) approximation scheme if there is some fixed constant  $\delta > 0$  such that for any  $\varepsilon > 0$ , a multiplicative  $(1 + \varepsilon)$ -approximation for  $\text{OPT}(\psi)$  can be computed in time  $O(m^{k-\delta})$ .

- **Theorem 17.** *Let  $\psi$  be an  $\text{OPTSP}_k$  graph formula. If  $3 \leq H_{\text{deg}}(\psi)$  or  $H_{\text{deg}}(\psi) = k$ , then there exists no efficient approximation scheme for  $\text{OPT}(\psi)$  assuming the Sparse  $\text{MAX-3-SAT}$  Hypothesis.*

### 3.5 Proving the Main Theorem

We can now put things together to prove Theorem 1.

**Proof of Theorem 1.** We only sketch the proof for a maximization problem  $\psi$  with  $k \geq 3$ ; the other cases are similar. We show how to classify  $\psi$  into one of the four stated regimes with a case distinction based on the hardness parameters  $H_{\text{and}}(\psi)$  and  $H_{\text{deg}}(\psi)$ . This case distinction can also be read off Figure 1.

- If  $H_{\text{deg}}(\psi) \leq 2$ : By Theorem 7,  $\psi$  is efficiently optimizable, so it lies in R1.
- Otherwise, it holds that  $3 \leq H_{\text{deg}}(\psi) \leq k$ . In this case, we make a further distinction:
  - $H_{\text{and}}(\psi) \leq 1$ : By Theorem 11,  $\psi$  admits an approximation scheme but by Theorem 17 not an efficient one, so it lies in R2.
  - $H_{\text{and}}(\psi) = 2$ : By Theorem 9,  $\psi$  admits an efficient constant-factor approximation but by Theorem 17, it does not admit an efficient approximation scheme. Thus, depending on whether  $\psi$  admits an approximation scheme or not, it lies in R2 or R3. (As mentioned below Theorem 1, this is the single case where we cannot place the formula in its precise regime, see also Open Problem 1.)
  - $3 \leq H_{\text{and}}(\psi) < k$ : By Theorem 9,  $\psi$  admits an efficient constant factor approximation but by Theorem 11 it has no approximation scheme, so it lies in R3.
  - $H_{\text{and}}(\psi) = k$ : By Theorem 13,  $\psi$  admits an efficient polynomial-factor approximation, and this is best possible, so it lies in R4. ◀

## 4 Discussion and Open Problems

Our investigation reveals all possible approximability types (in better-than-exhaustive-search time) for general classes of polynomial-time optimization problems, namely graph formulas in  $\text{MAXSP}_k$  and  $\text{MINSP}_k$ . Our results, which give an almost complete characterization, open up the following questions:

- Can we extend our classification beyond graph formulas, i.e., when we allow more binary relations, or even higher-arity relations? Such settings include, e.g., generalizations of Max  $k$ -XOR from  $\mathbb{F}_2$  to  $\mathbb{F}_q$ , or variants of the densest subgraph problem on hypergraphs.
- In our setting, each variable  $x_i$  ranges over a separate set  $X_i$ , also known as a *multichromatic* setting. We leave it open to transfer our results to the *monochromatic* setting (see e.g. [52]).
- While we consider running times expressed in the input size (as usual in database contexts), it would also be natural to consider parameterization in the number  $n$  of objects in the relational structure, see [68].

Besides these extensions, we ask whether one can close the remaining gap in our classification: Do formulas  $\psi$  with  $H_{\text{and}}(\psi) = 2$  and  $H_{\text{deg}}(\psi) \geq 3$  admit an approximation scheme? As a specific challenge, we give the following open problem:

► **Open Problem 1.** *Is there an approximation scheme for Maximum Exact-3-Cover (or its minimization variant)? Specifically, can we prove or rule out that for every  $\varepsilon > 0$ , there is some  $\delta > 0$  such that we can  $(1 + \varepsilon)$ -approximate Maximum Exact-3-Cover in time  $O(m^{3-\delta})$ ?*

It appears likely that showing existence of an approximation scheme for Maximum Exact-3-Cover would lead to a full characterization of  $\text{MAXSP}_k$ .

Finally, while we focused on the qualitative question whether or not exhaustive search can be beaten, a follow-up question is to determine precise *approximability-time tradeoffs*. In this vein, consider the well-studied Maximum  $k$ -Cover problem: A simple linear-time greedy approach is known to establish a  $(1 - 1/e)^{-1}$ -approximation [46]. Subsequent lower bounds show that this is conditionally best possible in polynomial-time [35] and even  $f(k)m^{o(k)}$ -time (under GAP-ETH) [31, 59]. On the other hand, for every fixed  $k$ , we show (1) existence of an approximation scheme, but (2) rule out an efficient one assuming SETH, i.e., for every  $\delta > 0$ , there is some  $\varepsilon > 0$  such that an  $(1 + \varepsilon)$ -approximation requires time  $\Omega(m^{k-\delta})$ .

