

Classification and Approximation of Geometric Location Problems

Sascha Brauer, M. Sc.

June 4, 2019

A dissertation submitted to the
Department of Computer Science
Paderborn University

for the degree of
Doktor der Naturwissenschaften
(doctor rerum naturalium)

accepted on the recommendation of

Prof. Dr. Johannes Blömer
Paderborn University

Prof. Dr. Heiko Röglin
University of Bonn

defended on
September 25, 2019

Acknowledgements

I would like to start by thanking my advisor Prof. Dr. Johannes Blömer. He gave me the opportunity to work in his group and provided a collaborative research environment, as well as the freedom to pursue my own research interests.

I conducted the research on this thesis while employed at Paderborn University. This was partly supported by the German Research Foundation (DFG) under grant BL 314/8-1. I am grateful for the funding I received.

Over the years, there were many colleagues that contributed to my research in one way or the other. My co-author Dr. Kathrin Bujna, the political action committee Jan Bobolz and Nils Löken, the gaming collective Fabian Eidens and Dr. Jakob Juhnke, and the old men Dr. Peter Günther and Dr. Gennadij Liske. I want to thank all of you for the countless fruitful discussions. I also want to thank Prof. Dr. Alexander Skopalik for providing motivation for some of my work and Pascal Bemann and Dr. Gareth Davies for proofreading parts of my thesis. Furthermore, I am thankful to Claudia Jahn for her invaluable administrative work and the professional social counselling.

There are a lot of people outside of my working environment which helped me by providing getaways. There is Dominik Buder as one of the constants throughout my life. Also, Robert Kreutzer and Patrick Schönlau from Team Lauch, who kept challenging me to always improve myself. To all the other guys and gals I cannot name here: Thank you for being there.

Last but not least I want to thank my family. This includes my parents, my brother Marc, and my in-laws for allowing me to grow up to be the person I am today. In particular however, my wife Anna for always being by my side and enduring all our decisions with me.

Abstract

In this thesis we explore aspects of computational complexity and algorithms with a guaranteed approximation ratio for several geometric location problems. The main focus is on the FUZZY k -MEANS problem, a problem that has so far not been subject of a substantial theoretical analysis.

The first part of the thesis focuses on hardness and impossibility results. We show that a variant of DISCRETE FUZZY k -MEANS on metric spaces, as well as a newly defined radius variant of that problem, are **NP**-hard. A popular algorithmic approach to these problems is the single-swap heuristic. This local search algorithm iteratively swaps a single representative for another candidate until it finds a stable solution. We prove that this algorithm is tightly **PLS**-complete for METRIC UNCAPACITATED FACILITY LOCATION, DISCRETE k -MEANS, and DISCRETE FUZZY k -MEANS. This means that, in the worst case, it takes an exponential number of iterations until a single-swap local search terminates. We conclude the first part by showing that optimal solutions of the general FUZZY k -MEANS problem are not solvable by radicals over the rational numbers. Thus, there are instances where optimal solutions cannot be represented by a finite concatenation of rational numbers using elementary arithmetic and extracting roots. In consequence, any algorithm trying to solve the problem can at most find an approximately optimal solution.

The second part of the thesis discusses algorithmic techniques. We present an algorithm approximating the FUZZY k -MEANS problem with a factor of $(1 + \epsilon)$. Our algorithm improves previously presented algorithms as its runtime is independent of any weight of the input points and the exponential dependence on the number of clusters is linear. Furthermore, we show that our radius variant of that problem can be solved in polynomial time for instances on the real line and with two clusters. We complement our algorithmic results by presenting a construction of small coresets for FUZZY k -MEANS.

Zusammenfassung

In dieser Arbeit untersuchen wir komplexitätstheoretische Aspekte und Algorithmen mit garantierter Approximationsgüte für verschiedene Probleme der geometrischen Platzierung. Der Fokus liegt dabei auf dem FUZZY k -MEANS Problem, welches bisher nicht substantiell theoretisch untersucht wurde.

Der erste Teil der Arbeit beschäftigt sich mit Härte- und Unmöglichkeitsergebnissen. Wir zeigen, dass eine Variante von DISCRETE FUZZY k -MEANS auf metrischen Räumen, sowie eine hier neu definierte Radiusvariante dieses Problems, **NP**-schwer sind. Eine beliebte algorithmische Herangehensweise an diese Probleme ist die single-swap Heuristik. Dieser lokale Suchalgorithmus tauscht iterativ einen Repräsentanten der aktuellen Lösung gegen einen anderen Kandidatenpunkt, bis eine stabile Lösung gefunden wird. Wir beweisen, dass dieser Algorithmus für die Probleme METRIC UNCAPACITATED FACILITY LOCATION, DISCRETE k -MEANS und DISCRETE FUZZY k -MEANS streng **PLS**-vollständig ist. Das heißt, dass im worst-case exponentiell viele Iterationen benötigt werden, bis eine single-swap Suche terminiert. Zum Abschluss des ersten Teils zeigen wir, dass optimale Lösungen des allgemeinen FUZZY k -MEANS Problems nicht durch Radikale über den rationalen Zahlen lösbar sind. Es gibt also Instanzen, deren optimale Lösungen nicht durch eine endliche Verkettung rationaler Zahlen durch die Grundrechenarten und das Ziehen von Wurzeln dargestellt werden können. Folglich kann jeder Algorithmus, der versucht, das Problem zu lösen, höchstens eine approximativ optimale Lösung finden.

Der zweite Teil der Arbeit diskutiert algorithmische Techniken. Wir präsentieren einen Algorithmus, der das FUZZY k -MEANS Problem mit Faktor $(1+\epsilon)$ approximiert. Dieser Algorithmus stellt eine Verbesserung bisher bekannter Algorithmen dar, da seine Laufzeit unabhängig von der Gewichtung der Eingabepunkte ist und die exponentielle Abhängigkeit von der Anzahl an Clustern linear ist. Des Weiteren zeigen wir, dass unsere Radiusvariante dieses Problems für eindimensionale Instanzen mit zwei Clustern in polynomieller Zeit gelöst werden kann. Wir ergänzen unsere algorithmischen Ergebnisse durch eine Konstruktion kleiner Kernmengen für FUZZY k -MEANS.

Contents

1	Introduction	1
1.1	Outline	2
1.2	Publication Notes	3
I	Fundamentals	5
2	Remarks on Notation	7
3	Problem Statements	11
3.1	Facility Location	11
3.2	Clustering	14
3.2.1	Hard Clustering	18
3.2.2	Soft Clustering	20
3.2.3	Discussion	24
3.2.4	Further Analysis	25
3.3	Related Work	27
II	Classification and Impossibility	29
4	Classical Hardness	31
4.1	Basics of Computational Complexity	32
4.2	Discrete Clustering Problems	34
4.3	Related Work	36
4.4	Towards a Classification of DFKM	36
4.4.1	Construction	38
4.4.2	Correctness	39
4.5	Parameterized Hardness of DFRKM	44
4.6	DKM is Hard in Fixed Dimensions	47
5	Complexity of Single-Swap Local Search	51
5.1	Polynomial Local Search	53
5.2	Related Work	58
5.3	Completeness of MUFL/SingleSwap	59
5.3.1	Construction	60

Contents

5.3.2	Correctness	60
5.3.3	Tightness	64
5.4	Completeness of DKM/SingleSwap	65
5.4.1	Construction	66
5.4.2	Correctness and Tightness	67
5.4.3	Embedding	72
5.5	Completeness of DFKM/SingleSwap	75
5.5.1	Construction	76
5.5.2	Correctness, Tightness, and Embedding	77
6	Non-Representability of Solutions	87
6.1	Algebraic Basics	88
6.2	Related Work	90
6.3	Unsolvability of Optimal FKM Solutions	90
6.4	Implementation of the Finite Taylor Grid	97
III	Approaching Hard Problems	101
7	Approximation Algorithms	103
7.1	Related Work	104
7.2	A PTAS for FKM With Fixed Number of Clusters	105
7.2.1	From Soft to Hard Clusters	106
7.2.2	A Sampling Based Approach	111
7.3	Solving Small FRKM Instances	114
8	Coresets	119
8.1	Related Work	120
8.2	Small Coresets for FKM	120
8.2.1	From Weak to Strong via Non-Negligibility	122
8.2.2	Weak Coresets for Non-Negligible Solutions	125
8.3	Applying FKM Coresets	132
8.4	Correctness of our Weak Coreset Algorithm	133
9	Future Research	143
	Bibliography	145

List of Abbreviations

UFL UNCAPACITATED FACILITY LOCATION

MUFL METRIC UNCAPACITATED FACILITY LOCATION

KM k -MEANS

RKM RADIUS k -MEANS

FKM FUZZY k -MEANS

FRKM FUZZY RADIUS k -MEANS

DKM DISCRETE k -MEANS

DRKM DISCRETE RADIUS k -MEANS

DFKM DISCRETE FUZZY k -MEANS

DFRKM DISCRETE FUZZY RADIUS k -MEANS

X3C EXACT COVER BY 3-SETS

TSP TRAVELLING SALESMAN PROBLEM

DS DOMINATING SET

P3NAESAT POSITIVE 3-NOTALLEQUAL-SAT

M2SAT MAX 2-SAT

PM2NAESAT POSITIVE MAX 2-NOTALLEQUAL-SAT

MC MAX CUT

PTAS Polynomial Time Approximation Scheme

Introduction

In the digital age, data is abundant, and new data is generated at an unprecedented rate. Already in 2016, the search engine Google reported that they processed more than two trillion searches in the previous twelve months¹. Assuming the average search query consists of only four letters (which is probably a conservative estimate), this yields roughly 22 gigabytes of data every day – just for the queries. The popular multimedia messaging app Snapchat currently claims that their users create three billion *snap*s per day². In its simplest form, a snap is a 1920×1080 pixel image with an accompanying short text. Provided they use some standard image compression, this is more than 1 petabyte of daily image data. Storing and distributing such amounts of data is already difficult. But what if these companies are interested in analyzing their data? Designing sophisticated and efficient algorithms for huge data sets is a daunting task. A popular technique for unsupervised analysis or pre-processing of data is *clustering*.

Clustering constitutes a wide field of tasks, which occur in different types of applications of data analysis. The core goal of a clustering is to identify groups of similar elements in the data set, and find a somehow suitable representative for each group. One popular class of clustering tasks are the *geometric location problems*. In these types of problems, elements of the data set are assumed to be points in some geometric sense, for example, elements of some general metric space or d -dimensional Euclidean space. These problems are divided into two different categories: the hard clustering and the soft clustering problems. In a hard clustering problem, the data points are partitioned into subgroups, the so-called *clusters*. Each point of the data set belongs to exactly one cluster. In a soft clustering problem, points

¹<https://www.blog.google/products/ads/ads-and-analytics-innovations-for-a-mobile-first-world/>, last access: 22.05.2019

²<https://forbusiness.snapchat.com/home>, last access: 22.05.2019

can be divided among clusters. That is, some fraction of a data point might belong to one cluster, while another fraction of the same point belongs to some other cluster. One popular soft clustering problem is FUZZY k -MEANS, which regularly finds applications in, for example, image segmentation and analysis of medical data. Despite its popularity in practice, the theory of the FUZZY k -MEANS problem is poorly understood. A simple heuristic called the FUZZY k -MEANS algorithm, which is usually applied to solve the problem, provides no guarantees in terms of runtime or quality of the produced solutions. The goal of this thesis is to obtain a foundation of the theory of FUZZY k -MEANS in terms of computational complexity and algorithms with provable approximation ratio. We further show that some of the techniques we develop can also be applied to other clustering problems, such as METRIC UNCAPACITATED FACILITY LOCATION and k -MEANS, and yield new insights into these problems, as well.

1.1 Outline

This thesis is organized in three parts. In [Part I](#), we present the notational and formal basis of our work. We introduce different problems which we analyze in the subsequent parts. Furthermore, we discuss how these problems are related to one another and provide some elementary analysis. In [Part II](#), we present negative results of our analysis. For some of our problems, we discuss hardness in the classical sense ([Chapter 4](#)) and hardness of a popular local search algorithm called the single-swap heuristic ([Chapter 5](#)). Afterwards, we prove that optimal solutions of the FUZZY k -MEANS problem can not be represented finitely ([Chapter 6](#)). In [Part III](#), we present positive algorithmic techniques. More specifically, we present two different types of results: algorithms with performance guarantees ([Chapter 7](#)) and the construction of small coresets ([Chapter 8](#)).

1.2 Publication Notes

Several of the results presented in this thesis were developed in cooperation with coauthors and have already been published:

(Blömer et al., 2016a)

Johannes Blömer, Sascha Brauer, and Kathrin Bujna
A Theoretical Analysis of the Fuzzy K-Means Problem
presented at the 16th International Conference on Data Mining,
ICDM 2016

(Brauer, 2017)

Sascha Brauer
*Complexity of Single-Swap Heuristics for Metric Facility Location
and Related Problems*
presented at the 10th International Conference on Algorithms and
Complexity, CIAC 2017

(Blömer et al., 2018)

Johannes Blömer, Sascha Brauer, and Kathrin Bujna
Coresets for Fuzzy K-Means with Applications
presented at the 29th International Symposium on Algorithms
and Computation, ISAAC 2018

(Brauer, 2019)

Sascha Brauer
*Complexity of Single-Swap Heuristics for Metric Facility Location
and Related Problems*
Theoretical Computer Science, Volume 754

Part I



Fundamentals

Remarks on Notation

A core concept used throughout this thesis is that of a data set – an unordered collection of elements. For most of the following, it is not important whether the elements in a data set are pairwise distinct. Collections allowing multiple instances of the same element are mostly known as *multisets* or *bags*. There is no generally-agreed-upon notation differentiating multisets from the well-established notion of a set, which of course contains no duplicate elements. In a disquisition about notations he adopted over the years, Edsger W. Dijkstra wrote,

“Not making superfluous distinctions should always be encouraged; it is bad enough that we don’t have a canonical representation for the unordered pair.”

[Dijkstra, 2000]

In this spirit, we abuse standard notation by not explicitly discriminating between the notation of sets and multisets. Consequently, all operations on sets have to be generalized accordingly. Let X, Y be two multisets and let u be an element from the underlying universe.

- The relation $u \in X$ is true if and only if element u occurs in X at least once.
- All operations iterating over a multiset (for example, $\forall x \in X$, $\sum_{x \in X}$, or $\prod_{x \in X}$) consider each element individually, respecting multiple occurrences of the same element.
- The size of a multiset is $|X| = \sum_{x \in X} 1$.
- The multiplicity of each element in X is $m_X(u) = |\{x \in X \mid x = u\}|$.
- In the union $X \cup Y$, the number of times each element occurs is equal to the maximum of the number of times it occurs in X and Y , that is $m_{X \cup Y}(u) = \max\{m_X(u), m_Y(u)\}$.

- In the intersection $X \cap Y$, the number of times each element occurs is equal to the minimum of the number of times it occurs in X and Y , that is $m_{X \cap Y}(u) = \min\{m_X(u), m_Y(u)\}$.
- In the sum $X + Y$ of two multisets (a generalization of the disjoint union), the number of times each element occurs is equal to the sum of the number of times it occurs in X and Y , that is $m_{X+Y}(u) = m_X(u) + m_Y(u)$.
- In the difference $X \setminus Y$, the number of times each element occurs is the maximum of 0 and the difference between the number of times it occurs in X and in Y , that is $m_{X \setminus Y}(u) = \max\{0, m_X(u) - m_Y(u)\}$.
- The relation $X \subseteq Y$ ($X \subset Y$) is true if each elements occurs in Y at least as often as (more often than) it occurs in X , that is if $m_X(u) \leq m_Y(u)$ ($m_X(u) < m_Y(u)$).
- Further multiset operation can be derived from the above, as usual.

The advantage of this is that we can use sets and multisets without having to jump back and forth between notations. Notice that for multisets where the multiplicity of each element is 1, this notation is consistent with classical set operations. To adhere to formal rigor, we need a specialized notation when dealing with multisets in combination with sets.

Definition 2.1 *Let X be a multiset, and Y be a set. We define the relation subset with replacements as*

$$X \Subset Y \Leftrightarrow \forall x \in X : x \in Y .$$

The difference to the standard subset relation is that subsets with replacements ignore multiplicity. This provides a compact notation when forming multisets with elements from a universe which is not a multiset, like \mathbb{N} or \mathbb{R} . Additionally, we introduce a variant of the difference operator, which removes all elements occurring in the subtrahend from the minuend.

Definition 2.2 *Let X, Y be multisets. We define the full difference as*

$$X \setminus\setminus Y := \{x \in X \mid x \notin Y\} .$$

Another important notion is that of *weighted* data sets. A weighted data set is a data set X together with a weight function $w : X \rightarrow \mathbb{N}$ assigning some positive integral weight to each element. There is a canonical equivalence of multisets and weighted sets.

Definition 2.3 For any multiset X , let Y be the set we obtain by removing all but one occurrence of each element from X . We call Y the underlying set of X and let

$$\begin{aligned} w : Y &\rightarrow \mathbb{N} \\ y &\mapsto m_X(y) . \end{aligned}$$

We say that (Y, w) is X as a weighted set.

For any weighted set (Y, w) let

$$X := \sum_{y \in Y} \sum_{i=1}^{w(y)} \{y\} .$$

We say that X is (Y, w) as a multiset.

This bijection shows that, conceptually, multisets are close to weighted sets. One can directly formulate a formal definition of multisets using weighted sets. In this case, one would call some set X the underlying set, and the weight function $w : X \rightarrow \mathbb{N}$ the multiplicity of each element. We do not use such a definition here because algorithmically speaking we differentiate the semantics of a multiset from that of a weighted set. An algorithm obtaining a multiset as an input is oblivious to the copies of an element. It processes each element in the same way, regardless of any potential duplicates. By giving an algorithm a weighted set, we make it aware of the presence of multiple copies. In essence, we allow the algorithm's treatment of an element to depend on its weight. An algorithm could transform a multiset to a weighted set, or vice versa, via the canonical mapping described in [Definition 2.3](#). We ignore this and expect algorithms to work with their input in the form provided to them. Reasons for this are, among others, that it declutters analysis by allowing a more concise notation and that, for some problems, algorithms are not able to profit from the knowledge of multiplicity in a meaningful way.

Notice that we can also associate some weight function $w : X \rightarrow \mathbb{N}$ with a multiset X . This is the most general form of data set we consider here. The mapping from [Definition 2.3](#) can canonically be generalized to transform a weighted multiset to a multiset or to a weighted set.

For the sake of brevity, we introduce the following succinct notations.

Definition 2.4 For all $i, n \in \mathbb{Z}$ with $i < n$, we denote

$$[n]_i := \{i, \dots, n\} \subset \mathbb{Z} .$$

For $i = 1$, we usually omit the subscript $[n] := [n]_1$.

Definition 2.5 Let X be a multiset and f be a function on X . For each $Y \subseteq X$, we denote

$$f[Y] := \{f(y) \mid y \in Y\}$$

and

$$f(Y) := \sum_{y \in Y} f(y).$$

Problem Statements

Contribution Summary We present a formal introduction of the problems that are the subject of this thesis. Aside from popular facility location and clustering problems, we introduce a soft variant of radius clustering. To the best of our knowledge, this formulation has not been considered in the literature, so far. Furthermore, we provide some basic analysis and discuss the relation of the presented problems to each other.

3.1 Facility Location	11
3.2 Clustering	14
3.2.1 Hard Clustering	18
3.2.2 Soft Clustering	20
3.2.3 Discussion	24
3.2.4 Further Analysis	25
3.3 Related Work	27

The central topic of this thesis are some of the so-called *geometric location problems*. More specifically, we study *facility location* and *clustering* problems. In this chapter, we formally introduce the relevant computational problems and present some elementary analysis.

3.1 Facility Location

At the core of the problems considered here lies the UNCAPACITATED FACILITY LOCATION (UFL) problem. UFL provides a general framework, which comprises a plethora of different optimization problems and has spawned a significant amount of research.

Problem 3.1.1 (UNCAPACITATED FACILITY LOCATION) *Let C be a finite multiset, F be a set, and $w : C \rightarrow \mathbb{N}$, $d : C \times F \rightarrow \mathbb{R}_{\geq 0}$, $f : F \rightarrow \mathbb{R}_{\geq 0}$ be functions. The goal is to find $O \subseteq F$ and $\kappa : C \times O \rightarrow [0, 1]$ minimizing*

$$\begin{aligned} \phi_{ufl}(C, w, f, d, O, \kappa) &:= \sum_{c \in C} \sum_{o \in O} \kappa(c, o) \cdot w(c) \cdot d(c, o) + f(O) \\ \text{subject to } \forall c \in C &: \sum_{o \in O} \kappa(c, o) = 1 . \end{aligned}$$

The elements in C are called *clients*, the ones in F are called *facilities*. The clients are weighted by the function w , which can, for example, be used to model importance of certain clients. In a solution, consisting of O and κ , we call O the *opened* facilities. The assignment κ models which opened facilities serve which fraction of each client's demand. The distance function d models the cost of each facility serving a specific client. Hence, the first sum of the objective function is usually called the *service cost* of the solution. One trivial solution, obviously minimizing the service cost, would be to simply open all available facilities. To counteract the opening of too many facilities, the problem associates an *opening cost* f with each facility. Consequently, the second part of the objective function is usually called the *opening cost* of the solution. Generally speaking, opening a lot of facilities incurs high opening cost, but small service cost, while opening few facilities incurs small opening cost at the expense of high service cost. The problem is called *uncapacitated* since each facility can serve any amount of demand for any number of clients. All problems discussed in this thesis are some special case or variant of UFL. In the following, we denote $n := |C|$ for all facility location problems.

We could restrict O to be a set. It is easy to see that there is always an optimal solution which does not open multiple of the same facility. One can also easily see that there is always an optimal solution where κ is a binary function assigning each client to exactly one facility. That is why we sometimes exclude κ from the optimization by choosing the optimal assignment implicitly

$$\phi_{ufl}(C, w, f, d, O) := \sum_{c \in C} w(c) \cdot \min_{o \in O} \{d(c, o)\} + f(O) .$$

We still explicitly formalize O as a multiset and incorporate the assignment as a variable of the general problem. This is used later on when we derive our soft clustering problems as variants of UFL.

3.1 Facility Location

When the input is unweighted, we simply omit the weight function w . For notational consistency we introduce the following auxiliary function.

Definition 3.1.2 Let $\mathbb{1}$ denote the function mapping every input to 1.

Using this, we write

$$\phi_{ufl}(C, f, d, O, \kappa) := \phi_{ufl}(C, \mathbb{1}, f, d, O, \kappa) .$$

Example We discuss a small toy example motivating the definition of UFL. Assume the town council of Paderborn decided to build new hospitals to increase health care coverage in the city. They need to figure out where and how many hospitals should be build. The clients C are given by street addresses in the city, the weight w by the number of people living at a particular address. The facilities F are addresses of possible locations for new hospitals, the opening cost f is the monetary cost of building a hospital at each specific location. The distance d is given by the number of minutes an ambulance needs to get from one street address to another.

Similar to our toy example, a lot of motivation for UFL comes from real-world problems asking for some sort of geographic placement of facilities. This leads us to METRIC UNCAPACITATED FACILITY LOCATION (MUFL), an important special case of UFL.

Definition 3.1.3 (Metric Space) Let U be a set. A function $d : U \times U \rightarrow \mathbb{R}_{\geq 0}$ is called metric if, for all $u, v, w \in U$,

1. $d(u, v) = d(v, u)$ (symmetry)
2. $d(u, v) = 0 \Leftrightarrow u = v$ (identity of indiscernibles)
3. $d(u, v) \leq d(u, w) + d(w, v)$ (triangle inequality).

For a multiset X , we call (X, d) a metric space if d is a metric on the underlying set of X .

Problem 3.1.4 (METRIC UNCAPACITATED FACILITY LOCATION) Let C be a finite multiset, F be a set, $w : C \rightarrow \mathbb{N}$, $f : F \rightarrow \mathbb{R}_{\geq 0}$ be functions, and $(C \cup F, d)$ be a metric space. The goal is to find $O \subseteq F$ minimizing

$$\text{mfl}(C, w, f, d, O) := \sum_{c \in C} w(c) \cdot \min_{o \in O} \{d(c, o)\} + f(O) .$$

There are three restrictions forming MUFL from UFL. The most important one is that we only allow instances where d forms a metric on $C \cup F$. The other two are the replacement of the assignment function κ by its optimal choice and the restriction of O to sets. We already discussed that the latter two restrictions have no influence on the optima of the objective function, and that their general form was only kept for syntactical reasons.

3.2 Clustering

In the following, we introduce *clustering* problems. These are variants of UFL where we impose no opening cost. Instead we fix an upper bound on the number of facilities which can be opened. In the context of clustering, we usually refer to the clients as *data points* and call an opened facility a *representative*.

Problem 3.2.1 (*k*-CLUSTERING) *Let $k \in \mathbb{N}$, X be a finite multiset, R be a set, and $w : C \rightarrow \mathbb{N}$, $d : X \times R \rightarrow \mathbb{R}_{\geq 0}$ be functions. The goal is to find $M \subseteq R$ and $\kappa : X \times M \rightarrow [0, 1]$ minimizing*

$$\begin{aligned} \phi_{clus}(X, w, d, M, \kappa) &:= \sum_{x \in X} \sum_{\mu \in M} \kappa(x, \mu) \cdot w(x) \cdot d(x, \mu) \\ &\text{subject to } |M| \leq k \\ &\forall x \in X : \sum_{\mu \in M} \kappa(x, \mu) = 1. \end{aligned}$$

We only consider reasonable instances for clustering problems, that is, $|R| > k$. Otherwise, an optimal solution is trivially given by choosing $M = R$. Similar to UFL, we observe that there is always an optimal solution consisting of exactly k distinct representatives and a binary assignment function κ . Moreover, for every fixed set of representatives, there is an optimal κ which simply assigns each point to its closest representative. The difficulty of finding optimal representatives given some fixed assignment heavily depends on the distance function. As we discuss later, for example, it is easy if d is the squared Euclidean distance, but optimal representatives are not finitely representable if d is the standard Euclidean distance.

We further introduce what we call *radius covering* problems. In this flavor of clustering problem, we replace the sums over all points and representatives by taking the maximum. This leads to a problem where we want to minimize the largest radius when covering the input by k balls.

3.2 Clustering

Problem 3.2.2 (RADIUS k -COVER) *Let $k \in \mathbb{N}$, X be a finite set, R be a set, and $w : C \rightarrow \mathbb{N}$, $d : X \times R \rightarrow \mathbb{R}_{\geq 0}$ be functions. The goal is to find $M \subseteq R$ and $\kappa : X \times M \rightarrow [0, 1]$ minimizing*

$$\begin{aligned} \phi_{rad}(X, w, d, M, \kappa) &:= \max_{x \in X} \left\{ \max_{\mu \in M} \{ \kappa(x, \mu) \cdot w(x) \cdot d(x, \mu) \} \right\} \\ &\text{subject to } |M| \leq k \\ &\forall x \in X : \sum_{\mu \in M} \kappa(x, \mu) = 1 . \end{aligned}$$

There are two apparent differences between facility location/clustering and radius covering problems. First, in facility location and clustering, which are sum-type problems, we can replace a weighted input set by its respective multiset, and vice versa. The objective functions on the weighted set and the multiset are the same. Any algorithm for either of these problems could simply pick the representation of the input set which is more convenient. This is not the case for radius covering problems. Since this is a max-type problem, there is a significant difference between considering a weighted set and considering its respective multiset. We can remove all but one copy of each point in the input set of a radius covering problem and obtain the same optimal solutions. In contrast, replacing w by $\mathbb{1}$ potentially has a large influence on optimal solutions. Second, radius covering does not necessarily have an optimal solution with a binary assignment function κ . However, it is still easy to compute an optimal assignment κ , with respect to a given set of representatives.

Lemma 3.2.3 *Let (k, X, R, w, d) be an instance of RADIUS k -COVER and fix some $M \subseteq R$ with $|M| \leq k$. $\phi_{rad}(X, w, d, M, \kappa)$ is minimized if, for all $x \in X$ and $\mu \in M$, it holds that*

$$\kappa(x, \mu) = \begin{cases} 0 & \text{if } x \neq \mu \text{ and } x \in M \\ \mathfrak{m}_M(x)^{-1} & \text{if } x = \mu \\ \frac{d(x, \mu)^{-1}}{\sum_{\mu' \in M} d(x, \mu')^{-1}} & \text{else.} \end{cases}$$

Proof. The first two cases of the claim are easy to see. If a point $x \in X$ coincides with any number of representatives, then we distribute the assignment uniformly over all these representatives to obtain a cost of 0 for this point. This is a valid assignment since

$$\sum_{\mu \in M} \kappa(x, \mu) = \sum_{\mu \in \{\mu' \in M \mid x = \mu'\}} \kappa(x, \mu) = \frac{\mathfrak{m}_M(x)}{\mathfrak{m}_M(x)} = 1 .$$

Distributing the assignment uniformly is completely arbitrary and can be replaced by any other valid assignment of the point to the representatives it coincides with. It is however important that the assignment to the representatives not coinciding with the point is 0.

For the third case of the claim, fix any $x \in X$ which is distinct from all representatives. We argue that if $\phi_{rad}(\{x\}, w, d, M, \kappa)$ is minimized, then there exists some $\alpha \in \mathbb{R}$ such that

$$\forall \mu \in M : \kappa(x, \mu) \cdot w(x) \cdot d(x, \mu) = \alpha .$$

Assume to the contrary that there exist $\mu, \mu' \in M$ such that

$$\kappa(x, \mu) \cdot d(x, \mu) > \kappa(x, \mu') \cdot d(x, \mu') .$$

We set

$$\epsilon := \frac{\kappa(x, \mu) \cdot d(x, \mu) - \kappa(x, \mu') \cdot d(x, \mu')}{d(x, \mu) + d(x, \mu')} ,$$

and observe that $0 < \epsilon < \kappa(x, \mu)$. Consider the modified assignment κ' , with $\kappa'(x, \mu) := \kappa(x, \mu) - \epsilon$, $\kappa'(x, \mu') := \kappa(x, \mu') + \epsilon$, and equal to κ for all other representatives. Notice that κ' is also a valid assignment. Furthermore, we obtain

$$\begin{aligned} & \kappa'(x, \mu) \cdot d(x, \mu) \\ &= (\kappa(x, \mu) - \epsilon) \cdot d(x, \mu) \\ &= \kappa(x, \mu) \cdot d(x, \mu) - \frac{\kappa(x, \mu) \cdot d(x, \mu)^2 - \kappa(x, \mu') \cdot d(x, \mu') \cdot d(x, \mu)}{d(x, \mu) + d(x, \mu')} \\ &= \frac{\kappa(x, \mu) \cdot d(x, \mu) \cdot d(x, \mu') + \kappa(x, \mu') \cdot d(x, \mu') \cdot d(x, \mu)}{d(x, \mu) + d(x, \mu')} \\ &= \kappa(x, \mu') \cdot d(x, \mu') + \frac{\kappa(x, \mu) \cdot d(x, \mu) \cdot d(x, \mu') - \kappa(x, \mu') \cdot d(x, \mu')^2}{d(x, \mu) + d(x, \mu')} \\ &= (\kappa(x, \mu') + \epsilon) \cdot d(x, \mu') = \kappa'(x, \mu') \cdot d(x, \mu') , \end{aligned}$$

which is strictly less than $\kappa(x, \mu) \cdot d(x, \mu)$. We can repeatedly apply this process until x has the same cost with respect to every representative. This never increases the overall cost of the assignment. In particular, it reduces the overall cost each time we apply this to a uniquely-most-expensive representative μ . Let κ be the assignment where the cost of x with respect to every representative is α . We have that

$$1 = \sum_{\mu \in M} \kappa(x, \mu) = \sum_{\mu \in M} \alpha \cdot w(x)^{-1} \cdot d(x, \mu)^{-1} ,$$

3.2 Clustering

hence

$$\alpha = \frac{w(x)}{\sum_{\mu \in M} d(x, \mu)^{-1}},$$

and thus

$$\forall \mu \in M : \kappa(x, \mu) = \frac{d(x, \mu)^{-1}}{\sum_{\mu' \in M} d(x, \mu')^{-1}}.$$

■

It is not inherently clear that optimal solutions of radius covering problems consist of k distinct representatives. In the facility location and clustering case, the argument is that a solution where a representative occurs multiple times has the same cost as the same solution where all but one of the identical representatives are removed. [Lemma 3.2.3](#) shows us that this is not the case for radius covering problems. A second of the same representative might actually reduce the overall cost. Additionally, just as with [Problem 3.2.1](#) we do not know of a universal way of finding an optimal set of representatives given some fixed assignment function.

[Problem 3.2.1](#) and [Problem 3.2.2](#) are still formulated broadly, encompassing many different practical problems. In the following, we introduce the special cases and variants of these problems that we focus on in this thesis. We distinguish two different classes of problems: *hard* and *soft* clustering problems. We denote $n := |X|$ for all clustering problems.

One thing all problems in both classes have in common is that the domain of the input points is a vector space over the reals. As it turns out, the *mean* or *center of gravity* of a weighted multiset of vectors is important for these types of problems. Hence, we introduce the following notation.

Definition 3.2.4 (Mean) *Let $d \in \mathbb{N}$, $X \subseteq \mathbb{R}^d$, and $f : X \rightarrow \mathbb{R}$ be some function. If $f(x) > 0$, then we denote by*

$$\mu_f(X) := \frac{\sum_{x \in X} f(x) \cdot x}{f(X)}$$

the mean of X weighted by f . As a shorthand notation for non-empty X , we denote

$$\mu(X) := \mu_{\mathbb{1}}(X) = \frac{1}{n} \cdot \sum_{x \in X} x.$$

3.2.1 Hard Clustering

The first class of clustering problems are hard clustering problems. The name stems from the hard, sometimes also called *crisp*, assignment of data points to clusters. In this context, a *cluster* is the subset of data points assigned to the same representative. That is, each point is assigned to exactly one of the representatives, similar to MUFL. We formalize the notion of a cluster as follows.

Definition 3.2.5 (Cluster) *Let X be a finite multiset, M be a set, and $d : X \times M \rightarrow \mathbb{R}_{\geq 0}$ be a function. For each $\mu \in M$ we call*

$$\mathcal{C}_{\mu}^{(X,M)}(d) := \left\{ x \in X \mid \mu = \arg \min_{\mu' \in M} \{d(x, \mu')\} \right\}$$

the cluster of μ in X . We often omit the argument d if the distance function is clear from context.

The k -MEANS (KM) problem, one of the most popular clustering problems to date, arises naturally as a special case of [Problem 3.2.1](#). We fix some dimension $d \in \mathbb{N}$, restrict the input to multisets of points in \mathbb{R}^d , set $R = \mathbb{R}^d$, and let the distance be the squared Euclidean distance.

Problem 3.2.6 (k -MEANS) *Let $k, d \in \mathbb{N}$, $X \subseteq \mathbb{R}^d$ be finite, and $w : X \rightarrow \mathbb{N}$ be a function. The goal is to find $M \subset \mathbb{R}^d$ minimizing*

$$\begin{aligned} \text{km}(X, w, M) &:= \sum_{x \in X} w(x) \cdot \min_{\mu \in M} \{ \|x - \mu\|_2^2 \} \\ &\text{subject to } |M| \leq k . \end{aligned}$$

For reasons already discussed, and as we have done with MUFL, we omit the assignment function in the problem definition and restrict M to sets. A core observation on the KM problem is that the representative of each cluster of an optimal solution is the cluster's mean – hence, the name k -MEANS. Moreover, using a representative different from the mean, increases the cost by the distance of that representative to the mean multiplied by the overall weight of the cluster.

Lemma 3.2.7 *Let $d \in \mathbb{N}$, $X \subseteq \mathbb{R}^d$ be finite, $f : X \rightarrow \mathbb{R}_{\geq 0}$ be a function, and $\mu \in \mathbb{R}^d$. If $f(X) > 0$, then we have*

$$\sum_{x \in X} f(x) \cdot \|x - \mu\|_2^2 = \sum_{x \in X} f(x) \cdot \|x - \mu_f(X)\|_2^2 + \|\mu - \mu_f(X)\|_2^2 \cdot f(X) .$$

3.2 Clustering

Proof. Let $\langle \cdot, \cdot \rangle$ denote the dot product and observe

$$\begin{aligned}
& \sum_{x \in X} f(x) \cdot \|x - \mu\|_2^2 \\
&= \sum_{x \in X} f(x) \cdot \|x - \mu_f(X) + \mu_f(X) - \mu\|_2^2 \\
&= \sum_{x \in X} f(x) \cdot \|x - \mu_f(X)\|_2^2 + \|\mu_f(X) - \mu\|_2^2 \cdot f(X) \\
&\quad + 2 \cdot \sum_{x \in X} f(x) \cdot \langle x - \mu_f(X), \mu_f(X) - \mu \rangle
\end{aligned}$$

where

$$\begin{aligned}
& \sum_{x \in X} f(x) \cdot \langle x - \mu_f(X), \mu_f(X) - \mu \rangle \\
&= \left\langle \underbrace{\sum_{x \in X} f(x) \cdot (x - \mu_f(X))}_{=\mathbf{0}_d}, \mu_f(X) - \mu \right\rangle = 0.
\end{aligned}$$

■

Lemma 3.2.7 implies that the optimal representative for each cluster is its mean. Furthermore, for each optimal solution M we have

$$\forall \mu \in M : \mu = \mu_w \left(\mathcal{C}_\mu^{(X, M)} \right) = \frac{\sum_{x \in \mathcal{C}_\mu^{(X, M)}} w(x) \cdot x}{w \left(\mathcal{C}_\mu^{(X, M)} \right)}.$$

As a shorthand notation, we denote the smallest single-mean KM cost of some multiset X with weight function w as

$$\text{km}(X, w) := \text{km}(X, w, \{\mu_w(X)\}).$$

A radius hard clustering does not arise as naturally from **Problem 3.2.2**. As we showed in **Lemma 3.2.3**, an optimal assignment is not necessarily assigning each point to exactly one representative. Hence, to formalize the RADIUS k -MEANS (RKM) problem we have to restrict κ to a function to $\{0, 1\}$ instead of $[0, 1]$. This leads to the desired effect that an optimal κ assigns each point to its closest representative.

Problem 3.2.8 (RADIUS k -MEANS) *Let $k, d \in \mathbb{N}$, $X \subset \mathbb{R}^d$ be finite, and $w : X \rightarrow \mathbb{N}$ be a function. The goal is to find $M \subset \mathbb{R}^d$ minimizing*

$$\begin{aligned}
& \text{rkkm}(X, w, M) := \max_{x \in X} \left\{ w(x) \cdot \min_{\mu \in M} \left\{ \|x - \mu\|_2^2 \right\} \right\} \\
& \text{subject to } |M| \leq k.
\end{aligned}$$

Finding an optimal set of representatives given some fixed clusters is not as easy as it was for KM. In computational geometry, this problem is known as the *smallest enclosing sphere* problem (in the unweighted case), or the *1-center* problem (in the general weighted case).

There is an interesting connection between KM and RKM. Similar to our first, general definition of these types of problems, the objective of RKM is the same as that of KM, only with the sum replaced by taking the maximum over all points. However, to obtain the RKM objective function from [Problem 3.2.2](#), we had to make the constraints that derived KM from [Problem 3.2.1](#) and additionally, explicitly restrict κ to binary assignments. In the next section we introduce a soft clustering and a soft radius problem in Euclidean space. There however, the radius problem arises naturally and we have to put in additional effort to arrive at a clustering problem which has the radius problem as a max-cost-per-point variant.

Choosing the Number of Clusters Finding an appropriate choice for the parameter k constitutes a whole research area of its own. If k is chosen too large or too small, then there is no real value in interpreting solutions of clustering problems. Additionally, choosing the *correct* k is highly dependent on the concrete application, and even the specific data set. In this thesis, we concern ourselves with the theoretical analysis of problems and algorithms, not with their actual application. For this reason, we simply assume k to be some value chosen in advance.

3.2.2 Soft Clustering

In a hard clustering, each point either belongs to some cluster, or it does not. In practice however, entities are not so black and white, but rather are a mixture of different influencing factors. To model this effect we consider soft clustering problems as an alternative to classical hard clusterings. In the following, we define FUZZY k -MEANS (FKM), a popular soft clustering problem in Euclidean space. Just as with KM we fix some $d \in \mathbb{N}$, restrict the input to multisets from \mathbb{R}^d , set $R = \mathbb{R}^d$, and choose the squared Euclidean distance. To distinguish this problem from KM we furthermore choose an $m \in \mathbb{R}_{>1}$ (called the *fuzzifier*) and raise the assignment κ to the m^{th} power. Thereby, we obtain that optimal solutions do not have a binary assignment of points to representatives (except for some trivial border cases).

3.2 Clustering

Problem 3.2.9 (FUZZY k -MEANS) *Let $k, d \in \mathbb{N}$, $m \in \mathbb{R}_{>1}$, $X \subseteq \mathbb{R}^d$ be finite, and $w : X \rightarrow \mathbb{N}$ be a function. The goal is to find $M \subseteq \mathbb{R}^d$ and $\kappa : X \times M \rightarrow [0, 1]$ minimizing*

$$\begin{aligned} \text{fkm}(X, w, M, \kappa) &:= \sum_{x \in X} \sum_{\mu \in M} \kappa(x, \mu)^m \cdot w(x) \cdot \|x - \mu\|_2^2 \\ &\text{subject to } |M| \leq k \\ &\forall x \in X : \sum_{\mu \in M} \kappa(x, \mu) = 1. \end{aligned}$$

Notice that we deliberately omitted the fuzzifier from the arguments of the objective function. It is not sensible to optimize for m since every non-trivial solution can always be improved by increasing m . Hence, unless we state some fixed value for m , we always consider m to be some (small) constant chosen in advance. Further observe that optimal solutions for FKM with $m = 1$ would coincide with optimal solutions for KM on the same instance (k, X, w) .

Applying [Lemma 3.2.7](#), we see that for any optimal solution (M, κ) we have

$$\forall \mu \in M : \mu = \mu_{\kappa(\cdot, \mu)^m \cdot w(\cdot)}(X) = \frac{\sum_{x \in X} \kappa(x, \mu)^m \cdot w(x) \cdot \|x - \mu\|_2^2}{\sum_{x \in X} \kappa(x, \mu)^m \cdot w(x)}.$$

In the case of KM it was easy to find an optimal assignment given some set M – simply assign each point to its nearest representative. For FKM this turns out to be a little more complicated, but still efficiently computable.

Lemma 3.2.10 *Let $k, d \in \mathbb{N}$, $X \subseteq \mathbb{R}^d$ be finite, $w : X \rightarrow \mathbb{N}$ be a function, and $M \subseteq \mathbb{R}^d$ with $|M| = k$. $\text{fkm}(X, w, M, \kappa)$ is minimized if, for all $x \in X$ and $\mu \in M$, it holds that*

$$\kappa(x, \mu) = \begin{cases} 0 & \text{if } x \neq \mu \text{ and } x \in M \\ \mathbf{m}_M(x)^{-1} & \text{if } x = \mu \\ \frac{\|x - \mu\|_2^{-\frac{2}{m-1}}}{\sum_{\mu' \in M} \|x - \mu'\|_2^{-\frac{2}{m-1}}} & \text{else.} \end{cases}$$

Proof. As in the proof of [Lemma 3.2.3](#), the first two cases are easy to see, and uniformly distributing the assignment is arbitrary.