► **Open Problem 2.** *Let  $k \geq 2$ . Can we determine, for every  $1 \leq \gamma \leq (1 - 1/e)^{-1}$ , the optimal exponent  $\alpha$  of the fastest  $\gamma$ -approximation for Maximum  $k$ -Cover running in time  $O(m^{\alpha \pm o(1)})$ , assuming plausible fine-grained hardness assumptions?*

Note that the extreme cases for  $\gamma = 1$  and  $\gamma = (1 - 1/e)^{-1}$  are already settled and that [59] shows that for all immediate cases,  $\alpha$  must have a linear dependence on  $k$ , assuming GAP-ETH.

---

## References

- 1 Amir Abboud and Arturs Backurs. Towards hardness of approximation for polynomial time problems. In *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science*, volume 67 of *ITCS '17*, pages 11:1–11:26. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- 2 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In *Proceedings of the 56th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '15, pages 59–78. IEEE Computer Society, 2015.

- 3 Amir Abboud and Aviad Rubinfeld. Fast and deterministic constant factor approximation algorithms for LCS imply new circuit lower bounds. In *Proceedings of the 9th Conference on Innovations in Theoretical Computer Science*, volume 94 of *ITCS '18*, pages 35:1–35:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- 4 Amir Abboud, Aviad Rubinfeld, and Ryan Williams. Distributed PCP theorems for hardness of approximation in P. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '17, pages 25–36. IEEE Computer Society, 2017.
- 5 Josh Alman, Timothy Chan, and Ryan Williams. Polynomial representations of threshold functions and algorithmic applications. In *Proceedings of the 57th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '16, pages 467–476. IEEE Computer Society, 2016.
- 6 Josh Alman, Timothy Chan, and Ryan Williams. Faster deterministic and Las Vegas algorithms for offline approximate nearest neighbors in high dimensions. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, pages 637–649. SIAM, 2020.
- 7 Josh Alman and Virginia Vassilevska Williams. OV graphs are (probably) hard instances. In *Proceedings of the 11th Conference on Innovations in Theoretical Computer Science*, volume 151 of *ITCS '20*, pages 83:1–83:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- 8 Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '21, pages 522–539. SIAM, 2021.
- 9 Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *Proceedings of the 47th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '06, pages 459–468. IEEE Computer Society, 2006.
- 10 Alexandr Andoni, Piotr Indyk, Huy L. Nguyen, and Ilya P. Razenshteyn. Beyond locality-sensitive hashing. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 1018–1028. SIAM, 2014.
- 11 Alexandr Andoni, Piotr Indyk, and Ilya P. Razenshteyn. Approximate nearest neighbor search in high dimensions. In *Proceedings of the International Congress of Mathematicians*, ICM '18, pages 3287–3318, 2018.
- 12 Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Polylogarithmic approximation for edit distance and the asymmetric query complexity. In *Proceedings of the 51st IEEE Annual Symposium on Foundations of Computer Science*, FOCS '10, pages 377–386. IEEE Computer Society, 2010.
- 13 Alexandr Andoni and Negev Shekel Nosatzki. Edit distance in near-linear time: it's a constant factor. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science*, FOCS '20, pages 990–1001. IEEE, 2020.
- 14 Alexandr Andoni and Krzysztof Onak. Approximating edit distance in near-linear time. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, STOC '09, pages 199–204. ACM, 2009.
- 15 Alexandr Andoni and Ilya P. Razenshteyn. Optimal data-dependent hashing for approximate near neighbors. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, STOC '15, pages 793–801. ACM, 2015.
- 16 Arturs Backurs, Liam Roditty, Gilad Segal, Virginia Vassilevska Williams, and Nicole Wein. Towards tight approximation bounds for graph diameter and eccentricities. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing*, STOC '18, pages 267–280. ACM, 2018.
- 17 Ziv Bar-Yossef, T. S. Jayram, Robert Krauthgamer, and Ravi Kumar. Approximating edit distance efficiently. In *Proceedings of the 45th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '04, pages 550–559. IEEE Computer Society, 2004.
- 18 Tugkan Batu, Funda Ergün, Joe Kilian, Avner Magen, Sofya Raskhodnikova, Ronitt Rubinfeld, and Rahul Sami. A sublinear algorithm for weakly approximating edit distance. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, STOC '03, pages 316–324. ACM, 2003.