For the third case, we introduce n Lagrange multipliers λ_x and consider the modified objective function

$$\text{fkm}'(X, w, M, \kappa) = \text{fkm}(X, w, M, \kappa) + \sum_{x \in X} \lambda_x \cdot \left(\sum_{\mu \in M} \kappa(x, \mu) - 1 \right).$$

Fix some $x \in X \setminus M$ and $\mu \in M$. We find a stationary point of the modified objective function in the direction of $\kappa(x, \mu)$, by setting the partial derivative to zero

$$\begin{aligned} \frac{\partial \text{fkm}'}{\partial \kappa(x, \mu)}(X, w, M, \kappa) &= m \cdot \kappa(x, \mu)^{m-1} \cdot w(x) \cdot \|x - \mu\|_2^2 + \lambda_x = 0 \\ \Leftrightarrow \kappa(x, \mu) &= \left(\frac{-\lambda_x}{m \cdot w(x) \cdot \|x - \mu\|_2^2} \right)^{\frac{1}{m-1}}. \end{aligned}$$

Using the constraint $\sum_{\mu' \in M} \kappa(x, \mu') = 1$ we obtain

$$-\lambda_x = \left(\sum_{\mu' \in M} \left(\frac{1}{m \cdot w(x) \cdot \|x - \mu'\|_2^2} \right)^{\frac{1}{m-1}} \right)^{-m+1},$$

and hence

$$\begin{aligned} \kappa(x, \mu) &= \left(\frac{\left(\sum_{\mu' \in M} \left(\frac{1}{m \cdot w(x) \cdot \|x - \mu'\|_2^2} \right)^{\frac{1}{m-1}} \right)^{-m+1}}{m \cdot w(x) \cdot \|x - \mu\|_2^2} \right)^{\frac{1}{m-1}} \\ &= \frac{\|x - \mu\|_2^{-\frac{2}{m-1}}}{\sum_{\mu' \in M} \|x - \mu'\|_2^{-\frac{2}{m-1}}}. \end{aligned}$$

■

Observe that the proof of [Lemma 3.2.10](#) does not actually use any properties of the squared Euclidean distance. Hence, the result trivially generalizes to any other distance function used in the FKM objective function.

For every $M \subseteq \mathbb{R}^d$ we call an assignment function κ chosen according to [Lemma 3.2.10](#) the assignment *induced* by M . Consequently, we call

$$\begin{aligned} \text{fkm}(X, w, M) &:= \sum_{x \in X \setminus M} \sum_{\mu \in M} \left(\frac{\|x - \mu\|_2^{-\frac{2}{m-1}}}{\sum_{\mu' \in M} \|x - \mu'\|_2^{-\frac{2}{m-1}}} \right)^m \cdot w(x) \cdot \|x - \mu\|_2^2 \\ &= \sum_{x \in X \setminus M} \frac{w(x)}{\left(\sum_{\mu' \in M} \|x - \mu'\|_2^{-\frac{2}{m-1}} \right)^{m-1}} \end{aligned}$$

the FKM cost induced by M .

3.2 Clustering

For every feasible assignment function $\kappa : X \times M \rightarrow [0, 1]$ we call the set

$$\{\mu_{\kappa(\cdot, \mu)^m, w(\cdot)}(X) \mid \mu \in M\}$$

the set of representatives induced by κ , and denote by

$$\text{fkm}(X, w, \kappa) := \text{fkm}(X, w, \{\mu_{\kappa(\cdot, \mu)^m, w(\cdot)}(X) \mid \mu \in M\}, \kappa)$$

the FKM cost induced by κ .

A pair (M, κ) is called a *stationary pair* if M is induced by κ and κ is induced by M . We already observed that every optimal solution to FKM is a stationary pair. The converse however is not true. In fact, a stationary pair can be arbitrarily worse than an optimal solution [Blömer et al., 2016a].

Choosing the Fuzzifier As already discussed, increasing m always reduces the cost of every non-trivial solution. By Lemma 3.2.10 we also know that when m gets larger, an optimal assignment gets closer to the uniform distribution. The corresponding optimal representatives all move closer to the mean of the data set $\mu_w(X)$. Hence, if m is chosen too large, then solutions lack any useful interpretation. We treat m the same way we treat k and just assume this to be some constant chosen from the outside.

Nevertheless, one important choice we want to explicitly discuss is $m = 2$, as this case has a nice geometric interpretation. The induced cost of any set of representatives for FKM with $m = 2$ is

$$\text{fkm}(X, w, M) = \sum_{x \in X \setminus M} \frac{w(x)}{\sum_{\mu \in M} \|x - \mu\|_2^{-2}},$$

which is, for each point x , the unnormalized harmonic mean of the k distances $\|x - \mu\|_2^2$. Some of the results presented here hold only for this choice of m . As we see next, $m = 2$ furthermore has an interesting parallel to radius clustering.

As our final problem statement we consider FUZZY RADIUS k -MEANS (FRKM). This is a special case of Problem 3.2.2 with input sets from \mathbb{R}^d , $R = \mathbb{R}^d$, and the squared Euclidean distance.

Problem 3.2.11 (FUZZY RADIUS k -MEANS) *Let $k, d \in \mathbb{N}$, $X \subset \mathbb{R}^d$ be finite, and $w : C \rightarrow \mathbb{N}$ be a function. The goal is to find $M \subseteq \mathbb{R}^d$ and $\kappa : X \times M \rightarrow [0, 1]$ minimizing*

$$\begin{aligned} \text{frkm}(X, w, M, \kappa) &:= \max_{x \in X} \left\{ \max_{\mu \in M} \left\{ \kappa(x, \mu) \cdot w(x) \cdot \|x - \mu\|_2^2 \right\} \right\} \\ &\text{subject to } |M| \leq k \\ &\quad \forall x \in X : \sum_{\mu \in M} \kappa(x, \mu) = 1. \end{aligned}$$

For any set of representatives M , [Lemma 3.2.3](#) tells us that a corresponding optimal assignment is

$$\forall x \in X \setminus M, \forall \mu \in M : \kappa(x, \mu) = \frac{\|x - \mu\|_2^{-2}}{\sum_{\mu' \in M} \|x - \mu'\|_2^{-2}},$$

which is also an optimal assignment for FKM with $m = 2$, for the same set of representatives. This implies that by substituting an optimal assignment into the objective function, we obtain

$$\text{frkm}(X, w, M) := \max_{x \in X \setminus M} \left\{ \frac{w(x)}{\sum_{\mu \in M} \|x - \mu\|_2^{-2}} \right\},$$

which is the per point maximum of the FKM cost $\text{fkm}(X, w, M)$.

3.2.3 Discussion

In the following, we discuss how KM, RKM, FKM (with $m = 2$), and FRKM relate to each other. We derived these four problems from the two base problems [Problem 3.2.1](#) and [Problem 3.2.2](#). The difference between these two base problems is that we replaced the sum over the data points and the sum over the representatives by taking the maximum over each. Similarly, if we plug an optimal assignment function into the FKM objective function, then the difference between KM/FKM and RKM/FRKM is that the clustering problems sum over the data points and the radius problems take the maximum. To arrive at this situation we had to make different changes to the constraints of the problems. For the hard clustering variants, we restricted the assignment function of the radius problem to take values from $\{0, 1\}$ instead of $[0, 1]$. For the soft clustering variants, we raised the assignment function in the objective function of the clustering problem to the second power. These are substantially different modifications to arrive at problems which

3.2 Clustering

have a similar relation to one another. This is particularly unexpected for the soft clustering case. The clustering problem takes the sum over the data points and squares the assignment function, while the radius problem takes the maximum over points and representatives. Still, given some set of representatives, the optimal assignment functions are the same for the two problems. These observations indicate that all four problems are closely related and that any result for either of them might lead to non-trivial insight into all of them.

3.2.4 Further Analysis

We supplement our arguments from the previous section by presenting some well-known formal results regarding the relationship of our problems to each other. First, we see that there is an upper bound on the difference of the cost of some set of representatives with respect to the soft and hard problems.

Lemma 3.2.12 *Let $k, d \in \mathbb{N}$, $m \in \mathbb{R}_{>1}$ $X \subseteq \mathbb{R}^d$ be finite, and $w : X \rightarrow \mathbb{N}$ be a function. For each $M \subseteq \mathbb{R}^d$ with $|M| \leq k$ we have*

$$\text{fkm}(X, w, M) \leq \text{km}(X, w, M) \leq k^{m-1} \cdot \text{fkm}(X, w, M) , \text{ and}$$

$$\text{frkm}(X, w, M) \leq \text{rkm}(X, w, M) \leq k \cdot \text{frkm}(X, w, M) .$$

Proof. The first inequalities are immediate. Setting $\kappa(x, \mu) = 1$ for $\mu = \arg \min_{\mu' \in M} \{\|x - \mu'\|_2\}$ is a feasible assignment function. The soft and hard objective functions are the same for this assignment. Hence, with respect to an optimal assignment function, the cost of the soft problem can not be larger than the respective cost of the hard problem.

For the clustering case we observe

$$\begin{aligned} \text{fkm}(X, w, M) &= \sum_{x \in X \setminus M} \frac{w(x)}{\left(\sum_{\mu \in M} \|x - \mu\|_2^{-\frac{2}{m-1}} \right)^{m-1}} \\ &\geq \sum_{x \in X \setminus M} \frac{w(x)}{\left(\sum_{\mu \in M} (\min_{\mu' \in M} \{\|x - \mu'\|_2\})^{-\frac{2}{m-1}} \right)^{m-1}} \\ &= \sum_{x \in X \setminus M} \frac{w(x)}{k^{m-1} \cdot (\min_{\mu' \in M} \{\|x - \mu'\|_2\})^{-2}} \\ &= \frac{1}{k^{m-1}} \cdot \text{km}(X, w, M) . \end{aligned}$$

Similarly, the radius case can be bounded as

$$\begin{aligned} \text{frkm}(X, w, M) &= \max_{x \in X \setminus M} \left\{ \frac{w(x)}{\sum_{\mu \in M} \|x - \mu\|_2^{-2}} \right\} \\ &\geq \max_{x \in X \setminus M} \left\{ \frac{w(x)}{\sum_{\mu \in M} (\min_{\mu' \in M} \{\|x - \mu'\|_2\})^{-2}} \right\} \\ &= \frac{1}{k} \cdot \text{rkm}(X, w, M) . \end{aligned}$$

■

Second, we show that the representatives of an optimal solution always lie in the interior of the data set.

Lemma 3.2.13 *If M is an optimal set of representatives for KM, FKM, RKM, or FRKM, then the representatives all lie inside the convex hull of the input data set X .*

Proof. Assume to the contrary that there exists some $M' \subseteq M$ such that all $\mu \in M'$ lie outside of the convex hull of X . If no point is assigned to any $\mu \in M'$, then we can remove those representatives from the solution without increasing the cost. Afterwards, we add $|M'|$ many points at the location of $|M'|$ different data points in X which have cost larger than 0, and assign these points to the newly added representative at their location. The cost of the other points still remains unchanged. However, the cost of the reassigned points decreases, and thus the initial set of representatives was not optimal.

Assume that, for each $\mu \in M'$, there is at least one point with non-trivial assignment to μ . If we move μ to its closest point in the convex hull of X , then the distance of all points in X to μ decreases. Thus, the overall cost of the solution decreases. Hence, the initial set of representatives was not optimal.

If some representatives outside the convex hull have trivial assignment and some not, then we apply these arguments to each representative individually. ■

Finally, we present iterative relocation: a simple scheme to find stationary points of the objective functions of KM and FKM. Iterative relocation, sometimes also called alternating optimization, describes a general framework to solve clustering problems, similar to the Expectation-Maximization algorithm for mixture models [Bishop, 2006]. The idea is to start with an arbitrarily chosen set of representatives M . Then we repeat the following two steps until we reach a stationary pair:

3.3 Related Work

1. Compute an optimal assignment of points to M . That is, for KM compute for each $\mu \in M$ the cluster $\mathcal{C}_\mu^{(X,M)}$, and for FKM compute an optimal assignment function κ according to [Lemma 3.2.10](#).
2. Compute a new set of representatives based on this assignment. That is, for KM replace μ by $\mu_w(\mathcal{C}_\mu^{(X,M)})$, and for FKM replace each μ by $\mu_{\kappa(\cdot,\mu)^m \cdot w(\cdot)}(X)$.

These iterative relocation schemes are also known as the k -MEANS algorithm [[Lloyd, 1982](#)] and the FUZZY k -MEANS algorithm [[Bezdek et al., 1984](#)], respectively. They have two significant downsides. There is no known polynomial upper bound on the number of update steps until convergence is reached. In fact, we show in [Chapter 6](#) that there are data sets for the FKM case where, assuming a model of computation with arbitrary precision, iterative relocation never reaches a stationary pair. Furthermore, the cost of a stationary pair can be arbitrarily bad relative to the cost of an optimal solution [[Kanungo et al., 2004](#); [Blömer et al., 2016a](#)]. From a theoretician's point of view, iterative relocation has no desirable properties: neither polynomial runtime, nor a guaranteed approximation ratio.

3.3 Related Work

The first mention of a facility location type problem goes back to Pierre de Fermat, who asked for a point minimizing the sum of the distances to the vertices of a given triangle [[de Fermat, 1891](#)]. Today, facility location is ubiquitous, especially in operations research, and comes in many different flavors. Among others, one can consider capacitated problems where there is a bound on the number of clients each facility can serve, multi-commodity problems where clients can demand different types of goods, or any combination of models (see, for example, [[Pirkul & Jayaraman, 1998](#)]). Most of these formulations come with their own set of issues and essentially constitute a research area on their own.

A solution to KM solves the problem of least sum of squared errors. Its iterative relocation scheme is one of the most important clustering algorithms to date. It was first introduced by [Lloyd \[1982\]](#) and has since been the topic of vast amounts of research. The runtime of the algorithm is bounded from above by the number of different Voronoi partitions with respect to k centers. That is, the k -MEANS algorithm stops after at most $\mathcal{O}(n^{d \cdot (k+1)})$ iterations [[Inaba et al., 1994](#)], which

is polynomial if d and k are considered to be constants. If however k is part of the input, then there are data sets in the plane (i.e. $d = 2$) where the k -means algorithm requires $2^{\Omega(n)}$ iterations [Vattani, 2011]. Furthermore, there are simple examples of data sets ($n = 4$, $k = 3$, and $d = 1$) where the algorithm gets stuck at an arbitrarily poor local minimum of the objective function [Kanungo et al., 2004]. Despite these pitfalls, the k -MEANS algorithm is still employed extensively in practice. This is due to its simplicity and empirically fast runtime. The fast convergence of the algorithm observed in practice finds theoretical foundation in smoothed analysis. In a nutshell, this means that the data sets for which the k -MEANS algorithm exhibits an exponentially large runtime are pathologic instances. If input points are just slightly perturbed by Gaussian noise, then the algorithm has an expected polynomial runtime [Arthur & Vassilvitskii, 2006; Arthur et al., 2009; Manthey & Röglin, 2009]. For a more thorough overview on theoretical aspects of the k -MEANS algorithm we refer the interested reader to the survey of Blömer et al. [2016b].

The literature on FKM is not as extensive as it is for KM. The FKM objective was first formalized by Dunn [1973]. Building on this work, Bezdek et al. [1984] generalized the objective function and presented some analysis of the problem as well as the iterative relocation scheme following the KM example. Continuing their work, Bezdek et al. [1987] showed that the FUZZY k -MEANS algorithm converges to a stationary point of the objective function. Examining whether the algorithm reached a local minimum or a saddle point received some attention of its own [Kim et al., 1988; Hoppner & Klawonn, 2003]. Just as for the k -MEANS algorithm, there are simple examples of data sets ($n = 4$, $k = 2$, and $d = 2$) where the algorithm gets stuck at an arbitrarily poor local minimum of the objective function [Blömer et al., 2016a].

Part II



Classification and Impossibility

Classical Hardness

Contribution Summary We present classical computational hardness results for discrete versions of our clustering problems. In [Section 4.4](#) we discuss the difficulty with reductions to our soft clustering problems and a hardness result of a variant of DISCRETE FUZZY k -MEANS for general metrics. Afterwards, in [Section 4.5](#) we present the parameterized hardness of the radius problem DISCRETE FUZZY RADIUS k -MEANS. Finally, the reduction presented for DISCRETE k -MEANS in [Section 4.6](#) is a technical adaptation of a proof originally presented by [Papadimitriou \[1981\]](#) for the Euclidean distance to squared Euclidean distance. These are previously unpublished results.

4.1	Basics of Computational Complexity	32
4.2	Discrete Clustering Problems	34
4.3	Related Work	36
4.4	Towards a Classification of DFKM	36
	4.4.1 Construction	38
	4.4.2 Correctness	39
4.5	Parameterized Hardness of DFRKM	44
4.6	DKM is Hard in Fixed Dimensions	47

We discuss the placement of a specific decision variant of some of our location problems into complexity classes. The complexity of hard clustering problems is, in general, well understood. There is a detailed classification of KM and RKM. However, to the best of our knowledge, so far there has been no work on the computational complexity of the fuzzy problems FKM and FRKM. Classifying these problems turns out to be quite complicated due to the difficult algebraic structure of their objective functions. To obtain a combinatorially manageable form of these problems we discuss their what we call *discrete* variants.

4.1 Basics of Computational Complexity

The driving force behind computational complexity theory is the question about the tractability of solving problems. This topic is so essential to computer science that we teach the foundations to undergraduate students as part of almost every computer science curriculum. The literature on complexity theory is vast and cannot be summarized in truly concise form. For this reason, we expect the reader to be familiar with the core concepts and terminology of the field. A from-scratch introduction can, for example, be found in the excellent work of Papadimitriou [1993]; Goldreich [2008]; Arora & Barak [2009]; Sipser [2012].

The most important classes of computational complexity are **P**, the languages decidable by a deterministic Turing machine in polynomial time, and **NP**, the languages decidable by a non-deterministic Turing machine in polynomial time. We present an equivalent definition of **NP** via polynomially balanced relations.

Definition 4.1.1 (P and NP) *The class **P** consists of all polynomial-time decidable languages.*

We say a relation \sim is polynomially balanced if there exists $k \in \mathbb{N}$ such that, for all $x \sim y$, we have $|y| \leq |x|^k$.

*The class **NP** consists of all languages L for which there exists a polynomial-time decidable and polynomially balanced relation \sim_L such that $L = \{x \mid \exists y : x \sim_L y\}$.*

For each $L \in \mathbf{NP}$ and $x \in L$ we call a y with $x \sim_L y$ a witness for x 's membership in L .

The question whether $\mathbf{P} \subset \mathbf{NP}$ or $\mathbf{P} = \mathbf{NP}$ has been puzzling computer scientists for decades. Even with occasional claims to the resolution of the question, we are still lacking a universally accepted proof.

Problems classified in **NP** are so-called *decision problems* as the computational goal is to decide whether a given x is a member of language L or not. There is no obvious analogue of **NP** for optimization problems where we want to compute the *best* (where the quality of a solution varies, depending on the problem) solution given some instance x . Without going into detail, usually **MaxSNP** is considered to complement **NP** for the classification of optimization problem [Papadimitriou & Yannakakis, 1991].

A central tool of computational complexity are *reductions*. These allow us to find a hierarchy of problems and identify the most difficult problems of a complexity class.

4.1 Basics of Computational Complexity

Definition 4.1.2 (Polynomial-Time Reduction) *We say a language A is polynomial-time reducible to language B , write $A \leq_p B$, if there exists a polynomial-time computable function f such that $x \in A \Leftrightarrow f(x) \in B$.*

*A language L is called **NP**-hard if all problems in **NP** are polynomial-time reducible to it. If additionally $L \in \mathbf{NP}$, then L is called **NP**-complete.*

Since the inception of Boolean Satisfiability as the first **NP**-complete problem [Cook, 1971; Levin, 1973], there has been a lot of work on finding more **NP**-hard problems. Being **NP**-hard still classifies a problem as one of the hardest computational problems known. We introduce three hard problems which we use throughout this chapter.

- Given a set $U = \{u_1, \dots, u_m\}$ and a collection of subsets $S = \{S_1, \dots, S_k\}$ such that, for all $i \in [k]$, we have $S_i \subset U$ and $|S_i| = 3$. The EXACT COVER BY 3-SETS (X3C) problem asks if there is a collection of pairwise disjoint subsets $C \subset S$ such that $\bigcup_{S_i \in C} S_i = U$? X3C is a special case of the well-know **NP**-complete SET COVER problem [Karp, 1972].
- Given an undirected graph $G = (V, E)$ and $k \in \mathbb{N}$. The DOMINATING SET (DS) problem asks if there is a set of vertices $C \subseteq V$ with $|C| = k$ such that, for all $v \in V \setminus C$, there exists a $c \in C$ with $\{v, c\} \in E$? Simply speaking, is there a set C of k vertices in the graph so that every vertex is either in C , or in the neighborhood of a vertex in C ?
- Given a Boolean formula in conjunctive normal form where each clause consists of 3 non-negated literals. The POSITIVE 3-NOTALLEQUAL-SAT (P3NAESAT) problem asks if there is a truth assignment of the variables such that, in each clause, at least one literal evaluates to true and at least one to false?

Lemma 4.1.3 (Garey & Johnson [1979]) *X3C, DS, and P3NAESAT are **NP**-complete.*

In general, the notion of hardness and completeness are the same for every complexity class. However, the reduction used to classify problems might vary. For example, defining **P**-completeness with respect to polynomial-time reductions is not sensible. The reduction itself is able to solve the problem, and hence, all non-trivial languages are reducible to each other. For this reason, we usually use logarithmic-space reductions to define **P**-completeness. Throughout this thesis, we present the placement of some of our problems in certain not so well-known complexity classes. We introduce these classes together with their respective reduction types in the corresponding sections.

4.2 Discrete Clustering Problems

Whenever we discuss classical computational complexity, we need to consider the *decision variant* of our optimization problems. That is, given an instance of one of our location problems and a cost bound B , is there a solution with cost at most B ? To show that one of our problems lies in **NP** we need to prove the existence of a witness with size polynomial in the size of the instance and B .

It is not difficult to see that KM lies in **NP**. A witness is simply the mapping of the n input points to the k clusters. By [Lemma 3.2.7](#), we know that for each cluster an optimal representative is given by the mean of the cluster, which we can compute in polynomial time. Hence, we can represent and evaluate the cost of an optimal solution in time polynomial in the size of the input.

Realizing that RKM lies in **NP** is a little more involved. Finding a solution with cost at most B is equivalent to covering the given data set by k balls with radius B . To find an optimal solution we can partition the input data into k clusters and find the smallest enclosing sphere for each cluster. Such a smallest enclosing sphere is a $d' \leq d$ dimensional ball which is uniquely determined by $d' + 1$ points of the input data set on its boundary [[Cheng et al., 2006](#)]. Thus, a witness for an RKM solution with cost at most B , provided it exists, can be represented by at most $k \cdot (d + 1)$ points of the input data set. For each cluster we obtain the $d' + 1$ boundary points of the smallest enclosing ball. The optimal set of representatives is then given as the midpoints of these k balls. Again, we can represent and evaluate the cost of an optimal solution in time polynomial in the size of the input.

It is not readily apparent why FKM should lie in **NP**. A solution to an FKM instance consists of a continuous assignment function and a set of means. As we discuss in [Chapter 6](#), optimal solutions to FKM are in general not solvable by radicals. This is evidence that the problem might actually not be a part of **NP**. Assuming that for a given instance there exists an optimal solution with cost less than B , we do not know whether we can approximate this solution using size polynomial in the instance and B , but still maintain cost of at most B . Although our unsolvability result does not transfer to FRKM, we neither know if the radius variant of FKM is part of **NP**.

To circumvent the difficult structure of FKM solutions, in the following, we consider a combinatorial variant of our location problems. Similar to UFL type problems, we discretize the solution space by expanding instances by a set of candidate representatives. More specifically, we only allow representatives to be taken from the input points

4.2 Discrete Clustering Problems

X . It is immediately clear that these so-called *discrete* problems lie in **NP**. We denote our discrete problems as DISCRETE k -MEANS (DKM), DISCRETE RADIUS k -MEANS (DRKM), DISCRETE FUZZY k -MEANS (DFKM), and DISCRETE FUZZY RADIUS k -MEANS (DFRKM).

There are two main reasons for us to chose the input points as candidate representatives. First, instances are the same as instances for the non-discrete variant since we do not actually add a set of candidates. Second, an optimal solution to the discrete problem is also always a *good* solution to the original problem.

Lemma 4.2.1 *An optimal solution to DKM, or DFKM, is a 2-approximation to KM, or FKM respectively, on the same instance.*

An optimal solution to DRKM, or DFRKM, is a 4-approximation to RKM, or FRKM respectively, on the same instance.

Proof. Let (X, w) be an instance, (M^*, κ^*) be an optimal solution to FKM, and (M_d^*, κ_d^*) an optimal solution to DFKM. Consider the set M where for each $\mu \in M^*$ we add the point $x \in X$ closest to μ to M . Let $f : M^* \rightarrow M$ be the function mapping each optimal representative to its replacement. We bound

$$\begin{aligned}
 & \text{fkm}(X, w, M_d^*, \kappa_d^*) \\
 & \leq \text{fkm}(X, w, M) && (M \text{ is a discrete solution}) \\
 & \leq \sum_{x \in X} \sum_{\mu \in M} \kappa^*(x, \mu)^m \cdot w(x) \cdot \|x - \mu\|_2^2 && (\kappa^* \text{ is not optimal for } M) \\
 & \leq \text{fkm}(X, w, M^*, \kappa^*) + \sum_{x \in X} \sum_{\mu \in M^*} \kappa^*(x, \mu)^m \cdot w(x) \cdot \underbrace{\|\mu - f(\mu)\|_2^2}_{\leq \|x - \mu\|_2^2} \\
 & && (\text{Lemma 3.2.7 and choice of } M) \\
 & \leq 2 \cdot \text{fkm}(X, w, M^*, \kappa^*) .
 \end{aligned}$$

The same holds for KM and DKM where we replace the fuzzy assignment functions by the respective binary assignments.

Now, let (M^*, κ^*) be an optimal solution to FRKM, and (M_d^*, κ_d^*) an optimal solution to DFRKM. Let M again be chosen as above. We obtain

$$\begin{aligned}
 & \text{frkm}(X, w, M_d^*, \kappa_d^*) \\
 & \leq \text{frkm}(X, w, M) && (M \text{ is a discrete solution}) \\
 & \leq \max_{x \in X} \left\{ \max_{\mu \in M} \left\{ \kappa^*(x, \mu) \cdot w(x) \cdot \|x - \mu\|_2^2 \right\} \right\} && (\kappa^* \text{ is not optimal for } M) \\
 & \leq \max_{x \in X} \left\{ \max_{\mu \in M^*} \left\{ \kappa^*(x, \mu) \cdot w(x) \cdot \left(2 \cdot \|x - \mu\|_2^2 + 2 \cdot \|\mu - f(\mu)\|_2^2 \right) \right\} \right\} \\
 & && (\text{approximate triangle inequality for each point})
 \end{aligned}$$

$$\leq 4 \cdot \text{frkm}(X, w, M^*, \kappa^*) . \quad (\text{by choice of } M)$$

Again, the same holds for RKM and DRKM where we replace the fuzzy assignment functions by the respective binary assignments. ■

Choosing X as the candidate set for representatives is not the best possible choice. For KM and FKM there are algorithms computing a candidate set containing a $(1 + \epsilon)$ -approximation to the original problem [Matoušek, 2000; Blömer et al., 2016a]. However, these candidate sets cannot be constructed in polynomial time.

4.3 Related Work

There are numerous classification results regarding the complexity of our facility location and hard clustering problems. MUFL is known to be **NP**-hard and **MaxSNP**-hard, and if there exists a polynomial-time approximation algorithm with an approximation factor better than 1.463, then $\mathbf{NP} \subseteq \text{DTIME}(n^{\log(\log(n))})$ [Guha & Khuller, 1999]. Since MUFL is just a special case, this also holds for the general UFL problem.

The radius problems RKM and DRKM are also known as the absolute, and respectively discrete, k -CENTER problem. The former is **NP**-hard even in the plane [Megiddo & Supowit, 1984] and is **NP**-hard to approximate with a factor better than 1.82 with respect to the Euclidean distance [Feder & Greene, 1988]. The latter is **NP**-hard to approximate with a factor strictly smaller than 2 for general metrics [Hochbaum & Shmoys, 1986].

Similar results have been obtained for the clustering problem KM. It is **NP**-hard for $k = 2$ [Dasgupta, 2008] and also for $d = 2$ [Mahajan et al., 2012]. That means it is hard to find an optimal solution as long as either of the two parameters is part of the input. If both k and d are part of the input, then it is even **NP**-hard to approximate the problem with a factor strictly better than 1.0013 [Lee et al., 2017]. The discrete variant DKM is **NP**-hard for $d = 2$ with respect to the Euclidean distance [Papadimitriou, 1981]. To the best of our knowledge, there are no classification results regarding the soft clustering variants FKM, DFKM, FRKM, and DFRKM.

4.4 Towards a Classification of DFKM

Placing DFKM into classical complexity classes seems to be a hard problem itself. In the following, we discuss some indications as to

4.4 Towards a Classification of DFKM

why standard reduction techniques appear not to be applicable for this problem. Afterwards, we present the **NP**-hardness of a variant of DFKM with $m = 2$ using a general metric instead of squared Euclidean distance.

Most of the well-known **NP**-hard problems belong to one of two classes, let us call them *local* and *global* problems. By local problems, we mean that the task is to find some subset of the input satisfying some structural or quantitative constraint. Some examples for such a local problem would be CLIQUE/INDEPENDENTSET, KNAPSACK, or SUBSETSUM/SUBSETPRODUCT. A witness for these types of problems can be analyzed isolated from the rest of the instance. Each subset of vertices either is a clique of a certain size, or it is not. Each set of items either fulfils the weight bound and the has the required value, or it does not. As it seems, reducing such a problem to DFKM should not work. The cost of a set of representatives does not only depend on some subset of the input points, but on all of them.

By global problems, we mean that a solution to the problem satisfies some constraint for *all* elements of the input. Examples for global problems include BOOLEANSATISFIABILITY, VERTEXCOVER/EDGECOVER, or INTEGERPROGRAMMING. To analyze a witness for these problems, we need to consider the whole problem instance and verify that the posed constraints are fulfilled globally. These type of problems are naturally more similar to DFKM. Furthermore, reductions for other types of clustering problems, like the ones we discussed previously for DKM and DFRKM, come from this category.

The underlying structure of reductions to clustering problems mostly follows one general approach. Start by finding a set of points with a certain minimum pairwise distance. This essentially imposes a minimum cost on each point. Then, make sure that if the instance for the problem we are reducing from has a solution, then there is a set of representatives such that each point (at least almost) attains its minimum cost. If the original instance does not have a solution, then we want that, for every set of representatives, there is at least one point which costs significantly more than its minimum cost. Finally, choose the cost bound to be (maybe slightly more than) the number of points multiplied by each points minimum cost. The main difficulty then lies in embedding a set of points with such distances into as few dimensions as possible, or being able to restrict the number of clusters k . So far, this has been the recipe for success for almost all clustering hardness results.

Assume we have a reduction following this approach. Consider the image of such a reduction given a no-instance of the original

problem. As we said, for every set of representatives, there is a point with high cost. Let us call this the *uncovered* point. Most of the times this uncovered point corresponds (under the reduction function) to an uncovered element of the original instance. For example, an unsatisfied clause, or an edge not adjacent to a vertex in the cover. This is an inherent property of such a global problem. A single uncovered element is enough to invalidate a solution. The main problem that arises is that this is not true for DFKM. We call this *overcovering* a point. A central characteristic of reductions to clustering problems which we described above is that a minimum distance between points dictates a minimum cost each point has in each solution. A point in a DFKM solution which has multiple representatives at minimum distance however has significantly lower cost than this minimum distance. This is not beneficial for an element corresponding to such a point in the original instance. For example, we cannot make a clause more true by having more than one true literal. What this leads to is that, although there is an uncovered point, solutions to DFKM instances arising from a reduction fall below the cost threshold by overcovering some other parts of the input. This cannot be solved by simply lowering the cost threshold because we still need to recognize yes-instances where a witness covers each element by exactly one representative.

To find a reduction showing the **NP**-hardness of DFKM we need to deal with overcovering. So far, we do not know whether DFKM actually is **NP**-hard. However, we are able to show a weaker result by proving the **NP**-hardness of unweighted DFKM with $m = 2$ for general metrics.

Theorem 4.4.1 *A variant of DISCRETE FUZZY k -MEANS with $m = 2$, $w = 1$, and using a general metric (instead of squared Euclidean distance) is **NP**-complete.*

Our proof of [Theorem 4.4.1](#) is a reduction from P3NAESAT and is separated into two parts: the construction and the correctness.

4.4.1 Construction

Let φ be a Boolean formula over n variables consisting of m clauses of size 3. Without loss of generality, we assume that no literal appears more than once in any given clause. Let $X := \emptyset$. We set $W = 6 \cdot m$ and add W copies of each of $2 \cdot n$ different points to X , two for each variable in φ . For each variable, one of these points corresponds to the positive literal, the other to the negation. Next, we add m points to X corresponding to the clauses in φ . Finally, we add m additional points to X . For each clause $(x \vee y \vee z)$ in φ , we add a point corresponding

4.4 Towards a Classification of DFKM

to the (in φ non-existent) clause $(\bar{x} \vee \bar{y} \vee \bar{z})$. For the sake of brevity we abuse notation and call these points *literals* and *clauses*, omitting that they only correspond to a literal/clause in φ . We define the distance function $d : X \times X \rightarrow \mathbb{R}_{\geq 0}$ for our DFKM instance as

$$d(p, q) = d(q, p) := \begin{cases} 0 & \text{if } p = q \\ 1 & \text{if } p \text{ and } q \text{ are a literal and its negation} \\ n & \text{if } p \text{ is a literal appearing in clause } q \\ n + 1 & \text{if } \bar{p} \text{ is a literal appearing in clause } q \\ n + 2 & \text{else.} \end{cases}$$

The central idea is that literals are close to their negation, clauses are closer to literals they contain than to the literal's negation, and all other points are far away from each other. The DFKM instance is (X, n) and we ask for the cost bound

$$L := \frac{n \cdot W}{1 + \frac{n-1}{n+2}} + \frac{m}{\frac{1}{n} + \frac{2}{n+1} + \frac{n-3}{n+2}} + \frac{m}{\frac{2}{n} + \frac{1}{n+1} + \frac{n-3}{n+2}}.$$

Lemma 4.4.2 *The function d is a metric.*

Proof. Non-negativity, symmetry, and identity of indiscernibles are immediate by definition of d . Establishing that d fulfils the triangle inequality requires a few observations. We consider which sums of two non-zero distances might occur on the right hand side of the triangle inequality.

The smallest of sum is 2. This occurs only as $d(x, \bar{x}) + d(\bar{x}, x)$ for some variable x . However, in that case the left hand side of the inequality is $d(x, x)$ or $d(\bar{x}, \bar{x})$, which are both 0.

The next larger sum is $n + 1$. This appears as $d(x, b) + d(x, \bar{x})$ where x is a literal appearing in b (x might be a negative literal – the argument is the same). However, by definition we have $d(\bar{x}, b) = n + 1$, and hence the triangle inequality is fulfilled.

All other sums are at least $n + 2$ and thus trivially fulfil the triangle inequality. ■

4.4.2 Correctness

To show correctness we have to carefully analyze the structure of solutions for X . An important notion we use in this analysis is the classification of a solution as a so-called *reasonable* solution.

Definition 4.4.3 (Reasonable Solutions) *Let $M \in X$. We call M reasonable if for each variable x in φ*

$$x \in M \vee \bar{x} \in M .$$

Observe that a reasonable solution contains a representative at either the positive or the negative literal of each variable x . Since solutions consist of n representatives and φ contains n variables, this means that a reasonable solution never contains a representative at a variable and at its negation. By this, we can define a one-to-one mapping between truth assignments for φ and reasonable solutions for X . A variable x is true if and only if $x \in M$ and x is false if and only if $\bar{x} \in M$. This allows us to express the cost of a reasonable solution in terms of the number of NAE-satisfied clauses. In the following, for each reasonable solution M , we denote by $S(M)$ the number of NAE-satisfied clauses and by $U(M)$ the number of NAE-unsatisfied clauses of the assignment corresponding to M .

Lemma 4.4.4 *If $M \in X$ is reasonable, then*

$$\text{fkm}(X, M) = \frac{n \cdot W}{1 + \frac{n-1}{n+2}} + (\Gamma_1 + \Gamma_2) \cdot S(M) + (\Gamma_0 + \Gamma_3) \cdot U(M)$$

where

$$\Gamma_i = \left(\frac{i}{n} + \frac{3-i}{n+1} + \frac{n-3}{n+2} \right)^{-1} .$$

Proof. First, consider the cost of the literals. Each of these points is either in M and has cost 0, or it has a representative at its negated literal at distance 1 and $n-1$ representatives at distance $n+2$ and hence has cost

$$\frac{1}{1 + \frac{n-1}{n+2}} .$$

Recall that there are W copies of each of these points. Second, consider the cost of the clauses. Since clauses are of length 3, each of these has $n-3$ representatives at distance $n+2$. The others are either at distance n or at distance $n+1$, dependent on whether the literals are true or false in the clause. Hence, the cost of such a clause is Γ_i where i is the number of true literals in that clause. Recall that a clause is NAE-satisfied if it contains at least one true and one false literal. Furthermore, recall that by construction, for each clause b of φ , we added a clause where all literals from b are negated and which does

4.4 Towards a Classification of DFKM

not occur in φ . We consider the cost of the clauses in pairs of points b, \bar{b} where one corresponds to a clause in φ and the other to the newly added clause. If b has i true literals, then \bar{b} has exactly $3 - i$ true literals. Thus, each of these pairs has either cost $\Gamma_1 + \Gamma_2$ or $\Gamma_0 + \Gamma_3$, dependent on whether b is NAE-satisfied, or not. We obtain the claimed cost. ■

The first direction of the correctness proof is fairly simple. Assume there exists an assignment of the variables NAE-satisfying φ and let M be the corresponding reasonable solution. By [Lemma 4.4.4](#) we have

$$\text{fkm}(X, M) = \frac{n \cdot W}{1 + \frac{n-1}{n+2}} + (\Gamma_1 + \Gamma_2) \cdot m = L ,$$

since all clauses are NAE-satisfied.

For the second direction we have spend more effort by further analyzing solutions for X . The central argument lies in showing that any solution with cost at most L is reasonable.

Lemma 4.4.5 *If $M \in X$, such that $\text{fkm}(X, M) \leq L$, then M is reasonable.*

Proof. Assume M is not reasonable. By definition, there exist two points corresponding to literals of the same variable $x, \bar{x} \notin M$. Since $|M| = n$ we have one of two cases: Either there exists a point corresponding to a clause $b \in M$, or there exist two points corresponding to literals of the same variable $y, \bar{y} \in M$ such that $x \neq y$. In the following, we only discuss cases where exactly one of these two applies, and where there is no location at which there are multiple representatives. The general claim can easily be derived as a number of combinations of the presented analyses.

Case 1 Assume there is a representative at a clause b . We obtain that there are $n + 1$ variables with no representative at their location. Since $x, \bar{x} \notin M$, we observe that all representatives are at distance at least n to x and \bar{x} . The other $n - 1$ literals with non-trivial cost have at most 1 representative at distance 1 (their negation) and $n - 1$ at distance at least n . Hence, we obtain

$$\text{fkm}(X, M) \geq \frac{2 \cdot W}{\frac{n}{n}} + \frac{(n-1) \cdot W}{1 + \frac{n-1}{n}} = 2 \cdot W + \frac{n \cdot (n-1) \cdot W}{2 \cdot n - 1} .$$

Observe that for all $i \in [3]_0$ and $n \in \mathbb{N}$ we have

$$\begin{aligned}\Gamma_i &= \left(\frac{i}{n} + \frac{3-i}{n+1} + \frac{n-3}{n+2} \right)^{-1} \\ &\leq \left(\frac{i}{n+2} + \frac{3-i}{n+2} + \frac{n-3}{n+2} \right)^{-1} \\ &= \frac{n+2}{n} \leq 3.\end{aligned}$$

Thus, by choice of W , we have

$$W > (\Gamma_1 + \Gamma_2) \cdot m.$$

Furthermore, we bound

$$\begin{aligned}W + \frac{n \cdot (n-1) \cdot W}{2 \cdot n - 1} &= \frac{(2 \cdot n^3 + 3 \cdot n^2 - n - 1) \cdot W}{(2 \cdot n - 1) \cdot (2 \cdot n + 1)} \\ &\geq \frac{(2 \cdot n^3 + 3 \cdot n^2 - 2n) \cdot W}{(2 \cdot n - 1) \cdot (2 \cdot n + 1)} \quad (\text{since } n \geq 1) \\ &= \frac{n \cdot (n+2) \cdot W}{2 \cdot n + 1} \\ &= \frac{n \cdot W}{1 + \frac{n-1}{n+2}}.\end{aligned}$$

Taking the sum of the two lower bounds we obtain the claim for this case.

Case 2 Assume there are representatives at literals y, \bar{y} . There are still n literals with no representative at their location. Similar to before, all representatives are at distance at least $n+2$ to x and \bar{x} . The other $n-2$ literals with non-trivial cost have at most 1 representative at distance 1 (their negation) and $n-1$ at distance at least $n+2$. We obtain

$$\text{fkm}(X, M) \geq \frac{2 \cdot W}{\frac{n}{n+2}} + \frac{(n-2) \cdot W}{1 + \frac{n-1}{n+2}} = \frac{2 \cdot (n+2) \cdot W}{n} + \frac{(n-2) \cdot (n+2) \cdot W}{2 \cdot n + 1}.$$

As before, our choice of W yields

$$\frac{(n+2) \cdot W}{n} > W > (\Gamma_1 + \Gamma_2) \cdot m.$$

4.4 Towards a Classification of DFKM

Also similar to the first case we have

$$\begin{aligned}
\frac{(n+2) \cdot W}{n} + \frac{(n-2) \cdot (n+2) \cdot W}{2 \cdot n + 1} &= \frac{(n^2 + 1) \cdot (n+2) \cdot W}{n \cdot (2 \cdot n + 1)} \\
&\geq \frac{n^2 \cdot (n+2) \cdot W}{n \cdot (2 \cdot n + 1)} \\
&= \frac{n \cdot (n+2) \cdot W}{2 \cdot n + 1} \\
&= \frac{n \cdot W}{1 + \frac{n-1}{n+2}}.
\end{aligned}$$

Once again, taking the sum yields the claim and concludes the proof. \blacksquare

Recall that we can uniquely map each reasonable solution to a truth assignment of the variables in φ . Since

$$\begin{aligned}
&\Gamma_0 + \Gamma_3 \\
&= \frac{2 \cdot n^{10} + 10 \cdot n^9 + 45 \cdot n^8 + 140 \cdot n^7 + 332 \cdot n^6 + 667 \cdot n^5 + 990 \cdot n^4 + 1190 \cdot n^3 + 1040 \cdot n^2 + 504 \cdot n + 96}{n^{10} + 3 \cdot n^9 + 21 \cdot n^8 + 49 \cdot n^7 + 161 \cdot n^6 + 283 \cdot n^5 + 538 \cdot n^4 + 668 \cdot n^3 + 696 \cdot n^2 + 516 \cdot n + 144} \\
&> \frac{2 \cdot n^{10} + 10 \cdot n^9 + 45 \cdot n^8 + 140 \cdot n^7 + 328 \cdot n^6 + 639 \cdot n^5 + 900 \cdot n^4 + 990 \cdot n^3 + 720 \cdot n^2 + 216 \cdot n}{n^{10} + 3 \cdot n^9 + 21 \cdot n^8 + 49 \cdot n^7 + 161 \cdot n^6 + 283 \cdot n^5 + 538 \cdot n^4 + 668 \cdot n^3 + 696 \cdot n^2 + 516 \cdot n + 144} \\
&= \Gamma_1 + \Gamma_2,
\end{aligned}$$

we conclude our proof by observing that the cost of a reasonable solution is at most L if and only if $U(M) = 0$, i.e. all clauses are NAE-satisfied.

Our **NP**-hardness reduction for the general metric variant of DFKM circumvented the overcovering problem by exploiting the special structure of P3NAESAT. In this variant of the satisfiability problem clauses can actually be overcovered because setting all variables to true falsifies a clause. By complementing each clause by a sort-of negated version of itself we are able to find an exact characterisation of optimal solutions for our instance. Due to the difficult algebraic structure of the DFKM objective function we used pretty coarse lower bounds for the cost of non-reasonable solutions. This is strong evidence that it might be possible to apply our technique together with a sophisticated embedding to obtain a hardness result for DFKM.

Conjecture 4.4.6 DISCRETE FUZZY k -MEANS with $m = 2$ and $w = 1$ is **NP**-complete.

4.5 Parameterized Hardness of DFRKM

We present results classifying the hardness of the fuzzy variant of radius covering DFRKM. While we show that the problem is hard for unweighted instances, we require d and k to be part of the input. In the following, we first introduce a special case of a metric.

Definition 4.5.1 ($(1, 2) - B$ Metric) *Let U be a set and $B \in \mathbb{N}$. A metric $d : U \times U \rightarrow \mathbb{R}_{\geq 0}$ is called a $(1, 2) - B$ metric if,*

1. *for all $u, v \in U$ with $u \neq v$, we have $d(u, v) \in \{1, 2\}$ and*
2. *for all $u \in U$, we have $|\{v \in U \mid d(u, v) = 1\}| \leq B$.*

These types of metrics arise naturally in graph theory as the 2-truncated metric [Deza & Deza, 2009]. That is, we consider the standard path (also called geodesic) metric of a graph, but replace any distance greater than 2 by 2. For a graph $G = (V, E)$ this then forms a $(1, 2) - \delta(G)$ metric where $\delta(G) := \max_{v \in V} \{\deg(v)\}$ is the maximum degree of the vertices in the graph. $(1, 2) - B$ metrics are particularly interesting as they are a very restrictive type of metric, and despite having such a simple structure, there are well-known problems which are still hard when they are restricted to such a metric. For example, there exists a constant B such that TRAVELLING SALESMAN PROBLEM (TSP) is still **MaxSNP**-hard if it is restricted to $(1, 2) - B$ metrics [Papadimitriou & Yannakakis, 1993]. This means that there is no Polynomial Time Approximation Scheme (PTAS) for the problem unless $\mathbf{P} = \mathbf{NP}$ [Arora et al., 1998].

One of the main difficulties when dealing with geometric problems is finding an embedding of points exhibiting the required distances between points. For our purposes in this section we use an embedding result for $(1, 2) - B$ metrics from coding theory.

Lemma 4.5.2 (Trevisan [2000]) *Let U be a set with $|U| = n$, $B \in \mathbb{N}$, and $d : U \times U \rightarrow \mathbb{R}_{\geq 0}$ be a $(1, 2) - B$ metric. There exists a polynomial time algorithm which given U and d computes an embedding $f : U \rightarrow \mathbb{R}^{n(B+1)}$ such that, for all $u, v \in U$,*

- *if $d(u, v) = 1$, then $\|f(u) - f(v)\|_2^2 = 2 \cdot B$ and*
- *if $d(u, v) = 2$, then $\|f(u) - f(v)\|_2^2 = 2 \cdot (B + 1)$.*

This embedding is based on Hadamard codes, actually maps the elements in U to the unit hypercube $\{0, 1\}^{n(B+1)}$, and was designed as an isometry for the Hamming distance. However, we can repurpose

4.5 Parameterized Hardness of DFRKM

this because Hamming distance and squared Euclidean distance coincide on binary vectors. Using this embedding technique we obtain a straightforward hardness proof for DFRKM.

Theorem 4.5.3 DISCRETE FUZZY RADIUS k -MEANS with $w = \mathbb{1}$ is **NP**-complete.

Proof. In the following, we present a reduction from DS to DFRKM.

Let $(G = (V, E), k)$ be an instance of DS. By Lemma 4.5.2, there exists a function $f : V \rightarrow \mathbb{R}^{n \cdot (\delta(G)+1)}$ such that

- if $\{u, v\} \in E$, then $\|f(u) - f(v)\|_2^2 = 2 \cdot \delta(G)$, and
- if $\{u, v\} \notin E$, then $\|f(u) - f(v)\|_2^2 = 2 \cdot (\delta(G) + 1)$.

Our reduction maps (G, k) to the instance $(f[V], k)$ for DFRKM and asks for the cost bound

$$L = 2 \cdot \delta(G) \cdot \frac{\delta(G) + 1}{\delta(G) \cdot k + 1}.$$

Assume $C \subseteq V$ is a size k dominating set in G . We analyze the cost of the set $f[C]$. Every point in $f[V]$ is either in $f[C]$, or has at least one point in $f[C]$ at distance $2 \cdot \delta(G)$. Hence, we can bound

$$\begin{aligned} \text{frkm}(f[V], f[C]) &= \max_{x \in f[V] \setminus f[C]} \left\{ \frac{1}{\sum_{\mu \in f[C]} \|x - \mu\|_2^{-2}} \right\} \\ &\leq \frac{1}{(2 \cdot \delta(G))^{-1} + (k-1) \cdot (2 \cdot (\delta(G) + 1))^{-1}} = L. \end{aligned}$$

Conversely assume that there is no dominating set of size k in G . Hence, for every $C \subseteq f[V]$ with $|C| = k$, there exists an $x \in f[V]$ such that, for every $c \in C$, we have $\{f^{-1}(x), f^{-1}(c)\} \notin E$. We conclude

$$\text{frkm}(f[V], f[C]) = \frac{1}{k \cdot (2 \cdot (\delta(G) + 1))^{-1}} = \frac{2 \cdot (\delta(G) + 1)}{k} > L.$$

■

In general DFRKM is **NP**-complete if d and k are part of the input. In the following, we discuss that we can at least bound the dimension d we allow our instance to have.

Corollary 4.5.4 DISCRETE FUZZY RADIUS k -MEANS with $w = \mathbb{1}$ and $d \in \mathcal{O}(n)$ is **NP**-complete.

Proof. Use the reduction from the proof of [Theorem 4.5.3](#) and observe that DS is still **NP**-complete when restricted to graphs G with $\delta(G) = 3$ [[Garey & Johnson, 1979](#)]. ■

Interestingly, we are able to obtain a stronger hardness result by relaxing the bound on the dimensionality of the input set. To show this we take a shallow dive into *parameterized complexity*. The central notion of this domain of computational complexity is to understand languages as sets of tuples (x, k) where x is the *instance* and k the *parameter*. Such a parameterized language is called *fixed-parameter tractable* if there exists a computable function f and an algorithm deciding the language in time $f(k) \cdot \text{poly}(|x|)$. The idea of this being that the problem is polynomial-time solvable if we fix the parameter to a constant. The class **FPT** then consists of all fixed-parameter tractable parameterized languages. Above **FPT** there exists an infinite hierarchy of classes of hard parameterized languages called the **W**-hierarchy. Without going into the definition of the **W** classes, the hierarchy is arranged as follows

$$\mathbf{FPT} \subseteq \mathbf{W}[1] \subseteq \mathbf{W}[2] \subseteq \dots \subseteq \mathbf{W}[\text{Sat}] \subseteq \mathbf{W}[\mathbf{P}] .$$

A collapse of just the first inclusion of the **W**-hierarchy would also refute the *exponential time hypothesis*. That is, if $\mathbf{FPT} = \mathbf{W}[1]$, then $\mathbf{NP} \subseteq \text{DTIME}(2^{o(n)})$ [[Downey & Fellows, 2012](#)]. Problems in this hierarchy are related to one another using *parameterized reductions*.

Definition 4.5.5 (Parameterised Reduction) *We say that a parameterized language A is parameterized reducible to parameterized language B if there exist computable functions f, g , and h such that*

1. $f(x, k)$ is computable in time $h(k) \cdot \text{poly}(|x|)$, and
2. $(x, k) \in A \Leftrightarrow (f(x, k), g(k)) \in B$.

This very brief outline of the foundation of parameterized complexity does no justice to this extensive and elegant subfield of computational complexity. However, it is sufficient for the results presented here. The interested reader can find a detailed exposition of the theory of parameterized complexity in the seminal work by [Downey & Fellows \[2012, 2013\]](#).

Usefully, DS is one of the most important problems of parameterized hardness.

Theorem 4.5.6 ([Downey & Fellows \[2013\]](#)) **DOMINATING SET** with parameter k (size of the dominating set) is **W**[2]-complete.

4.6 DKM is Hard in Fixed Dimensions

Once again, consider the reduction presented in the proof of [Theorem 4.5.3](#). We can easily recognize that this is also a parameterized reduction. g is the identity (k remains unchanged) and computing the embedding of the graph is independent of k . Furthermore, observe that for every graph $G = (V, E)$ we have $\delta(G) \leq |V|$.

Corollary 4.5.7 DISCRETE FUZZY RADIUS k -MEANS with $w = \mathbb{1}$, $d \in \mathcal{O}(n^2)$, and parameter k (number of representatives) is $\mathbf{W}[2]$ -complete.

4.6 DKM is Hard in Fixed Dimensions

For the problems discussed in this thesis we obtain the strongest result, with respect to classical complexity, for DKM. We show that DKM is \mathbf{NP} -hard even for unweighted instances in the plane (i.e. $w = \mathbb{1}$ and $d = 2$). The proof is just a slight adaptation of a proof [Papadimitriou \[1981\]](#) used to establish \mathbf{NP} -hardness of what he called k -MEDIAN, which in terms of our notation is just DKM using the standard Euclidean distance. In his proof he assumed that the input points lie on an integral lattice and he assumed the distance between two points to be rounded down to the nearest integer. This is necessary to avoid comparing sums of radicals, which, to the best of our knowledge, is not known to be solvable in deterministic polynomial time. A detailed discussion of this long-standing open problem can for example be found in [[Blömer, 1993](#)]. However, radicals do not arise in DKM where we use the squared Euclidean distance. In the following, we establish \mathbf{NP} -hardness of DKM in the plane, essentially mirroring [Papadimitriou's](#) exposition. We present an explicit embedding of a point set into Euclidean space which exhibits basically the same structure as the point set constructed in the original proof, and outline the correctness arguments.

Theorem 4.6.1 DISCRETE k -MEANS with $w = \mathbb{1}$ and $d = 2$ is \mathbf{NP} -complete.

Proof. We show $\mathbf{X3C} \leq_p \text{DKM}$ where the reduction function constructs a point set in \mathbb{R}^2 . Let (U, S) be an instance of $\mathbf{X3C}$ and call $m := |U| = 3 \cdot n$ for some $n \in \mathbb{N}$. If no such n exists, then (U, S) is a no-instance, and we output some no-instance of DKM. We choose some $\epsilon > 0$ such that $\delta := 12/5 \cdot \epsilon - \epsilon^2 \leq k^{-4}$ and $X := \emptyset$.

We add n^2 copies of each of k rows R_1, \dots, R_k to X . Each row R_l consists of $2 \cdot m^2 + 6 \cdot m + 2$ points. First, a grid of points $6 \cdot m + 2$ points $p_{i,j,l}$, with $i \in [3 \cdot m + 1]$ and $j \in [2]$. We let $p_{i,j,l} \in \mathbb{R}^2$ and denote

$p_{i,j,l} = (p_{i,j,l}^x, p_{i,j,l}^y)$. For an arbitrarily chosen $p_{1,2,l}$, we let

$$p_{i,1,l} = \left(p_{i,2,l}^x, p_{i,2,l}^y + \frac{1}{5} \right) \text{ and}$$

$$p_{i+1,j,l} = \left(p_{i,j,l}^x + 1, p_{i,j,l}^y \right).$$

Second, m^2 copies of each of the two points

$$b_{1,l} = \left(p_{1,2,l}^x - \frac{6}{5}, p_{1,2,l}^y + \frac{1}{10} \right) \text{ and}$$

$$b_{2,l} = \left(p_{3 \cdot m + 1, 2, l}^x + \frac{6}{5} - \epsilon, p_{3 \cdot m + 1, 2, l}^y + \frac{1}{10} \right)$$

so that each row consists of points in $6 \cdot m + 4$ different locations. Figure 4.6.2 sketches one of the rows.




Figure 4.6.2: Sketch of a row R_l .

These k rows are arranged in parallel to each other. We choose an arbitrary location for the point $p_{1,2,1}$. Then, for each row $l \in [k]_2$, we choose $p_{1,2,l} = (p_{1,1,l-1}^x, p_{1,1,l-1}^y + 4)$.

Next, we add n^2 copies of $6 \cdot n \cdot (k - 1)$ points $q_{i,j,l}$, with $i \in [m]$, $j \in [2]$, and $l \in [k - 1]$. We call each these the q -pairs and locate them at

$$q_{i,1,l} = \left(p_{3 \cdot i, 2, l+1}^x - \frac{1}{2}, p_{3 \cdot i, 2, l+1}^y - \frac{9}{6} \right) \text{ and}$$

$$q_{i,2,l} = \left(p_{3 \cdot i, 2, l}^x - \frac{1}{2}, p_{3 \cdot i, 2, l}^y + \frac{9}{6} \right).$$

For each q -pair, we consider four possible locations

$$x_{i,l} = \left(q_{i,1,l}^x, q_{i,1,l}^y + \frac{5}{6} \right)$$

$$y_{i,l} = \left(q_{i,1,l}^x + \frac{1}{2}, q_{i,1,l}^y + \frac{4}{6} \right)$$

$$w_{i,l} = \left(q_{i,2,l}^x + \frac{1}{2}, q_{i,2,l}^y - \frac{4}{6} \right)$$

4.6 DKM is Hard in Fixed Dimensions

$$z_{i,l} = \left(q_{i,2,l}^x, q_{i,2,l}^y - \frac{5}{6} \right).$$

We add a single point of each of these under certain conditions. For each q -pair $(q_{i,1,l}, q_{i,2,l})$, we have

- $x_{i,l} \in X \Leftrightarrow u_i \notin S_{l+1}$,
- $y_{i,l} \in X \Leftrightarrow u_i \in S_{l+1}$,
- $w_{i,l} \in X \Leftrightarrow u_i \in S_l$, and
- $z_{i,l} \in X \Leftrightarrow u_i \notin S_l$.

Figure 4.6.3 sketches such a configuration of points.




Figure 4.6.3: Sketch of a q -pair configuration.

Our DKM instance is then given as $(X, k \cdot (m + 2) + m \cdot (k - 1))$ and we ask for the cost bound $L = n^2 \cdot (k \cdot (m \cdot 4.12 + 2.9) - n \cdot 2 \cdot \delta + m \cdot (k - 1)) + n \cdot (25/6) \cdot (k - 1)$.

We only outline the correctness of this reduction as the arguments are the same as in the original proof for the Euclidean distance [Papadimitriou, 1981]. The core idea is that we can precisely characterize the cheapest and the second cheapest solutions of a row R_l using $m + 2$ representatives. The cost bound L of the instance X is fulfilled if and

only if $m + 2$ representatives are assigned to each row and one to each q -pair. Furthermore, we need to have exactly n rows using a cheapest configuration of representatives and the rest a second cheapest configuration. Each row R_l in the cheapest configuration then corresponds to the set S_l being part of an exact cover. ■

Complexity of Single-Swap Local Search

Contribution Summary We present local search hardness of the single-swap neighborhood for METRIC UNCAPACITATED FACILITY LOCATION, DISCRETE k -MEANS, and DISCRETE FUZZY k -MEANS. The results and proofs discussed in [Section 5.3](#), [Section 5.4](#), and [Section 5.5](#) were published in [Brauer, 2019], with a preliminary version presented in [Brauer, 2017].

5.1	Polynomial Local Search	53
5.2	Related Work	58
5.3	Completeness of MUFL/SingleSwap	59
5.3.1	Construction	60
5.3.2	Correctness	60
5.3.3	Tightness	64
5.4	Completeness of DKM/SingleSwap	65
5.4.1	Construction	66
5.4.2	Correctness and Tightness	67
5.4.3	Embedding	72
5.5	Completeness of DFKM/SingleSwap	75
5.5.1	Construction	76
5.5.2	Correctness, Tightness, and Embedding	77

A popular approach to dealing with hard optimization problems is a metaheuristic called *local search*. The underlying principle is to define a small *neighborhood* of each solution. A neighborhood usually contains solutions that are in some sense similar to each other – hence the name *local search*. Having defined such a neighborhood, we immediately obtain a simple approximation algorithm for our problem. Start by selecting any feasible solution and repeatedly replace the current solution by a better solution from its neighborhood. This fairly general framework exhibits several hyperparameters, each having

more of less significant impact on the quality of the local search. The most prominent is the definition of the neighborhood, which has three major impacts on the algorithm. First, we need to be able to check if there is a better solution in the current neighborhood in polynomial time. This is often achieved by defining a small enough neighborhood so that we can simply check all solutions in the neighborhood. Second, aside from the polynomial runtime of a single improvement, we also need to bound the number of improvement steps a local search takes before terminating. Finally, the neighborhood defines the quality of the final solution. A local search terminates when it has found a solution with no better solution in its neighborhood. This in itself is no guarantee on the overall quality of the solution. Showing some global guarantee of a locally optimal solution is the hard part of analyzing local search neighborhoods and much harder to predict than just the runtime aspect of the neighborhood.

In this chapter, we examine the *single-swap* neighborhood of MUFL, DKM, and DFKM. While similar to each other, the neighborhood for MUFL slightly differs from the other two. The neighborhood of some MUFL solution O consists of all solutions opening at most one facility which is closed in O , and having at most one facility closed which is open in O . Formally, we set

$$N_{ocs}(O) = \{O' \subseteq F \mid |O \setminus O'| \leq 1 \wedge |O' \setminus O| \leq 1\} .$$

We say that the MUFL neighborhood either *opens* a facility, *closes* a facility, or *swaps* an open facility. Recall that solutions for DKM and DFKM consists of at most k representatives and that using less than k is always worse than using exactly k . Hence, for these problems, we disallow the open and close operations. The neighborhood then consists of the solutions swapping exactly one representative. Formally, we set

$$N_s(O) = \{O' \in X \mid |O \setminus O'| = 1 \wedge |O' \setminus O| = 1\} .$$

For MUFL and DKM it is known that a single-swap local search outputs a constant factor approximation of the global optimum [Arya et al., 2004; Kanungo et al., 2004]. However, no approximation bound has so far been shown for DFKM. Nevertheless, we expect such an analysis to be possible.

Conjecture 5.1 *A single-swap local search finds a constant factor approximation for DISCRETE FUZZY k -MEANS.*

5.1 Polynomial Local Search

The computational theory of local search is captured in the class **PLS**. As one expects from a meaningful complexity class, **PLS** has its own type of reductions, for which it has complete problems. In the following, we present an introduction to the theory of **PLS** and some of the relation of local search complexity to classical computational complexity. The definitions and results presented in this section are due to Johnson et al. [1988]; Papadimitriou et al. [1990]; Schäffer & Yannakakis [1991].

In contrast to classical computational complexity, where we classify decision problems, **PLS** classifies optimization problems.

Definition 5.1.1 (Combinatorial Optimization Problem) *A combinatorial optimization problem $\Pi = (g, D, F, \text{cost})$ consists of*

- *an optimization goal g (min or max),*
- *a set of instances D ,*
- *for each instance $I \in D$, a finite set of feasible solutions $F(I)$, and*
- *an objective function cost mapping an instance I and a feasible solution $s \in F(I)$ to a numerical (usually positive real) value.*

Given $I \in D$, the goal of the problem Π is to find $s \in F(I)$ with

$$\text{cost}(I, s) = g \{ \text{cost}(I, s') \mid s' \in F(I) \} .$$

If we augment such a combinatorial optimization problem by a neighborhood and show the existence of several polynomial-time algorithms, then we obtain the class **PLS**.

Definition 5.1.2 (Polynomial Local Search (**PLS**)) *Let $\Pi = (g, D, F, \text{cost})$ be a combinatorial optimization problem and*

$$\mathcal{N}_\Pi = \{ N(I, s) \subseteq F(I) \mid I \in D, s \in F(I) \}$$

be a collection of subsets of feasible solutions of Π . We call $N(I, s)$ the neighborhood of the solution s with respect to \mathcal{N}_Π and I .

We call a pair of solutions $(s, s') \in F(I) \times F(I)$ a local improvement of s if $s' \in N(I, s)$ and $\text{cost}(I, s')$ is strictly better than $\text{cost}(I, s)$ where better depends on the optimization goal g . A solution s is called locally optimal if there is no local improvement of s . We further call a solution s' reachable from solution s if there exists a sequence of solutions

$$(s = s_1, \dots, s_n = s') \in F(I)^n$$

such that, for all $i \in [n - 1]$, the pair (s_i, s_{i+1}) is a local improvement of s_i .

The class **PLS** consists of all problems Π/\mathcal{N}_Π for which there exist three polynomial-time algorithms A , B , and C such that

- A , given $I \in D$, outputs some solution $s \in F(I)$,
- B , given $I \in D$ and $s \in F(I)$, outputs $\text{cost}(I, s)$, and
- C , given $I \in D$ and $s \in F(I)$, outputs s' such that (s, s') is a local improvement of s . If no local improvement of s exists, then C reports that s is locally optimal.

The task of a **PLS** problem Π/\mathcal{N}_Π is, given an instance $I \in D$, to find a locally optimal solution. We observe that such a **PLS** problem always lies in **TFNP**, the class of all function problems of the form

Given some x and a polynomial-time decidable relation \sim , where there exists a y with $x \sim y$, output z with $x \sim z$.

Due to the definition using total relations, **TFNP** is a natural functional analog to $\mathbf{NP} \cap \mathbf{coNP}$. Hence, no problem in **TFNP**, and thus no problem in **PLS**, can be **NP**-hard unless $\mathbf{NP} = \mathbf{coNP}$ [Megiddo & Papadimitriou, 1991].

The motivation for the definition of **PLS**, as well as its sibling classes **PPA** and **PPP**, was a little atypical. The goal was to group problems by the technique used to prove their membership of **TFNP**. For problems in **PLS** the proof relies on the fact that every finite directed acyclic graph has a sink. We can find this argument in the definition of the so-called *standard algorithm*. To solve a **PLS** problem one can always use the algorithms A and C , which are provided by definition. Run A to obtain some solution s and repeatedly update s using C until C reports that s is locally optimal. Since the set of feasible solutions is finite, this algorithm always terminates with a locally optimal solution. This standard algorithm traverses a path on the transition graph.

Definition 5.1.3 (Transition Graph) *Let Π/\mathcal{N}_Π be a **PLS** problem and $I \in D$. The transition graph $\text{TG}_{\Pi/\mathcal{N}_\Pi}(I) = (F(I), E(I))$ is a directed graph such that $(s, s') \in E(I)$ if and only if (s, s') is a local improvement of s with respect to I .*

The transition graph captures the local structure of the **PLS** problem. Each vertex is a feasible solution, and the edges represent improving steps the algorithm C might take. The transition graph is acyclic, and thus, we find the reason why **PLS** problems lie in **TFNP**.

5.1 Polynomial Local Search

While we suspect that **PLS** problems are not **NP**-hard, we have no universal recipe to finding locally optimal solutions except for the standard algorithm. Thus, the second driving task in the definition of **PLS** was the analysis of the runtime of the standard algorithm. This led to the important notion of the *standard algorithm problem*:

Given a **PLS** problem Π/\mathcal{N}_Π , an instance $I \in D$, and a solution $s \in F(I)$, find a locally optimal solution $s^* \in F(I)$ which is reachable from s .

Notice that the standard algorithm problem is independent of the concrete algorithm C . It covers all possible tiebreaks and improvement choices an algorithm C might take. As we discuss later, solving the standard algorithm problem might be significantly harder than solving the corresponding **PLS** problem.

Definition 5.1.4 (PLS-reduction) *Let Π/\mathcal{N}_Π and $\Lambda/\mathcal{N}_\Lambda$ be **PLS** problems. We say Π/\mathcal{N}_Π is **PLS**-reducible to $\Lambda/\mathcal{N}_\Lambda$, write $\Pi/\mathcal{N}_\Pi \leq_{\text{PLS}} \Lambda/\mathcal{N}_\Lambda$, if there are polynomial-time computable functions Φ and Ψ such that*

- Φ maps an instance $I \in D_\Pi$ to an instance $\Phi(I) \in D_\Lambda$,
- Ψ maps an instance $I \in D_\Pi$ and a solution $s \in F_\Lambda(\Phi(I))$ to a solution $\Psi(I, s) \in F_\Pi(I)$, and
- if $s \in F_\Lambda(\Phi(I))$ is locally optimal, then $\Psi(I, s)$ is locally optimal.

These reductions behave exactly as one expects from a sound reduction. **PLS**-reductions are reflexive, transitive, and there are **PLS**-complete problems. However, being **PLS**-complete does not tell us anything about the complexity of the respective standard algorithm problem. To obtain such a classification, we need to consider a restriction of **PLS**-reductions to so-called *tight **PLS**-reductions*.

Definition 5.1.5 (Tight PLS-reduction) *Let Π/\mathcal{N}_Π and $\Lambda/\mathcal{N}_\Lambda$ be **PLS** problems and let (Φ, Ψ) be a **PLS**-reduction from Π/\mathcal{N}_Π to $\Lambda/\mathcal{N}_\Lambda$. We call (Φ, Ψ) tight if for every $I \in D_\Pi$ there is a set of reasonable solutions $\mathcal{R} \subseteq F_\Lambda(\Phi(I))$ with the following properties*

- \mathcal{R} contains all locally optimal solutions of $\Phi(I)$,
- there is a polynomial-time computable function $h : F_\Pi(I) \rightarrow \mathcal{R}$ such that, for all $s \in F_\Pi(I)$, we have $s = \Psi(I, h(s))$, and
- if $s \rightsquigarrow s'$ be a path in $\text{TG}_{\Lambda/\mathcal{N}_\Lambda}(\Phi(I))$ with $s, s' \in \mathcal{R}$ such that all internal path vertices are solutions in $F_\Lambda(\Phi(I)) \setminus \mathcal{R}$, then either $\Psi(I, s) = \Psi(I, s')$ or $(\Psi(I, s), \Psi(I, s'))$ is an edge in $\text{TG}_{\Pi/\mathcal{N}_\Pi}(I)$.

We call a **PLS** problem tightly **PLS**-complete if all **PLS** problems are reducible to it via a tight reduction.

Once again observe that tight **PLS**-reductions are reflexive and transitive. While the arguments are a little technical, this can be obtained with a straightforward check of the required properties and will not be presented here.

The third constraint to a tight **PLS**-reduction might seem a bit daunting and technical. However, it is a technical formulation of an intuitive idea: we require that tight **PLS**-reductions preserve paths of local improvement. The problem that the tight reduction maps to might introduce additional intermediate solutions to improving paths in the original problem. This makes in some sense sure that the standard algorithm of the reduced-to problem cannot take new shortcuts that the original standard algorithm could not have taken. This means that we cannot speed up a standard algorithm by first reducing our problem to a different one. Tight **PLS**-reductions are of special interest since they open a connection to classical complexity theory and provide a lower bound for the worst case runtime of the standard algorithm.

Theorem 5.1.6 *Let Π/\mathcal{N}_Π be a tightly **PLS**-complete problem.*

- *The standard algorithm problem of Π/\mathcal{N}_Π is **PSPACE**-complete.*
- *There is no polynomial p such that, for each instance $I \in D$ and each solution $s \in F(I)$, there is a locally optimal solution of I which is reachable from s with less than $p(|I|)$ local improvements.*

This not only means that the standard algorithm problem of a tightly **PLS**-complete problem is hard, but also, that there always is a solution which is superpolynomially many steps away from every local optimum.

In the following, we show that **MUFL/SingleSwap**, **DKM/SingleSwap**, and **DFKM/SingleSwap** are tightly **PLS**-complete problems.

Local Search for Boolean Satisfiability For most complexity classes, there is a variant of Boolean satisfiability which is a hard problem within the class, and **PLS** poses no exception to this. We formally present the **MAX 2-SAT (M2SAT)** and **POSITIVE MAX 2-NOTALLEQUAL-SAT (PM2NAESAT)** problems, as well as the Flip neighborhood. Their corresponding local search problems **M2SAT/Flip** and **PM2NAESAT/Flip** play an important role in the the analysis of **PLS**.

5.1 Polynomial Local Search

Let B be a finite set of clauses (i.e. Boolean disjunctions) over the set of variables X . As before, we denote occurrences of variables in a clause, either positive or negated, as literals. For a literal x we denote by $B(x)$ the subset of clauses in B in which x appears. For a weight function $w : B \rightarrow \mathbb{N}$ on the clauses we denote $w_{max}(B) = \max_{b \in B} \{w(b)\}$ and $w_{min}(B) = \min_{b \in B} \{w(b)\}$.

Let T be a truth assignment of the variables X . We denote by $B_t(T)$ the set of clauses in B which are satisfied by T in the classical sense. That is, a clause is satisfied if at least one of the literals evaluates to true. Analogously, we denote by $B_t^{NAE}(T)$ the set of clauses which are NAE-satisfied by T . As a reminder: a clause is NAE-satisfied if there is at least one literal evaluating to true and at least one to false. Furthermore, we denote by $B_f(T) = B \setminus B_t(T)$, and $B_f^{NAE}(T) = B \setminus B_t^{NAE}(T)$ respectively, the clauses which are not (NAE-)satisfied. We formally introduce two optimization problems.

- Given a set of clauses B where each clause consists of exactly two literals and a weight function $w : B \rightarrow \mathbb{N}$. The M2SAT problem asks for a truth assignment maximizing

$$w(B_t(T)) .$$

- Given a set of clauses B where each clause consists of exactly two positive literals and a weight function $w : B \rightarrow \mathbb{N}$. The PM2NAESAT problem asks for a truth assignment maximizing

$$w(B_t^{NAE}(T)) .$$

Notice that PM2NAESAT is equivalent to the well-known MAX CUT (MC) problem. The MC problem is, given an undirected graph $G = (V, E)$ and a weight on the edges $w : E \rightarrow \mathbb{N}$, find a partition of the vertices into two sets maximizing the sum of the weights of the edges going over the cut (i.e. from one set of vertices to the other). We map each variable to a vertex and draw edges between two vertices if and only if they occur together in a clause. By interpreting a partition of the vertices as assigning a truth value to the corresponding variable, one can easily see that the value of the cut and the weight of the NAE-satisfied clauses coincide. Furthermore, observe that it is not important which of the two subsets of vertices is assigned true since the NAE weight of a truth assignment is the same as the NAE weight of its conjugate assignment.

An intuitive local search approach to approximating MAX SAT-type problems is the Flip neighborhood. Given a truth assignment T , its

Flip neighborhood is the set of all truth assignments differing from T in the value of exactly one variable – the *flipped* variable. Notice that M2SAT/Flip and PM2NAESAT/Flip are both **PLS** problems.

- The algorithm A can simply output any fixed truth assignment.
- The algorithm B computes the cost function in polynomial time simply by evaluating each clause and summing up the weights of (NAE-)satisfied clauses.
- The algorithm C uses algorithm B to compute the cost of every assignment in the neighborhood and outputs either the best one or reports that the current assignment is locally optimal. This is also a polynomial-time algorithm since the Flip neighborhood contains only $|X|$ assignments.

Theorem 5.1.7 *MAX 2-SAT/Flip and POSITIVE MAX 2-NOTALLEQUAL-SAT/Flip are tightly **PLS**-complete problems.*

A peculiarity of **PLS**-complete problems is that the restriction to unweighted instances is usually not tightly **PLS**-complete. That is because a local improvement step always strictly reduces the cost of the current solution. For most **PLS** problems this means that there is at least a constant improvement in every step, independent of the actual instance. An unweighted instance (or one with a polynomial bound on the weights, for that matter) has a polynomial upper and lower (typically 0) bound on the total cost of a solution. Consequently, the local search will terminate after a polynomial number of improving steps. We know that a tightly **PLS**-complete problem has solutions which are superpolynomially many steps away from every local optimum. From this we conclude that even the restriction of these types of problem to polynomially bounded weight functions cannot be tightly **PLS**-complete.

5.2 Related Work

Swap-based local search heuristics for facility location and hard clustering have already been analyzed in terms of their runtime and approximation guarantee. A single-swap local search yields a 3-approximation for MUFL [Arya et al., 2004]. Furthermore, UFL/SingleSwap for general distance functions is known to be **PLS**-complete [Kochetov & Ivanenko, 2005].

5.3 Completeness of MUFL/SingleSwap

For DKM, single-swap also yields a constant-factor approximation [Kanungo et al., 2004]. The authors claim the approximation factor to be 25, however, there is a minor error in their calculations such that the actual factor of their analysis is 81. They also do not show an upper bound on the number of iterations. They instead show that a relaxation of the single-swap heuristic where we impose a lower bound on the cost improvement of a swap yields an algorithm with polynomial runtime but a slightly worse approximation ratio. As a remark: this result is also applicable to a more general formulation of the problem where an instance consists of a set of points and a set of potential representatives. Our problem DKM is a special case where the latter is the data set itself. Furthermore, recall that due to Lemma 4.2.1, this also yields a constant factor approximation for KM.

A straightforward generalization of single-swap is the multi-swap heuristic. There we are allowed, for some constant l , to swap up to l representatives in a single step. This heuristic yields a PTAS for variants of DKM in Euclidean space with fixed dimension [Cohen-Addad et al., 2016] and in metric spaces with bounded doubling dimension [Friggstad et al., 2016].

5.3 Completeness of MUFL/SingleSwap

The first tight **PLS**-reduction we present shows the completeness of the metric case of facility location together with the single-swap neighborhood. In particular, we show that this still holds for unweighted instances, i.e. for the weight function $w = \mathbb{1}$ on the clients. This appears to be in conflict with our earlier discussion on tightly **PLS**-complete problems always being weighted. Recall that we argued that a local improvement for problems discussed so far always improves the cost by a constant. For example, a local improvement of M2SAT always fulfils an additional clause hence the objective increases by at least 1. This is not true for MUFL where improvements depend on the metric. Specifically, our reduction encodes the weights of clauses into the distance between points in the metric space.

Theorem 5.3.1 METRIC UNCAPACITATED FACILITY LOCATION *with $w = \mathbb{1}$ and with single-swap neighborhood is tightly **PLS**-complete.*

The proof of Theorem 5.3.1 consists of three parts. First, we present the construction of the reduction functions Φ and Ψ , second we argue the correctness of the reduction, and finally show that the reduction is tight.

5.3.1 Construction

Let $(B, w : B \rightarrow \mathbb{N})$ be an instance of M2SAT over the set of variables X and denote $n := |X|$, $m := |B|$, and

$$W := m \cdot w_{\max}(B) .$$

Without loss of generality, we assume $m \geq 2$. We construct the function Φ mapping (B, w) to an instance (C, F, f, d) of MUFL as follows: Similar to the proof of [Theorem 4.4.1](#) we add $2 \cdot n$ points to F – two for each variable in X . For each variable, one of these corresponds to the positive literal, the other to the negation. For the clients we initially set $C = F$ and then add m additional points corresponding to the clauses in B . Simply speaking, there is a client corresponding to each clause and each literal, and facilities can be opened at the location of each client corresponding to a literal. We set the opening cost function to

$$\forall o \in F : f(o) := 2 .$$

We define the distance function $d : C \cup F \times C \cup F \rightarrow \mathbb{R}_{\geq 0}$ as

$$d(p, q) = d(q, p) := \begin{cases} 0 & \text{if } p = q \\ 1 & \text{if } p \text{ and } q \text{ are a literal and its negation} \\ 1 + \frac{w(q)}{W} & \text{if } p \text{ is a literal appearing in clause } q \\ 1 + \frac{c \cdot w(q)}{W} & \text{if } \bar{p} \text{ is a literal appearing in clause } q \\ 2 & \text{else} \end{cases}$$

where $1 < c < 2$ is some arbitrarily chosen constant.

It is easy to see that d is a metric. We immediately obtain non-negativity, symmetry, and identity of indiscernibles by definition. To recognize that d fulfils the triangle inequality, observe that the sum of any two non-trivial distances is always at least 2 while the distance between two points is itself at most 2.

Next, we construct the function Ψ mapping a solution of $\Phi(B, w)$ back to a solution of (B, w) . Let $O \subseteq F$ be a set of opened facilities. For each variable $x \in X$ we set x to true if $x \in O$ and to false otherwise. That is, a variable is true if there is an open facility at the point corresponding to its positive literal. There is no *only if* part to this statement. A variable is also set to true if both corresponding literals are opened.

5.3.2 Correctness

In the following, fix any M2SAT instance (B, w) , denote the image of the reduction as $(C, F, 2, d) := \Phi(B, w)$, and for each set of opened facilities

5.3 Completeness of MUFL/SingleSwap

$O \subseteq F$ denote by $T_O := \Psi(B, w, O)$ the truth assignment obtained as the image of Ψ . One can easily verify that Φ and Ψ are polynomial-time computable functions. It remains to show that T_O is locally optimal for (B, w) with respect to the flip neighborhood if O is locally optimal for $\Phi(B, w)$ with respect to the single-swap neighborhood. Since we fix the metric and the opening cost, for the remainder of this section, we introduce the following shorthand notation

$$\text{mfl}(C, O) := \text{mfl}(C, 2, d, O) .$$

Observe that $\Phi(B, w)$ has significantly more feasible solutions than (B, w) . Thus, Ψ cannot be injective, which we already discussed as the missing *only if* part in the definition of Ψ . In general this makes it difficult to relate the cost of a solution to $\Phi(B, w)$ to the cost of its image under Ψ . We tackle this problem by using a generalization of the set of reasonable solutions we introduced in [Definition 4.4.3](#).

Definition 5.3.2 (Reasonable Solutions) *Let $O \subseteq F$. We call O reasonable if $|O| = n$ and for each variable x*

$$x \in O \vee \bar{x} \in O .$$

We deliberately call this set of solutions *reasonable*, a term we previously introduced as the name of the set \mathcal{R} required to establish tightness of a **PLS**-reduction. This is because these solutions constitute the set \mathcal{R} in the tightness proof of this reduction as it is presented in [Section 5.3.3](#). Reasonable solutions of $\Phi(B, w)$ have several useful properties, which we establish in the following.

Recall the earlier discussion about reasonable solutions, which applies in this case, as well. A reasonable solutions contains exactly n open facilities and for each of the n variables occurring in B , either the facility at its positive or at its negative literal is opened. Hence, the restriction of Ψ to reasonable solutions is a bijection. We obtain a one-to-one mapping between truth assignments and reasonable solutions where a variable x is true if and only if $x \in O$ and x is false if and only if $\bar{x} \in O$. This allows us to express the cost of a reasonable solution O in terms of the cost of the corresponding truth assignment T_O .

Lemma 5.3.3 *If $O \subseteq F$ is reasonable, then*

$$\text{mfl}(C, O) = 3 \cdot n + m + \frac{c}{W} \cdot w(B) - \frac{c-1}{W} \cdot w(B_t(T_O)) .$$

Proof. O incurs opening cost of $2 \cdot n$ since it is reasonable and thus opens exactly n facilities. Each client corresponding to a literal either

has a facility opened at its location and has cost 0, or there is an opened facility at the location of its negated literal and it has cost 1. Overall the service cost of clients corresponding to literals is n . By definition of the point set and Ψ , we obtain that a client corresponding to a clause $b \in B_t(T_O)$ has at least one open facility at distance $1 + w(b)/W$. Analogously, a client corresponding to a clause $b \in B_f(T_O)$ has two facilities at distance $1 + c \cdot w(b)/W$. For both types of clauses, all other facilities are at distance 2. We sum this up to obtain

$$\begin{aligned}
 \text{mfl}(C, O) &= 2 \cdot n + n + \sum_{b \in B_t(T_O)} \left(1 + \frac{w(b)}{W}\right) + \sum_{b \in B_f(T_O)} \left(1 + \frac{c \cdot w(b)}{W}\right) \\
 &= 3 \cdot n + m + \sum_{b \in B_t(T_O)} \frac{w(b)}{W} + \sum_{b \in B_f(T_O)} \frac{c \cdot w(b)}{W} \\
 &= 3 \cdot n + m + \frac{c}{W} \cdot w(B) - \frac{c-1}{W} \cdot \underbrace{\sum_{b \in B_t(T_O)} w(b)}_{=w(B_t(T_O))} .
 \end{aligned}$$

■

Using this, we obtain that a cost improvement in the M2SAT instance (B, w) corresponds to a cost improvement of reasonable solutions in the MUFL instance $\Phi(B, w)$ via Ψ .

Corollary 5.3.4 *If $O, O' \subseteq F$ are reasonable, then*

$$\text{mfl}(C, O) > \text{mfl}(C, O') \text{ if and only if } w(B_t(T_O)) < w(B_t(T_{O'})) .$$

Proof. Using [Lemma 5.3.3](#) we obtain that the cost of a reasonable solution is a constant only depending on the instance minus a fraction of the M2SAT cost of T_O . Since $c > 1$, we obtain the claim. ■

The previous corollary illustrates that the M2SAT instance (B, w) has the same notion of a *better* solution as the reasonable solutions of its image under Φ . Hence, the bijection Ψ uniquely maps local improvements of (B, w) to local improvements between two reasonable solutions of $\Phi(B, w)$, and vice versa. However, there is no relation of this kind if any of the MUFL solutions involved is not reasonable. We show that this is irrelevant since all locally optimal solutions of instances in the image of Φ are, in fact, reasonable.

Lemma 5.3.5 *If $O \subseteq F$ is locally optimal with respect to the single-swap neighborhood, then O is reasonable.*

5.3 Completeness of MUFL/SingleSwap

Proof. We differentiate two cases in contraposition of the claim. First, we argue that no locally optimal solution can open a facility corresponding to a literal and the facility corresponding to the respective negated literal. Second, we show that every locally optimal solution contains a facility corresponding to either the positive or the negative literal of each of the variables.

Case 1 Assume there is a variable $x \in X$ such that $x, \bar{x} \in O$. By definition we have $|B(x)| \leq m$ and $0 < c - 1 < 1$, and obtain

$$1 \geq \frac{|B(x)|}{m} \geq \sum_{b \in B(x)} \frac{w(b)}{w_{\max}(B) \cdot m} > \sum_{b \in B(x)} \frac{(c-1) \cdot w(b)}{W}.$$

By definition of the metric d , no client in $\tilde{C} := C \setminus (B(x) \cup \{x\})$ is closer to x than it is to \bar{x} . Thus, the service cost of these clients does not change when we close the facility at x .

$$\begin{aligned} \text{mfl}(C, O) &= \underbrace{\sum_{c \in \tilde{C}} \min_{o \in O} \{d(c, o)\}}_{\text{service cost of } \tilde{C}} + \underbrace{\sum_{b \in B(x)} \min_{o \in O} \{d(b, o)\}}_{\text{service cost of } B(x)} + \underbrace{2 \cdot |O|}_{\text{opening cost}} \\ &= \sum_{c \in \tilde{C}} \min_{o \in O \setminus \{x\}} \{d(c, o)\} + \sum_{b \in B(x)} \left(1 + \frac{w(b)}{W}\right) + 2 + 2 \cdot (|O| - 1) \\ &> \sum_{c \in \tilde{C}} \min_{o \in O \setminus \{x\}} \{d(c, o)\} + \sum_{b \in B(x)} \left(1 + \frac{w(b)}{W}\right) + \sum_{b \in B(x)} \frac{(c-1) \cdot w(b)}{W} \\ &\quad + 1 + 2 \cdot (|O| - 1) \qquad \qquad \qquad \text{(Definition of } c \text{ and } W) \\ &= \sum_{c \in \tilde{C}} \min_{o \in O \setminus \{x\}} \{d(c, o)\} + \sum_{b \in B(x)} \underbrace{\left(1 + \frac{c \cdot w(b)}{W}\right)}_{d(b, \bar{x})} + \underbrace{1}_{d(x, \bar{x})} + 2 \cdot (|O| - 1) \\ &\geq \text{mfl}(C, O \setminus \{x\}). \end{aligned}$$

Case 2 Assume there is a variable $x \in X$ such that $x, \bar{x} \notin O$. All opened facilities are at distance 2 to the clients located at x and \bar{x} . Opening a facility at x is sufficient to reduce the overall cost.

$$\begin{aligned}
 \text{mfl}(C, O) &= \sum_{c \in C \setminus \{x, \bar{x}\}} \min_{o \in O} \{d(c, o)\} + \underbrace{\sum_{c \in \{x, \bar{x}\}} \min_{o \in O} \{d(c, o)\}}_{=4} + 2 \cdot |O| \\
 &> \sum_{c \in C \setminus \{x, \bar{x}\}} \min_{o \in O} \{pd(c, o)\} + \underbrace{1}_{d(x, \bar{x})} + 2 \cdot (|O| + 1) \\
 &\geq \text{mfl}(C, O \cup \{x\}) .
 \end{aligned}$$

■

By combining the facts that local improvements of (B, w) correspond to improvements between reasonable solutions and that all locally optimal solutions are reasonable, we can conclude correctness of our reduction.