- 19 Tugkan Batu, Funda Ergün, and Süleyman Cenk Sahinalp. Oblivious string embeddings and edit distance approximations. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '06, pages 792–801. SIAM, 2006.
- 20 Allan Borodin, Rafail Ostrovsky, and Yuval Rabani. Subquadratic approximation algorithms for clustering problems in high dimensional spaces. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, STOC '99, pages 435–444. ACM, 1999.
- 21 Joshua Brakensiek and Aviad Rubinfeld. Constant-factor approximation of near-linear edit distance in near-linear time. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing*, STOC '20, pages 685–698. ACM, 2020.
- 22 Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '15, pages 661–670. IEEE Computer Society, 2014.
- 23 Karl Bringmann, Alejandro Cassis, Nick Fischer, and Marvin Künnemann. Fine-grained completeness for optimization in P. In *APPROX-RANDOM*, volume 207 of *LIPICs*, pages 9:1–9:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- 24 Karl Bringmann, Nick Fischer, and Marvin Künnemann. A fine-grained analogue of Schaefer's theorem in P: Dichotomy of  $\exists^k\forall$ -quantified first-order graph properties. In *Proceedings of the 34th Computational Complexity Conference*, volume 137 of *CCC '19*, pages 31:1–31:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- 25 Karl Bringmann, Marvin Künnemann, and Karol Wegrzycki. Approximating APSP without scaling: Equivalence of approximate min-plus and exact min-max. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing*, STOC '19, pages 943–954. ACM, 2019.
- 26 Diptarka Chakraborty, Debarati Das, Elazar Goldenberg, Michal Koucký, and Michael E. Saks. Approximating edit distance within constant factor in truly sub-quadratic time. In *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '18, pages 979–990. IEEE Computer Society, 2018.
- 27 Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan. From gap-ETH to FPT-inapproximability: Clique, dominating set, and more. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '17, pages 743–754. IEEE Computer Society, 2017.
- 28 Lijie Chen. On the hardness of approximate and exact (bichromatic) maximum inner product. In *Proceedings of the 33th Computational Complexity Conference*, volume 102 of *CCC '18*, pages 14:1–14:45. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- 29 Lijie Chen, Shafi Goldwasser, Kaifeng Lyu, Guy N. Rothblum, and Aviad Rubinfeld. Fine-grained complexity meets IP = PSPACE. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pages 1–20. SIAM, 2019.
- 30 Lijie Chen and Ryan Williams. An equivalence class for orthogonal vectors. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pages 21–40. SIAM, 2019.
- 31 Vincent Cohen-Addad, Anupam Gupta, Amit Kumar, Euiwoong Lee, and Jason Li. Tight FPT approximations for  $k$ -median and  $k$ -means. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming*, volume 132 of *ICALP '19*, pages 42:1–42:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- 32 Nadia Creignou. A dichotomy theorem for maximum generalized satisfiability problems. *J. Comput. Syst. Sci.*, 51(3):511–522, 1995.
- 33 Martin Dietzfelbinger, Philipp Schlag, and Stefan Walzer. A subquadratic algorithm for 3XOR. In *Proceedings of the 43rd International Symposium on Mathematical Foundations of Computer Science*, volume 117 of *MFCs '18*, pages 59:1–59:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- 34 Ran Duan and Seth Pettie. Linear-time approximation for maximum weight matching. *J. ACM*, 61(1):1:1–1:23, 2014.