Corollary 5.3.6 *If $O \subseteq F$ is locally optimal for $\Phi(B, w)$ with respect to the single-swap neighborhood, then T_O is locally optimal for (B, w) with respect to the Flip neighborhood.*

Proof. Assume to the contrary that T_O is not locally optimal. If O is not reasonable, then it is not locally optimal by [Lemma 5.3.5](#). Therefore assume that O is, in fact, reasonable. Since T_O is not locally optimal, we know that there exists a variable x such that

$$w(B_t(T_O^{\bar{x}})) > w(B_t(T_O))$$

where $T_O^{\bar{x}}$ denotes T_O with a flipped assignment of the variable x . Since $O^{\bar{x}} := (O \setminus \{x\}) \cup \{\bar{x}\}$ is reasonable, $\Psi(B, w, O^{\bar{x}}) = T_O^{\bar{x}}$, and by [Corollary 5.3.4](#) we obtain that

$$\text{mfl}(C, O^{\bar{x}}) < \text{mfl}(C, O) .$$

We conclude that O is not locally optimal because $O^{\bar{x}}$ is in the single-swap neighborhood of O . ■

5.3.3 Tightness

We validate that (Φ, Ψ) fulfils the properties required from a tight **PLS**-reduction. Let \mathcal{R} be the set of all reasonable solutions for the MUFL instance $\Phi(B, w)$. In our analysis we focus on $\Phi|_{\mathcal{R}}$, the restriction of Ψ to the set of reasonable solutions. We already discussed that, by

5.4 Completeness of DKM/SingleSwap

Corollary 5.3.4, the single-swap neighborhood behaves, with respect to reasonable solutions, exactly the same way the Flip neighborhood behaves on (B, w) . If we inspect our previous analysis a little closer, then we can see that once single-swap has reached a reasonable solution it will always stay at a reasonable solution.

Lemma 5.3.7 *If $O \in \mathcal{R}$ and $O' \notin \mathcal{R}$, then $(O, O') \notin \text{TG}(\Phi(B, w))$.*

Proof. Assume $O \in \mathcal{R}$ and let $x \in O$ be any (possibly negated) literal opened by the solution. Recall the arguments we have seen in the proof of **Lemma 5.3.5**. The solutions $O \setminus \{x\}$ and $O \cup \{\bar{x}\}$ are more expensive than the current solution. A swap, let $y \in O$ and choose $(O \cup \{\bar{y}\}) \setminus \{x\}$ as the new solution, is even more expensive than either only opening \bar{y} or only closing x . ■

Finally, we conclude this section with the tightness proof of our reduction.

Lemma 5.3.8 *The PLS-reduction (Φ, Ψ) is tight.*

Proof. We need to verify that (Φ, Ψ) and \mathcal{R} fulfil the properties required by **Definition 5.1.5**. \mathcal{R} contains all local optima of $\Phi(B, w)$ by **Lemma 5.3.5**. As the function h we choose $(\Psi|_{\mathcal{R}})^{-1}$. This is well-defined since $\Psi|_{\mathcal{R}}$ is a bijection. Furthermore, we can compute this in polynomial-time by simply adding the respective literals to the solution.

For the third property, let $O \rightsquigarrow O'$ be a path in $\text{TG}(\Phi(B, w))$ with $O, O' \in \mathcal{R}$ and all internal path vertices outside of \mathcal{R} . Applying **Lemma 5.3.7** we obtain that either $O = O'$ or that the path consists of a single edge (O, O') . In the former case we have $\Psi(B, w, O) = \Psi(B, w, O')$.

In the latter case (O, O') is a local improvement of O under the single-swap neighborhood. That is, there exists a variable x such that $O' = O^{\bar{x}}$ and $\text{mfl}(C, O) > \text{mfl}(C, O^{\bar{x}})$. Since $\Psi(B, w, O^{\bar{x}}) = T_{O^{\bar{x}}} = T_{O'}$ and by applying **Corollary 5.3.4**, we have $w(B_t(T_O)) < w(B_t(T_{O'}))$. We conclude by observing that $(T_O, T_{O'}) \in \text{TG}(B, w)$. ■

5.4 Completeness of DKM/SingleSwap

In this section, we present our second tight **PLS**-reduction showing the completeness of DKM together with the single-swap neighborhood. Just like our previous reduction, this still holds for $w = \mathbb{1}$. However, a reduction to DKM poses two difficulties, which did not occur in the MUFL case. In our definition of DKM, each input point is a potential

representative. In contrast, a facility location instance contains a dedicated set of potential locations for facilities. Furthermore, there is a stronger restriction on distances we can choose between points. A DKM instance is a point set in \mathbb{R}^d where the distance between two points is the squared Euclidean distance. Thus, we have to ensure that there exists a set of points in \mathbb{R}^d exhibiting the required pairwise interpoint distances. The MUFL reduction was able to freely assign distances between points, only making sure that the final distance function is a metric. There also is no opening cost for the representatives as there was for the facilities. We extensively used this to dominate any improvement in the service cost by the additional opening cost.

Despite these additional constraints, we maintain the general approach and obtain conceptually the same intermediate results.

Theorem 5.4.1 DISCRETE k -MEANS with $w = \mathbb{1}$ and with single-swap neighborhood is tightly **PLS**-complete.

The construction of the reduction (Φ, Ψ) is similar to the proof of [Theorem 5.3.1](#). We start by pointing out modifications required to adjust it to DKM. The ensuing proof of correctness and tightness is based on an abstract definition of the point set via specified interpoint distances. Subsequently, we argue on the existence of a point set in \mathbb{R}^d exhibiting the required squared Euclidean distances. The section is concluded by a complementing result proving that our reduction cannot be embedded in significantly less dimensions than the embedding we provide.

5.4.1 Construction

As before, let $(B, w : B \rightarrow \mathbb{N})$ be an instance of M2SAT over the set of variables X and denote $n := |X|$ and $m := |B|$. The choice of W is changed to

$$W := 2 \cdot m \cdot w_{\max}(B) .$$

We construct the function Φ mapping (B, w) to an instance (X, k) to DKM. Abstractly define the point set X to contain $2 \cdot n + m$ points – one corresponding to the positive and negative literal of each variable and one corresponding to each clause. This is the same as the set C in the proof of [Theorem 5.3.1](#). Let $1 < c < 2$ be some arbitrarily chosen constant and

$$\epsilon := \frac{1}{4 \cdot n + 2 \cdot m} .$$

5.4 Completeness of DKM/SingleSwap

As the distance between points in X we choose

$$\|p - q\|_2^2 := \begin{cases} 0 & \text{if } p = q \\ 1 & \text{if } p \text{ and } q \text{ are a literal and its negation} \\ 1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w(q)}{W}\right) & \text{if } p \text{ is a literal appearing in clause } q \\ 1 + \epsilon \cdot \left(\frac{3}{2} + \frac{c \cdot w(q)}{W}\right) & \text{if } \bar{p} \text{ is a literal appearing in clause } q \\ 1 + 2 \cdot \epsilon & \text{else.} \end{cases}$$

The number of clusters is set to $k := n$.

Also as before, the function Ψ , given some set $M \subseteq X$, sets each variable x to true if $x \in M$ and to false otherwise. Notice the still missing *only if* part.

5.4.2 Correctness and Tightness

In the following, fix any M2SAT instance (B, w) , denote the image of the reduction as $(X, n) := \Phi(B, w)$, and, for each set of representatives $M \subseteq X$, denote by $T_M := \Psi(B, w, M)$ the truth assignment obtained under the image of Ψ . Again we use the concept of reasonable solutions to tackle the non-injective function Ψ . However, in the DKM case the condition $|M| = n$ is trivially fulfilled for every feasible solution (which means, reasonable solutions in the sense of [Definition 4.4.3](#)). Using this we obtain the same proof structure as the correctness proof of the MUFL reduction. However, we need to adapt to the changed distance function and to the larger set of potential locations for representatives.

Lemma 5.4.2 *If $M \subseteq X$ is reasonable, then*

$$\text{km}(X, M) = n + \left(1 + \frac{3 \cdot \epsilon}{2}\right) \cdot m + \frac{c \cdot \epsilon}{W} \cdot w(B) - \frac{(c-1) \cdot \epsilon}{W} \cdot w(B_t(T_M)) .$$

Proof. We follow the same arguments as the proof of [Lemma 5.3.3](#) where we simply insert the changed distances between points and omit the opening cost. Recall that by definition of Ψ and reasonable solutions, a variable x is assigned true if and only if $x \in M$ and false if and only if $x \notin M$. The cost of a point corresponding to a literals is hence either 0 if the literal is part of M or its negated literal is in M and its cost is 1. Thus, the total cost of these points is n . By definition of the point set, each point corresponding to a clause $b \in B_t(T_M)$ has at least one representative at distance $1 + \epsilon \cdot (3/2 + w(b)/W)$. A point corresponding to a clause $b \in B_f(T_M)$ has two points in M at distance

$1 + \epsilon \cdot (3/2 + c \cdot w(b)/W)$. For both types of clauses all other representatives are at distance $1 + 2 \cdot \epsilon$. Summing this up we obtain

$$\begin{aligned}
 \text{km}(X, M) &= n + \sum_{b \in B_t(T_M)} \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w(q)}{W} \right) \right) \\
 &\quad + \sum_{b \in B_f(T_M)} \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{c \cdot w(q)}{W} \right) \right) \\
 &= n + \left(1 + \frac{3}{2} \cdot \epsilon \right) \cdot m + \sum_{b \in B_t(T_M)} \frac{\epsilon \cdot w(b)}{W} + \sum_{b \in B_f(T_M)} \frac{c \cdot \epsilon \cdot w(b)}{W} \\
 &= n + \left(1 + \frac{3}{2} \cdot \epsilon \right) \cdot m + \frac{c \cdot \epsilon}{W} \cdot w(B) - \frac{(c-1) \cdot \epsilon}{W} \cdot \underbrace{\sum_{b \in B_t(T_M)} w(b)}_{=w(B_t(T_M))} .
 \end{aligned}$$

■

As before, we conclude that cost improvements in the M2SAT instance (B, w) correspond to cost improvements of reasonable solutions in the DKM instance $\Phi(B, w)$.

Corollary 5.4.3 *If $M, M' \subseteq X$ are reasonable, then*

$$\text{km}(X, M) > \text{km}(X, M') \text{ if and only if } w(B_t(T_M)) < w(B_t(T'_M)) .$$

Proof. Apply the proof of [Corollary 5.3.4](#) with [Lemma 5.4.2](#) instead of [Lemma 5.3.3](#). ■

Most of the additional work for our DKM correctness is found in the proof of the following lemma. In this case, there are no facilities that can simply be opened or closed. Previously, we could use this because the opening cost dominated any potential increases or decreases in the service cost. Furthermore, we now have to ensure that locally optimal solutions contain no representatives at the locations corresponding to clauses, which were simply not available as a facility in the MUFL case.

Lemma 5.4.4 *If $M \subseteq X$ is locally optimal with respect to the single-swap neighborhood, then M is reasonable.*

Proof. Each point corresponding to some clause b has exactly two points at distance $1 + \epsilon \cdot (3/2 + w(b)/W)$ and two points at distance $1 + \epsilon \cdot (3/2 + c \cdot w(b)/W)$. The former are the points corresponding to the literals in b , the latter are the points corresponding to the respective negated literals. We call these the points *adjacent* to b . All other points

5.4 Completeness of DKM/SingleSwap

in X have distance $1 + 2 \cdot \epsilon$ to b , and thus, are strictly farther away than the adjacent points. Moreover, for all $b \in B$, we have

$$\frac{3}{2} < \frac{3}{2} + \frac{w(b)}{W} < \frac{3}{2} + \frac{c \cdot w(b)}{W} < 2 .$$

Assume in contraposition that M is not reasonable. By definition, there exists a variable x such that neither of the literals corresponding to x are in M , i.e. $x, \bar{x} \notin M$. Since all feasible solutions contain exactly n representatives, we can differentiate two cases: either there exists a point corresponding to a clause b with $b \in M$ or there exists a variable $y \neq x$ such that both of the corresponding literals are part of the solution $y, \bar{y} \in M$. In the following, we show that, in both cases, M is not locally optimal.

Case 1 Assume there is clause $b \in M$ and that, without loss of generality, $b = (z \vee v)$. If any of the literals in the clause are negative, then exchange the role of positive and negative literals in the following arguments. An important observation is that if we exchange b for some other representative, then only its own cost and the cost of its adjacent points can increase. All other points in X which might have b as their closest representative are at distance $1 + 2 \cdot \epsilon$ anyway and can be reassigned to any other representative for at most the same cost. Formally, we have

$$\text{km}(X \setminus \{b, z, \bar{z}, v, \bar{v}\}, M) = \text{km}(X \setminus \{b, z, \bar{z}, v, \bar{v}\}, M \setminus \{b\}) .$$

Recall that adding an additional mean to M can never increase the cost of the solution. Hence, for any $u \in X$,

$$\text{km}(X \setminus \{b, z, \bar{z}, v, \bar{v}\}, M) \geq \text{km}(X \setminus \{b, z, \bar{z}, v, \bar{v}\}, (M \setminus \{b\}) \cup \{u\}) . \quad (5.4.5)$$

We differentiate three subcases depending on which of b 's adjacent points are already part of the current solution M .

Case 1.1 Assume none of the adjacent points are part of the solution $z, \bar{z}, v, \bar{v} \notin M$. None of the points adjacent to b has a representative at distance 1, and hence, has cost at least $1 + \epsilon \cdot (3/2 + w_{\min}(B)/W)$. We

obtain

$$\begin{aligned}
 \text{km}(\{b, z, \bar{z}, v, \bar{v}\}, M) &\geq 4 \cdot \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w_{\min}(B)}{W}\right)\right) \\
 &> 4 + 6 \cdot \epsilon \\
 &> \underbrace{1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w(b)}{W}\right)}_{=\|b-z\|_2^2} + \underbrace{1}_{\|\bar{z}-z\|_2^2} + \underbrace{2 + 4 \cdot \epsilon}_{\|v-z\|_2^2 + \|\bar{v}-z\|_2^2} \\
 &= \text{km}(\{b, z, \bar{z}, v, \bar{v}\}, \{z\}) .
 \end{aligned}$$

By combining this with (5.4.5), we observe that M is not locally optimal for the single-swap neighborhood since

$$\text{km}(X, M) > \text{km}(X, (M \setminus \{b\}) \cup \{z\}) .$$

Case 1.2 Assume at least one of the adjacent points $z \in M$ or $\bar{z} \in M$ but $v, \bar{v} \notin M$. Removing b from M does not increase the cost of z and \bar{z} as they either both have cost 0 or one of them has cost 0 and the other cost 1

$$\text{km}(\{z, \bar{z}\}, M) = \text{km}(\{z, \bar{z}\}, M \setminus \{b\}) .$$

Similar to the previous case we obtain

$$\begin{aligned}
 \text{km}(\{b, v, \bar{v}\}, M) &\geq 2 \cdot \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w_{\min}(B)}{W}\right)\right) \\
 &> 2 + 3 \cdot \epsilon \\
 &> \underbrace{1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w(b)}{W}\right)}_{=\|b-v\|_2^2} + \underbrace{1}_{\|\bar{v}-v\|_2^2} \\
 &= \text{km}(\{b, v, \bar{v}\}, \{v\}) .
 \end{aligned}$$

Again in combination with (5.4.5), we observe that M is not locally optimal for the single-swap neighborhood as

$$\text{km}(X, M) > \text{km}(X, (M \setminus \{b\}) \cup \{v\}) .$$

Case 1.3 Assume at least two of the adjacent points are part of the solution where at least one is located at the literals corresponding to z and one at the literals corresponding to v . That is, $z \in M \vee \bar{z} \in M$ and

5.4 Completeness of DKM/SingleSwap

$v \in M \vee \bar{v} \in M$. Removing b from M , similar to the previous case, does not affect the cost of the four adjacent points

$$\text{km}(X \setminus \{b\}, M) = \text{km}(X \setminus \{b\}, M \setminus \{b\}) .$$

However, recall that still $x \notin M$ and $\bar{x} \notin M$, and hence,

$$\begin{aligned} \text{km}(\{b, x, \bar{x}\}, M) &\geq 2 \cdot \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w_{\min}(B)}{W} \right) \right) \\ &> 2 + 3 \cdot \epsilon \\ &> \underbrace{1 + \epsilon \cdot \left(\frac{3}{2} + \frac{c \cdot w(b)}{W} \right)}_{=\|b-\bar{x}\|_2^2} + \underbrace{1}_{\|\bar{x}-x\|_2^2} \\ &= \text{km}(\{b, v, \bar{v}\}, \{x\}) . \end{aligned}$$

Combining these two observations we again conclude that M is not locally optimal for the single-swap neighborhood

$$\text{km}(X, M) > \text{km}(X, (M \setminus \{b\}) \cup \{x\}) .$$

Case 2 Assume there is no point corresponding to a clause in M but there exists a variable $y \neq x$ with $y, \bar{y} \in M$. We observe that, similar to the first case, if we remove y from M , then only the cost of the point y itself and the cost of points corresponding to clauses $B(y)$ can increase. Adding any representative to a solution never increases the cost. Thus,

$$\begin{aligned} \text{km}(X \setminus (B(y) \cup \{y, x, \bar{x}\}), M) \\ \geq \text{km}(X \setminus (B(y) \cup \{y, x, \bar{x}\}), (M \setminus \{y\}) \cup \{x\}) . \end{aligned}$$

The points x and \bar{x} have distance $1 + 2 \cdot \epsilon$ to any representative in M since there is not representative at a location corresponding to a clause. Additionally, all points corresponding to clauses in $B(y)$ have distance $1 + \epsilon \cdot (3/2 + c \cdot w(B)/W)$ to $\bar{y} \in M$ and $|B(y)| \leq m$. We obtain

$$\begin{aligned} \text{km}(B(y) \cup \{y, \bar{y}, x, \bar{x}\}, M) \\ &= 2 + 4 \cdot \epsilon + \sum_{b \in B(y)} \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w(b)}{W} \right) \right) \\ &> 2 + \sum_{b \in B(y)} \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w(b)}{W} \right) \right) + \frac{(c-1) \cdot \epsilon}{2} \quad \text{(Definition of } c) \\ &\geq 2 + \sum_{b \in B(y)} \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{w(b)}{W} \right) \right) + \frac{(c-1) \cdot \epsilon}{W} \underbrace{\sum_{b \in B(y)} w(b)}_{\leq m \cdot w_{\max}(B)} \end{aligned}$$

$$\begin{aligned}
 &= \underbrace{1}_{\|y-\bar{y}\|_2^2} + \underbrace{1}_{\|\bar{x}-x\|_2^2} + \sum_{b \in B(y)} \left(1 + \epsilon \cdot \left(\frac{3}{2} + \frac{c \cdot w(b)}{W} \right) \right) \\
 &= \text{km}(B(y) \cup \{y, \bar{y}, x, \bar{x}\}, \{x, \bar{y}\}) .
 \end{aligned}$$

We conclude the proof as, also in this case, M is not locally optimal with respect to the single-swap neighborhood

$$\text{km}(X, M) > \text{km}(X, (M \setminus \{y\}) \cup \{x\}) .$$

■

Since the intermediate results are so similar to our MUFL reduction, we obtain the correctness and the tightness of the DKM reduction without requiring any additional arguments.

Lemma 5.4.6 (Φ, Ψ) is a tight **PLS**-reduction.

Proof. Substitute [Corollary 5.4.3](#) and [Lemma 5.4.4](#) into the proofs of [Corollary 5.3.6](#), [Lemma 5.3.7](#), and [Lemma 5.3.8](#). ■

5.4.3 Embedding

So far we treated X as an abstract point set purely defined by its interpoint distances. In the following, we show that there is isometric embedding of X into Euclidean space. That is, there exists a point set in \mathbb{R}^d exhibiting exactly the required interpoint distances as squared Euclidean distance.

We represent point sets with fixed interpoint distances a so-called *distance matrix*. That is, for a set P of size n and a distance function $d : P \times P \rightarrow \mathbb{R}_{\geq 0}$, chose some ordering (p_1, \dots, p_n) of the points and let D_P be the matrix

$$D_P = (D_P)_{i,j} := d(p_i, p_j) .$$

Throughout this section we denote by D_X the distance matrix of the point set X in $\Phi(B, w)$, by $\vec{\mathbf{1}}_d$ the d -dimensional vector of ones, and by δ_{ij} the Kronecker delta, which is 1 if $i = j$ and 0 otherwise. We use the following classic result to show that X can be isometrically embedded into squared Euclidean space.

Theorem 5.4.7 (Schoenberg [1938]) Let $D \in \mathbb{R}^{d \times d}$ be a distance matrix. A point set $P \subseteq \mathbb{R}^d$ with squared Euclidean distance matrix $D_P = D$ exists if and only if

$$\forall u \in \mathbb{R}^d \text{ with } u \cdot \vec{\mathbf{1}}_d = 0 \text{ it holds that } u^T D u \leq 0 .$$

5.4 Completeness of DKM/SingleSwap

Omitting further details: there is a well-know technique called MULTIDIMENSIONAL SCALING, which finds points with the desired squared Euclidean distances.

Theorem 5.4.8 (Torgerson [1952]) *Let $D \in \mathbb{R}^{d \times d}$ be a distance matrix which is embeddable into squared Euclidean space. There is a polynomial-time algorithm which computes a matrix $P \in \mathbb{R}^{d \times d}$ whose rows form a set $\{p_i\} \subset \mathbb{R}^d$ with $(D)_{ij} = \|p_i - p_j\|_2^2$.*

Independent of the order of points chosen, observe that for our point set X we have

$$(D_X)_{i,j} = \|x_i - x_j\|_2^2 = 1 - \delta_{ij} + (\|x_i - x_j\|_2^2 - 1) \cdot (1 - \delta_{ij})$$

where the second summand is always non-negative.

Lemma 5.4.9 *The point set X from $\Phi(B, w)$ can be isometrically embedded into squared Euclidean space using $2 \cdot n + m$ dimensions.*

Proof. We apply [Theorem 5.4.7](#) with $d := 2 \cdot n + m$. For any $u \in \mathbb{R}^d$ with $u \cdot \vec{\mathbf{1}}_d = 0$, we obtain

$$\begin{aligned} & u^T D_X u \\ &= \sum_{i=1}^d \sum_{j=1}^d (D_X)_{i,j} \cdot u_i \cdot u_j \\ &= \sum_{i=1}^d \sum_{j=1}^d \left(u_i \cdot u_j - \delta_{ij} \cdot u_i \cdot u_j + (\|x_i - x_j\|_2^2 - 1) \cdot (1 - \delta_{ij}) \cdot u_i \cdot u_j \right) \\ &= \left(\sum_{i=1}^d u_i \right)^2 - \sum_{i=1}^d u_i^2 + \sum_{i=1}^d \sum_{j=1}^d (\|x_i - x_j\|_2^2 - 1) \cdot (1 - \delta_{ij}) \cdot u_i \cdot u_j \\ &= -\|u\|_2^2 + \sum_{i=1}^d \sum_{j=1}^d (\|x_i - x_j\|_2^2 - 1) \cdot (1 - \delta_{ij}) \cdot u_i \cdot u_j \quad (u \cdot \vec{\mathbf{1}}_d = 0) \\ &\leq -\|u\|_2^2 + \sum_{i=1}^d \sum_{j=1}^d (\|x_i - x_j\|_2^2 - 1) \cdot |u_i| \cdot |u_j| \\ &\leq -\|u\|_2^2 + 2 \cdot \epsilon \cdot \left(\sum_{i=1}^d |u_i| \right)^2 \\ &\leq -\|u\|_2^2 + 2 \cdot \epsilon \cdot d \cdot \|u\|_2^2 \quad \text{(Cauchy-Schwarz inequality)} \\ &= 0. \quad \text{(Definition of } d \text{ and } \epsilon) \end{aligned}$$

■

A downside of MULTIDIMENSIONAL SCALING is that, in general, the algorithm requires $2 \cdot n + m$ dimensions to exactly preserve the distances of the $2 \cdot n + m$ points in our set X . We complement our embedding result by showing that there cannot be another embedding of X which uses asymptotically less than $2 \cdot n + m$ dimensions.

Lemma 5.4.10 *If $P \subset \mathbb{R}^d$ is a set of pairwise equidistant points, then $|P| \leq d + 1$.*

Proof. We prove the claim by induction on the number of dimensions.

The base case $d = 1$ holds since there are at most 2 pairwise equidistant points on the real line.

Assume the claim holds for any fixed $d \in \mathbb{N}$. Let $c \in \mathbb{R}_{>0}$ be some constant and assume there exists a set $P \subset \mathbb{R}^{d+1}$ of $d + 1$ points with pairwise distance c . We want to bound how many points $q \in \mathbb{R}^{d+1}$ satisfy

$$\forall p \in P : \|p - q\|_2^2 = c.$$

The points in P span a d -dimensional hyperplane \mathcal{H} in \mathbb{R}^{d+1} . By induction hypothesis, \mathcal{H} cannot contain a point q with distance c to all points in P . Nevertheless, observe that the mean $\mu(P) \in \mathcal{H}$ has the same distance to all points in P . We have that

$$\|p - q\|_2^2 = \|p\|_2^2 + \|q\|_2^2 - 2 \cdot \langle p, q \rangle,$$

and thus, for all $p, q \in P$ with $p \neq q$,

$$\langle p, q \rangle = \frac{\|p\|_2^2 + \|q\|_2^2 - c}{2}.$$

For any fixed $p \in P$, we have

$$\begin{aligned} & \|p - \mu(P)\|_2^2 \\ &= \langle p - \mu(P), p - \mu(P) \rangle \\ &= \frac{1}{(d+1)^2} \cdot \sum_{\substack{q \in P \\ q \neq p}} \sum_{\substack{r \in P \\ r \neq p}} \langle p - q, p - r \rangle \\ &= \frac{1}{(d+1)^2} \cdot \sum_{\substack{q \in P \\ q \neq p}} \|p - q\|_2^2 + \frac{1}{(d+1)^2} \cdot \sum_{\substack{q \in P \\ q \neq p}} \sum_{\substack{r \in P \\ r \neq p, r \neq q}} \langle p - q, p - r \rangle \\ &= \frac{1}{(d+1)^2} \cdot \left(c \cdot d + \sum_{\substack{q \in P \\ q \neq p}} \sum_{\substack{r \in P \\ r \neq p, r \neq q}} \|p\|_2^2 - \langle p, q \rangle - \langle p, r \rangle + \langle q, r \rangle \right) \end{aligned}$$

5.5 Completeness of DFKM/SingleSwap

$$\begin{aligned}
&= \frac{1}{(d+1)^2} \cdot \left(c \cdot d + \sum_{\substack{q \in P \\ q \neq p}} \sum_{\substack{r \in P \\ r \neq p, r \neq q}} \frac{c}{2} \right) \\
&= \frac{c \cdot d}{(d+1)^2} \cdot \left(1 + \frac{d-1}{2} \right).
\end{aligned}$$

This means that the points in P all lie on the surface of a ball centered at $\mu(P)$. Therefore, the only point in \mathcal{H} that has the same distance to all points in P is $\mu(P)$.

Let $q \in \mathbb{R}^{d+1}$ be a point with distance c to all points in P . Since the squared Euclidean distance is the sum of the squared distances per dimension, we know that the orthogonal projection of q onto \mathcal{H} also has the same distance to all points in P . Thus, q has to lie on a line which is orthogonal to \mathcal{H} and intersects the hyperplane at $\mu(P)$. There are exactly two points on this line having distance exactly c to all points in P – one on each side of \mathcal{H} . It is easy to see that the distance between these points is larger than c . Thus, we can only add one additional point from \mathbb{R}^{d+1} and still maintain pairwise distance c between all points in P . ■

Observe that the set of positive/negative literals each form a set of n equidistant points in X , and the set of clauses forms a set of m equidistant points in X . Consequently, by [Lemma 5.4.10](#), there is no embedding of our point set into less than $\max\{n, m\} - 1$ dimensions. Thus, proving **PLS**-hardness of DKM with bounded, or even constant, number of dimensions requires a substantially new reduction.

5.5 Completeness of DFKM/SingleSwap

The third and final reduction presented in this chapter is a tight **PLS**-reduction showing the completeness of DFKM with $m = 2$ and with the single-swap neighborhood. This reduction is, once again, similar to the two previously presented reductions. However, in addition to the difficulties of the DKM reduction we now also have to find a reduction which circumvents the problems with overcovering (see [Section 4.4](#)). Furthermore, due to the hard to analyze objective function of our fuzzy problems, this result does not hold for unweighted instances. Instead of somehow encoding the clause weights into the distances of the points, here, we require a non-trivial weight function on the points. Recall that for the tightly **PLS**-hard Boolean satisfiability problems

there cannot be a polynomial upper bound on the weights. Hence, we are also unable to emulate these weights by copies of points.

Theorem 5.5.1 DISCRETE FUZZY k -MEANS *with $m = 2$ and with single-swap neighborhood is tightly PLS-complete.*

The construction of (Φ, Ψ) is similar to the one presented in the proofs of [Theorem 5.3.1](#) and [Theorem 5.4.1](#). However, to avoid problems with overcovering we reduce from PM2NAESAT instead of M2SAT. As before, we present the reduction based on an abstract point set purely defined by interpoint distances and argue embeddability afterwards.

5.5.1 Construction

Let $(B, w : B \rightarrow \mathbb{N})$ be an instance of PM2NAESAT over the set of variables X and denote $n := |X|$. We extend B to a larger set of clauses B' . For each clause $(x \vee y) \in B$, we add two clauses to B' :

$$b = (x \vee y) \text{ and } b' = (\bar{x} \vee \bar{y})$$

where, by definition of PM2NAESAT, b' does not appear in B . Let $m := |B'|$ and choose

$$W := 4 \cdot n^2 \cdot m \cdot w_{\max}(B) .$$

We construct the function Φ mapping (B, w) to an instance (X, k) to DFKM. As before, abstractly define X to contain $2 \cdot n + m$ points – one corresponding to the positive and negative literal of each variable and one corresponding to each clause in B' . This is the same construction as the one used in the DKM reduction but for the extended clause set B' . Let $1 < c < 2$ be some arbitrarily chosen constant and

$$\epsilon := \frac{m - 1}{9 \cdot n^2 \cdot m} .$$

As the distance between points in X we choose

$$\|p - q\|_2^2 := \begin{cases} 0 & \text{if } p = q \\ 1 & \text{if } p \text{ and } q \text{ are a literal and its negation} \\ 1 + \epsilon & \text{if } p \text{ is a literal appearing in clause } q \\ 1 + c \cdot \epsilon & \text{if } \bar{p} \text{ is a literal appearing in clause } q \\ 1 + 2 \cdot \epsilon & \text{else.} \end{cases}$$

5.5 Completeness of DFKM/SingleSwap

Observe that we still have $\epsilon \leq 1/(4 \cdot n + 2 \cdot m)$ since each clause is a unique combination of two different positive or negative literals, and hence, $m \leq 2 \cdot n \cdot (n - 1)$.

We let the weight function $w : X \rightarrow \mathbb{N}$ be defined as

$$w(x) := \begin{cases} w(b) & \text{if } x \text{ is a clause} \\ W & \text{else.} \end{cases}$$

The number of clusters is again set to $k := n$. Also as before, the function Ψ , given some set $M \subseteq X$, sets each variable x to true if $x \in M$ and to false otherwise. Once again notice the still missing *only if* part.

Before proceeding to the correctness proof we observe a basic property of these types of point sets. For each set of representatives $M \subseteq X$, we can lower bound the assignment of each point to each representative.

Lemma 5.5.2 *If $M \subseteq X$ with $|M| = n$ and κ is an optimal assignment function with respect to M , then*

$$\forall x \in X \setminus M \forall \mu \in M : \kappa(x, \mu) > \frac{1}{2 \cdot n} .$$

Proof. Fix any $x \in X \setminus M$ and $\mu \in M$. Using [Lemma 3.2.10](#) we conclude

$$\kappa(x, \mu) = \frac{\|x - \mu\|_2^{-2}}{\sum_{\mu' \in M} \|x - \mu'\|_2^{-2}} \geq \frac{(1 + 2 \cdot \epsilon)^{-1}}{\sum_{\mu' \in M} 1} > \frac{1}{2 \cdot n} .$$

■

Furthermore, recall that also by [Lemma 3.2.10](#), for all $M \subseteq X$,

$$\text{fkm}(X, w, M) = \sum_{x \in X \setminus M} \frac{w(x)}{\sum_{\mu \in M} \|x - \mu\|_2^{-2}} .$$

5.5.2 Correctness, Tightness, and Embedding

In the following, fix any PM2NAESAT instance (B, w) , denote the image of the reduction as $(X, w, n) := \Phi(B, w)$, and for each set of representatives $M \subseteq X$ denote by $T_M := \Psi(B, w, M)$ the truth assignment obtained under the image of Ψ . The structure of the following arguments is similar to the DKM reduction. Since Ψ is not injective, we use the concept of reasonable solutions to characterize the cost of solutions for the DFKM instance (X, w, n) .

Lemma 5.5.3 *If $M \in X$ is reasonable, then*

$$\begin{aligned} \text{fkm}(X, w, M) &= \frac{n \cdot W \cdot (1 + 2 \cdot \epsilon)}{n + 2 \cdot \epsilon} + (\Gamma_0 + \Gamma_2) \cdot w(B) \\ &\quad - (\Gamma_0 + \Gamma_2 - 2 \cdot \Gamma_1) \cdot w^{\text{NAE}}(B_t(T_M)) \end{aligned}$$

where

$$\Gamma_i = \left(\frac{i}{1 + \epsilon} + \frac{2 - i}{1 + c \cdot \epsilon} + \frac{n - 2}{1 + 2 \cdot \epsilon} \right)^{-1}.$$

Proof. Recall that, by the definitions of Ψ and reasonable solutions, a variable x is assigned true if and only if $x \in M$ and false if and only if $x \notin M$. Consider the cost of points corresponding to literals. Each of these either is in M and has cost 0 or it has a representative at its negated literal at distance 1 and $n - 1$ representatives at distance $1 + 2 \cdot \epsilon$. Since each of these points has weight W , we obtain an overall cost of

$$\frac{W}{1 + \frac{n-1}{1+2\epsilon}} = \frac{W \cdot (1 + 2 \cdot \epsilon)}{n + 2 \cdot \epsilon}.$$

Consider the cost of points corresponding to clauses. Each clause has length 2 and has $n - 2$ representatives at distance $1 + 2 \cdot \epsilon$. The other two are either at distance $1 + \epsilon$ or $1 + c \cdot \epsilon$, dependent on whether the literals in the clause are true or false. Hence, the cost of such a clause $b \in B'$ is $w(b) \cdot \Gamma_i$ where i is the number of true literals in that clause. Recall that a clause is NAE-satisfied if it contains exactly one true and one false literal. Furthermore, for each clause $b \in B$ we added a clause b' to B' , which contains the negated literals of the variables in b . Thus, if b has i true literals, then b' has $2 - i$ true literals. For each clause $b \in B$ we consider the pair of points $b, b' \in B'$. If $b \in B_t^{\text{NAE}}(T_M)$, then this pair has cost $2 \cdot w(b) \cdot \Gamma_1$. If $b \in B_f^{\text{NAE}}(T_M)$, then this pair has cost $w(b) \cdot (\Gamma_0 + \Gamma_2)$. Taking the sum we obtain

$$\begin{aligned} \text{fkm}(X, w, M) &= \frac{n \cdot W \cdot (1 + 2 \cdot \epsilon)}{n + 2 \cdot \epsilon} + \sum_{b \in B_t^{\text{NAE}}(T_M)} 2 \cdot w(b) \cdot \Gamma_1 \\ &\quad + \sum_{b \in B_f^{\text{NAE}}(T_M)} w(b) \cdot (\Gamma_0 + \Gamma_2) \\ &= \frac{n \cdot W \cdot (1 + 2 \cdot \epsilon)}{n + 2 \cdot \epsilon} + 2 \cdot \Gamma_1 \sum_{b \in B_t^{\text{NAE}}(T_M)} w(b) \\ &\quad + (\Gamma_0 + \Gamma_2) \cdot \sum_{b \in B_f^{\text{NAE}}(T_M)} w(b) \end{aligned}$$

5.5 Completeness of DFKM/SingleSwap

$$\begin{aligned}
&= \frac{n \cdot W \cdot (1 + 2 \cdot \epsilon)}{n + 2 \cdot \epsilon} + 2 \cdot (\Gamma_0 + \Gamma_2) \cdot w(B) \\
&\quad - (\Gamma_0 + \Gamma_2 - 2 \cdot \Gamma_1) \cdot \underbrace{\sum_{b \in B_t^{\text{NAE}}(T_M)} w(b)}_{=w^{\text{NAE}}(B_t(T_M))} .
\end{aligned}$$

■

Again we conclude that the cost improvements in the PM2NAESAT instance (B, w) correspond to cost improvements of reasonable solutions in the DFKM instance $\Phi(B, w)$.

Corollary 5.5.4 *If $M, M' \in X$ are reasonable, then*

$$\text{fkm}(X, w, M) > \text{fkm}(X, w, M') \text{ if and only if } w^{\text{NAE}}(B_t(T_M)) < w^{\text{NAE}}(B_t(T'_M)) .$$

Proof. Observe that

$$\begin{aligned}
&\Gamma_0 + \Gamma_2 - 2 \cdot \Gamma_1 \\
&= \frac{2 \cdot \left(\frac{1}{1+c\epsilon} - \frac{1}{1+\epsilon} \right)^2}{\left(\frac{2}{1+\epsilon} + \frac{n-2}{1+2\epsilon} \right) \cdot \left(\frac{2}{1+c\epsilon} + \frac{n-2}{1+2\epsilon} \right) \cdot \left(\frac{1}{1+\epsilon} + \frac{1}{1+c\epsilon} + \frac{n-2}{1+2\epsilon} \right)} > 0
\end{aligned}$$

and apply the proof of [Corollary 5.3.4](#) with [Lemma 5.5.3](#) instead of [Lemma 5.3.3](#). ■

In the following, we present a result analogous to [Lemma 5.3.5](#) and [Lemma 5.4.4](#). The main challenge that we have to deal with in this case is the analytical difficulty of the fuzzy objective function. This means that we have to analyze the cost of solutions significantly closer than before.

Lemma 5.5.5 *If $M \in X$ is locally optimal with respect to the single-swap neighborhood, then M is reasonable.*

Proof. The structure of the following proof is similar to that of the proof of [Lemma 5.4.4](#). However, it is technically more demanding due to the more involved objective function. Recall that each point corresponding to a clause $b \in B'$ (where b is either a clause from B or one of the newly added clauses) has four adjacent points: two at distance $1 + \epsilon$ and two at distance $1 + c \cdot \epsilon$. All other points in X have distance $1 + 2 \cdot \epsilon$ to b .

Assume in contraposition that M is not reasonable. Recall that, by definition, there exists a variable x such that $x, \bar{x} \notin M$. Again we differentiate two cases: either there exists a point corresponding to a clause b with $b \in M$ or there exists a variable $y \neq x$ such that $y, \bar{y} \in M$. We show that in both cases M is not locally optimal.

Case 1 Assume there is a clause $b \in M$ and that, without loss of generality, $b = (z \vee v)$. We observe that if we exchange b for some other representative, then only its own cost and the cost of its adjacent points can increase. All other points in X are at distance $1 + 2 \cdot \epsilon$ from b , thus, no location swapped in for b can be farther away from these points than b . If we use a respective optimal assignment function κ , then the cost of the other points cannot increase with respect to the swap. In the following, let \tilde{M} be a set of representatives obtainable from M by a single swap removing b . That is, there exists some $u \in X$ such that

$$\tilde{M} = (M \setminus \{b\}) \cup \{u\} .$$

Formally, we have

$$\text{fkm}(X \setminus \{b, z, \bar{z}, v, \bar{v}\}, w, M) \geq \text{fkm}(X \setminus \{b, z, \bar{z}, v, \bar{v}\}, w, \tilde{M}) ,$$

and hence,

$$\begin{aligned} & \text{fkm}(X, w, M) - \text{fkm}(X, w, \tilde{M}) \\ & \geq \text{fkm}(\{b, z, \bar{z}, v, \bar{v}\}, w, M) - \text{fkm}(\{b, z, \bar{z}, v, \bar{v}\}, w, \tilde{M}) . \end{aligned}$$

Let κ be an optimal assignment function with respect to M and $\tilde{\kappa}$ be an optimal assignment function with respect to \tilde{M} . As a shorthand notation, for each $u \in X$ and $\mu \in M$, we introduce

$$\begin{aligned} K(u, \mu) &= \sum_{\mu' \in M \setminus \{\mu\}} \kappa(u, \mu')^2 \cdot \|u - \mu'\|_2^2 \quad \text{and} \quad (5.5.6) \\ \tilde{K}(u, \mu) &= \sum_{\mu' \in M \setminus \{\mu\}} \tilde{\kappa}(u, \mu')^2 \cdot \|u - \mu'\|_2^2 . \end{aligned}$$

That is, $K(u, \mu)$ is the cost contribution of the representatives different from μ to the total cost of u . We differentiate three subcases, depending on which of b 's adjacent points are already part of the current solution M . In each subcase, we analyze the difference of the cost of M to the cost of a solution from the single-swap neighborhood of M .

Case 1.1 Assume none of the adjacent points are part of the solution $z, \bar{z}, v, \bar{v} \notin M$. We choose $\tilde{M} := (M \setminus \{b\}) \cup \{z\}$. Recall that the cost of a solution can only increase if we use a non-optimal assignment function, and since M and \tilde{M} agree on all representatives except for b

5.5 Completeness of DFKM/SingleSwap

and z , we have, for all $u \in X \setminus \{b, z\}$, that $K(u, z) = K(u, b)$. Thus, we obtain

$$\begin{aligned}
& \text{fkm}(\{b, z, \bar{z}, v, \bar{v}\}, w, M) - \text{fkm}(\{b, z, \bar{z}, v, \bar{v}\}, w, \tilde{M}) \\
&= \sum_{u \in \{z, \bar{z}, v, \bar{v}\}} w(u) \cdot \left(\kappa(u, b)^2 \cdot \|u - b\|_2^2 + K(u, b) \right) \\
&\quad - \sum_{u \in \{b, \bar{z}, v, \bar{v}\}} w(u) \cdot \left(\tilde{\kappa}(u, z)^2 \cdot \|u - z\|_2^2 + \tilde{K}(u, z) \right) \\
&\geq \sum_{u \in \{z, \bar{z}, v, \bar{v}\}} w(u) \cdot \left(\kappa(u, b)^2 \cdot \|u - b\|_2^2 + K(u, b) \right) \\
&\quad - \sum_{u \in \{b, \bar{z}, v, \bar{v}\}} w(u) \cdot \left(\kappa(u, b)^2 \cdot \|u - z\|_2^2 + K(u, b) \right) .
\end{aligned}$$

Since $b \in M$, there is an optimal κ with $\kappa(b, b) = 1$, and hence,

$$\begin{aligned}
& \text{fkm}(\{b, z, \bar{z}, v, \bar{v}\}, w, M) - \text{fkm}(\{b, z, \bar{z}, v, \bar{v}\}, w, \tilde{M}) \\
&\geq w(z) \cdot \left(\kappa(z, b)^2 \cdot \|z - b\|_2^2 + \underbrace{K(z, b)}_{\geq 0} \right) \\
&\quad - w(b) \cdot \left(\kappa(b, b)^2 \cdot \|b - z\|_2^2 + \underbrace{K(b, b)}_{=0} \right) \\
&\quad + \sum_{u \in \{\bar{z}, v, \bar{v}\}} w(u) \cdot \kappa(u, b)^2 \cdot \left(\|u - b\|_2^2 - \|u - z\|_2^2 \right) \\
&\geq (W \cdot \kappa(z, b)^2 - w(b)) \cdot \|z - b\|_2^2 + \underbrace{W \cdot \kappa(\bar{z}, b)^2 \cdot c \cdot \epsilon}_{\geq 0} \\
&\quad + W \cdot (\epsilon - 2 \cdot \epsilon) \cdot (\kappa(v, b)^2 + \kappa(\bar{v}, b)^2) \\
&\geq \left(W \cdot \frac{1}{4 \cdot n^2} - w(b) \right) \cdot \|z - b\|_2^2 \quad \text{(Lemma 5.5.2)} \\
&\quad - W \cdot \epsilon \cdot 2 \quad (\kappa(v, b), \kappa(\bar{v}, b) \leq 1) \\
&= (m \cdot w_{\max}(B) - w(b)) \cdot (1 + \epsilon) - 8 \cdot n^2 \cdot m \cdot w_{\max}(B) \cdot \frac{m - 1}{9 \cdot n^2 \cdot m} \\
&\geq (m - 1) \cdot w_{\max}(B) - (m - 1) \cdot w_{\max}(B) \cdot \frac{8}{9} \\
&> 0 .
\end{aligned}$$

Case 1.2 Assume at least one of the adjacent points $z \in M$ or $\bar{z} \in M$ but $v, \bar{v} \notin M$. Without loss of generality assume $\bar{z} \in M$. Removing b from M does not increase the cost of \bar{z} as the point has cost 0 anyway.

We choose $\tilde{M} := (M \setminus \{b\}) \cup \{v\}$. As in the previous case, we have, for all $u \in X \setminus \{b, v\}$, that $K(u, v) = K(u, b)$ and that there is an optimal κ with $\kappa(b, b) = 1$. Hence, we obtain

$$\begin{aligned}
 & \text{fkm}(\{b, z, \bar{z}, v, \bar{v}\}, w, M) - \text{fkm}(\{b, z, \bar{z}, v, \bar{v}\}, w, \tilde{M}) \\
 & \geq \sum_{u \in \{z, v, \bar{v}\}} w(u) \cdot \left(\kappa(u, b)^2 \cdot \|u - b\|_2^2 + K(u, b) \right) \\
 & \quad - \sum_{u \in \{b, z, \bar{v}\}} w(u) \cdot \left(\kappa(u, b)^2 \cdot \|u - v\|_2^2 + K(u, b) \right) \\
 & \geq w(v) \cdot \left(\kappa(v, b)^2 \cdot \|v - b\|_2^2 + \underbrace{K(v, b)}_{\geq 0} \right) \\
 & \quad - w(b) \cdot \left(\kappa(b, b)^2 \cdot \|b - v\|_2^2 + \underbrace{K(b, b)}_{=0} \right) \\
 & \quad + \sum_{u \in \{z, \bar{v}\}} w(u) \cdot \kappa(u, b)^2 \cdot \left(\|u - b\|_2^2 - \|u - v\|_2^2 \right) \\
 & \geq (W \cdot \kappa(v, b)^2 - w(b)) \cdot \|v - b\|_2^2 + \underbrace{W \cdot \kappa(\bar{v}, b)^2 \cdot c \cdot \epsilon}_{\geq 0} \\
 & \quad + W \cdot (\epsilon - 2 \cdot \epsilon) \cdot \kappa(z, b)^2 \\
 & \geq \left(W \cdot \frac{1}{4 \cdot n^2} - w(b) \right) \cdot \|v - b\|_2^2 \quad \text{(Lemma 5.5.2)} \\
 & \quad - W \cdot \epsilon \quad \quad \quad (\kappa(z, b) \leq 1) \\
 & \geq (m - 1) \cdot w_{\max}(B) - (m - 1) \cdot w_{\max}(B) \cdot \frac{4}{9} \\
 & > 0.
 \end{aligned}$$

Case 1.3 Assume at least two of the adjacent points are part of the solution where at least one is located at the literals corresponding to z and one at the literals corresponding to v . Without loss of generality assume $\bar{z}, \bar{v} \in M$. Similar to the previous case, removing b from M does not increase the cost of \bar{z} and \bar{v} . Recall that $x, \bar{x} \notin M$ and choose $\tilde{M} := (M \setminus \{b\}) \cup \{x\}$. Once again we have, for all $u \in X \setminus \{b, x\}$, that $K(u, x) = K(u, b)$ and that there is an optimal κ with $\kappa(b, b) = 1$. Analogously to the previous cases, we obtain

$$\begin{aligned}
 & \text{fkm}(\{b, z, \bar{z}, v, \bar{v}, x, \bar{x}\}, w, M) - \text{fkm}(\{b, z, \bar{z}, v, \bar{v}, x, \bar{x}\}, w, \tilde{M}) \\
 & \geq \sum_{u \in \{z, v, x, \bar{x}\}} w(u) \cdot \left(\kappa(u, b)^2 \cdot \|u - b\|_2^2 + K(u, b) \right)
 \end{aligned}$$

5.5 Completeness of DFKM/SingleSwap

$$\begin{aligned}
& - \sum_{u \in \{b, z, v, \bar{x}\}} w(u) \cdot \left(\kappa(u, b)^2 \cdot \|u - x\|_2^2 + K(u, b) \right) \\
& \geq w(x) \cdot \left(\kappa(x, b)^2 \cdot \|x - b\|_2^2 + \underbrace{K(x, b)}_{\geq 0} \right) \\
& \quad - w(b) \cdot \left(\kappa(b, b)^2 \cdot \|b - x\|_2^2 + \underbrace{K(b, b)}_{=0} \right) \\
& \quad + \sum_{u \in \{z, v, \bar{x}\}} w(u) \cdot \kappa(u, b)^2 \cdot \left(\|u - b\|_2^2 - \|u - x\|_2^2 \right) \\
& \geq (W \cdot \kappa(x, b)^2 - w(b)) \cdot \|x - b\|_2^2 + \underbrace{W \cdot \kappa(\bar{x}, b)^2 \cdot 2 \cdot \epsilon}_{\geq 0} \\
& \quad + W \cdot (\epsilon - 2 \cdot \epsilon) \cdot (\kappa(z, b)^2 + \kappa(v, b)^2) \\
& \geq \left(W \cdot \frac{1}{4 \cdot n^2} - w(b) \right) \cdot \|x - b\|_2^2 \quad \text{(Lemma 5.5.2)} \\
& \quad - W \cdot \epsilon \cdot 2 \quad (\kappa(z, b), \kappa(v, b) \leq 1) \\
& \geq (m - 1) \cdot w_{\max}(B) - (m - 1) \cdot w_{\max}(B) \cdot \frac{8}{9} \\
& > 0.
\end{aligned}$$

Case 2 Assume there is no point corresponding to a clause in M , but there exists a variable $y \neq x$ with $y, \bar{y} \in M$. We choose $\tilde{M} := (M \setminus \{y\}) \cup \{x\}$. Observe that, similar to before, if we remove y from M , then only the cost of the point y itself and the cost of points corresponding to clauses $B(y)$ can increase. All other points in X have distance $1 + 2 \cdot \epsilon$ to y , and thus, no location swapped in for y can be farther away from these points than y . Hence, their cost cannot increase if we use a respective optimal assignment function. We obtain

$$\text{fkm}(X \setminus (B(y) \cup \{y, x, \bar{x}\}), M) \geq \text{fkm}(X \setminus (B(y) \cup \{y, x, \bar{x}\}), \tilde{M}).$$

Furthermore, since there is no representative at a location corresponding to a clause, $y \in M$, and $x \in \tilde{M}$, we obtain

$$\begin{aligned}
\text{fkm}(\{x, \bar{x}, y\}, w, M) &= W \cdot \frac{2 + 4 \cdot \epsilon}{n} \text{ and} \\
\text{fkm}(\{x, \bar{x}, y\}, w, \tilde{M}) &= W \cdot \frac{2 + 4 \cdot \epsilon}{n + 2 \cdot \epsilon}.
\end{aligned}$$

Let κ be an optimal assignment function with respect to M , $K(u, \mu)$

be defined as in (5.5.6), and recall that $|B(y)| \leq m$. We observe

$$\begin{aligned} & \text{fkm}(B(y) \cup \{x, \bar{x}, y\}, w, \tilde{M}) \\ & \leq \text{fkm}(\{x, \bar{x}, y\}, w, \tilde{M}) + \text{fkm}(B(y), w, \tilde{M}, \kappa) \\ & = W \cdot \frac{2 + 4 \cdot \epsilon}{n + 2 \cdot \epsilon} + \sum_{b \in B(y)} w(b) \cdot \left(\kappa(b, y)^2 \cdot \|b - x\|_2^2 + K(b, y) \right). \end{aligned}$$

Recall that, by definition of the distances, we have $\forall y \in B(y) : \|b - y\|_2^2 = 1 + \epsilon$ and $\forall x \in X : \|b - x\|_2^2 \leq 1 + 2 \cdot \epsilon$. We conclude

$$\begin{aligned} & \text{fkm}(B(y) \cup \{x, \bar{x}, y\}, w, \tilde{M}) \\ & \leq W \cdot \frac{2 + 4 \cdot \epsilon}{n + 2 \cdot \epsilon} + \sum_{b \in B(y)} w(b) \cdot \left(\kappa(b, y)^2 \cdot (\|b - y\|_2^2 + \epsilon) + K(b, y) \right) \\ & = W \cdot \frac{2 + 4 \cdot \epsilon}{n + 2 \cdot \epsilon} + \sum_{b \in B(y)} w(b) \cdot \kappa(b, y)^2 \cdot \epsilon + \text{fkm}(B(y), w, M) \\ & \leq W \cdot \frac{2 + 4 \cdot \epsilon}{n + 2 \cdot \epsilon} + \epsilon \cdot M \cdot w_{\max}(B) + \text{fkm}(B(y), w, M) \quad (\kappa(b, y) \leq 1) \\ & = W \cdot \frac{2 + 4 \cdot \epsilon}{n + 2 \cdot \epsilon} + \epsilon \cdot \frac{W}{4 \cdot n^2} + \text{fkm}(B(y), w, M) \\ & < W \cdot \frac{2 + 4 \cdot \epsilon}{n + 2 \cdot \epsilon} + W \cdot \frac{\epsilon/n}{n + 2 \cdot \epsilon} + \text{fkm}(B(y), w, M) \\ & < W \cdot \frac{2 + 4 \cdot \epsilon + 2 \cdot \epsilon \cdot (2 + 4 \cdot \epsilon)/n}{n + 2 \cdot \epsilon} + \text{fkm}(B(y), w, M) \\ & = W \cdot \frac{2 + 4 \cdot \epsilon}{n} + \text{fkm}(B(y), w, M) \\ & = \text{fkm}(B(y) \cup \{x, \bar{x}, y\}, w, M). \end{aligned}$$

Thus, M is not locally optimal for the single-swap neighborhood. \blacksquare

Fortunately, our careful analysis once again led to intermediate results which are so similar to the ones of the previous two reductions that we obtain correctness and tightness without any additional arguments.

Lemma 5.5.7 (Φ, Ψ) is a tight **PLS**-reduction.

Proof. Substitute Corollary 5.5.4 and Lemma 5.5.5 into the proofs of Corollary 5.3.6, Lemma 5.3.7, and Lemma 5.3.8. \blacksquare

This also holds for the embedding of X .

5.5 Completeness of DFKM/SingleSwap

Lemma 5.5.8 *The point set X can be embedded into squared Euclidean space using $2 \cdot n + m$ dimensions, and there is no embedding using less than $\max\{n, m\} - 1$ dimensions.*

Proof. We can apply the proof of [Lemma 5.4.9](#) as is. Furthermore, we apply [Lemma 5.4.10](#) since in X the set of positive/negative literals each form a set of n equidistant points and the set of clauses forms a set of m equidistant points. ■

Non-Representability of Solutions

Contribution Summary We present the unsolvability by radicals of optimal solutions of FUZZY k -MEANS. The result discussed in [Section 6.3](#) has been published in [Blömer et al., 2016a], however, the proof is not part of that publication.

6.1 Algebraic Basics	88
6.2 Related Work	90
6.3 Unsolvability of Optimal FKM Solutions	90
6.4 Implementation of the Finite Taylor Grid	97

So far, we focused on classifying the hardness of discrete versions of fuzzy clustering problems. These types of geometric problems have a finite solution space and thus enable us to analyze them using combinatorial methods. One might think that the KM problem does not have a finite solution space as representatives can be placed arbitrarily in \mathbb{R}^d . However, recall that there is a binary assignment of points to representatives, and that, by [Lemma 3.2.7](#), we can characterize optimal representatives as a linear combination of the points assigned to them. Thus, there is a finite subspace of solutions which we know to contain all global optima, namely the set of all means of subsets of the input points. We do not know of the existence of a similar subspace for FKM. In the fuzzy variant of clustering we are allowed to freely place representatives in \mathbb{R}^d and to choose a continuous assignment of points to the solution. This has led to the FKM problem resisting any attempt to combinatorially express solutions in terms of the input points without either fixing the representatives or the assignment function. In the following, we show that no closed form combinatorial expression can exist since optimal solutions to FKM are, in general, algebraically difficult.

6.1 Algebraic Basics

We start by formalizing our notion of *algebraically difficult*. Our main result is that, in general, optimal solutions of FKM are not *solvable by radicals*. The mathematical foundation of unsolvability lies in the theory of Galois groups of polynomials. It is not expedient to actually take a deep dive into Galois theory as a part of this thesis. The results we need to apply to prove our claim can be understood without any knowledge of Galois groups. However, to provide some context, we review core definitions and concepts of unsolvability. In the following, we assume that the reader has basic knowledge of algebra. The definitions and results presented in this section are taken from [Herstein, 1975; Hungerford, 2003] where one can also find a thorough discussion of the topic.

Let K and F be fields. If $K \subset F$, then we call K a subfield of F , respectively F an extension field of K . Let $u_1, \dots, u_n \in F$. We denote by $K(u_1, \dots, u_n)$ the subfield of F generated by $K \cup \{u_1, \dots, u_n\}$, that is, the intersection of all subfields of F that contain $K \cup \{u_1, \dots, u_n\}$ (this is also the smallest subfield of F containing $K \cup \{u_1, \dots, u_n\}$). If $F = K(u_1, \dots, u_n)$, then we call F a finitely generated extension.

Definition 6.1.1 *A finitely generated extension $K(u_1, \dots, u_n) = F$ is called radical extension if*

1. *there exists a $p_1 \in \mathbb{N}$ such that $u_1^{p_1} \in K$, and,*
2. *for all $i \in [n]_2$, there exists a $p_i \in \mathbb{N}$ such that $u_i^{p_i} \in K(u_1, \dots, u_{i-1})$.*

Our main concern are radical extensions of \mathbb{Q} . On the one hand, if $x \in \mathbb{Q}(r_1, \dots, r_n)$, for some $r_1, \dots, r_n \in \mathbb{R}$ such that $\mathbb{Q}(r_1, \dots, r_n)$ is a radical extension, then we can represent x as a finite concatenation of rational numbers using addition, multiplication, and taking roots. On the other hand, if $x \in \mathbb{R}$ does not lie in any radical extension of \mathbb{Q} , then we cannot represent x in the described manner. This has several algorithmic consequences when x is the solution of some problem. Mainly, there is no exact algorithm solving this problem in general, assuming a model of computation using standard arithmetics and root extraction to express roots of algebraic equations [Bajaj, 1988]. This is independent of any runtime constraints when trying to compute a solution. Numbers which do not lie in any radical extension of \mathbb{Q} cannot be explicitly represented in finite memory using the standard model of Turing or RAM machines. As we see in the following, there are such numbers which can be represented symbolically as “ x is the

6.1 Algebraic Basics

root of the following (finite) polynomial”. However, it is not clear how to, in general, efficiently compute the value of some objective function given a solution which is implicitly represented as the root of some polynomial.

The natural follow-up question is how to actually show that some number $x \in \mathbb{R}$ does not lie in any radical extension of \mathbb{Q} . To this end, consider some non-constant polynomial $p(x)$ with coefficients in K (we also say $p \in K[x]$). A finitely generated extension $F = K(u_1, \dots, u_n)$ is called a splitting field over K of the polynomial p if p splits over F and u_1, \dots, u_n are the roots of p in F . That is, there exists a $c \in K$ such that

$$p(x) = c \cdot \prod_{i=1}^n (x - u_i) .$$

Splitting fields are unique up-to isomorphism. Consequently, the equation $p(x) = 0$ is called *solvable by radicals over K* if the splitting field over K of p is a radical extension.

We introduce the *Galois group* of a polynomial. Let $p \in K[x]$ be some irreducible polynomial and F be the splitting field over K of p . The Galois group of the polynomial p is the set of all automorphisms on F fixing K . That is, the set of all isomorphisms $\varphi : F \rightarrow F$ such that $\varphi|_K$ is the identity on K . This set together with the function composition \circ constitutes a group.

Lemma 6.1.2 *Let $p \in K[x]$ be an irreducible polynomial. If $p(x) = 0$ is solvable by radicals over K , then the Galois group of p is solvable.*

We do not discuss the concept of solvable groups, which is tied to finite sequences of normal subgroups with abelian factor groups. Instead, we only introduce a family of unsolvable groups which we later use to prove our main result of this chapter. Let S be some finite set and consider the set $P(S)$ of all bijections $\varphi : S \rightarrow S$. Again, together with the function composition \circ this forms a group. An element of this group is called a permutation, and thus, $(P(S), \circ)$ is called the *permutation group* of S . Assume $S = [n]$ for some natural n . The permutation group of $[n]$ is called the *Symmetric group of n symbols* and is usually denoted as S_n . Symmetric groups are important in the study of finite groups. Most notably for us as cases of unsolvable groups.

Lemma 6.1.3 *The Symmetric group S_n is not solvable for $n \geq 5$.*

We conclude this section by presenting a characterisation of the Galois group of certain polynomials.

Lemma 6.1.4 (Bajaj [1988]) *Let $p \in \mathbb{Q}[x]$ be a polynomial with $n := \deg(p) > 2$ and $n \equiv 0 \pmod{2}$. We call a prime number good for p if it does not divide the discriminant of p .*

If there are good prime numbers p_1, p_2, p_3 for p such that

1. *$p \pmod{p_1}$ is an irreducible polynomial of degree n ,*
2. *$p \pmod{p_2}$ factors into a linear term and an irreducible polynomial of degree $n - 1$, and*
3. *$p \pmod{p_3}$ factors into a linear term, an irreducible polynomial of degree 2, and an irreducible polynomial of degree $n - 3$,*

then the Galois group of p is isomorphic to the Symmetric group S_n .

6.2 Related Work

The Weber problem asks for an optimal representative for the sum of Euclidean distances error, also called the *geometric median*. That is, given $X \subseteq \mathbb{R}^d$ find $\mu \in \mathbb{R}^d$ minimizing $\sum_{x \in X} \|x - \mu\|_2$. It is a straightforward generalization of Fermat's problem, which asks for the same point but only considering 3 points in the plane. Famously, Bajaj [1988] resolved the Weber problem by showing that the solution for $n \geq 5$ points in $d \geq 2$ dimensions is, in general, unsolvable by radicals over \mathbb{Q} . The proof consists of the construction of an instance with 5 points in the plane with integral coordinates. Using a symmetry argument one can show that an optimal solution of this instance has to lie on the y -axis. Thus, finding an optimal solution reduces to finding a root of the derivative of a univariate polynomial. As it turns out, the roots of this polynomial are unsolvable by radicals over \mathbb{Q} .

6.3 Unsolvability of Optimal FKM Solutions

We state and prove the main result of this chapter.

Theorem 6.3.1 *Optimal solutions to the FUZZY k -MEANS problem are, in general, not solvable by radicals over \mathbb{Q} .*

A solution to FKM is not solvable by radicals if neither the representatives nor all values of the assignment function are solvable by radicals. Recall that optimal solutions are stationary pairs, and in a stationary pair each representative is a rational expression in the values of the assignment function, and vice versa. Hence, it is easy to

6.3 Unsolvability of Optimal FKM Solutions

see that either both representatives and the assignment function are solvable or they are both not solvable. If all optimal solutions of an instance are unsolvable by radicals, then there are two implications for the FKM problem. First, we cannot finitely represent an optimal solution to the problem, and thus, can only ever approximate the optimum of the instance. Second, the iterative relocation scheme for FKM never terminates if it is initialized with a rational solution (assuming arbitrary precision).

Over the remainder of this chapter we show that the unique optimal solution to the FKM problem with $m = 2$, $k = 2$, and

$$X = \{-3, -2, -1, 1, 2, 3\}$$

is not solvable by radicals over \mathbb{Q} . Specifically note that this instance is a set of integral points on the real line. Also note that

$$\text{fkm}(X, \mathbb{1}, \{-2, 2\}) = \frac{242}{65},$$

and hence, that any optimal solution has cost at most $242/65$.

We start by showing that there is a unique (up to renumbering of representatives) optimal solution for FKM on our instance.

Lemma 6.3.2 *Let $M = \{\mu_1, \mu_2\}$ be an optimal solution of our instance and assume without loss of generality that $\mu_1 \leq \mu_2$. It holds that*

$$\mu_1 \in (-3, -1) \text{ and } \mu_2 \in (1, 3).$$

Proof. First, we show that the signs of μ_1 and μ_2 are different. To the contrary assume that $\mu_1, \mu_2 \leq 0$. Then we obtain

$$\text{fkm}(X, \mathbb{1}, M) \geq \text{fkm}(\{3\}, \mathbb{1}, M) \geq \text{fkm}(\{3\}, \mathbb{1}, \{0, 0\}) = \frac{1}{1/9 + 1/9} = 4.5$$

where $\mu_1, \mu_2 \geq 0$ is equivalent due to the symmetry of X .

Recall that optimal solutions always lie in the convex hull of the input points. Hence, we can already conclude

$$\mu_1 \in [-3, 0] \text{ and } \mu_2 \in [0, 3].$$

Second, assume that $\mu_1 \in [-1, 0]$. We differentiate two subcases. If $\mu_2 \in [0, 1]$, then

$$\begin{aligned} \text{fkm}(X, \mathbb{1}, M) &\geq \text{fkm}(\{-3, 3\}, \mathbb{1}, M) \\ &\geq \text{fkm}(\{-3\}, \mathbb{1}, \{-1, 0\}) + \text{fkm}(\{3\}, \mathbb{1}, \{0, 1\}) \\ &= \frac{2}{1/4 + 1/9} = \frac{72}{13} > \frac{242}{65}. \end{aligned}$$

If $\mu_2 \in [1, 3]$, then

$$\begin{aligned} \text{fkm}(X, \mathbb{1}, M) &\geq \text{fkm}(\{-2, -3\}, \mathbb{1}, M) \\ &\geq \text{fkm}(\{-2, -3\}, \mathbb{1}, \{-1, 1\}) \\ &= \frac{1}{1 + 1/9} + \frac{1}{1/4 + 1/16} = \frac{41}{10} > \frac{242}{65}. \end{aligned}$$

Hence, we obtain $\mu_1 \in [-3, -1]$ and by symmetry also $\mu_2 \in (1, 3]$.

Third, assume that $\mu_1 = -3$. We also differentiate two subcases. If $\mu_2 \in [1, 2]$, then

$$\begin{aligned} \text{fkm}(X, \mathbb{1}, M) &\geq \text{fkm}(\{-2, -1, 3\}, \mathbb{1}, M) \\ &\geq \text{fkm}(\{-2, -1\}, \mathbb{1}, \{-3, 1\}) + \text{fkm}(\{3\}, \mathbb{1}, \{-3, 2\}) \\ &= \frac{1}{1 + 1/9} + \frac{1}{1/4 + 1/4} + \frac{1}{1/81 + 1} = \frac{797}{205} > \frac{242}{65}. \end{aligned}$$

If $\mu_2 \in [2, 3]$, then

$$\begin{aligned} \text{fkm}(X, \mathbb{1}, M) &\geq \text{fkm}(\{-2, -1, 1\}, \mathbb{1}, M) \\ &\geq \text{fkm}(\{-2, -1, 1\}, \mathbb{1}, \{-3, 2\}) + \text{fkm}(\{1\}, \mathbb{1}, \{-3, 2\}) \\ &= \frac{1}{1 + 1/16} + \frac{1}{1/4 + 1/9} + \frac{1}{1/16 + 1} = \frac{1028}{221} > \frac{242}{65}. \end{aligned}$$

Hence, we obtain $\mu_1 \in (-3, -1)$ and by symmetry also $\mu_2 \in (1, 3)$. \blacksquare

We narrowed our solution space down to a small square in which all optimal solutions of FKM on our instance X lie. In the next step, we argue that the objective function is strictly convex on the square $(-3, -1) \times (1, 3)$. This means that the objective function has a single stationary point in the square and that this stationary point describes the unique global optimum of the objective function (still, up to renumbering of the two representatives). As it turns out, classical manual analysis is infeasible, which is why we resort to a computer-assisted proof.

Lemma 6.3.3 *The function $\text{fkm}(X, \mathbb{1}, \{\mu_1, \mu_2\})$ is convex for $\mu_1 \in (-3, -1)$ and $\mu_2 \in (1, 3)$.*

Proof. Throughout this proof, we use the shorthand notation

$$\text{fkm}(\mu_1, \mu_2) := \text{fkm}(X, \mathbb{1}, \{\mu_1, \mu_2\})$$

for our objective function.

6.3 Unsolvability of Optimal FKM Solutions

To prove that the objective function is strictly convex, it is sufficient to show that its Hessian matrix is positive definite [Bertsekas, 1999]. This is equivalent to showing that both partial derivatives and the determinant of the Hessian matrix are positive. That is, for all $\mu_1 \in (-3, -1)$ and $\mu_2 \in (1, 3)$, we have

- $\frac{\partial^2 \text{fkm}}{\partial \mu_1^2}(\mu_1, \mu_2) > 0$,
- $\frac{\partial^2 \text{fkm}}{\partial \mu_2^2}(\mu_1, \mu_2) > 0$, and
- $\left(\frac{\partial^2 \text{fkm}}{\partial \mu_1^2} \cdot \frac{\partial^2 \text{fkm}}{\partial \mu_2^2} - \frac{\partial^2 \text{fkm}}{\partial \mu_1 \partial \mu_2} \cdot \frac{\partial^2 \text{fkm}}{\partial \mu_2 \partial \mu_1} \right) (\mu_1, \mu_2) > 0$.

Each of these second derivatives is a fraction of two high-degree polynomials. Showing the positiveness of these quantities manually seems to be infeasible. This is why we propose a computer-assisted approach.

A core ingredient to our proof is the well-known Taylor Theorem. More specifically, we use a variant of the theorem with the mean-value form of the remainder.

Lemma 6.3.4 (Forster [2010]) *Let $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be a bivariate function and $\epsilon_x, \epsilon_y > 0$ be constants. For each $x, y \in \mathbb{R}$ such that f is continuously differentiable on $[x, x + \epsilon_x] \times [y, y + \epsilon_y]$, there exists a $\xi \in [0, 1]$ such that*

$$\begin{aligned} & f(x + \epsilon_x, y + \epsilon_y) \\ &= f(x, y) + \frac{\partial f}{\partial x}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \cdot \epsilon_x + \frac{\partial f}{\partial y}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \cdot \epsilon_y . \end{aligned}$$

We use Lemma 6.3.4 to prove that if our function is positive at some point in the desired interval, then it is also positive in a small neighborhood around this point. To this end, let $\epsilon > 0$ be some constant, $x, y \in \mathbb{R}$ such that f is continuously differentiable on $[x, x + \epsilon] \times [y, y + \epsilon]$, and set

$$\Delta_\epsilon := \frac{f(x, y)}{\rho_x + \rho_y}$$

where

$$\begin{aligned} \rho_x &:= \max_{(x_0, y_0) \in [x, x + \epsilon] \times [y, y + \epsilon]} \left\{ \left| \frac{\partial f}{\partial x}(x_0, y_0) \right| \right\} \\ \rho_y &:= \max_{(x_0, y_0) \in [x, x + \epsilon] \times [y, y + \epsilon]} \left\{ \left| \frac{\partial f}{\partial y}(x_0, y_0) \right| \right\} . \end{aligned}$$

Assume $f(x, y) > 0$, let $0 \leq \epsilon_x, \epsilon_y \leq \min\{\epsilon, \Delta_\epsilon\}$, and obtain $\xi \in [0, 1]$ from [Lemma 6.3.4](#) applied to $f(x + \epsilon_x, y + \epsilon_y)$. We bound

$$\begin{aligned}
 & f(x + \epsilon_x, y + \epsilon_y) \\
 &= f(x, y) + \frac{\partial f}{\partial x}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \cdot \epsilon_x + \frac{\partial f}{\partial y}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \cdot \epsilon_y \\
 &\geq f(x, y) - \left| \frac{\partial f}{\partial x}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \right| \cdot \epsilon_x - \left| \frac{\partial f}{\partial y}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \right| \cdot \epsilon_y \\
 &\geq f(x, y) - \left(\left| \frac{\partial f}{\partial x}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \right| + \left| \frac{\partial f}{\partial y}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \right| \right) \cdot \Delta_\epsilon \\
 &= f(x, y) \cdot \left(1 - \frac{\left| \frac{\partial f}{\partial x}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \right| + \left| \frac{\partial f}{\partial y}(x + \xi \cdot \epsilon_x, y + \xi \cdot \epsilon_y) \right|}{\rho_x + \rho_y} \right) \\
 &\geq 0. \qquad \qquad \qquad \text{(since } \epsilon_x, \epsilon_y \leq \epsilon)
 \end{aligned}$$

Thus, if f is positive at (x, y) , then f is also positive on the whole square $[x, x + \min\{\epsilon, \Delta_\epsilon\}] \times [y, y + \min\{\epsilon, \Delta_\epsilon\}]$. Based on this observation, we propose the following general algorithm to verify the positiveness of a function on a given area.

Algorithm 6.3.5: POSITIVENESSBYFINITE TAYLOR GRID

Input: Area $A = [x_0, x_1] \times [y_0, y_1] \subset \mathbb{R}^2$, a continuously differentiable $f : A \rightarrow \mathbb{R}$, and $\epsilon > 0$

```

1 if  $f(x_0, y_0) \leq 0$  then
2   | return false
3  $x \leftarrow x_0$ 
4  $y \leftarrow y_0$ 
5  $\epsilon_{\min} \leftarrow \epsilon$ 
6 while  $y \leq y_1$  do
7   |  $\delta \leftarrow \min \left\{ \epsilon, \frac{f(x, y)}{\rho_x + \rho_y} \right\}$ 
8   |  $\epsilon_{\min} \leftarrow \min\{\delta, \epsilon_{\min}\}$ 
9   |  $x \leftarrow x + \delta$ 
10  | if  $x > x_1$  then
11  |   |  $x \leftarrow x_0$ 
12  |   |  $y \leftarrow y + \epsilon_{\min}$ 
13  |   |  $\epsilon_{\min} \leftarrow \epsilon$ 
14 return true
    
```

The core idea is to cover the area A by a finite grid. This grid is chosen small enough such that we can use [Lemma 6.3.4](#). If the

6.3 Unsolvability of Optimal FKM Solutions

function is positive on a corner of a cell of the grid, then it is positive on the whole cell. The algorithm checks the function value at a corner of A for positiveness and then adaptively chooses the correct grid size at runtime. Thus, if the algorithm covers A by these small grid cells, then f is positive on the whole area A .

We cannot give an upper bound on the runtime of [Algorithm 6.3.5](#). In fact, if f is negative at some point in the interior of A , then the algorithm goes into an infinite loop. However, if [Algorithm 6.3.5](#) terminates, then f is positive on A . There are two more issues we have to resolve to accept a halting computation of the algorithm as a formally sound proof of positiveness. First, we have to make sure that we compute with enough precision to represent δ and ϵ_{\min} without rounding errors. Second, we have to figure out how to compute ρ_x and ρ_y . In the following, we discuss our approaches to solving these two issues for our functions.

The core observation here is that our second derivatives of the FKM objective function are fractions of polynomials. This has several important implications. Each of them has continuous derivatives, and if you evaluate them at a rational input, then the function value is also rational. Hence, we can solve the precision problem of [Algorithm 6.3.5](#) by starting at rational endpoints of our area, that is, checking $[-3, -1] \times [1, 3]$. All computations of the algorithm also stay within rational numbers. Thus, we avoid rounding errors by representing everything as a fraction of integers. We do not know how to compute ρ_x and ρ_y exactly. However, we can easily find an upper bound on these quantities. First, we compute a minimum/maximum value, by plugging in the smallest/largest possible value for all terms with a positive coefficient and the largest/smallest value for all terms with a negative coefficient at the numerator, and vice versa at the denominator. We assume that coefficients are adjusted for any sign changes due to the sign of the arguments, for example, a positive coefficient which is evaluated at a negative argument has to be treated as a negative coefficient. Second, we obtain an upper bound on ρ_x and ρ_y as the maximum of the absolute values of the upper and lower bounds we just computed.

We implemented [Algorithm 6.3.5](#) in Sage¹, a general purpose computer algebra system based on Python². One of the strengths of this language is that it inherently uses arbitrary size integers to represent rational numbers as numerator and denominator. The algo-

¹<https://www.sagemath.org>, last access: 22.05.2019

²<https://www.python.org>, last access: 22.05.2019

rithm terminates for $\frac{\partial^2 \text{fkm}}{\partial \mu_1^2}$, $\frac{\partial^2 \text{fkm}}{\partial \mu_2^2}$, and $\left(\frac{\partial^2 \text{fkm}}{\partial \mu_1^2} \cdot \frac{\partial^2 \text{fkm}}{\partial \mu_2^2} - \frac{\partial^2 \text{fkm}}{\partial \mu_1 \partial \mu_2} \cdot \frac{\partial^2 \text{fkm}}{\partial \mu_2 \partial \mu_1} \right)$ on $[-3, -1] \times [1, 3]$. Our implementation can be found in [Section 6.4](#). ■

The previous two lemmata tell us that all optimal solutions for FKM on X lie in $[-3, -1] \times [1, 3]$ and that the objective function is convex on this square. Hence, if we find a stationary point of the objective function in that area, then this is the unique, globally optimal solution of the problem. Let

$$f(\mu) := \text{fkm}(X, \mathbb{1}, \{-\mu, \mu\})$$

be the restriction of the FKM objective function on X to solutions which are symmetric around 0, and consider its derivative

$$\frac{df}{d\mu}(\mu) = \frac{2 \cdot \mu \cdot (3 \cdot \mu^{12} + 84 \cdot \mu^{10} + 490 \cdot \mu^8 - 292 \cdot \mu^6 - 8981 \cdot \mu^4 - 17640 \cdot \mu^2 - 11664)}{(\mu^2 + 9)^2 \cdot (\mu^2 + 4)^2 \cdot (\mu^2 + 1)^2}.$$

Since the denominator is strictly positive, we only have to find roots of the numerator. The trivial root $\mu = 0$ does not lie in our desired interval and is thus not the global optimum we seek. Let

$$g(\mu) := 3 \cdot \mu^{12} + 84 \cdot \mu^{10} + 490 \cdot \mu^8 - 292 \cdot \mu^6 - 8981 \cdot \mu^4 - 17640 \cdot \mu^2 - 11664,$$

and observe that $g(2) = -20864$ and $g(3) = 8658576$. Further notice that the roots of the polynomial g are square roots of the roots of the polynomial

$$h(\mu) := 3 \cdot \mu^6 + 84 \cdot \mu^5 + 490 \cdot \mu^4 - 292 \cdot \mu^3 - 8981 \cdot \mu^2 - 17640 \cdot \mu - 11664.$$

Thus, the unique global optimum of FKM on X is of the form $\{-\mu^*, \mu^*\}$ where μ^* is the root of the polynomial h in the interval $(2, 3)$.

Proof of Theorem 6.3.1. We apply [Lemma 6.1.4](#) to the polynomial h . The discriminant of h is

$$3086428236130279646902930636800 = 2^{31} \cdot 3^7 \cdot 5^2 \cdot 7^3 \cdot 76637866514129.$$

Thus, 11, 17, and 89 are good prime numbers for h . We factor the polynomial modulo these primes to obtain

$$h(\mu) = 3 \cdot (\mu^6 + 6 \cdot \mu^5 + 2 \cdot \mu^4 + 9 \cdot \mu^3 + 2 \cdot \mu^2 + 5 \cdot \mu + 6) \pmod{11},$$

$$h(\mu) = 3 \cdot (\mu^5 + 3 \cdot \mu^4 + 9 \cdot \mu^3 + 12 \cdot \mu^2 + 10 \cdot \mu + 7) \cdot (\mu + 8) \pmod{17}, \text{ and}$$

$$h(\mu) = 3 \cdot (\mu^3 + 17 \cdot \mu^2 + 50 \cdot \mu + 17) \cdot (\mu^2 + 9 \cdot \mu + 27) \cdot (\mu + 2) \pmod{89}.$$

Thus, the Galois group of h is isomorphic to S_6 , which yields the claim by applying [Lemma 6.1.3](#). ■

6.4 Implementation of the Finite Taylor Grid

```

1 # Since we are only computing fractions and evaluating polynomials,
2 # we can work with the rational numbers and receive exact
3 # computation results.
4 R.<x,y> = PolynomialRing(QQ, 'x,y')
5
6 # Set our objective function ...
7 obj = (((3 - x)^(-2) + (3 - y)^(-2))^( -1) +
8         ((2 - x)^(-2) + (2 - y)^(-2))^( -1) +
9         ((1 - x)^(-2) + (1 - y)^(-2))^( -1) +
10        ((1 + x)^(-2) + (1 + y)^(-2))^( -1) +
11        ((2 + x)^(-2) + (2 + y)^(-2))^( -1) +
12        ((3 + x)^(-2) + (3 + y)^(-2))^( -1))
13 obj(x,y) = obj # Transform expression to function
14
15 # ... and compute the second partial derivatives
16 oxx = derivative(derivative(obj,x),x)
17 oyy = derivative(derivative(obj,y),y)
18 # Omit oxy, since dobj^2/dxdy = dobj^2/dydx
19 oyy = derivative(derivative(obj,y),y)
20 r = oxx*oyy - oxy^2
21
22 # We need to show that, on the square [-3,-1]x[1,3],
23 # oxx > 0, oyy > 0, and oxx*oyy - oxy^2 > 0.
24
25 # Next, we adjust the signs in each polynomial. Since x is always
26 # negative, we flip the sign of all terms with an odd powered
27 # occurrence of x. Afterwards, we can work with the square [1,3]^2.
28 def flipOddX(p):
29     res(x,y) = [0,0]
30     p = p.numerator_denominator()
31     for i in [0,1]:
32         for o in (expand(p[i])).operands():
33             degx = int(o.degree(x))
34             if degx > 0 and degx % 2 == 1:
35                 res[i] -= o
36             else:
37                 res[i] += o
38     return res[0]/res[1]
39
40 oxx = flipOddX(oxx)
41 oyy = flipOddX(oyy)
42 r = flipOddX(r)
43
44 # Finally, we check all function values for a
45 # sufficiently small grid.
46
47 # Given a polynomial p, find an upper bound on the absolute value
48 # in the area XIxYI.
49 def ubAbsInInterval(p, XI, YI):

```

```

50 # We compute a lower and an upper bound, and then take the
51 # maximum of the absolute values
52 lb, ub = [0, 0], [0, 0]
53 p = p.numerator_denominator()
54 for i in [0, 1]:
55     for o in (expand(p[i])).operands():
56         if o.leading_coefficient(x).leading_coefficient(y) > 0:
57             lb[i] += o(x = XI[i], y = YI[i])
58             ub[i] += o(x = XI[1-i], y = YI[1-i])
59         else:
60             lb[i] += o(x = XI[1-i], y = YI[1-i])
61             ub[i] += o(x = XI[i], y = YI[i])
62 return max(abs(ub[0]/ub[1]), abs(lb[0]/lb[1]))
63
64 # We use the first order Taylor expansion to find a
65 # sufficiently small grid.
66 def checkGrid(p):
67     if p(1, 1) <= 0:
68         print('Found a function value <= 0 at 1 1')
69         sys.exit()
70     maxstep = 1/16
71     dx = derivative(p, x)
72     dy = derivative(p, y)
73     curx, cury = 1, 1
74     epsmin = maxstep
75     while cury <= 3:
76         fvalue = p(curx, cury)
77         if fvalue <= 0:
78             # This actually cannot happen.
79             # We verify it anyways as a simple sanity check.
80             print('Found a function value <= 0 at '+
81                 str(curx)+' '+str(cury))
82             sys.exit()
83         eps = min(maxstep,
84                 fvalue/
85                 (ubAbsInInterval(dx, [curx, curx+maxstep],
86                                 [cury, cury+maxstep])
87                  + ubAbsInInterval(dy, [curx, curx+maxstep],
88                                     [cury, cury+maxstep])))
89         epsmin = min(eps, epsmin)
90         curx += eps
91         if curx > 3:
92             curx = 1
93             cury += epsmin
94             epsmin = maxstep
95             print(str(round((cury-1)*100/(2+maxstep)))+ '%... ')
96
97 print('Checking df^2/dx^2 ... ')
98 checkGrid(oxx)
99 print('Checking df^2/dy^2 ... ')
100 checkGrid(oyy)

```

6.4 Implementation of the Finite Taylor Grid

```
101 print( 'Checking  $df^2/dx^2*df^2/dy^2 - (df/(dxdy))^2 \dots$  ' )  
102 checkGrid(r)  
103 print( 'All checks successful' )
```


Part III



Approaching Hard Problems

Approximation Algorithms

Contribution Summary We present algorithms trying to solve FUZZY k -MEANS and FUZZY RADIUS k -MEANS. The algorithm discussed in [Section 7.2](#) is an improvement over the algorithm published in [\[Blömer et al., 2018\]](#), utilizing a variant of the soft-to-hard lemma published in [\[Blömer et al., 2016a\]](#). In [Section 7.3](#) we present a previously unpublished algorithm, solving FUZZY RADIUS k -MEANS for small instances.