- 35 Uriel Feige. A threshold of  $\ln n$  for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
- 36 Andreas Emil Feldmann, C. S. Karthik, Euiwoong Lee, and Pasin Manurangsi. A survey on approximation in parameterized complexity: Hardness and algorithms. *Electronic Colloquium on Computational Complexity (ECCC)*, 27:86, 2020.
- 37 Harold N. Gabow. Scaling algorithms for network problems. *J. Comput. Syst. Sci.*, 31(2):148–168, 1985.
- 38 Jiawei Gao. On the fine-grained complexity of least weight subsequence in multitrees and bounded treewidth DAGs. In *Proceedings of the 14th International Symposium on Parameterized and Exact Computation*, volume 148 of *IPEC '19*, pages 16:1–16:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- 39 Jiawei Gao and Russell Impagliazzo. The fine-grained complexity of strengthenings of first-order logic. *Electronic Colloquium on Computational Complexity (ECCC)*, 26:9, 2019.
- 40 Jiawei Gao, Russell Impagliazzo, Antonina Kolokolova, and Ryan Williams. Completeness for first-order properties on sparse structures with algorithmic applications. *ACM Trans. Algorithms*, 15(2):23:1–23:35, 2019.
- 41 Aristides Gionis, Piotr Indyk, and Rajeev Motwani. Similarity search in high dimensions via hashing. In *Proceedings of the 25th International Conference on Very Large Data Bases, VLDB '99*, pages 518–529. Morgan Kaufmann, 1999.
- 42 Ashish Goel, Piotr Indyk, and Kasturi R. Varadarajan. Reductions among high dimensional proximity problems. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '01, pages 769–778. SIAM, 2001.
- 43 Elazar Goldenberg, Aviad Rubinfeld, and Barna Saha. Does preprocessing help in fast sequence comparisons? In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing*, STOC '20, pages 657–670. ACM, 2020.
- 44 Sarel Har-Peled, Piotr Indyk, and Rajeev Motwani. Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory of Computing*, 8(1):321–350, 2012.
- 45 Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- 46 Dorit S. Hochbaum. *Approximating Covering and Packing Problems: Set Cover, Vertex Cover, Independent Set, and Related Problems*, pages 94–143. PWS Publishing Co., 1996.
- 47 Piotr Indyk. Dimensionality reduction techniques for proximity problems. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '00, pages 371–378. SIAM, 2000.
- 48 Piotr Indyk. Better algorithms for high-dimensional proximity problems via asymmetric embeddings. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03, pages 539–545. SIAM, 2003.
- 49 Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, STOC '98, pages 604–613. ACM, 1998.
- 50 Zahra Jafargholi and Emanuele Viola. 3SUM, 3XOR, triangles. *Algorithmica*, 74(1):326–343, 2016.
- 51 C. S. Karthik, Bundit Laekhanukit, and Pasin Manurangsi. On the parameterized complexity of approximating dominating set. *J. ACM*, 66(5):33:1–33:38, 2019.
- 52 Karthik C. S. and Pasin Manurangsi. On closest pair in euclidean metric: Monochromatic is as hard as bichromatic. In *Proceedings of the 10th Conference on Innovations in Theoretical Computer Science*, volume 124 of *ITCS '19*, pages 17:1–17:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- 53 Sanjeev Khanna, Madhu Sudan, Luca Trevisan, and David P. Williamson. The approximability of constraint satisfaction problems. *SIAM J. Comput.*, 30(6):1863–1920, 2000.
- 54 Donald E. Knuth. Dancing links. *CoRR*, abs/cs/0011047, 2000. [arXiv:0011047](https://arxiv.org/abs/cs/0011047).
- 55 Michal Koucký and Michael E. Saks. Constant factor approximations to edit distance on far input pairs in nearly linear time. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing*, STOC '20, pages 699–712. ACM, 2020.

- 56 Gad M. Landau, Eugene W. Myers, and Jeanette P. Schmidt. Incremental string comparison. *SIAM J. Comput.*, 27(2):557–582, 1998.
- 57 Bingkai Lin. Constant approximating  $k$ -clique is  $w[1]$ -hard. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing*, STOC '21, pages 1749–1756. ACM, 2021.
- 58 Andrea Lincoln, Virginia Vassilevska Williams, and Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, pages 1236–1252. SIAM, 2018.
- 59 Pasin Manurangsi. Tight running time lower bounds for strong inapproximability of maximum  $k$ -coverage, unique set cover and related problems (via  $t$ -wise agreement testing theorem). In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, pages 62–81. SIAM, 2020.
- 60 Christos H. Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. Syst. Sci.*, 43(3):425–440, 1991.
- 61 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, STOC '13, pages 515–524. ACM, 2013.
- 62 Aviad Rubinfeld. Hardness of approximate nearest neighbor search. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing*, STOC '18, pages 1260–1268. ACM, 2018.
- 63 Noah Stephens-Davidowitz and Vinod Vaikuntanathan. SETH-hardness of coding problems. In *Proceedings of the 60th IEEE Annual Symposium on Foundations of Computer Science*, FOCS '19, pages 287–301. IEEE Computer Society, 2019.
- 64 Gregory Valiant. Finding correlations in subquadratic time, with applications to learning parities and the closest pair problem. *J. ACM*, 62(2), 2015.
- 65 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the International Congress of Mathematicians*, ICM '18, pages 3447–3487, 2018.
- 66 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005.
- 67 Ryan Williams. *Algorithms and resource requirements for fundamental problems*. ProQuest LLC, Ann Arbor, MI, 2007. Thesis (Ph.D.)—Carnegie Mellon University.
- 68 Ryan Williams. Faster decision of first-order graph properties. In *Proceedings of the Joint Meeting of the 23rd EACSL Annual Conference on Computer Science Logic (CSL) and the 29th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, CSL-LICS '14. ACM, 2014.
- 69 Ryan Williams. On the difference between closest, furthest, and orthogonal pairs: Nearly-linear vs barely-subquadratic complexity. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, pages 1207–1215. SIAM, 2018.
- 70 Uri Zwick. All pairs shortest paths using bridging sets and rectangular matrix multiplication. *J. ACM*, 49(3):289–317, 2002.