7.1 Related Work	104
7.2 A PTAS for FKM With Fixed Number of Clusters	105
7.2.1 From Soft to Hard Clusters	106
7.2.2 A Sampling Based Approach	111
7.3 Solving Small FRKM Instances	114

The first part of this thesis focused on complexity theory. We have seen several hardness results for the discrete variants of our clustering problems. Furthermore, in [Chapter 6](#) we discussed that, in general, no conventional algorithm is able to solve FKM exactly. Consequently, the best we can expect is a polynomial-time algorithm, approximating FKM up to a factor of $(1 + \epsilon)$, for an arbitrarily chosen $\epsilon > 0$. In the following, we present an algorithm solving the FKM problem up to a factor of $(1 + \epsilon)$. The runtime of our algorithm is polynomial in the number of points and the dimension, and exponential in the number of clusters k . This makes it a PTAS for any constant k . To the best of our knowledge, it is the first algorithm of this type whose runtime is independent of the weight function on the input points. Furthermore, it is also the first algorithm where the exponential dependence on k is linear. Afterwards, we show that the radius variant FRKM can be solved for two clusters on instances on the real line.

The FKM algorithm we present in this chapter is not an improvement to the fastest $(1 + \epsilon)$ -approximation algorithm known, which was presented in Blömer et al. [2016a]. However, the independence of the weight function makes it a key ingredient for our combination with fuzzy coresets, which we present in Chapter 8. This combination then yields an approximation algorithm for FKM which is faster than all previously presented algorithms with the same approximation guarantee.

7.1 Related Work

Since most of the problems we discuss in this thesis are computationally hard, they have been subject to substantial research in approximation algorithms.

Besides the approximation guarantee of local search algorithms (see Section 5.2), MUFL can be approximated using many different algorithmic techniques. A rounded solution of an LP formulation of the problem yields a factor 4 approximation [Shmoys et al., 1997]. This was improved by utilizing primal-dual and filtering techniques, to achieve a 3-approximation [Jain & Vazirani, 2001]. Furthermore, there is a 1.61-approximative greedy algorithm [Jain et al., 2002]. Following this, a lot of algorithms with constant approximation factor have been presented, steadily closing in on the 1.463-inapproximability bound (see Section 4.3). Currently, the best known algorithm has an expected 1.488-approximation ratio [Li, 2013].

Our radius covering problem RKM can be approximated with a factor of 2 by a fairly simple farthest-first traversal algorithm [Gonzalez, 1985]. This bound is tight assuming $\mathbf{P} \neq \mathbf{NP}$ (see Section 4.3).

Although the approximation factor is not constant, one of the most popular approximation algorithms for KM is the k -means++ seeding technique [Arthur & Vassilvitskii, 2007], due to its simplicity and fast runtime. For some time, local search heuristics were the best constant-factor approximation algorithms known for KM. Currently, the best algorithm achieves a $(6.357 + \epsilon)$ -approximation ratio using a primal-dual approach [Ahmadian et al., 2017]. The latter also yields a $(9 + \epsilon)$ -approximation for the KM-problem on general metric spaces. $(1 + \epsilon)$ -approximation algorithms for KM assume either a constant number of clusters or constant dimension since there is probably no PTAS for KM (see Section 4.3). For constant k , there has been a string of research of successively improving runtimes based on coresets [Bădoiu et al., 2002; Har-Peled & Kushal, 2006] (and many more) with the

currently fastest known algorithm due to [Feldman et al. \[2007\]](#). [Kumar et al. \[2010\]](#) present an algorithm based on sampling candidate means, which has a similar runtime to the fastest coresets-based algorithms.

So far, the only algorithm with guaranteed approximation ratio for FKM have been presented by [Blömer et al. \[2016a\]](#) and [Blömer et al. \[2018\]](#). They have presented several different approaches, which all yielded $(1 + \epsilon)$ -approximations. The fastest of these has runtime $\mathcal{O}(d \cdot k \cdot n) + (d \cdot \log(n))^{\tilde{\mathcal{O}}(k^2/\epsilon)}$ (where $\tilde{\mathcal{O}}$ hides logarithmic factors).

7.2 A PTAS for FKM With Fixed Number of Clusters

In the following, we present a $(1 + \epsilon)$ -approximation algorithm for FKM, whose runtime is polynomial in the number of points n and the dimensionality of the input set d . Notably, the runtime of the algorithm does not depend on the weight function on the points. This makes the algorithm particularly useful in combination with the coresets we present in [Chapter 8](#). However, the algorithm's runtime is still exponential in the number of clusters k and ϵ .

Algorithm 7.2.1: DERANDOMIZED SAMPLING

Input: $X \in \mathbb{R}^d$, $w : X \rightarrow \mathbb{N}$, $k \in \mathbb{N}$, $\epsilon \in (0, 1)$

- 1 $\mathcal{T} \leftarrow \{\mu(S) \mid S \subseteq X, |S| = 33/\epsilon\}$
- 2 $M \leftarrow \arg \min_{T \in \mathcal{T}, |T|=k} \{\text{fkm}(X, w, T)\}$
- 3 **return** M

Notice that [Algorithm 7.2.1](#) ignores the weight function for the computation of the candidate representatives in \mathcal{T} . The algorithm computes the unweighted mean for each multiset of $33/\epsilon$ points from X where points can occur multiple times and are counted with multiplicity. One might think of such a set S as drawing points from X with replacement. However, despite ignoring the weight function, \mathcal{T} contains a set of k representatives which form a $(1 + \epsilon)$ -approximation to the weighted problem. In addition to being independent of the weight function, this constitutes the first $(1 + \epsilon)$ -approximation algorithm for FKM where the exponential runtime dependence on k is only linear.

Theorem 7.2.2 *Let (X, w, k) be an FKM instance and M^* be an optimal solution. [Algorithm 7.2.1](#) computes a set $M \subseteq \mathbb{R}^d$ with $|M| = k$ and*

$$\text{fkm}(X, w, M) \leq (1 + \epsilon) \cdot \text{fkm}(X, w, M^*) .$$

in time $d \cdot n^{\mathcal{O}(k/\epsilon)}$.

We prove [Theorem 7.2.2](#) in two steps. First, we present a variation of the soft-to-hard lemma, which has already been a key ingredient in the proof of previously presented approximation algorithms for FKM [[Blömer et al., 2016a](#)]. Our slightly modified variant of the lemma shows that, for each representative in an FKM solution, there exists, under a certain condition, a hard cluster (i.e. a subset of input points) which has characteristics similar to certain statistics with respect to the fuzzy representative. Second, we show how to use a sampling technique, which is well-known from algorithms for the KM problem, to generate a small number of candidate means. In this candidate set, there is a multiset of representatives with the desired approximation ratio.

7.2.1 From Soft to Hard Clusters

Once again, the main difficulty when analyzing algorithms intended to solve FKM is the missing combinatorial structure of optimal solutions. We circumvent this problem by proving the existence of similar hard clusters. Similarity of clusters is expressed in three statistics: weight, mean, and cost contribution. For a hard cluster, i.e. some set $C \subseteq \mathbb{R}^d$ with weight $w : C \rightarrow \mathbb{N}$, these are $w(C)$, $\mu_w(C)$ and $\text{km}(C, w)$.

Definition 7.2.3 (Fuzzy Cluster Statistics) *Let $X \subseteq \mathbb{R}^d$, $w : X \rightarrow \mathbb{N}$, $M \subseteq \mathbb{R}^d$, and $\kappa : X \times M \rightarrow [0, 1]$. For a representative $\mu \in M$ we denote its cluster statistics*

- $\kappa(X, w, \mu) := \sum_{x \in X} \kappa(x, \mu)^m \cdot w(x)$ *the weight,*
- $\mu_{\kappa(\cdot, \mu)^m \cdot w(\cdot)}(X)$ *the mean, and*
- $\text{fkm}(X, w, \{\mu\}, \kappa)$ *the cost contribution.*

For the following analysis we require several elementary observations from probability theory. Instead of presenting a self-contained overview of the field, we expect the reader to be familiar with the basic terms and only outline core notation and results we apply later on. A full introduction to probability theory, including proofs of the results we present in the next paragraph, can, for example, be found in the book by [Mitzenmacher & Upfal \[2005\]](#).

Basic Probability Theory Let R be a random variable. We call $E[R]$ its *expected value* and $\text{Var}(R)$ its *variance*. For any non-negative random

7.2 A PTAS for FKM With Fixed Number of Clusters

variable R and $a \in \mathbb{R}_{>0}$, *Markov's inequality* states

$$\Pr(R \geq a \cdot \mathbb{E}[R]) \leq \frac{1}{a}.$$

For any random variable R with finite expected value and non-zero variance and every $a \in \mathbb{R}_{>0}$, *Chebyshev's inequality* states

$$\Pr\left(|R - \mathbb{E}[R]| \geq a \cdot \sqrt{\text{Var}(R)}\right) \leq \frac{1}{a^2}.$$

Fix some multiset $X \subseteq \mathbb{R}^d$, weight function $w : X \rightarrow \mathbb{N}$, $k \in \mathbb{N}$, and stationary pair (M, κ) for FKM on (X, w, k) . Technically speaking, the following results do not require a stationary pair but just some arbitrary feasible assignment function and its induced set of representatives. However, choosing a stationary pair is sufficient for our application and eases notation.

We introduce a random process which creates, for each $\mu \in M$, a subset $C_\mu \subseteq X$. Each point $x \in X$ is independently added to C_μ with probability $\kappa(x, \mu)^m$.

Fix some $\mu \in M$ and let $\chi_{C_\mu} : X \rightarrow \{0, 1\}$ be the indicator function of the set C_μ . For each $x \in X$, $\chi_{C_\mu}(x)$ is a binary random variable with

$$\begin{aligned} \Pr(\chi_{C_\mu}(x) = 1) &= \kappa(x, \mu)^m, \\ \mathbb{E}[\chi_{C_\mu}(x)] &= \kappa(x, \mu)^m, \text{ and} \\ \text{Var}(\chi_{C_\mu}(x)) &= \kappa(x, \mu)^m \cdot (1 - \kappa(x, \mu)^m). \end{aligned}$$

Using χ_{C_μ} we derive concentration bounds on the cluster statistics of C_μ where each statistic is a random variable depending on the process generating C_μ . Notice that the random variables $\chi_{C_\mu}(x)$ are independent of the weight function w .

First, we bound the difference of the weights of the hard and fuzzy clusters. Observe that

$$w(C_\mu) = \sum_{x \in C_\mu} w(x) = \sum_{x \in X} \chi_{C_\mu}(x) \cdot w(x),$$

and thus, we obtain

$$\mathbb{E}[w(C_\mu)] = \sum_{x \in X} \kappa(x, \mu)^m \cdot w(x) = \kappa(X, w, \mu).$$

Applying Chebyshev's inequality yields

$$\Pr\left(|w(C_\mu) - \kappa(X, w, \mu)| \geq \sqrt{4} \cdot \sqrt{\text{Var}(w(C_\mu))}\right) \leq \frac{1}{4}. \quad (7.2.4)$$

Second, we bound the squared Euclidean distance of the means of the hard clusters to their respective mean μ . Let

$$D_\mu := \left\| \sum_{x \in X} \chi_{C_\mu}(x) \cdot w(x) \cdot (x - \mu) \right\|_2^2$$

be a random variable depending on χ_{C_μ} , as well. Similar to the weights, we apply Markov's inequality to obtain

$$\Pr(D_\mu \geq 4 \cdot \mathbb{E}[D_\mu]) \leq \frac{1}{4}. \quad (7.2.5)$$

Finally, we bound the cost in the same fashion. Observe that

$$\text{km}(C_\mu, w) \leq \text{km}(C_\mu, w, \{\mu\}) = \sum_{x \in X} \chi_{C_\mu}(x) \cdot w(x) \cdot \|x - \mu\|_2^2,$$

and since $\|x - \mu\|_2^2$ is independent of the random variables $\chi_{C_\mu}(x)$,

$$\mathbb{E}[\text{km}(C_\mu, w, \{\mu\})] = \sum_{x \in X} \kappa(x, \mu)^m \cdot w(x) \cdot \|x - \mu\|_2^2 = \text{fkm}(X, w, \{\mu\}, \kappa).$$

Applying Markov's inequality one more time we obtain

$$\begin{aligned} \Pr(\text{km}(C_\mu, w) \geq 4 \cdot \text{fkm}(X, w, \{\mu\}, \kappa)) \\ \leq \Pr(\text{km}(C_\mu, w) \geq 4 \cdot \mathbb{E}[\text{km}(C_\mu, w)]) \leq \frac{1}{4}. \end{aligned} \quad (7.2.6)$$

We formulate our soft-to-hard lemma by combining these concentration bounds into a single argument. In the following, we denote $w_{\max}(X) := \max_{x \in X} \{w(x)\}$.

Lemma 7.2.7 (Soft-To-Hard) *Let $k, d \in \mathbb{N}$, $X \subseteq \mathbb{R}^d$, $w : X \rightarrow \mathbb{N}$, (M, κ) a stationary pair for FKM on (X, w, k) , and $\epsilon \in (0, 1)$. For each $\mu \in M$ with*

$$\kappa(X, w, \mu) \geq \frac{16 \cdot w_{\max}(X)}{\epsilon},$$

there exists a set $C_\mu \subseteq X$ such that

$$w(C_\mu) \geq \frac{\kappa(X, w, \mu)}{2}, \quad (7.2.8)$$

$$\|\mu_w(C_\mu) - \mu\|_2^2 \leq \frac{\epsilon}{\kappa(X, w, \mu)} \cdot \text{fkm}(X, w, \{\mu\}, \kappa), \quad \text{and} \quad (7.2.9)$$

$$\text{km}(C_\mu, w) \leq 4 \cdot \text{fkm}(X, w, \{\mu\}, \kappa). \quad (7.2.10)$$

7.2 A PTAS for FKM With Fixed Number of Clusters

Proof. Let

$$\begin{aligned} \hat{w} : X &\rightarrow (0, 1] \\ x &\mapsto \frac{w(x)}{w_{\max}(X)} \end{aligned}$$

be the normalized weight function, fix some $\mu \in M$ with $\kappa(X, w, \mu) \geq 16 \cdot w_{\max}(X)/\epsilon$, and let χ_{C_μ} be the previously described random variable. We apply the probabilistic method. Using the union bound we obtain that the probability for the events (7.2.4), (7.2.5), and (7.2.6) (applied to \hat{w} instead of w) to fail simultaneously is at least $1 - 3/4 > 0$. Hence, there exists a set $C_\mu \subseteq X$ with

$$\begin{aligned} |\hat{w}(C_\mu) - \kappa(X, \hat{w}, \mu)| &\leq 2 \cdot \sqrt{\text{Var}(\hat{w}(C_\mu))}, \\ D_\mu &\leq 4 \cdot \mathbb{E}[D_\mu], \text{ and} \\ \text{km}(C_\mu, \hat{w}) &\leq 4 \cdot \text{fkm}(X, \hat{w}, \{\mu\}, \kappa). \end{aligned}$$

Observe that by our precondition we have

$$2 = \frac{\sqrt{16}}{2} \leq \frac{\sqrt{\epsilon \cdot \kappa(X, w, \mu)/w_{\max}(X)}}{2} = \frac{\sqrt{\epsilon \cdot \kappa(X, \hat{w}, \mu)}}{2} \leq \frac{\sqrt{\kappa(X, \hat{w}, \mu)}}{2}.$$

Furthermore, since the $\chi_{C_\mu}(x)$ are drawn independent of each other,

$$\begin{aligned} \sqrt{\text{Var}(\hat{w}(C_\mu))} &= \sqrt{\sum_{x \in X} \hat{w}(x)^2 \cdot \kappa(x, \mu)^m \cdot (1 - \kappa(x, \mu))^m} \\ &\leq \sqrt{\sum_{x \in X} \hat{w}(x) \cdot \kappa(x, \mu)^m} \quad (0 < \hat{w}(x), \kappa(x, \mu) \leq 1) \\ &= \sqrt{\kappa(X, \hat{w}, \mu)}. \end{aligned}$$

Combining these inequalities, we conclude

$$\begin{aligned} |w(C_\mu) - \kappa(X, w, \mu)| &= w_{\max}(X) \cdot |\hat{w}(C_\mu) - \kappa(X, \hat{w}, \mu)| \\ &\leq w_{\max}(X) \cdot 2 \cdot \sqrt{\text{Var}(\hat{w}(C_\mu))} \\ &\leq w_{\max}(X) \cdot \frac{\kappa(X, \hat{w}, \mu)}{2} = \frac{\kappa(X, w, \mu)}{2}, \end{aligned}$$

and thus, (7.2.8).

Next, we compute the expectation $\mathbb{E}[D_\mu]$. We expand the term

$$\begin{aligned} D_\mu &= \left\| \sum_{x \in X} \chi_{C_\mu}(x) \cdot \hat{w}(x) \cdot (x - \mu) \right\|_2^2 \\ &= \sum_{x \in X} \sum_{y \in X} \chi_{C_\mu}(x) \cdot \chi_{C_\mu}(y) \cdot \hat{w}(x) \cdot \hat{w}(y) \cdot \langle x - \mu, y - \mu \rangle. \end{aligned}$$

Since the $\chi_{C_\mu}(x)$ are independent binary random variables, we have $\mathbb{E}[\chi_{C_\mu}(x)^2] = \kappa(x, \mu)^m$ and, for all $x \neq y$, that $\mathbb{E}[\chi_{C_\mu}(x) \cdot \chi_{C_\mu}(y)] = \kappa(x, \mu)^m \cdot \kappa(y, \mu)^m$. We obtain

$$\begin{aligned}
 \mathbb{E}[D_\mu] &= \sum_{x \in X} \sum_{y \in X} \mathbb{E}[\chi_{C_\mu}(x) \cdot \chi_{C_\mu}(y)] \cdot \hat{w}(x) \cdot \hat{w}(y) \cdot \langle x - \mu, y - \mu \rangle \\
 &= \sum_{x \in X} \left(\kappa(x, \mu)^m \cdot \hat{w}(x)^2 \cdot \|x - \mu\|_2^2 \right. \\
 &\quad \left. + \sum_{\substack{y \in X \\ x \neq y}} \kappa(x, \mu)^m \cdot \kappa(y, \mu)^m \cdot \hat{w}(x) \cdot \hat{w}(y) \cdot \langle x - \mu, y - \mu \rangle \right) \\
 &= \sum_{x \in X} \left((\kappa(x, \mu)^m - \kappa(x, \mu)^{2m}) \cdot \hat{w}(x)^2 \cdot \|x - \mu\|_2^2 \right. \\
 &\quad \left. + \sum_{y \in X} \kappa(x, \mu)^m \cdot \kappa(y, \mu)^m \cdot \hat{w}(x) \cdot \hat{w}(y) \cdot \langle x - \mu, y - \mu \rangle \right) \\
 &= \sum_{x \in X} \left(\kappa(x, \mu)^m \cdot (1 - \kappa(x, \mu)^m) \cdot \hat{w}(x)^2 \cdot \|x - \mu\|_2^2 \right. \\
 &\quad \left. + \kappa(x, \mu)^m \cdot \hat{w}(x) \cdot \langle x - \mu, \underbrace{\sum_{y \in X} \kappa(y, \mu)^m \cdot \hat{w}(y) \cdot (y - \mu)}_{= 0 \text{ since } (M, \kappa) \text{ stationary}} \rangle \right) \\
 &\leq \sum_{x \in X} \kappa(x, \mu)^m \cdot \hat{w}(x) \cdot \|x - \mu\|_2^2 \quad (0 < \hat{w}(x), \kappa(x, \mu) \leq 1) \\
 &= \text{fkm}(X, w, \{\mu\}, \kappa) / w_{\max}(X) .
 \end{aligned}$$

To bound the difference between the representatives, we rewrite

$$\|\mu_{\hat{w}(C_\mu)} - \mu\|_2^2 = \left\| \frac{\sum_{x \in X} \chi_{C_\mu}(x) \cdot \hat{w}(x) \cdot (x - \mu)}{\hat{w}(C_\mu)} \right\|_2^2 = \frac{D_\mu}{\hat{w}(C_\mu)^2} .$$

Recall that $\hat{w}(C_\mu)^2 \geq \kappa(X, \hat{w}, \mu)/4$, $\epsilon \cdot \kappa(X, \hat{w}, \mu)/4 \geq 4$, and $\mu_{\hat{w}(C_\mu)} = \mu_w(C_\mu)$. Thus, we obtain (7.2.9) by

$$\begin{aligned}
 \|\mu_w(C_\mu) - \mu\|_2 &= \frac{D_\mu}{\hat{w}(C_\mu)^2} \leq \frac{4 \cdot \text{fkm}(X, w, \{\mu\}, \kappa) / w_{\max}(X)}{\kappa(X, \hat{w}, \mu)^2 / 4} \\
 &\leq \frac{\epsilon \cdot \kappa(X, \hat{w}, \mu)}{4} \cdot \frac{4 \cdot \text{fkm}(X, w, \{\mu\}, \kappa)}{\kappa(X, \hat{w}, \mu)^2 \cdot w_{\max}(X)} \\
 &= \frac{\epsilon \cdot \text{fkm}(X, w, \{\mu\}, \kappa)}{\kappa(X, w, \mu)} .
 \end{aligned}$$

Finally, we obtain (7.2.10) immediately from

$$\begin{aligned} \text{km}(C_\mu, w) &= w_{\max}(X) \cdot \text{km}(C_\mu, \hat{w}) \\ &\leq w_{\max}(X) \cdot 4 \cdot \text{fkm}(X, \hat{w}, \{\mu\}, \kappa) = 4 \cdot \text{fkm}(X, w, \{\mu\}, \kappa) . \end{aligned}$$

■

The soft-to-hard lemma shows us that, for each fuzzy representative with a large enough assignment volume, there is a subset of input points with similar cluster statistics. In the following, we use the probabilistic method again, this time on a well-know sampling technique, to show the correctness of [Algorithm 7.2.1](#).

7.2.2 A Sampling Based Approach

We do not have any structural knowledge about the subsets of input points we are trying to find. The only information available to us is a lower bound on the number of points in each subset. To approximate the means of such an unstructured subset, we use a simple uniform sampling technique which is oblivious to any underlying structure of the points. More specifically, that the mean of a uniformly drawn subset is, with high probability, a good approximation of the mean of the whole set.

Lemma 7.2.11 (Inaba et al. [1994]) *Let $P \subseteq \mathbb{R}^d$, $n \in \mathbb{N}$, and $\delta \in (0, 1)$. If S is a set of n points drawn uniformly at random from P , then*

$$\Pr \left(\|\mu(S) - \mu(P)\|_2^2 \leq \frac{1}{\delta \cdot n} \cdot \frac{\text{km}(P, \mathbb{1})}{|P|} \right) \geq 1 - \delta .$$

We use the probabilistic method to conclude the existence of a constant size subset whose mean approximates the mean of an arbitrary subset of X well.

Corollary 7.2.12 *Let $P \subseteq \mathbb{R}^d$, $w : X \rightarrow \mathbb{N}$, $k \in \mathbb{N}$, and $\epsilon \in (0, 1)$. If $|P| > 0$, then there exists a set $S \subseteq P$ with $|S| = 33/(32 \cdot \epsilon)$ such that*

$$\|\mu(S) - \mu_w(P)\|_2^2 \leq \epsilon \cdot \frac{\text{km}(P, w)}{w(P)} .$$

Proof. Let $S \subseteq P$ with $|S| = 33/(32 \cdot \epsilon)$ be uniformly sampled from P with replacement. The sampling treats all points $p \in P$ as $w(p)$ copies of points at the same location. Applying [Lemma 7.2.11](#), we obtain

$$\Pr \left(\|\mu(S) - \mu_w(P)\|_2^2 \leq \epsilon \cdot \frac{\text{km}(P, w)}{w(P)} \right) \geq 1/33 > 0 .$$

Hence, there exists a set $S \subseteq P$ with the claimed property. ■

We combine this result with the soft-to-hard lemma to finally proof the correctness of [Algorithm 7.2.1](#).

Correctness We analyze the result M of [Algorithm 7.2.1](#). Consider the point set X_c which contains

$$c := \left\lceil \frac{64 \cdot w_{max}(X)}{\epsilon \cdot \min_{\mu \in M} \{\kappa^*(X, w, \mu)\}} \right\rceil$$

copies of each point $x \in X$. Notice that this is purely analytical as we can not actually compute c . For all sets of representatives M and all assignment functions κ , we have

$$\text{fkm}(X_c, w, M, \kappa) = c \cdot \text{fkm}(X, w, M, \kappa) .$$

We extend κ^* to an assignment function for X_c by setting, for each $\mu \in M$, $x \in X$, and $y \in X_c$ with $y = x$, $\kappa^*(y, \mu) := \kappa^*(x, \mu)$. Then, (M^*, κ^*) is also optimal for the instance (X_c, w, k) . Furthermore, for all $\mu \in M$, we have

$$\begin{aligned} \kappa^*(X_c, w, \mu) &\geq \frac{64 \cdot w_{max}(X)}{\epsilon \cdot \min_{\mu \in M} \{\kappa^*(X, w, \mu)\}} \cdot \sum_{x \in X} \kappa^*(x, \mu)^m \cdot w(x) \\ &\geq \frac{64 \cdot w_{max}(X)}{\epsilon} . \end{aligned}$$

Hence, we can apply [Lemma 7.2.7](#) to X_c , w , and $\epsilon/4$, for each representative $\mu \in M^*$. There exists a collection of k non-empty subsets

$$\{C_\mu \subseteq X_c \mid \mu \in M^*\} ,$$

such that, for each $\mu \in M^*$,

$$\begin{aligned} w(C_\mu) &\geq \frac{\kappa^*(X_c, w, \mu)}{2} , \\ \|\mu_w(C_\mu) - \mu\|_2 &\leq \frac{\epsilon}{4 \cdot \kappa^*(X_c, w, \mu)} \cdot \text{fkm}(X_c, w, \{\mu\}, \kappa^*) , \text{ and} \\ \text{km}(C_\mu, w) &\leq 4 \cdot \text{fkm}(X_c, w, \{\mu\}, \kappa^*) . \end{aligned}$$

Next, we apply [Corollary 7.2.12](#) to w and $\epsilon/32$, for each set C_μ . We obtain that there exists another collection of k non-empty subsets

$$\{S_\mu \subseteq C_\mu \subseteq X_c \mid \mu \in M^*\}$$

such that, for each $\mu \in M^*$, we have $|S_\mu| = 33/\epsilon$ and

$$\|\mu(S_\mu) - \mu_w(C_\mu)\|_2^2 \leq \frac{\epsilon}{32} \cdot \frac{\text{km}(C_\mu, w)}{w(C_\mu)} .$$

7.2 A PTAS for FKM With Fixed Number of Clusters

Recall that X_c consists only of copies of points from X . Hence, we have, for each $\mu \in M^*$, that $S_\mu \subseteq X$. Simply speaking, we can also obtain the sets S_μ from X by allowing each point to occur multiple times independent of its actual multiplicity. Since [Algorithm 7.2.1](#) computes the means of all sets $S \subseteq X$ with $|S| = 33/\epsilon$ and chooses the best k representatives among those, we have

$$\text{fkm}(X, w, M) \leq \text{fkm}(X, w, \{\mu(S_\mu) \mid \mu \in M^*\}) .$$

Plugging all of this together, we bound the cost of M as follows.

$$\begin{aligned} & \text{fkm}(X, w, M) \\ & \leq \text{fkm}(X, w, \{\mu(S_\mu) \mid \mu \in M^*\}) \\ & = \frac{1}{c} \cdot \text{fkm}(X_c, w, \{\mu(S_\mu) \mid \mu \in M^*\}) \\ & \leq \frac{1}{c} \cdot \text{fkm}(X_c, w, \{\mu(S_\mu) \mid \mu \in M^*\}, \kappa^*) \\ & = \frac{1}{c} \cdot \sum_{x \in X_c} \sum_{\mu \in M^*} \kappa^*(x, \mu)^m \cdot w(x) \cdot \|x - \mu(S_\mu)\|_2^2 \\ & = \text{fkm}(X, w, M^*, \kappa^*) + \frac{1}{c} \cdot \sum_{x \in X_c} \sum_{\mu \in M^*} \kappa^*(x, \mu)^m \cdot w(x) \cdot \|\mu - \mu(S_\mu)\|_2^2 \\ & \hspace{20em} \text{(Lemma 3.2.7)} \\ & \leq \text{fkm}(X, w, M^*, \kappa^*) + \frac{2}{c} \cdot \sum_{x \in X_c} \sum_{\mu \in M^*} \kappa^*(x, \mu)^m \cdot w(x) \cdot \|\mu - \mu_w(C_\mu)\|_2^2 \\ & \quad + \frac{2}{c} \cdot \sum_{x \in X_c} \sum_{\mu \in M^*} \kappa^*(x, \mu)^m \cdot w(x) \cdot \|\mu_w(C_\mu) - \mu(S_\mu)\|_2^2 \\ & \hspace{10em} \text{(2-approximate triangle inequality)} \\ & \leq \text{fkm}(X, w, M^*, \kappa^*) + \frac{\epsilon}{2 \cdot c} \sum_{\mu \in M^*} \text{fkm}(X_c, w, \{\mu\}, \kappa^*) \\ & \quad + \frac{\epsilon}{c \cdot 16} \sum_{x \in X_c} \sum_{\mu \in M^*} \kappa^*(x, \mu)^m \cdot w(x) \cdot \frac{\text{km}(C_\mu, w)}{w(C_\mu)} \\ & = \left(1 + \frac{\epsilon}{2}\right) \cdot \text{fkm}(X, w, M^*, \kappa^*) + \frac{\epsilon}{c \cdot 16} \sum_{\mu \in M^*} \frac{\kappa^*(X_c, w, \mu) \cdot \text{km}(C_\mu, w)}{w(C_\mu)} \\ & \leq \left(1 + \frac{\epsilon}{2}\right) \cdot \text{fkm}(X, w, M^*, \kappa^*) + \frac{\epsilon}{c \cdot 8} \sum_{\mu \in M^*} \text{km}(C_\mu, w) \\ & \leq \left(1 + \frac{\epsilon}{2}\right) \cdot \text{fkm}(X, w, M^*, \kappa^*) + \frac{\epsilon}{c \cdot 2} \sum_{\mu \in M^*} \text{fkm}(X_c, w, \{\mu\}, \kappa^*) \\ & = (1 + \epsilon) \cdot \text{fkm}(X, w, M^*, \kappa^*) . \end{aligned}$$

Bounding the runtime is straightforward. The algorithm evaluates the cost of $|\mathcal{T}|^k$ different FKM solutions where each evaluation costs time $\mathcal{O}(d \cdot k \cdot n)$. Hence, the total runtime is bounded by

$$\mathcal{O}\left(d \cdot k \cdot n \cdot |\mathcal{T}|^k\right) = \mathcal{O}\left(d \cdot k \cdot n \cdot \left(n^{33/\epsilon}\right)^k\right) = d \cdot n^{\mathcal{O}(k/\epsilon)}.$$

7.3 Solving Small FRKM Instances

To the best of our knowledge, there is still no polynomial-time algorithm with a constant approximation guarantee for the FKM problem. One of the main motivations behind the definition of FRKM was to gain additional algorithmic insight into FKM-type problems. As it turns out however, just as with FKM, we still do not know of the existence of a polynomial-time constant-factor approximation algorithm for the radius problem, either. Albeit the similarly hard to analyze objective function, at least we are able to solve FRKM for two clusters on unweighted instances on the real line.

Algorithm 7.3.1: FUZZYRADIUSOLVER

Input: $X \subset \mathbb{R}$

- 1 **if** $|X| \leq 2$ **then**
- 2 **return** X
- 3 $x_{max} \leftarrow \max_{x \in X} \{x\}$
- 4 $x_{min} \leftarrow \min_{x \in X} \{x\}$
- 5 $x_a \leftarrow \arg \min_{x \in X} \left\{ \left| x - \frac{x_{max} + x_{min}}{2} \right| \right\}$
- 6 **return** $\left\{ \pm \sqrt{\left(x_a^2 + x_{max}^2 + \sqrt{x_a^4 + 14 \cdot x_a^2 \cdot x_{max}^2 + x_{max}^4} \right) / 6} \right\}$

Theorem 7.3.2 *Algorithm 7.3.1 optimally solves the FUZZY RADIUS k -MEANS problems with $d = 1$, $k = 2$, and $w = \mathbb{1}$ in time $\mathcal{O}(n)$.*

Proof. Let x_{max} , x_{min} , and x_a be defined as in Algorithm 7.3.1. Without loss of generality assume that the midpoint $(x_{max} + x_{min})/2$ of X is 0, i.e. $x_{max} = -x_{min}$.

We start by showing that x_{max} and x_{min} are among the most expensive points in each optimal solution. Let $M = \{\mu_1, \mu_2\}$ and assume, without loss of generality, that $x_{min} \leq \mu_1 \leq \mu_2 \leq x_{max}$ (recall that, by Lemma 3.2.13, optimal solutions have to lie in the convex hull of the input points). Assume x_{max} does not have the maximum cost among

7.3 Solving Small FRKM Instances

points in X . That is, there exists a point $x \in X$ with

$$\text{frkm}(\{x_{max}\}, M) < \text{frkm}(\{x\}, M) .$$

All points $\mu_2 \leq y \leq x_{max}$ are at most as far away from both representatives as x_{max} . Thus, their cost is at most the cost of x_{max} and we obtain $x < \mu_2$. Choose $\epsilon > 0$ small enough such that for $M' = \{\mu_1, \mu_2 - \epsilon\}$ we still have

$$\begin{aligned} & x < \mu_2 - \epsilon , \\ & \nexists z \in X : \mu_2 - \epsilon \leq z < \mu_2 \text{ and} \\ & \text{frkm}(\{x_{max}\}, M') < \text{frkm}(\{x\}, M') . \end{aligned}$$

Such an ϵ exists since the objective function is continuous. From this we obtain

$$\text{frkm}(\{z \in X \mid z < \mu_2 - \epsilon\}, M') < \text{frkm}(\{z \in X \mid z < \mu_2 - \epsilon\}, M)$$

since μ_2 got closer to these points. However, we still have

$$\text{frkm}(\{z \in X \mid z \geq \mu_2\}, M') \leq \text{frkm}(\{x_{max}\}, M') .$$

Since there are no points between $\mu_2 - \epsilon$ and μ_2 , we have that the overall cost of M' is strictly smaller than the cost of M , and thus M is not optimal.

Analogously, we can conclude that M is not optimal if x_{min} does not have the maximum cost among points in X . Combining these two observations, we obtain that if M is an optimal set of representatives, then

$$\text{frkm}(X, M) = \text{frkm}(\{x_{max}\}, M) = \text{frkm}(\{x_{min}\}, M) .$$

From this we can already conclude $x_{min} < \mu_1 \leq 0 \leq \mu_2 < x_{max}$. If μ_1 and μ_2 were both strictly larger (or smaller) than 0, then they were both closer to x_{max} than to x_{min} (or vice versa). In that case the cost of the two points would not be equal. If either $\mu_1 = x_{min}$ or $\mu_2 = x_{max}$, then both have to be true so that the overall cost is 0. In this case we have $|X| \leq 2$ because any third point in X would have cost strictly larger than 0. Thus, [Algorithm 7.3.1](#) would correctly answer by choosing the points in X as the optimal set of representatives.

For the sake of brevity we denote $u := x_{max}$ and $-u := x_{min} = -x_{max}$.

Observe that

$$\begin{aligned}
 0 &= \text{frkm}(\{-u\}, M) - \text{frkm}(\{u\}, M) \\
 &= \frac{1}{(-u - \mu_1)^{-2} + (-u - \mu_2)^{-2}} - \frac{1}{(u - \mu_1)^{-2} + (u - \mu_2)^{-2}} \\
 &= \frac{4 \cdot u \cdot (\mu_1 + \mu_2) \cdot (u^4 - 2 \cdot u^2 \cdot \mu_1 \cdot \mu_2 + \mu_1^3 \cdot \mu_2 + \mu_1 \cdot \mu_2^3 - \mu_1^2 \cdot \mu_2^2)}{((u - \mu_1)^2 + (u - \mu_2)^2) \cdot ((u + \mu_1)^2 + (u + \mu_2)^2)}.
 \end{aligned}$$

We differentiate the three cases for which the numerator vanishes. First, if $u = 0$, then $|X| = 1$, which is once again a trivial special case.

Second, let $u^4 - 2 \cdot u^2 \cdot \mu_1 \cdot \mu_2 + \mu_1^3 \cdot \mu_2 + \mu_1 \cdot \mu_2^3 - \mu_1^2 \cdot \mu_2^2 = 0$, and fix any u and $0 \leq \mu_2 < u$. Consider the polynomial as a function of μ_1

$$f(\mu_1) = \mu_2 \cdot \mu_1^3 - \mu_2^2 \cdot \mu_1^2 + (\mu_2^3 - 2 \cdot u^2 \cdot \mu_2) \cdot \mu_1 + u^4.$$

The discriminant of f is

$$\begin{aligned}
 \Delta &= - (3 \cdot \mu_2^4 - 14 \cdot \mu_2^2 \cdot u^2 + 27 \cdot u^4) \cdot (x + \mu_2)^2 \cdot (x - \mu_2)^2 \cdot \mu_2^2 \\
 &< - (3 \cdot \mu_2^4 + 13 \cdot u^4) \cdot (x + \mu_2)^2 \cdot (x - \mu_2)^2 \cdot \mu_2^2 \quad (\mu_2 < u) \\
 &< 0.
 \end{aligned}$$

Thus, f has one real and two non-real roots [Bronštein et al., 2008]. Furthermore, observe that

$$\begin{aligned}
 \lim_{\mu_1 \rightarrow \infty} f(\mu_1) &= \infty \text{ and} \\
 f(-u) &= u^4 + u^3 \cdot \mu_2 - u^2 \cdot \mu_2^2 - u \cdot \mu_2^3 > u^4 + u^3 \cdot \mu_2 - u^4 - u^3 \cdot \mu_2 = 0.
 \end{aligned}$$

We conclude that $\forall \mu_1 \geq x_{min} : f(\mu_1) > 0$, thus, is not 0 for any optimal solution.

This leaves us with $(\mu_1 + \mu_2) = 0$ as the third case. We obtain that, for all $x_{min} < \mu_1 \leq 0 \leq \mu_2 < x_{max}$ with $\text{frkm}(\{x_{max}\}, M) = \text{frkm}(\{x_{min}\}, M)$, we have $\mu_1 = -\mu_2$.

We established that all optimal solutions are of the form $M = \{\mu, -\mu\}$ for some $\mu \in \mathbb{R}_{\geq 0}$, $\mu < x_{max}$. Next, we examine for which value of μ the cost is minimized. As before, we still have that, for all $x \in X$ with $x_{min} < x \leq -\mu$ or $\mu \leq x < x_{max}$,

$$\text{frkm}(\{x\}, M) < \text{frkm}(X, M)$$

as these points are closer to both means than the two extremal points.

Fix any $\mu \in \mathbb{R}_{\geq 0}$ and consider the derivative of the objective function in the direction of x for any $x \neq \mu$, $x \neq -\mu$

$$\frac{\partial \text{frkm}}{\partial x}(\{x\}, M) = \frac{x \cdot (x - \mu) \cdot (x + \mu) \cdot (x^2 + 3 \cdot \mu^2)}{(x^2 + \mu^2)^2},$$

7.3 Solving Small FRKM Instances

which is negative for all $0 \leq x < \mu$ and positive for all $-\mu < x \leq 0$. Hence, points in X get more expensive the closer they are to 0 (in general: the closer they are to the midpoint of X). Recall the definition of x_a . If $x_a < \mu$, then, for all $x \in X$ with $x_a \leq |x| < \mu$, we have

$$\text{frkm}(\{x\}, M) < \text{frkm}(\{x_a\}, M) .$$

Due to symmetry, the actual sign of the closest point to 0 in X does not matter. We already know that $\text{frkm}(\{x_a\}, M)$ cannot be larger than $\text{frkm}(\{x_{max}\}, M)$. Assume that $\text{frkm}(\{x_a\}, M) < \text{frkm}(\{x_{max}\}, M)$ and consider the solution $M = \{\mu + \epsilon, -\mu - \epsilon\}$ for $\epsilon > 0$ small enough such that still

$$\text{frkm}(\{x_a\}, M') < \text{frkm}(\{x_{max}\}, M') .$$

Consider the derivative of the objective function for x_{max} in the direction of μ

$$\frac{\partial \text{frkm}}{\partial \mu}(\{x_{max}\}, M) = \frac{\mu \cdot (\mu - x_{max}) \cdot (\mu + x_{max}) \cdot (\mu^2 + 3 \cdot x_{max}^2)}{(\mu^2 + x_{max}^2)^2} ,$$

which is negative for all $0 \leq \mu < x_{max}$. From this we conclude that

$$\text{frkm}(\{x_{max}\}, M') < \text{frkm}(\{x_{max}\}, M) .$$

Hence, we obtain that, for optimal M ,

$$\text{frkm}(X, M) = \text{frkm}(\{x_{max}\}, M) = \text{frkm}(\{x_{min}\}, M) = \text{frkm}(\{x_a\}, M) .$$

Similar to before we consider

$$\begin{aligned} 0 &= \text{frkm}(\{x_a\}, M) - \text{frkm}(\{x_{max}\}, M) \\ &= \frac{1}{(x_a - \mu)^{-2} + (x_a + \mu)^{-2}} - \frac{1}{(x_{max} - \mu)^{-2} + (x_{max} + \mu)^{-2}} \\ &= \frac{(3 \cdot \mu^4 - \mu^2 \cdot x_a^2 - \mu^2 \cdot x_{max}^2 - x_a^2 \cdot x_{max}^2) \cdot (x_a + x_{max}) \cdot (x_{max} - x_a)}{2 \cdot (\mu^2 + x_a^2) \cdot (\mu^2 + x_{max}^2)} . \end{aligned}$$

The solutions to this are the roots of

$$g(\mu) = 3 \cdot \mu^4 - (x_a^2 + x_{max}^2) \cdot \mu^2 - x_a^2 \cdot x_{max}^2$$

because x_{max} and x_a are distinct and positive real numbers. Since g is biquadratic, its roots are

$$\mu = \pm \sqrt{\frac{x_a^2 + x_{max}^2 \pm \sqrt{(x_a^2 + x_{max}^2)^2 + 12 \cdot x_a^2 \cdot x_{max}^2}}{6}} .$$

Two of these are non-real because

$$\sqrt{(x_a^2 + x_{max}^2)^2 + 12 \cdot x_a^2 \cdot x_{max}^2} > x_a^2 + x_{max}^2 .$$

The other two are symmetric around 0, and hence, define the same solution M , which is also the set [Algorithm 7.3.1](#) outputs. Since this is the unique stationary point of the objective function with $x_{min} < -\mu \leq 0 \leq \mu < x_{max}$ and we know that there is an optimal solution with those properties, we obtain that M is a unique global minimum. ■

The downside of the previous proof is its rather explicit character. We closely examine the objective function using tools from real analysis which use the order on the real numbers to provide strong symmetry arguments. It is not clear if and how this could be generalized to a larger number of clusters or dimensions.

Coresets

Contribution Summary We present an algorithm constructing small coresets for FUZZY k -MEANS. The coreset construction in this chapter is based on a construction by [Chen \[2009\]](#) and is published in [\[Blömer et al., 2018\]](#). The application of coresets to approximation algorithms presented in [Section 8.3](#) yields an improvement over [\[Blömer et al., 2018\]](#) due to the improvement of the algorithm presented in [Section 7.2](#).

8.1	Related Work	120
8.2	Small Coresets for FKM	120
	8.2.1 From Weak to Strong via Non-Negligibility	122
	8.2.2 Weak Coresets for Non-Negligible Solutions	125
8.3	Applying FKM Coresets	132
8.4	Correctness of our Weak Coreset Algorithm	133

A central technique for dealing with clustering problems on huge data sets are *coresets*. Essentially, a coreset is a representation of a data set that preserves certain properties of the original data [\[Har-Peled & Mazumdar, 2004\]](#). In the context of this thesis, we formalize this by, given some number of clusters k , requiring the cost of any set of at most k representatives on the coreset to be close to the cost of the same set of representatives on the original data. Furthermore, the coresets we discuss in the following are weighted subsets of the original data. Such a coreset has two main applications. First, assume the original data set to be so large, that it does not fit into main memory. Constantly transferring data from main memory to hard drive, and vice versa, is time consuming. Thus, we usually want to process such a data set in a streaming setting, where we parse the data set only once. If we are able to maintain a coreset while parsing the stream and the resulting coreset is small enough so that it fits into main

memory after reading the whole stream, then we can apply standard algorithms to the coreset to find approximative solutions to the original data. Second, consider an approximation algorithm whose runtime is independent of any weight function on the data set, but only depends on the number of distinct point locations. For a small enough coreset, such an algorithm runs significantly faster on the coreset than on the original data. However, a good solution for the coreset is also a good solution on the original data. In the following, we present an efficient construction of small coresets for FKM. We also discuss how these can be applied in the two previously described scenarios (specifically, how our construction reduces the runtime of [Algorithm 7.2.1](#)).

8.1 Related Work

Coreset for KM have initially been introduced by [Bădoiu et al. \[2002\]](#). Since then, there has been a long string of research producing continuously improving results using different techniques, and constructing smaller and smaller coresets, for example, [Har-Peled & Mazumdar \[2004\]](#); [Har-Peled & Kushal \[2006\]](#); [Chen \[2009\]](#); [Feldman et al. \[2013\]](#), to name just a few. Currently, the smallest and most general coreset is due to [Sohler & Woodruff \[2018\]](#). They present a coreset of size $\mathcal{O}(\log(k/\epsilon))$ for any clustering and subspace approximation problem with, for any $p \geq 1$, the sum of Euclidean distances to the p^{th} power as the objective function.

There is a coreset construction for a large class of general soft clustering problems based on μ -similar Bregman divergences with size polynomial in k and d but independent of n [[Lucic et al., 2016](#)]. However, this construction can not be applied to FKM, and so far, there have been no other constructions of small coresets for FKM.

8.2 Small Coresets for FKM

We begin by formalizing our notion of FKM coresets.

Definition 8.2.1 (Coreset) *Let (X, w, k) be an FKM instance. A set $S \subseteq \mathbb{R}^d$ together with a weight function $w_S : S \rightarrow \mathbb{N}$ is called an ϵ -coreset (for the FKM problem) if, for all $M \subseteq \mathbb{R}^d$ with $|M| \leq k$,*

$$(1 - \epsilon) \cdot \text{fkm}(X, w, M) \leq \text{fkm}(S, w_S, M) \leq (1 + \epsilon) \cdot \text{fkm}(X, w, M) .$$

We often omit the explicit ϵ and denote an ϵ -coreset simply as a coreset. Furthermore, we sometimes refer to a coreset as a strong coreset.

8.2 Small Coresets for FKM

As it turns out, it is quite difficult to directly prove the existence of small coresets for FKM. The main pitfall being, once again, the algebraic structure of the FKM objective function. To still achieve our goal, we take a slight detour over so-called *weak coresets*. Consider some weighted set of points together with a set of candidate solutions. This triple (consisting of the set of points, weight function, and candidate solution set) is a weak coreset if the set of candidate solutions contains a good approximation to the original problem and the coreset property is fulfilled for all candidate solutions. We consider a set of candidate solutions instead of a set of candidate representatives, as it is done in the definition of weak coresets for KM [Feldman et al., 2007]. This slight generalization allows us to characterize solutions more precisely.

Definition 8.2.2 (Weak Coreset) *Let (X, w, k) be an FKM instance. A set $S \subseteq \mathbb{R}^d$ together with a weight function $w_S : S \rightarrow \mathbb{N}$ and a set of solutions $\Theta \subseteq \{\theta \in \mathbb{R}^d \mid |\theta| \leq k\}$ is called a weak ϵ -coreset (for the FKM problem) if there exists $M \in \Theta$ with*

$$\text{fkm}(X, w, M) \leq (1 + \epsilon) \cdot \text{fkm}(X, w, M^*)$$

where M^* is an optimal solution for FKM on (X, w, k) , and, for all $M \in \Theta$,

$$(1 - \epsilon) \cdot \text{fkm}(X, w, M) \leq \text{fkm}(S, w_S, M) \leq (1 + \epsilon) \cdot \text{fkm}(X, w, M).$$

As before, we often omit the explicit ϵ and denote a weak ϵ -coreset simply as a weak coreset.

In the following, we show how to efficiently construct small coreset for FKM, with high probability. To this end, we present two independent results. First, we argue that a weak coreset, with respect to so-called non-negligible solutions, which forms a subset of the original data and has an integral weight function is already a strong coreset. Second, we present an adaptation of Chen's coreset construction for KM [Chen, 2009] which computes a weak coreset with the desired properties, with high probability.

Chen's algorithm is neither the most efficient nor the smallest coreset construction there is for KM. However, we still use it since it is the best purely sampling based approach. Other techniques, such as ϵ -nets [Har-Peled & Kushal, 2006] or subspace techniques [Feldman et al., 2013] heavily rely on partitioning the input set. The issues with this is that the assignment function of FKM effectively introduces an unknown weighting of the input points. We do not know how to prevent the introduction of an error factor on the order $k^{\mathcal{O}(1)}$ to the cost estimation if we partition the input set (or project it into some subspace) in the way these algorithms do.

Theorem 8.2.3 *Let $d, k \in \mathbb{N}$, $X \subseteq \mathbb{R}^d$, and $\epsilon, \delta \in (0, 1)$. There is an algorithm which, given X, k, ϵ , and δ , outputs a set $S \subseteq X$ and $w_S : S \rightarrow \mathbb{N}$ such that (S, w_S) is a coresets for FKM on $(X, \mathbb{1}, k)$, with probability at least $1 - \delta$. The size of the coresets is*

$$|S| \in \mathcal{O}(d \cdot k^{4m} \cdot \log(n) \cdot \log(\log(n))^2 \cdot \epsilon^{-3} \cdot \log(\delta^{-1}))$$

and the algorithm runs in time $\mathcal{O}(d \cdot k \cdot n \cdot \log(\delta^{-1}) + |S|)$.

The algorithm constructing the coresets inherently expects an unweighted input set. We can trivially obtain a coresets for weighted input sets by giving the algorithm the weighted set as a multiset. However, all occurrences of n in the size of the coresets and the runtime are then replaced by $w(X)$. For this reason, in the following, we always assume X to be unweighted.

8.2.1 From Weak to Strong via Non-Negligibility

In contrast to hard clustering problems, an optimal FKM cluster is never empty. That is, if the assignment function is induced by the set of representatives, then every input point (which is different from all representatives) is always assigned to each representative by a non-trivial amount. Thus, given a set of representatives, each representative contributes non-trivially to the overall cost. It turns out that bounding the cost incurred by a representative with a small total assignment mass is rather difficult. We introduce the notion of *negligible* clusters to circumvent this problem.

Definition 8.2.4 (Negligible Fuzzy Cluster) *Let (X, w, k) be an FKM instance, $M \subseteq \mathbb{R}^d$, and $\kappa : X \times M \rightarrow [0, 1]$ be a feasible assignment function. We say that the cluster of a representative $\mu \in M$ is (k, ϵ) -negligible if, for all $x \in X$,*

$$\kappa(x, \mu) \leq \frac{\epsilon}{2 \cdot m \cdot k + \epsilon}.$$

We omit the parameters (k, ϵ) if they are clear from context.

We do not know whether optimal FKM solutions are in any way connected to non-negligible clusters. That is, we do not know if one could refute that there is a data set where all optimal solutions contain at least one negligible cluster. However, the total assignment mass of a negligible cluster is so small that we can simply remove the respective representative from the solution without significantly increasing the cost of the solution.

8.2 Small Coresets for FKM

Lemma 8.2.5 *Let (X, w, k) be an FKM instance, $M \subseteq \mathbb{R}^d$ with $|M| \leq k$, and $\epsilon \in (0, 1)$. There exists a set $M' \subseteq M$ with*

$$\text{fkm}(X, w, M') \leq (1 + \epsilon) \cdot \text{fkm}(X, w, M)$$

such that there is an optimal assignment function $\kappa' : X \times M' \rightarrow [0, 1]$ for FKM on (X, w) that contains no negligible cluster.

Proof. We present a somewhat backward inductive argument. Let $\tilde{M} \subseteq \mathbb{R}^d$ with $|\tilde{M}| \leq k$, and let $\tilde{\kappa} : X \times \tilde{M} \rightarrow [0, 1]$ be an optimal assignment function for FKM on (X, w) . Assume that the cluster of $\mu \in \tilde{M}$ is (k, ϵ) -negligible. Set $\hat{M} := \tilde{M} \setminus \{\mu\}$ and define an assignment function

$$\begin{aligned} \hat{\kappa} : X \times \hat{M} &\rightarrow [0, 1] \\ (x, \hat{\mu}) &\mapsto \tilde{\kappa}(x, \hat{\mu}) + \frac{\tilde{\kappa}(x, \mu) \cdot \tilde{\kappa}(x, \hat{\mu})}{1 - \tilde{\kappa}(x, \mu)}. \end{aligned}$$

Observe that this is well-defined since $\tilde{\kappa}(x, \mu) < 1$ because of $\tilde{\kappa}$'s negligibility. For all $x \in X$, we have

$$\begin{aligned} \sum_{\hat{\mu} \in \hat{M}} \hat{\kappa}(x, \hat{\mu}) &= \sum_{\hat{\mu} \in \hat{M}} \frac{(1 - \tilde{\kappa}(x, \mu)) \cdot \tilde{\kappa}(x, \hat{\mu}) + \tilde{\kappa}(x, \mu) \cdot \tilde{\kappa}(x, \hat{\mu})}{1 - \tilde{\kappa}(x, \mu)} \\ &= \sum_{\hat{\mu} \in \hat{M}} \frac{\tilde{\kappa}(x, \hat{\mu})}{1 - \tilde{\kappa}(x, \mu)} = \frac{\sum_{\hat{\mu} \in \hat{M}} \tilde{\kappa}(x, \hat{\mu})}{\sum_{\tilde{\mu} \in \tilde{M} \setminus \{\mu\}} \tilde{\kappa}(x, \tilde{\mu})} = 1, \end{aligned}$$

and hence, $\hat{\kappa}$ is a feasible assignment function. Furthermore, we can upper bound our new assignment function showing that the assignment of each point to the remaining representatives has not increased significantly. For all $x \in X$ and $\hat{\mu} \in \hat{M}$, we have

$$\begin{aligned} \hat{\kappa}(x, \hat{\mu}) &= \tilde{\kappa}(x, \hat{\mu}) + \frac{\tilde{\kappa}(x, \mu) \cdot \tilde{\kappa}(x, \hat{\mu})}{1 - \tilde{\kappa}(x, \mu)} \leq \tilde{\kappa}(x, \hat{\mu}) + \frac{\epsilon}{2 \cdot m \cdot k + \epsilon} \cdot \frac{\tilde{\kappa}(x, \hat{\mu})}{1 - \frac{\epsilon}{2 \cdot m \cdot k + \epsilon}} \\ &= \tilde{\kappa}(x, \hat{\mu}) \cdot \left(1 + \frac{\epsilon}{2 \cdot m \cdot k + \epsilon} \cdot \frac{2 \cdot m \cdot k + \epsilon}{2 \cdot m \cdot k}\right) = \left(1 + \frac{\epsilon}{2 \cdot m \cdot k}\right) \cdot \tilde{\kappa}(x, \hat{\mu}). \end{aligned}$$

Hence, we can bound the cost of \hat{M} as follows

$$\begin{aligned} \text{fkm}(X, w, \hat{M}) &\leq \sum_{x \in X} \sum_{\hat{\mu} \in \hat{M}} \hat{\kappa}(x, \hat{\mu})^m \cdot w(x) \cdot \|x - \hat{\mu}\|_2^2 \\ &\leq \sum_{x \in X} \sum_{\hat{\mu} \in \hat{M}} \left(1 + \frac{\epsilon}{2 \cdot m \cdot k}\right)^m \cdot \tilde{\kappa}(x, \hat{\mu})^m \cdot w(x) \cdot \|x - \hat{\mu}\|_2^2 \\ &< \left(1 + \frac{\epsilon}{2 \cdot m \cdot k}\right)^m \cdot \text{fkm}(X, w, \tilde{M}). \quad (\hat{M} = \tilde{M} \setminus \{\mu\}) \end{aligned}$$

At first, we set $\tilde{M} := M$. We repeatedly remove representatives with (k, ϵ) -negligible clusters, setting $\tilde{M} := \tilde{M}$ after each removal, until we arrive at a set M' which has no negligible clusters. The set M' is not empty because if there is only a single representative remaining, then that cluster cannot be negligible. The previously presented cost bound holds for each removal of a representative independently and we remove less than k representatives. Thus,

$$\text{fkm}(X, w, M') \leq \left(1 + \frac{\epsilon}{2 \cdot m \cdot k}\right)^{m \cdot k} \cdot \text{fkm}(X, w, M) \leq (1 + \epsilon) \cdot \text{fkm}(X, w, M) .$$

■

Recall that, given some set of representatives, an optimal assignment of a point depends only on the location of that point relative to the representatives. More specifically, it neither depends on the weight of the point nor on the location of other points in the data set. This means that if a cluster is negligible with respect to some data set, then it is negligible with respect to all subsets of that data set. Simply speaking, being negligible is preserved when we take a subset of the data. This is the key observation in the proof of our weak-to-strong lemma.

Lemma 8.2.6 (Weak-To-Strong) *Let (X, w, k) be an FKM instance, $\epsilon \in (0, 1)$, and*

$$\Theta_{(k, \epsilon)}(X) := \left\{ M \in \mathbb{R}^d \mid \begin{array}{l} |M| \leq k \text{ and } M \text{ has no negligible cluster} \\ \text{with respect to } X \end{array} \right\} .$$

If $S \subseteq X$ and $w_S : S \rightarrow \mathbb{N}$ are such that $(S, w_S, \Theta_{(k, \epsilon)}(X))$ is weak ϵ -coreset for FKM on (X, w, k) , then (S, w_S) is a strong $(3 \cdot \epsilon)$ -coreset for FKM on (X, w, k) .

Proof. We have to verify that the coreset property holds with a factor of $3 \cdot \epsilon$, for all $M \subseteq \mathbb{R}^d$ with $|M| \leq k$. Since $(S, w_S, \Theta_{(k, \epsilon)}(X))$ is a weak coreset, this trivially holds for $M \in \Theta_{(k, \epsilon)}(X)$. Thus, we only need to check solutions $M \notin \Theta_{(k, \epsilon)}(X)$. Consequently, fix some solution $M \subseteq \mathbb{R}^d$ with $M \notin \Theta_{(k, \epsilon)}(X)$.

From [Lemma 8.2.5](#), we obtain that there exists $M' \in \Theta_{(k, \epsilon)}(X)$ with $M' \subseteq M$ and

$$\text{fkm}(X, w, M') \leq (1 + \epsilon) \cdot \text{fkm}(X, w, M) .$$

8.2 Small Coresets for FKM

We immediately conclude the upper bound of the coreset property.

$$\begin{aligned}
 \text{fkm}(S, w_S, M) &\leq \text{fkm}(S, w_S, M') && (M' \subseteq M) \\
 &\leq (1 + \epsilon) \cdot \text{fkm}(X, w, M') && \text{(weak coreset property)} \\
 &\leq (1 + \epsilon)^2 \cdot \text{fkm}(X, w, M) \leq (1 + 3 \cdot \epsilon) \cdot \text{fkm}(X, w, M) .
 \end{aligned}$$

The lower bound is more involved. Again, we apply [Lemma 8.2.5](#), but to S . There exists $M'_S \in \Theta_{(k, \epsilon)}(S)$ with $M'_S \subseteq M$ and

$$\text{fkm}(S, w_S, M'_S) \leq (1 + \epsilon) \cdot \text{fkm}(S, w_S, M) .$$

A negligible cluster is also negligible for all subsets of the data set. That is, if there is a representative μ whose cluster is negligible with respect to X , then there is no point in X whose assignment to μ is more than some constant. Thus, there is also no point in $S \subseteq X$ whose assignment to μ is more than that constant. This means that the cluster of μ is also negligible with respect to S . Since $M' \in \Theta_{(k, \epsilon)}(X)$, it holds that the clusters of all representatives in $M \setminus M'$ are also negligible with respect to S , and thus, $M'_S \subseteq M'$. We conclude the proof by observing

$$\begin{aligned}
 \text{fkm}(S, w_S, M) &\geq \frac{1}{1 + \epsilon} \cdot \text{fkm}(S, w_S, M'_S) \\
 &\geq \frac{1}{1 + \epsilon} \cdot \text{fkm}(S, w_S, M') && (M'_S \subseteq M') \\
 &\geq \frac{1 - \epsilon}{1 + \epsilon} \cdot \text{fkm}(X, w, M') && \text{(weak coreset property)} \\
 &\geq \frac{1 - \epsilon}{1 + \epsilon} \cdot \text{fkm}(X, w, M) && (M' \subseteq M) \\
 &\geq (1 - 3 \cdot \epsilon) \cdot \text{fkm}(X, w, M) .
 \end{aligned}$$

■

Due to the weak-to-strong lemma, we only need to construct weak coresets for solutions without negligible clusters. This helps us with the initially mentioned problem: bounding the cost of representatives which are far away from all input points.

8.2.2 Weak Coresets for Non-Negligible Solutions

In the following, we show how to adapt Chen's coreset construction for KM [[Chen, 2009](#)] to obtain a pair (S, w_S) which fulfils the requirements of [Lemma 8.2.6](#). The core idea is easily explained. We take Chen's

original construction and adapt some parameters with two main goals. First, we need to compensate for the factor $1/k^{m-1}$ incurred by using KM bounds to bound the FKM objective function (see [Lemma 3.2.12](#)). Second, we use non-negligibility to obtain a large lower bound for representatives which are far away from all input points. However, as usual with FKM, actually showing that this is sufficient requires careful analysis.

Preliminaries Before proceeding to the actual construction, we introduce the following notations and results. For some $c \in \mathbb{R}^d$ and $r \in \mathbb{R}$ denote by $\mathcal{B}(c, r) := \{x \in \mathbb{R}^d \mid \|x - c\|_2 \leq r\}$ the Euclidean ball with radius r around c . Let (X, w, k) be an instance of any of our clustering problems. We call a set of representatives $M \subseteq \mathbb{R}^d$ an (α, β) -*bicriteria approximation* if the cost of M are at most a factor α worse than the optimal cost and $|M| \leq \beta \cdot k$. Furthermore, we apply the following concentration bound on the average function value of uniform sample sets.

Lemma 8.2.7 (Haussler [1992]) *Let $P \subseteq \mathbb{R}^d$, $\epsilon, \delta \in (0, 1)$, $f : X \rightarrow \mathbb{R}$, and $F \in \mathbb{R}$ such that, for all $x \in X$, we have $0 \leq f(x) \leq F$. If S is a set of at least $(1/2) \cdot \epsilon^{-2} \cdot \ln(2/\delta)$ points drawn uniformly at random from P , then*

$$\Pr \left(\left| \frac{f(P)}{|P|} - \frac{f(S)}{|S|} \right| \leq \epsilon \cdot F \right) \geq 1 - \delta.$$

Along the lines of the original proof, we first show how to construct a set $S \subseteq X$ and a weight function $w_S : S \rightarrow \mathbb{N}$ which fulfil the coresets property for a finite number of solutions.

Notice that [Algorithm 8.2.8](#) does not need to know for which solutions the coresets property has to hold but only for how many. Strictly speaking, the weight function w_S , as the algorithm defines it, is not integral, as is. However, this is just a minor technical issue. For each $X_{a,j}$ we take an arbitrary subset whose size is a multiple of q and sample from among these. This leaves strictly less than q points which are not considered during sampling. After sampling q points we add the points which were not considered to $S_{a,j}$. Trivially, the coresets property is fulfilled for all these points, and we only have to verify that this also holds for the points we actually sampled from. This approach increases the size of the coresets by at most a factor 2. Hence, in the following, we assume that $|X_{a,j}|$ is always some multiple of q .

8.2 Small Coresets for FKM

Algorithm 8.2.8: CHEN'S SAMPLING

Input: $X \in \mathbb{R}^d$, $k, \gamma \in \mathbb{N}, \alpha, \beta \in \mathbb{R}_{\geq 1}$, an (α, β) -bicriteria approximation $A \subset \mathbb{R}^d$ of KM on X , $\epsilon, \delta \in (0, 1)$

- 1 $F \leftarrow \lceil \frac{1}{2} \cdot \log(\alpha \cdot n) \rceil$
- 2 $R \leftarrow \sqrt{\text{km}(X, \mathbb{1}, A) / (\alpha \cdot n)}$
- 3 $q \leftarrow \mathfrak{q} \cdot (\alpha \cdot k^{m-1} / \epsilon)^2 \cdot \ln(4 \cdot \beta \cdot k \cdot F \cdot (\gamma + 1)^k / \delta)$
/* where \mathfrak{q} is a sufficiently large constant */
- 4
- 5 **for** $a \in A$ **and** $j \in [F]_0$ **do**
- 6 **if** $j = 0$ **then**
- 7 $L_{a,j} \leftarrow \mathcal{B}(a, R)$
- 8 **else**
- 9 $L_{a,j} \leftarrow \mathcal{B}(a, 2^j \cdot R) \setminus \mathcal{B}(a, 2^{j-1} \cdot R)$
- 10 $X_{a,j} \leftarrow L_{a,j} \cap \mathcal{C}_a^{(X,A)}$
- 11 **if** $X_{a,j} \neq \emptyset$ **then**
- 12 $S_{a,j} \leftarrow \emptyset$
- 13 **for** $i \in [q]$ **do**
- 14 Sample x uniformly at random from $X_{a,j}$
- 15 $S_{a,j} \leftarrow S_{a,j} + \{x\}$
- 16 $w_S(x) \leftarrow |X_{a,j}| / q$
- 17 **return** $(\bigcup_{a \in A, j \in [F]_0} S_{a,j}, w_S)$

Lemma 8.2.9 For each $\Gamma \in \mathbb{R}^d$ with $|\Gamma| \leq \gamma$, the output $S \in X$, $w_S : S \rightarrow \mathbb{N}$ of Algorithm 8.2.8 fulfils $w_S(S) = n$ and, with probability at least $1 - \delta$, for all $M \in \Gamma$ with $|M| \leq k$,

$$(1 - \epsilon) \cdot \text{fkm}(X, \mathbb{1}, M) \leq \text{fkm}(S, w_S, M) \leq (1 + \epsilon) \cdot \text{fkm}(X, \mathbb{1}, M)$$

$$\left(1 - \frac{\epsilon}{k^{m-1}}\right) \cdot \text{km}(X, \mathbb{1}, M) \leq \text{km}(S, w_S, M) \leq \left(1 + \frac{\epsilon}{k^{m-1}}\right) \cdot \text{km}(X, \mathbb{1}, M).$$

Proof. We start by arguing that the sets $X_{a,j}$ form a partition of X . Observe that the $X_{a,j}$ are pairwise disjoint subsets of X since $X_{a,j} \subseteq \mathcal{C}_a^{(X,A)} \subseteq X$ and the clusters $\mathcal{C}_a^{(X,A)}$ are pairwise disjoint by definition. Additionally, we have to show that the union $\bigcup_{a \in A, j \in [F]_0} X_{a,j}$ contains X . If $y \in \mathbb{R}^d$ but $y \notin \bigcup_{a \in A, j \in [F]_0} X_{a,j}$, then

$$\min_{a \in A} \{\|y - a\|_2\} > 2^F \cdot R = \sqrt{\alpha \cdot n} \cdot \sqrt{\text{km}(X, \mathbb{1}, A) / (\alpha \cdot n)} = \sqrt{\text{km}(X, \mathbb{1}, A)}.$$

However, for all $x \in X$, we have

$$\begin{aligned} \min_{a \in A} \left\{ \|x - a\|_2^2 \right\} &= \text{km}(\{x\}, \mathbb{1}, A) \\ &\leq \text{km}(X, \mathbb{1}, A) . \end{aligned}$$

Hence, $y \notin X$ and $X \subseteq \bigcup_{a \in A, j \in [F]_0} X_{a,j}$. Similarly, since the $S_{a,j}$ are sampled from $X_{a,j}$, they are pairwise disjoint and thus form a partition of S . Further notice that each $S_{a,j}$ contains q points with weight $|X_{a,j}|/q$, and thus, $w_S(S_{a,j}) = |X_{a,j}|$ and $w_S(S) = n$.

Fix an arbitrary $\Gamma \in \mathbb{R}^d$ with $|\Gamma| \leq \gamma$ and some $M \in \Gamma$ with $|M| \leq k$. Since the $X_{a,j}$ form a partition, we apply the triangle inequality to obtain

$$\begin{aligned} &|\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(S, w_S, M)| \\ &\leq \sum_{a \in A} \sum_{j=0}^F |\text{fkm}(X_{a,j}, \mathbb{1}, M) - \text{fkm}(S_{a,j}, w_S, M)| . \end{aligned}$$

We analyze each of these summands individually. To this end, fix some $a \in A$ and $j \in [F]_0$. Again by definition of the weights, we observe

$$\begin{aligned} \frac{1}{|X_{a,j}|} \cdot \text{fkm}(S_{a,j}, w_S, M) &= \frac{1}{|X_{a,j}|} \cdot \sum_{s \in S_{a,j}} w_S(s) \cdot \text{fkm}(\{s\}, \mathbb{1}, M) \\ &= \frac{1}{q} \cdot \text{fkm}(S_{a,j}, \mathbb{1}, M) . \end{aligned}$$

Using this, we rewrite the cost difference as a difference between the scaled cost of two unweighted sets

$$\begin{aligned} &|\text{fkm}(X_{a,j}, \mathbb{1}, M) - \text{fkm}(S_{a,j}, w_S, M)| \\ &= |X_{a,j}| \left| \frac{1}{|X_{a,j}|} \cdot \text{fkm}(X_{a,j}, \mathbb{1}, M) - \frac{1}{|X_{a,j}|} \cdot \text{fkm}(S_{a,j}, w_S, M) \right| \\ &= |X_{a,j}| \cdot \left| \frac{1}{|X_{a,j}|} \cdot \text{fkm}(X_{a,j}, \mathbb{1}, M) - \frac{1}{q} \cdot \text{fkm}(S_{a,j}, \mathbb{1}, M) \right| . \end{aligned}$$

Consider $\text{fkm}(\{\cdot\}, \mathbb{1}, M)$ as a function in a single point $y \in \mathbb{R}^d$. For each point $y \in \mathbb{R}^d$ we have $\text{fkm}(\{y\}, \mathbb{1}, M) \geq 0$. For

$$x_{a,j} := \arg \min_{x \in X_{a,j}} \left\{ \min_{\mu \in M} \{ \|x - \mu\|_2 \} \right\}$$

8.2 Small Coresets for FKM

and each $x \in X_{a,j}$, we can bound

$$\begin{aligned}
\text{fkm}(\{y\}, \mathbb{1}, M) &\leq \min_{\mu \in M} \left\{ \|x - \mu\|_2^2 \right\} \\
&\leq 2 \cdot \left(\min_{\mu \in M} \left\{ \|x_{a,j} - \mu\|_2^2 \right\} + \|x - x_{a,j}\|_2^2 \right) \\
&\hspace{10em} \text{(2-approximate triangle inequality)} \\
&\leq 4 \cdot \left(\min_{\mu \in M} \left\{ \|x_{a,j} - \mu\|_2^2 \right\} + \|x - a\|_2^2 + \|a - x_{a,j}\|_2^2 \right) \\
&\hspace{10em} \text{(2-approximate triangle inequality)} \\
&\leq 4 \cdot \left(\min_{\mu \in M} \left\{ \|x_{a,j} - \mu\|_2^2 \right\} + 2^{2 \cdot j+1} \cdot R^2 \right) . \\
&\hspace{10em} (X_{a,j} \subseteq L_{a,j} \subseteq \mathcal{B}(a, 2^j \cdot R))
\end{aligned}$$

Algorithm 8.2.8 sets the number of uniform samples q such that we can apply **Lemma 8.2.7** to $X_{a,j}$, $S_{a,j}$,

$$\epsilon' := \frac{\epsilon}{\mathfrak{e} \cdot \alpha \cdot k^{m-1}} , \text{ and } \delta' := \frac{\delta}{2 \cdot \beta \cdot k \cdot F \cdot (\gamma + 1)^k} ,$$

for some large enough constant \mathfrak{e} , which yields that, with probability at least $1 - \delta'$,

$$\begin{aligned}
&|\text{fkm}(X_{a,j}, \mathbb{1}, M) - \text{fkm}(S_{a,j}, w_S, M)| \\
&\leq 4 \cdot \epsilon' \cdot |X_{a,j}| \cdot \left(\min_{\mu \in M} \left\{ \|x_{a,j} - \mu\|_2^2 \right\} + 2^{2 \cdot j+1} \cdot R^2 \right) .
\end{aligned}$$

We expand the term on the right hand side and bound the two summands separately.

By definition of $x_{a,j}$ we obtain

$$|X_{a,j}| \cdot \min_{\mu \in M} \left\{ \|x_{a,j} - \mu\|_2^2 \right\} \leq \sum_{x \in X_{a,j}} \min_{\mu \in M} \left\{ \|x - \mu\|_2^2 \right\} = \text{km}(X_{a,j}, \mathbb{1}, M) .$$

For the second term, we differentiate two different cases depending on j . First, let $j = 0$ and observe

$$\begin{aligned}
|X_{a,j}| \cdot 2^{2 \cdot j+1} \cdot R^2 &= 2 \cdot |X_{a,j}| \cdot R^2 = \frac{2}{\alpha} \cdot \frac{|X_{a,j}|}{n} \cdot \text{km}(X, \mathbb{1}, A) \\
&\leq 2 \cdot \frac{|X_{a,j}|}{n} \cdot \text{km}(X, \mathbb{1}, A) . \hspace{10em} (\alpha \geq 1)
\end{aligned}$$

Second, let $j \geq 1$ and recall that $X_{a,j} \subseteq L_{a,j} = \mathcal{B}(a, 2^j \cdot R) \setminus \mathcal{B}(a, 2^{j-1} \cdot R)$. Hence, for all $x \in X_{a,j}$, we observe

$$2^{2 \cdot j-2} \cdot R^2 = (2^{j-1} \cdot R)^2 \leq \|x - a\|_2^2 = \text{km}(\{x\}, \mathbb{1}, A) ,$$

and thus,

$$|X_{a,j}| \cdot 2^{2 \cdot j+1} \cdot R^2 \leq 8 \cdot \text{km}(X_{a,j}, \mathbb{1}, A) .$$

We combine the three upper bounds to obtain

$$\begin{aligned} & |\text{fkm}(X_{a,j}, \mathbb{1}, M) - \text{fkm}(S_{a,j}, w_S, M)| \\ & \leq 4 \cdot \epsilon' \cdot (\text{km}(X_{a,j}, \mathbb{1}, M) \\ & \quad + 2 \cdot \frac{|X_{a,j}|}{n} \cdot \text{km}(X, \mathbb{1}, A) + 8 \cdot \text{km}(X_{a,j}, \mathbb{1}, A)) , \end{aligned} \quad (8.2.10)$$

still with probability at least $1 - \delta'$.

We apply the union bound to each event for $a \in A$ and $j \in [F]_0$. That is, with probability at least $1 - \delta/(2 \cdot (\gamma + 1)^k)$, the upper bound (8.2.10) holds simultaneously for all $X_{a,j}$. Recall that the $X_{a,j}$ form a partition, and thus, we can take the sum to obtain

$$\begin{aligned} & \sum_{a \in A} \sum_{j=0}^F |\text{fkm}(X_{a,j}, \mathbb{1}, M) - \text{fkm}(S_{a,j}, w_S, M)| \\ & \leq 4 \cdot \epsilon' \cdot (\text{km}(X, \mathbb{1}, M) + 10 \cdot \text{km}(X, \mathbb{1}, A)) . \end{aligned}$$

Recall that, since $|M| \leq k$,

$$\begin{aligned} \text{km}(X, \mathbb{1}, M) & \leq k^{m-1} \cdot \text{fkm}(X, \mathbb{1}, M) \text{ and} \\ \text{km}(X, \mathbb{1}, A) & \leq \alpha \cdot \text{km}(X, \mathbb{1}, M) \leq \alpha \cdot k^{m-1} \cdot \text{fkm}(X, \mathbb{1}, M) . \end{aligned}$$

We conclude

$$\begin{aligned} |\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(S, w_S, M)| & \leq 4 \cdot \epsilon' \cdot k^{m-1} \cdot (1 + 10 \cdot \alpha) \cdot \text{fkm}(X, \mathbb{1}, M) \\ & \leq \epsilon \cdot \text{fkm}(X, \mathbb{1}, M) . \end{aligned}$$

We apply the union bound once more, this time to each set $M \in \Gamma$ with $|M| \leq k$. There are less than $(\gamma + 1)^k$ of these sets and hence our upper bound holds simultaneously for each of these sets of representatives with probability at least $1 - \delta/2$.

Following the same line of arguments, we can show that, also with probability at least $1 - \delta/2$, we have, for all $M \in \Gamma$ with $|M| \leq k$,

$$|\text{km}(X, \mathbb{1}, M) - \text{km}(S, w_S, M)| \leq \frac{\epsilon}{k^{m-1}} \cdot \text{km}(X, \mathbb{1}, M) .$$

Using the union bound one final time, we conclude the proof by observing that the claimed properties hold simultaneously, with probability at least $1 - \delta$. \blacksquare

8.2 Small Coresets for FKM

Next, we present the algorithm which computes a weak coreset fulfilling the requirements of [Lemma 8.2.6](#). Basically, we apply [Algorithm 8.2.8](#) to appropriately chosen parameters. To this end, we fix any (α, β) -bicriteria approximation for KM with $\alpha, \beta \in \mathcal{O}(1)$, a significantly reduced ϵ , and a large enough γ . The main technical challenge is to prove the existence of a bounded size set of representatives such that if the coreset property holds for solutions chosen from among these representatives, then it holds for all solutions without negligible clusters.

Algorithm 8.2.11: FUZZY k -MEANS CORESET

Input: $X \subseteq \mathbb{R}^d$, $k \in \mathbb{N}$, $\epsilon, \delta \in (0, 1)$

- 1 Apply the algorithm presented by [Aggarwal et al. \[2009\]](#) to compute, with probability $1 - \delta/2$, a $(20, 32)$ -bicriteria approximation $A \subset \mathbb{R}^d$ of KM on $(X, \mathbb{1}, k)$.
- 2 $\tilde{\epsilon} \leftarrow \epsilon / (\mathbf{a} \cdot k^{m-1})$
/* where \mathbf{a} is a sufficiently large constant */
- 3
- 4 $\gamma \leftarrow k \cdot (1/2 \cdot \log(\mathbf{b} \cdot n / (\tilde{\epsilon}^2 \cdot (\epsilon / (2 \cdot m \cdot k + \epsilon))^m) + 1)) \cdot (\mathbf{c} / \tilde{\epsilon})^d$
/* where \mathbf{b} and \mathbf{c} are sufficiently large constants */
- 5
- 6 $(S, w_S) \leftarrow$ [Algorithm 8.2.8](#) applied to X , k , γ , 20, 32, $\tilde{\epsilon}$, and $\delta/2$
- 7 **return** (S, w_S)

Lemma 8.2.12 *Let $(X, \mathbb{1}, k)$ be an FKM instance. [Algorithm 8.2.11](#) computes a set $S \subseteq X$ and a function $w_S : S \rightarrow \mathbb{N}$ such that $(S, w_S, \Theta_{(k, \epsilon)}(X))$ is a weak ϵ -coreset for the FKM problem on $(X, \mathbb{1}, k)$, with probability at least $1 - \delta$.*

To increase readability we moved the, rather technical, proof of [Lemma 8.2.12](#) to [Section 8.4](#).

Proof of [Theorem 8.2.3](#). We invoke [Algorithm 8.2.11](#). [Lemma 8.2.12](#) gives us the correctness. This leaves us with arguing the size of the coreset and analyzing the runtime.

First, we bound $|S|$ in terms of $\tilde{\epsilon}$ and γ . S is the union of $32 \cdot k \cdot F$ sets of size at most $2 \cdot q$. By definition, we have $F \in \mathcal{O}(\log(n))$, and thus,

$$\begin{aligned} q &\in \mathcal{O}\left(\left(k^{m-1}/\tilde{\epsilon}\right)^2 \cdot \log\left(k \cdot \log(n) \cdot \gamma^k/\delta\right)\right) \\ &\subseteq \mathcal{O}\left(k^{2 \cdot m} \cdot \tilde{\epsilon}^{-2} \cdot \log(\gamma) \cdot \log(\log(n)) \cdot \log(\delta^{-1})\right) . \end{aligned}$$

Second, we analyze our choice of parameters. We set $\tilde{\epsilon} \in \mathcal{O}(\epsilon/k^{m-1})$ and

$$\begin{aligned} \log(\gamma) &\in \mathcal{O}(\log(k) + \log(\log(n) + \log(\tilde{\epsilon}^{-1}) + \log(m \cdot k/\epsilon)) + d \cdot \log(\tilde{\epsilon}^{-1})) \\ &\subseteq \mathcal{O}(d \cdot \log(k) \cdot \log(k/\epsilon) \cdot \log(\log(k/\epsilon)) \cdot \log(\log(n))) \end{aligned}$$

Putting this together we obtain the claimed size.

For the runtime, note that the $(20, 32)$ -bicriteria algorithm proposed by [Aggarwal et al. \[2009\]](#) takes time $\mathcal{O}(d \cdot k \cdot n \cdot \log(\delta^{-1}))$ to achieve the desired success probability. Next, we analyze the runtime of [Algorithm 8.2.8](#). Given some $x \in X$, we can determine $j \in [F]_0$ such that $x \in L_{a,j}$ by simply computing $\lceil \log(\|x - A_x\|_2/R) \rceil$. Hence, computing all $X_{a,j}$ takes time $\mathcal{O}(d \cdot k \cdot n)$. Afterwards, it is possible to sample the $|S|$ points in linear time. ■

8.3 Applying FKM Coresets

There are two main applications for coresets. One is the streaming setting where we want to maintain a coreset of an insertion-only stream of points (this is essentially equivalent to parsing the input once). There is a general framework by [Feldman et al. \[2013\]](#), which maintains a coreset of a stream, given a coreset construction fulfilling some formal properties. One can easily check that our construction fulfils these properties. We will not go into details on how the framework works since we are not really interested in this slightly more practical approach. Furthermore, we do not need to add anything to the already well-researched framework to make it work for our case.

We are more interested in the application of a coreset to speed up approximation algorithms for FKM – more specifically, [Algorithm 7.2.1](#). Recall that this algorithm's runtime is independent of the weight function on the points. Since our FKM coreset significantly reduces the number of different points in the input set, we obtain the, so far, fastest known $(1 + \epsilon)$ -approximation algorithm for FKM on unweighted input sets by combining [Algorithm 7.2.1](#) with our coreset construction.

Theorem 8.3.1 *There exists an algorithm which, given $k \in \mathbb{N}$, $X \subseteq \mathbb{R}^d$, and $\epsilon \in (0, 1)$, computes a $(1 + \epsilon)$ -approximation of the FUZZY k -MEANS problem, with constant probability, and in time*

$$\mathcal{O}(d \cdot k \cdot n) + (d \cdot k \cdot \log(n)/\epsilon)^{\mathcal{O}(k/\epsilon)} .$$

Proof. We call [Algorithm 8.2.11](#) with k , X , and $\epsilon/3$ to obtain, with constant probability, an $\epsilon/3$ -coreset (S, w_S) of $(X, \mathbb{1})$. Let M be the

8.4 Correctness of our Weak Coreset Algorithm

output of [Algorithm 7.2.1](#), given S , w_S , and $\epsilon/3$. Let M_S^* be an optimal set of representatives with respect to S , and M_X^* be an optimal set of representatives with respect to X . We bound

$$\begin{aligned} \text{fkm}(S, w_S, M) &\leq (1 + \epsilon/3) \cdot \text{fkm}(S, w_S, M_S^*) \leq (1 + \epsilon/3) \cdot \text{fkm}(S, w_S, M_X^*) \\ &\leq (1 + \epsilon/3)^2 \cdot \text{fkm}(X, \mathbb{1}, M_X^*) \leq (1 + \epsilon) \cdot \text{fkm}(X, \mathbb{1}, M_X^*) . \end{aligned}$$

The overall runtime is

$$\begin{aligned} &\mathcal{O}(d \cdot k \cdot n + |S|) + d \cdot |S|^{\mathcal{O}(k/\epsilon)} \\ &\subseteq \mathcal{O}(d \cdot k \cdot n) + d \cdot (d \cdot k^{4 \cdot m} \cdot \log(n) \cdot \log(\log(n))^2 \cdot \epsilon^{-3})^{\mathcal{O}(k/\epsilon)} \\ &\subseteq \mathcal{O}(d \cdot k \cdot n) + (d \cdot k \cdot \log(n)/\epsilon)^{\mathcal{O}(k/\epsilon)} \end{aligned}$$

■

Notice that the algorithm from [Theorem 8.3.1](#) can also be applied to weighted data sets. We argued that our PTAS [Algorithm 7.2.1](#) has runtime independent of any weight function. However, this is not true for our coreset construction. Again, to compute our coreset for a weighted set, we need to consider the weighted input as a set, and hence, all occurrences of n in the runtime are replaced by $w(X)$.

Strictly speaking, applying [Algorithm 7.2.1](#) directly on X is faster if $d \in \Omega(n)$. However, in that case we simply apply the [Johnson & Lindenstrauss \[1984\]](#) Lemma to replace d by $\log(n)/\epsilon^2$.

8.4 Correctness of our Weak Coreset Algorithm

In the following, we proof [Lemma 8.2.12](#). Assume we want to compare the cost of a set representatives with respect to a coreset computed by [Algorithm 8.2.8](#) to the cost of the same set with respect to the original input data set X . Since $w_S(S) = n$, we can treat treat the weighted set (S, w_S) as a set and compare the cost of two equally sized sets. We present a result showing that the difference in cost is mainly bounded by the distance of the points in the sets to each other.

Lemma 8.4.1 *Let $X, Y \in \mathbb{R}^d$ with $|X| = |Y| = n$, $M \in \mathbb{R}^d$, $\epsilon \in (0, 1)$, and $f : X \rightarrow Y$ be any bijection between points. It holds that*

$$\begin{aligned} &|\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(Y, \mathbb{1}, M)| \\ &\leq (1 + \epsilon^{-1}) \cdot \sum_{x \in X} \|x - f(x)\|_2^2 + \epsilon \cdot \min\{\text{fkm}(X, \mathbb{1}, M), \text{fkm}(Y, \mathbb{1}, M)\} . \end{aligned}$$

Proof. Without loss of generality assume $\text{fkm}(X, \mathbb{1}, M) \geq \text{fkm}(Y, \mathbb{1}, M)$. Let $\kappa_Y^* : Y \times M \rightarrow [0, 1]$ be an optimal assignment function. We bound

$$\begin{aligned}
& \text{fkm}(X, \mathbb{1}, M) - \text{fkm}(Y, \mathbb{1}, M) \\
& \leq \text{fkm}(X, \mathbb{1}, M, \kappa_Y^* \circ f) - \text{fkm}(Y, \mathbb{1}, M, \kappa_Y^*) \\
& = \sum_{x \in X} \sum_{\mu \in M} \kappa_Y^*(f(x), \mu)^m \cdot (\|x - \mu\|_2^2 - \|f(x) - \mu\|_2^2) \\
& \leq \sum_{x \in X} \sum_{\mu \in M} \kappa_Y^*(f(x), \mu)^m \cdot (\|x - f(x)\|_2^2 + 2 \cdot \|x - f(x)\|_2 \cdot \|f(x) - \mu\|_2) \\
& \hspace{20em} \text{(squared triangle inequality)} \\
& \leq \sum_{x \in X} \|x - f(x)\|_2^2 \hspace{10em} (\sum_{\mu \in M} \kappa_Y^*(x, \mu)^m \leq 1) \\
& \quad + \sum_{x \in X} \sum_{\mu \in M} 2 \cdot \kappa_Y^*(f(x), \mu)^m \cdot \|x - f(x)\|_2 \cdot \|f(x) - \mu\|_2 \\
& \leq \sum_{x \in X} \|x - f(x)\|_2^2 \\
& \quad + \sum_{x \in X} \sum_{\mu \in M} \kappa_Y^*(f(x), \mu)^m \cdot \left(\epsilon^{-1} \cdot \|x - f(x)\|_2^2 + \epsilon \cdot \|f(x) - \mu\|_2^2 \right) \\
& \hspace{10em} (\forall a, b \in \mathbb{R}_{\geq 0} : 0 \leq (a/\sqrt{\epsilon} - \sqrt{\epsilon} \cdot b)^2 = a^2/\epsilon - 2 \cdot a \cdot b + \epsilon \cdot b) \\
& \leq (1 + \epsilon^{-1}) \sum_{x \in X} \|x - f(x)\|_2^2 \hspace{10em} (\sum_{\mu \in M} \kappa_Y^*(x, \mu)^m \leq 1) \\
& \quad + \epsilon \cdot \sum_{x \in X} \sum_{\mu \in M} \kappa_Y^*(f(x), \mu)^m \cdot \|f(x) - \mu\|_2^2 .
\end{aligned}$$

■

In the following, we prove [Lemma 8.2.12](#) by analyzing a single run of [Algorithm 8.2.11](#). Assume that the first step succeeded. That is, $A \subset \mathbb{R}^d$ is a $(20, 32)$ -bicriteria approximation of KM on $(X, \mathbb{1}, k)$. Furthermore, fix some solution $M \in \Theta_{(k, \epsilon)}(X)$.

Overview We have to verify that the coreset property holds for M . To this end, we pursue the following general approach, which closely mirrors Chen's original exposition augmented with technical elaborations for FKM. We define a subspace of \mathbb{R}^d we call the *grid*. This grid is the union of large balls around the representatives of the bicriteria approximation A . As the name suggests, the grid contains many grid *cells*. The sidelength of each cell depends on the distance to its nearest representative in A – the closer the cell is to A , the smaller the sidelength. We differentiate two cases of M : either all representatives lie inside

8.4 Correctness of our Weak Coreset Algorithm

the grid or at least one is outside. For each case we independently establish that M fulfils the coreset property. More specifically, if M is contained in the grid, then we snap each representative to a corner of its surrounding grid cell. This does not significantly increase the cost of the solution. Hence, we only have to verify the coreset property for a finite number of solutions – the number of corners in the grid. To this end, we use [Algorithm 8.2.8](#).

The Grid The grid consists of a subspace $\mathcal{G} \subset \mathbb{R}^d$ and a discretizing function $g : \mathcal{G} \rightarrow \mathcal{G}$, i.e. g has a finite image. We set

$$\mathcal{G} := \bigcup_{a \in A} \mathcal{B}(a, 2^E \cdot R)$$

where $E := \lfloor 1/2 \cdot \log(\mathfrak{b} \cdot n / (\tilde{\epsilon}^2 \cdot (\epsilon / (2 \cdot m \cdot K + \epsilon))^m)) \rfloor$ and, as before, $R := \sqrt{\text{km}(X, \mathbb{1}, A) / (\alpha \cdot n)}$. We also partition \mathcal{G} into annuli of exponentially increasing radius. For each $j \in [E]_0$ and $a \in A$ we set

$$\mathcal{G}_{a,j} := \begin{cases} \mathcal{B}(a, R) & \text{if } j = 0 \text{ and} \\ \mathcal{B}(a, 2^j \cdot R) \setminus \mathcal{B}(a, 2^{j-1} \cdot R) & \text{else.} \end{cases}$$

We further partition each annulus into cells. Assume that each $\mathcal{G}_{a,j}$ consists of an axis-parallel grid with sidelength

$$\tilde{\epsilon} \cdot \frac{2^j \cdot R}{\sqrt{d}}.$$

The function g maps each point \mathcal{G} to the south-west corner of its surrounding grid cell. In fact, any fixed point in the same grid cell is fine, we simply choose one of the corners.

We bound the size of $g[\mathcal{G}]$, which is the set of corners in the grid. Applying the volume argument provided by [Chen \[2009\]](#), one can prove that

$$|g[\mathcal{G}]| \leq |A| \cdot (E + 1) \cdot \left(\frac{\mathfrak{c}}{\tilde{\epsilon}}\right)^d = \gamma.$$

By applying [Lemma 8.2.9](#), we know that the result (S, w_S) of our algorithm fulfils the coreset property for all $M \in g[\mathcal{G}]$ with $|M| \leq k$ for FKM with accuracy $\tilde{\epsilon}$ and for KM with accuracy $\tilde{\epsilon}/k^{m-1}$, with probability at least $1 - \delta/2$. In the following, we assume that the call to [Algorithm 8.2.8](#) successfully computed such an (S, w_S) . Note that the union bound over the call to the algorithm by [Aggarwal et al. \[2009\]](#)

and the call to [Algorithm 8.2.8](#), each with probability $1 - \delta/2$, gives us $1 - \delta$ as the success probability of [Algorithm 8.2.11](#).

The only modifications we make, compared to Chen's original algorithm, is that we modified some of the parameters. Hence, we can apply the original proof to obtain that (S, w_S) is already a strong $\tilde{\epsilon}/k^{m-1}$ -coreset for the KM problem.

Sets of Representatives Outside the Grid We start with the simpler case. Recall that all representatives have a non-negligible cluster. Hence, if there is a representative outside of the grid, then the cost of that representative is already so large that going from X to the coreset has no significant impact on the cost.

Lemma 8.4.2 *If there is a $\mu \in M$ with $\mu \notin \mathcal{G}$, then*

$$|\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(S, w_S, M)| \leq \epsilon \cdot \text{fkm}(X, \mathbb{1}, M) .$$

Proof. Recall the partitions $X = \bigcup_{a \in A, j \in [F]_0} X_{a,j}$ and $S = \bigcup_{a \in A, j \in [F]_0} S_{a,j}$ from [Algorithm 8.2.8](#). Where, for each $a \in A$ and $j \in [F]_0$, $S_{a,j}$ consists of q points sampled from $X_{a,j}$, each with weight $|X_{a,j}|/q \in \mathbb{N}$. Hence, if we treat (S, w_S) as a set, then we can find a function $s : X \rightarrow X$ such that $S_{a,j} = \{s(x) \mid x \in X_{a,j}\}$ and, for each $x \in S$, the pre-image has size $|s^{-1}(x)| = w_S(x)$. This function s essentially assigns each point in X to a surrogate in the coreset. For every $x \in X$, the point x and its surrogate $s(x)$ are both contained in the same annulus $L_{a,j}$. Due to the 2-approximate triangle inequality, any two points $x, y \in L_{a,j}$ have distance at most

$$\|x - y\|_2^2 \leq 2 \cdot (\|x - a\|_2^2 + \|a - y\|_2^2) \leq 4 \cdot (2^j \cdot R)^2 ,$$

and if $j \geq 1$, then the distance to the nearest bicriteria representative is at least

$$\|x - a\|_2^2 \geq (2^{j-1} \cdot R)^2 .$$

We combine these observations to obtain that, for each $x \in X_{a,j}$,

$$\begin{aligned} \|x - s(x)\|_2^2 &\leq 4 \cdot (2^j \cdot R)^2 \\ &= 4 \cdot 4 \cdot (2^{j-1} \cdot R)^2 \\ &\leq 4 \cdot \max\{4 \cdot \|x - a\|_2^2, R^2\} . \\ &\quad \text{(differentiate between } j \geq 1 \text{ and } j = 0) \end{aligned}$$

8.4 Correctness of our Weak Coreset Algorithm

By taking the sum over all data points, we obtain

$$\begin{aligned}
\sum_{x \in X} \|x - s(x)\|_2^2 &\leq \sum_{x \in X} 4 \cdot \max\{4 \cdot \min_{a \in A} \{\|x - a\|_2^2\}, R^2\} \\
&\leq \sum_{x \in X} 4 \cdot (4 \cdot \min_{a \in A} \{\|x - a\|_2^2\} + R^2) \\
&= 16 \cdot \text{km}(X, \mathbb{1}, A) + 4 \cdot n \cdot \frac{\text{km}(X, \mathbb{1}, A)}{20 \cdot n} \\
&\leq 17 \cdot \text{km}(X, \mathbb{1}, A). \tag{8.4.3}
\end{aligned}$$

Next, we lower bound the distance μ has to the representatives of the bicriteria approximation. Since $\mu \notin \mathcal{G}$, we have

$$\begin{aligned}
\min_{a \in A} \{\|\mu - a\|_2\} &> 2^E \cdot R \\
&\geq \sqrt{\frac{\mathfrak{b} \cdot n}{\tilde{\epsilon}^2 \cdot (\epsilon/(2 \cdot m \cdot K + \epsilon))^m}} \cdot \sqrt{\frac{\text{km}(X, \mathbb{1}, A)}{20 \cdot n}} \\
&= \sqrt{\frac{\mathfrak{b} \cdot \text{km}(X, \mathbb{1}, A)}{20 \cdot \tilde{\epsilon}^2 \cdot (\epsilon/(2 \cdot m \cdot K + \epsilon))^m}}.
\end{aligned}$$

Thus, we can lower bound the distance of each point $x \in X_{a,j}$ to μ by

$$\begin{aligned}
\|x - \mu\|_2 &\geq \|\mu - a\|_2 - \|a - x\|_2 \\
&\geq \sqrt{\frac{\mathfrak{b} \cdot \text{km}(X, \mathbb{1}, A)}{20 \cdot \tilde{\epsilon}^2 \cdot (\epsilon/(2 \cdot m \cdot K + \epsilon))^m}} - \sqrt{\text{km}(X, \mathbb{1}, A)} \\
&\geq \left(\sqrt{\frac{\mathfrak{b}}{20 \cdot \tilde{\epsilon}^2 \cdot (\epsilon/(2 \cdot m \cdot K + \epsilon))^m}} - 1 \right) \cdot \sqrt{\text{km}(X, \mathbb{1}, A)} \\
&\geq \sqrt{\frac{17}{\tilde{\epsilon}^2 \cdot (\epsilon/(2 \cdot m \cdot K + \epsilon))^m}} \cdot \sqrt{\text{km}(X, \mathbb{1}, A)}.
\end{aligned}$$

(for sufficiently large \mathfrak{b})

Let κ be an optimal assignment with respect to M . Because of μ 's non-negligibility, we have

$$\begin{aligned}
\text{fkm}(X, \mathbb{1}, M) &\geq \sum_{x \in X} \kappa(x, \mu)^m \cdot \|x - \mu\|_2^2 \\
&\geq \frac{17}{\tilde{\epsilon}^2 \cdot (\epsilon/(2 \cdot m \cdot K + \epsilon))^m} \cdot \text{km}(X, \mathbb{1}, A) \cdot \sum_{x \in X} \kappa(x, \mu)^m \\
&\geq \frac{17}{\tilde{\epsilon}^2} \cdot \text{km}(X, \mathbb{1}, A)
\end{aligned}$$

$$\geq \frac{1}{\tilde{\epsilon}^2} \cdot \sum_{x \in X} \|x - s(x)\|_2^2 .$$

If we treat (S, w_S) as a set, then s is a bijection between points in X and S . Thus, we apply [Lemma 8.4.1](#) with $\tilde{\epsilon}$ to conclude

$$\begin{aligned} & |\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(S, w_S, M)| \\ & \leq (1 + \tilde{\epsilon}^{-1}) \cdot \sum_{x \in X} \|x - s(x)\|_2^2 + \tilde{\epsilon} \cdot \min\{\text{fkm}(X, \mathbb{1}, M), \text{fkm}(S, w_S, M)\} \\ & \leq (\tilde{\epsilon} + \tilde{\epsilon}^2) \cdot \text{fkm}(X, \mathbb{1}, M) + \tilde{\epsilon} \cdot \text{fkm}(X, \mathbb{1}, M) \\ & \leq \epsilon \cdot \text{fkm}(X, \mathbb{1}, M) . \end{aligned} \tag{choice of \tilde{\epsilon}}$$

■

Sets of Representatives Inside the Grid Next, assume that $M \in \mathcal{G}$. Recall that we snap points $x \in \mathcal{G}$ to one of the corners of their containing grid cell via function g . We bound the distance between points and their image under g in terms of their KM cost with respect to the bicriteria approximation A . For the sake of brevity, we introduce the following shorthand notation: for a multiset $M \in \mathbb{R}^d$ and a point $x \in \mathbb{R}^d$, we denote

$$M_x := \min_{\mu \in M} \{\|x - \mu\|_2\} .$$

Lemma 8.4.4 *If $x \in \mathcal{G}$, then*

$$\|x - g(x)\|_2 \leq 2 \cdot \tilde{\epsilon} \cdot (\min\{\|x - A_x\|_2, \|g(x) - A_{g(x)}\|_2\} + R) .$$

Proof. Let $X_{a,j}$ be the grid cell containing both x and $g(x)$. By definition, we have, similar to before, that $\|x - g(x)\|_2 \leq \tilde{\epsilon} \cdot 2^j \cdot R$. If $j \geq 1$, then $\min\{\|x - A_x\|_2, \|g(x) - A_{g(x)}\|_2\} \geq 2^{j-1} \cdot R$. Again as before, we differentiate between $j = 0$ and $j \geq 1$ and upper bound the maximum between the two by taking the sum. This yields

$$\|x - g(x)\|_2 \leq \tilde{\epsilon} \cdot 2^j \cdot R \leq 2 \cdot \tilde{\epsilon} \cdot (\min\{\|x - A_x\|_2, \|g(x) - A_{g(x)}\|_2\} + R) .$$

■

We use this to bound the cost increase from snapping representatives to the grid.

Lemma 8.4.5 *Let $X \in \mathbb{R}^d$ be any multiset of points with $|X| = n$ and $g[M]$. It holds that*

$$\begin{aligned} & |\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(X, \mathbb{1}, g[M])| \\ & \leq 32 \cdot \tilde{\epsilon} \cdot (\text{km}(X, \mathbb{1}, M) + \text{km}(X, \mathbb{1}, A) + n \cdot R^2) . \end{aligned}$$

8.4 Correctness of our Weak Coreset Algorithm

Proof. Due to the triangle inequality, we can, once again, upper bound the cost point-wise

$$\begin{aligned} & |\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(X, \mathbb{1}, g[M])| \\ & \leq \sum_{x \in X} |\text{fkm}(\{x\}, \mathbb{1}, M) - \text{fkm}(\{x\}, \mathbb{1}, g[M])| . \end{aligned}$$

In the following, fix some $x \in X$ and let $\kappa : X \times M \rightarrow [0, 1]$ and $\kappa_g : X \times g[M] \rightarrow [0, 1]$ be optimal assignment functions. If $\text{fkm}(\{x\}, \mathbb{1}, M) \geq \text{fkm}(\{x\}, \mathbb{1}, g[M])$, then

$$\begin{aligned} & \text{fkm}(\{x\}, \mathbb{1}, M) - \text{fkm}(\{x\}, \mathbb{1}, g[M]) \\ & \leq \text{fkm}(\{x\}, \mathbb{1}, M, \kappa_g \circ g) - \text{fkm}(\{x\}, \mathbb{1}, g[M], \kappa_g) \\ & = \sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot (\|x - \mu\|_2^2 - \|x - g(\mu)\|_2^2) \\ & \leq \sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot (\|\mu - g(\mu)\|_2^2 + 2 \cdot \|\mu - g(\mu)\|_2 \cdot \|x - g(\mu)\|_2) . \end{aligned}$$

(squared triangle inequality)

Analogously, if $\text{fkm}(\{x\}, \mathbb{1}, M) \leq \text{fkm}(\{x\}, \mathbb{1}, g[M])$, then

$$\begin{aligned} & \text{fkm}(\{x\}, \mathbb{1}, g[M]) - \text{fkm}(\{x\}, \mathbb{1}, M) \\ & \leq \sum_{\mu \in M} \kappa(x, \mu)^m \cdot (\|\mu - g(\mu)\|_2^2 + 2 \cdot \|\mu - g(\mu)\|_2 \cdot \|x - \mu\|_2) . \end{aligned}$$

For both cases, we derive upper bounds for each of the two summands individually. Afterwards, we obtain the claim as the maximum of the upper bounds for the two cases.

First, observe that

$$\begin{aligned} & \|x - (g[M])_x\|_2^2 \\ & \leq \|x - g(M_x)\|_2^2 \\ & \leq \|x - M_x\|_2^2 + \|M_x - g(M_x)\|_2^2 + 2 \cdot \|x - M_x\|_2 \cdot \|M_x - g(M_x)\|_2 \\ & \hspace{15em} \text{(squared triangle inequality)} \\ & \leq \|x - M_x\|_2^2 + (2 \cdot \tilde{\epsilon} \cdot (\|M_x - A_x\|_2 + R))^2 \\ & \quad + 2 \cdot \|x - M_x\|_2 \cdot (2 \cdot \tilde{\epsilon} \cdot (\|M_x - A_x\|_2 + R)) \quad \text{(Lemma 8.4.4)} \\ & \leq \|x - M_x\|_2^2 + (2 \cdot \tilde{\epsilon} \cdot (\|M_x - x\|_2 + \|x - A_x\|_2 + R))^2 \\ & \quad + 4 \cdot \tilde{\epsilon} \cdot \|x - M_x\|_2 \cdot (\|M_x - x\|_2 + \|x - A_x\|_2 + R) \\ & \hspace{15em} \text{(triangle inequality)} \\ & \leq \|x - M_x\|_2^2 + 12 \cdot \tilde{\epsilon}^2 \cdot (\|x - M_x\|_2^2 + \|x - A_x\|_2^2 + R^2) \end{aligned}$$

$$\begin{aligned}
& + 2 \cdot \tilde{\epsilon} \cdot (\|x - M_x\|_2^2 + (\|x - M_x\|_2 + \|x - A_x\|_2 + R)^2) \\
& \leq \|x - M_x\|_2^2 + 20 \cdot \tilde{\epsilon} \cdot (\|x - M_x\|_2^2 + \|x - A_x\|_2^2 + R^2) \\
& \leq 2 \cdot \|x - M_x\|_2^2 + \|x - A_x\|_2^2 + R^2 .
\end{aligned} \tag{8.4.6}$$

We apply this to obtain

$$\begin{aligned}
& \sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot \|\mu - g(\mu)\|_2^2 \\
& \leq \sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot (2 \cdot \tilde{\epsilon} \cdot (\|g(\mu) - x\|_2 + \|x - A_x\|_2 + R))^2 \\
& \quad \text{(Lemma 8.4.4 and triangle inequality)} \\
& \leq 12 \cdot \tilde{\epsilon}^2 \cdot \sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot (\|x - g(\mu)\|_2^2 + \|x - A_x\|_2^2 + R^2) \\
& \leq 12 \cdot \tilde{\epsilon}^2 \cdot \left(\sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot \|x - g(\mu)\|_2^2 + \|x - A_x\|_2^2 + R^2 \right) \\
& \quad (\sum_{\mu \in M} \kappa_g(x, g(\mu))^m \leq 1) \\
& \leq 12 \cdot \tilde{\epsilon}^2 \cdot (\|x - (g[M])_x\|_2^2 + \|x - A_x\|_2^2 + R^2) \\
& \leq 12 \cdot \tilde{\epsilon}^2 \cdot (2 \cdot \|x - M_x\|_2^2 + \|x - A_x\|_2^2 + R^2 + \|x - A_x\|_2^2 + R^2) \\
& \quad \text{(by (8.4.6))} \\
& = 24 \cdot \tilde{\epsilon}^2 \cdot (\|x - M_x\|_2^2 + \|x - A_x\|_2^2 + R^2) ,
\end{aligned}$$

and similarly,

$$\begin{aligned}
& \sum_{\mu \in M} \kappa(x, g(\mu))^m \cdot \|\mu - g(\mu)\|_2^2 \\
& \leq 12 \cdot \tilde{\epsilon}^2 \cdot \left(\sum_{\mu \in M} \kappa(x, \mu)^m \cdot \|x - \mu\|_2^2 + \|x - A_x\|_2^2 + R^2 \right) \\
& \leq 12 \cdot \tilde{\epsilon}^2 \cdot (\|x - M_x\|_2^2 + \|x - A_x\|_2^2 + R^2) .
\end{aligned}$$

Using similar arguments, we bound the mixed terms

$$\begin{aligned}
& \sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot \|\mu - g(\mu)\|_2 \cdot \|x - g(\mu)\|_2 \\
& \leq 2 \cdot \tilde{\epsilon} \cdot \sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot (\|x - g(\mu)\|_2 + \|x - A_x\|_2 + R) \cdot \|x - g(\mu)\|_2 \\
& \quad \text{(Lemma 8.4.4 and triangle inequality)} \\
& \leq \tilde{\epsilon} \cdot \sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot (4 \cdot \|x - g(\mu)\|_2^2 + 3 \cdot \|x - A_x\|_2^2 + 3 \cdot R^2)
\end{aligned}$$

8.4 Correctness of our Weak Coreset Algorithm

$$\begin{aligned}
&\leq 4 \cdot \tilde{\epsilon} \cdot \left(\sum_{\mu \in M} \kappa_g(x, g(\mu))^m \cdot \|x - g(\mu)\|_2^2 + \|x - A_x\|_2^2 + R^2 \right) \\
&\leq 8 \cdot \tilde{\epsilon} \cdot \left(\|x - M_x\|_2^2 + \|x - A_x\|_2^2 + R^2 \right), \quad \text{(by (8.4.6))}
\end{aligned}$$

and

$$\begin{aligned}
&\sum_{\mu \in M} \kappa(x, \mu)^m \cdot \|\mu - g(\mu)\|_2 \cdot \|x - g(\mu)\|_2 \\
&\leq 2 \cdot \tilde{\epsilon} \cdot \sum_{\mu \in M} \kappa(x, g(\mu))^m \cdot (\|x - \mu\|_2 + \|x - A_x\|_2 + R) \cdot \|x - \mu\|_2 \\
&\quad \text{(Lemma 8.4.4 and triangle inequality)} \\
&\leq 4 \cdot \tilde{\epsilon} \cdot \left(\|x - M_x\|_2^2 + \|x - A_x\|_2^2 + R^2 \right).
\end{aligned}$$

Taking the maximum of both sums we obtain the claimed upper bound. \blacksquare

We conclude the correctness proof with the coreset property for all sets of representatives inside the grid.

Lemma 8.4.7 *If $M \in \mathcal{G}$, then*

$$|\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(S, w_S, M)| \leq \epsilon \cdot \text{fkm}(X, \mathbb{1}, M).$$

Proof. As before, let $g[M] := \{g(\mu) \mid \mu \in M\}$. By the triangle inequality, we have

$$\begin{aligned}
&|\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(S, w_S, M)| \\
&\leq |\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(X, \mathbb{1}, g[M])| \\
&\quad + |\text{fkm}(X, \mathbb{1}, g[M]) - \text{fkm}(S, w_S, g[M])| \\
&\quad + |\text{fkm}(S, w_S, g[M]) - \text{fkm}(S, w_S, M)|.
\end{aligned}$$

We apply [Lemma 8.4.5](#) to the first summand to obtain

$$\begin{aligned}
&|\text{fkm}(X, \mathbb{1}, M) - \text{fkm}(X, \mathbb{1}, g[M])| \\
&\leq 32 \cdot \tilde{\epsilon} \cdot (\text{km}(X, \mathbb{1}, M) + \text{km}(X, \mathbb{1}, A) + n \cdot R^2) \\
&\leq 32 \cdot \tilde{\epsilon} \cdot (\text{km}(X, \mathbb{1}, M) + (1 + 1/20) \cdot \text{km}(X, \mathbb{1}, A)) \quad \text{(definition of } R) \\
&\leq 704 \cdot \tilde{\epsilon} \cdot \text{km}(X, \mathbb{1}, M) \quad \text{(} A \text{ is bicriteria approximation)} \\
&\leq 704 \cdot k^{m-1} \cdot \tilde{\epsilon} \cdot \text{fkm}(X, \mathbb{1}, M) \\
&\leq \epsilon/3 \cdot \text{fkm}(X, \mathbb{1}, M). \quad \text{(definition of } \tilde{\epsilon})
\end{aligned}$$

Recall that w_S is an integral weight function with $w_S(S) = n$. Furthermore, recall the function $s : X \rightarrow X$, mapping input points to points in the coresets. If we consider S as a set, then we obtain an unweighted set of size n and can apply [Lemma 8.4.5](#) again to obtain

$$\begin{aligned}
& |\text{fkm}(S, w_S, g[M]) - \text{fkm}(S, w_S, M)| \\
& \leq 32 \cdot \tilde{\epsilon} \cdot (\text{km}(S, w_S, M) + \text{km}(S, w_S, A) + n \cdot R^2) \\
& \leq 32 \cdot \tilde{\epsilon} \cdot ((2 + \tilde{\epsilon}/k^{m-1}) \cdot \text{km}(X, \mathbb{1}, M) + \text{km}(S, w_S, A)) \\
& \quad \text{(S is an } \tilde{\epsilon}/k^{m-1}\text{-coreset for KM and definition of } R\text{)} \\
& \leq 32 \cdot \tilde{\epsilon} \cdot ((2 + \tilde{\epsilon}/k^{m-1}) \cdot \text{km}(X, \mathbb{1}, M) \\
& \quad + 2 \cdot \text{km}(X, \mathbb{1}, A) + 2 \cdot \sum_{x \in X} \|x - s(x)\|_2^2) \\
& \quad \text{(2-approximate triangle inequality)} \\
& \leq 32 \cdot \tilde{\epsilon} \cdot ((2 + \tilde{\epsilon}/k^{m-1}) \cdot \text{km}(X, \mathbb{1}, M) + 36 \cdot \text{km}(X, \mathbb{1}, A)) \quad \text{(by (8.4.3))} \\
& \leq 32 \cdot k^{m-1} \cdot \tilde{\epsilon} \cdot (722 + \tilde{\epsilon}/k^{m-1}) \cdot \text{fkm}(X, \mathbb{1}, M) \\
& \leq \epsilon/3 \cdot \text{fkm}(X, \mathbb{1}, M) . \quad \text{(definition of } \tilde{\epsilon}\text{)}
\end{aligned}$$

We previously assumed that [Algorithm 8.2.8](#) was successful in finding an $\tilde{\epsilon}$ -coreset for all solutions $M \in g[\mathcal{G}]$. Hence, we obtain

$$\begin{aligned}
|\text{fkm}(X, \mathbb{1}, g[M]) - \text{fkm}(S, w_S, g[M])| & \leq \tilde{\epsilon} \cdot \text{fkm}(X, \mathbb{1}, M) \\
& \leq \epsilon/3 \cdot \text{fkm}(X, \mathbb{1}, M) .
\end{aligned}$$

■

Future Research

The theory of geometric location problems, and particularly FKM, is far from being fully understood. While we show that the radius variant DFRKM is **NP**-hard in general, our result only holds if the dimension d and the number of clusters k are part of the input. Moreover, we show that a variant of FKM in general metric spaces with $m = 2$ is **NP**-hard. For KM it is known that the problem is **NP**-hard even if d or k is fixed to 2. Hence, an open question is whether FKM is **NP**-hard (see [Conjecture 4.4.6](#)), and if its radius variant DFRKM (and given **NP**-hardness, also FKM) is still hard if we fix the dimension or the number of clusters to some constant. Furthermore, one might ask if FKM lies in **NP**. Our unsolvability result for optimal solutions of FKM does not answer this question. The class **NP** consists of decision problems, thus, one only needs to find a solution with cost less than some given (finitely represented) bound. It is possible that there is no instance where a solution fulfilling this bound is necessarily an optimal solution. Instead we might have that, for every finite cost bound, if there is solution fulfilling this cost bound, then there also is a solution fulfilling the bound which can be represented using size polynomial in the input.

Similarly, our **PLS**-completeness results for DKM and DFKM assume d and k to be part of the input. It is also unknown whether we can fix either of the parameters and still obtain a hard local search problem. We show a lower bound of $\Omega(n)$ for the embedding of our reduction. Thus, any hardness result for bounded dimension requires a fundamentally different approach to the reduction. Additionally, in contrast to the MUFL and DKM reductions, our DFKM reduction introduces a non-trivial weight function on the data points. It would be interesting to explore whether we can encode these weights into the distances between points, similar to the technique we used in the DKM construction.

Generally, we could improve a lot of our results by showing a tighter

relation between KM and FKM. We show that the difference between the KM and FKM cost of a set of representatives is at most k^{m-1} . This bound is indeed tight (consider the set of k representatives located at the mean of the data set). However, we suspect that this can be improved if we only consider optimal solutions. If one could show some structural properties of optimal (or at least almost optimal) FKM solutions, then we might be able to obtain a tighter bound on the cost difference to the solution with respect to KM.

The algorithm we present for FKM still lacks practicality. It is the fastest algorithm (with a guaranteed approximation ratio) known for the problem. However, even if we consider k to be a constant, it is not reasonable to apply the algorithm to a practically relevant data set. There are still many unexplored avenues in terms of constant factor approximation of FKM. We do not know of the existence of a bicriteria approximation algorithm with constant α and β . And the quality of the single-swap heuristic, which was successfully applied to many different location problems, has not yet been analyzed for FKM (see [Conjecture 5.1](#)). Furthermore, our coresset construction is based on a KM technique which is much worse than the most efficient KM coresset constructions currently known. Applying any of the subspace approximation based coresset results for KM to FKM probably yields significantly faster approximation algorithms.

Bibliography

- Aggarwal, A., Deshpande, A., and Kannan, R. (2009). Adaptive Sampling for k-Means Clustering. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 15–28. Springer Berlin Heidelberg. [131](#), [132](#), [135](#)
- Ahmadian, S., Norouzi-Fard, A., Svensson, O., and Ward, J. (2017). Better Guarantees for k-Means and Euclidean k-Median by Primal-Dual Algorithms. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*. [104](#)
- Arora, S. and Barak, B. (2009). *Computational Complexity: A Modern Approach*. Cambridge University Press. [32](#)
- Arora, S., Lund, C., Motwani, R., Sudan, M., and Szegedy, M. (1998). Proof Verification and the Hardness of Approximation Problems. *Journal of the ACM*, 45(3):501–555. [44](#)
- Arthur, D., Manthey, B., and Röglin, H. (2009). k-Means Has Polynomial Smoothed Complexity. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*. [28](#)
- Arthur, D. and Vassilvitskii, S. (2006). Worst-case and Smoothed Analysis of the ICP Algorithm, with an Application to the k-Means Method. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science*. [28](#)
- Arthur, D. and Vassilvitskii, S. (2007). k-Means++: The Advantages of Careful Seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1027–1035. [104](#)
- Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., and Pandit, V. (2004). Local Search Heuristics for k-Median and Facility Location Problems. *SIAM Journal on Computing*, 33(3):544–562. [52](#), [58](#)
- Bădoiu, M., Har-Peled, S., and Indyk, P. (2002). Approximate Clustering via Core-Sets. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*. [104](#), [120](#)

Bibliography

- Bajaj, C. (1988). The Algebraic Degree of Geometric Optimization Problems. *Discrete & Computational Geometry*, 3(2):177–191. [88](#), [90](#)
- Bertsekas, D. P. (1999). *Nonlinear Programming*. Athena Scientific. [93](#)
- Bezdek, J. C., Ehrlich, R., and Full, W. (1984). FCM: The Fuzzy c-Means Clustering Algorithm. *Computers & Geosciences*, 10(2-3):191–203. [27](#), [28](#)
- Bezdek, J. C., Hathaway, R. J., Sabin, M. J., and Tucker, W. T. (1987). Convergence Theory for Fuzzy c-Means: Counterexamples and Repairs. *IEEE Transactions on Systems, Man, and Cybernetics*, 17(5):873–877. [28](#)
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer-Verlag New York Inc. [26](#)
- Blömer, J. (1993). *Simplifying Expressions Involving Radicals*. PhD thesis, Freie Universität Berlin. [47](#)
- Blömer, J., Brauer, S., and Bujna, K. (2016a). A Theoretical Analysis of the Fuzzy k-Means Problem. In *2016 IEEE 16th International Conference on Data Mining*. IEEE. [3](#), [23](#), [27](#), [28](#), [36](#), [87](#), [103](#), [104](#), [105](#), [106](#)
- Blömer, J., Brauer, S., and Bujna, K. (2018). Coresets for Fuzzy k-Means with Applications. In *29th International Symposium on Algorithms and Computation*, volume 123, pages 46:1–46:12. [3](#), [103](#), [105](#), [119](#)
- Blömer, J., Lammersen, C., Schmidt, M., and Sohler, C. (2016b). Theoretical Analysis of the k-Means Algorithm – A Survey. In *Algorithm Engineering*, pages 81–116. Springer International Publishing. [28](#)
- Brauer, S. (2017). Complexity of Single-Swap Heuristics for Metric Facility Location and Related Problems. In *Lecture Notes in Computer Science*, pages 116–127. Springer International Publishing. [3](#), [51](#)
- Brauer, S. (2019). Complexity of Single-Swap Heuristics for Metric Facility Location and Related Problems. *Theoretical Computer Science*, 754:88–106. [3](#), [51](#)
- Bronštein, I. N., Semendjajew, K. A., Musiol, G., and Mühlig, H. (2008). *Taschenbuch der Mathematik*. Verlag Harri Deutsch. (in German). [116](#)

- Chen, K. (2009). On Coresets for k-Median and k-Means Clustering in Metric and Euclidean Spaces and Their Applications. *SIAM Journal on Computing*, 39(3):923–947. 119, 120, 121, 125, 135
- Cheng, D., Hu, X., and Martin, C. (2006). On the Smallest Enclosing Balls. *Communications in Information and Systems*, 6(2):137–160. 34
- Cohen-Addad, V., Klein, P. N., and Mathieu, C. (2016). Local Search Yields Approximation Schemes for k-Means and k-Median in Euclidean and Minor-Free Metrics. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science*. 59
- Cook, S. A. (1971). The Complexity of Theorem-Proving Procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*. ACM Press. 33
- Dasgupta, S. (2008). The Hardness of k-Means Clustering. Technical Report CS2008-0916, Department of Computer Science and Engineering, University of California, San Diego. 36
- de Fermat, P. (1891). *Oeuvres de Fermat*. Gauthier-Villars et fils. 27
- Deza, E. and Deza, M. M. (2009). *Encyclopedia of Distances*. Springer Berlin Heidelberg. 44
- Dijkstra, E. W. (2000). The Notational Conventions I Adopted, and Why. <http://www.cs.utexas.edu/~EWD/transcriptions/EWD13xx/EWD1300.html>, last access: 22.05.2019. 7
- Downey, R. G. and Fellows, M. R. (2012). *Parameterized Complexity*. Springer New York. 46
- Downey, R. G. and Fellows, M. R. (2013). *Fundamentals of Parameterized Complexity*. Springer London. 46
- Dunn, J. C. (1973). A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters. *Journal of Cybernetics*, 3(3):32–57. 28
- Feder, T. and Greene, D. (1988). Optimal Algorithms for Approximate Clustering. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*. 36
- Feldman, D., Monemizadeh, M., and Sohler, C. (2007). A PTAS for k-Means Clustering Based on Weak Coresets. In *Proceedings of the*

Bibliography

- twenty-third annual symposium on Computational geometry.* 105, 121
- Feldman, D., Schmidt, M., and Sohler, C. (2013). Turning Big Data Into Tiny Data: Constant-Size Coresets for k-Means, PCA and Projective Clustering. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms.* 120, 121, 132
- Forster, O. (2010). *Analysis 2.* Vieweg + Teubner. 93
- Friggstad, Z., Rezapour, M., and Salavatipour, M. R. (2016). Local Search Yields a PTAS for k-Means in Doubling Metrics. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science.* 59
- Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness (A Series of Books in the Mathematical Sciences).* W. H. Freeman. 33, 46
- Goldreich, O. (2008). *Computational Complexity.* Cambridge University Press. 32
- Gonzalez, T. F. (1985). Clustering to Minimize the Maximum Intercluster Distance. *Theoretical Computer Science*, 38:293–306. 104
- Guha, S. and Khuller, S. (1999). Greedy Strikes Back: Improved Facility Location Algorithms. *Journal of Algorithms*, 31(1):228–248. 36
- Har-Peled, S. and Kushal, A. (2006). Smaller Coresets for k-Median and k-Means Clustering. *Discrete & Computational Geometry*, 37(1):3–19. 104, 120, 121
- Har-Peled, S. and Mazumdar, S. (2004). On Coresets for k-Means and k-Median Clustering. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing.* 119, 120
- Haussler, D. (1992). Decision Theoretic Generalizations of the PAC Model for Neural Net and Other Learning Applications. *Information and Computation*, 100(1):78–150. 126
- Herstein, I. N. (1975). *Topics in Algebra.* Wiley. 88
- Hochbaum, D. S. and Shmoys, D. B. (1986). A Unified Approach to Approximation Algorithms for Bottleneck Problems. *Journal of the ACM*, 33(3):533–550. 36

- Hoppner, F. and Klawonn, F. (2003). A Contribution to Convergence Theory of Fuzzy c-Means and Derivatives. *IEEE Transactions on Fuzzy Systems*, 11(5):682–694. [28](#)
- Hungerford, T. W. (2003). *Algebra*. Springer. [88](#)
- Inaba, M., Katoh, N., and Imai, H. (1994). Applications of Weighted Voronoi Diagrams and Randomization to Variance-Based k-Clustering. In *Proceedings of the tenth annual symposium on Computational geometry*. [27](#), [111](#)
- Jain, K., Mahdian, M., and Saberi, A. (2002). A New Greedy Approach for Facility Location Problems. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*. [104](#)
- Jain, K. and Vazirani, V. V. (2001). Approximation Algorithms for Metric Facility Location and k-Median Problems Using the Primal-Dual Schema and Lagrangian Relaxation. *Journal of the ACM*, 48(2):274–296. [104](#)
- Johnson, D. S., Papadimitriou, C. H., and Yannakakis, M. (1988). How Easy Is Local Search? *Journal of Computer and System Sciences*, 37(1):79–100. [53](#)
- Johnson, W. B. and Lindenstrauss, J. (1984). Extensions of Lipschitz Mappings Into a Hilbert Space. *Contemporary mathematics*, 26:189–206. [133](#)
- Kanungo, T., Mount, D. M., Netanyahu, N. S., Piatko, C. D., Silverman, R., and Wu, A. Y. (2004). A Local Search Approximation Algorithm for k-Means Clustering. *Computational Geometry*, 28(2-3):89–112. [27](#), [28](#), [52](#), [59](#)
- Karp, R. M. (1972). Reducibility Among Combinatorial Problems. In *Complexity of Computer Computations*, pages 85–103. Springer US. [33](#)
- Kim, T., Bezdek, J. C., and Hathaway, R. J. (1988). Optimality Tests for Fixed Points of the Fuzzy c-Means Algorithm. *Pattern Recognition*, 21(6):651–663. [28](#)
- Kochetov, Y. and Ivanenko, D. (2005). Computationally Difficult Instances for the Uncapacitated Facility Location Problem. In *Metaheuristics: Progress as Real Problem Solvers*, pages 351–367. Springer-Verlag. [58](#)

Bibliography

- Kumar, A., Sabharwal, Y., and Sen, S. (2010). Linear-Time Approximation Schemes for Clustering Problems in Any Dimensions. *Journal of the ACM*, 57(2):1–32. 105
- Lee, E., Schmidt, M., and Wright, J. (2017). Improved and Simplified Inapproximability for k-Means. *Information Processing Letters*, 120:40–43. 36
- Levin, L. A. (1973). Universal Sequential Search Problems. *Problems of Information Transmission*, 9(3):115–116. (in Russian). 33
- Li, S. (2013). A 1.488 Approximation Algorithm for the Uncapacitated Facility Location Problem. *Information and Computation*, 222:45–58. 104
- Lloyd, S. (1982). Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137. 27
- Lucic, M., Bachem, O., and Krause, A. (2016). Strong Coresets for Hard and Soft Bregman Clustering with Applications to Exponential Family Mixtures. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 1–9. 120
- Mahajan, M., Nimbhorkar, P., and Varadarajan, K. (2012). The Planar k-Means Problem is NP-Hard. *Theoretical Computer Science*, 442:13–21. 36
- Manthey, B. and Röglin, H. (2009). Improved Smoothed Analysis of the k-Means Method. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 461–470. 28
- Matoušek, J. (2000). On Approximate Geometric k-Clustering. *Discrete & Computational Geometry*, 24(1):61–84. 36
- Megiddo, N. and Papadimitriou, C. H. (1991). On Total Functions, Existence Theorems and Computational Complexity. *Theoretical Computer Science*, 81(2):317–324. 54
- Megiddo, N. and Supowit, K. J. (1984). On the Complexity of Some Common Geometric Location Problems. *SIAM Journal on Computing*, 13(1):182–196. 36
- Mitzenmacher, M. and Upfal, E. (2005). *Probability and Computing*. Cambridge University Press. 106

- Papadimitriou, C. H. (1981). Worst-Case and Probabilistic Analysis of a Geometric Location Problem. *SIAM Journal on Computing*, 10(3):542–557. [31](#), [36](#), [47](#), [49](#)
- Papadimitriou, C. H. (1993). *Computational Complexity*. Pearson. [32](#)
- Papadimitriou, C. H., Schäffer, A. A., and Yannakakis, M. (1990). On the Complexity of Local Search. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. ACM Press. [53](#)
- Papadimitriou, C. H. and Yannakakis, M. (1991). Optimization, Approximation, and Complexity Classes. *Journal of Computer and System Sciences*, 43(3):425–440. [32](#)
- Papadimitriou, C. H. and Yannakakis, M. (1993). The Traveling Salesman Problem with Distances One and Two. *Mathematics of Operations Research*, 18(1):1–11. [44](#)
- Pirkul, H. and Jayaraman, V. (1998). A Multi-Commodity, Multi-Plant, Capacitated Facility Location Problem: Formulation and Efficient Heuristic Solution. *Computers & Operations Research*, 25(10):869–878. [27](#)
- Schoenberg, I. J. (1938). Metric Spaces and Positive Definite Functions. *Transactions of the American Mathematical Society*, 44(3):522–522. [72](#)
- Schäffer, A. A. and Yannakakis, M. (1991). Simple Local Search Problems That Are Hard to Solve. *SIAM Journal on Computing*, 20(1):56–87. [53](#)
- Shmoys, D. B., Tardos, É., and Aardal, K. (1997). Approximation Algorithms for Facility Location Problems. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*. [104](#)
- Sipser, M. (2012). *Introduction to the Theory of Computation*. Course Technology. [32](#)
- Sohler, C. and Woodruff, D. P. (2018). Strong Coresets for k-Median and Subspace Approximation: Goodbye Dimension. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science*. [120](#)
- Torgerson, W. S. (1952). Multidimensional Scaling: I. Theory and Method. *Psychometrika*, 17(4):401–419. [73](#)

Bibliography

- Trevisan, L. (2000). When Hamming Meets Euclid: The Approximability of Geometric TSP and Steiner Tree. *SIAM Journal on Computing*, 30(2):475–485. [44](#)
- Vattani, A. (2011). k-Means Requires Exponentially Many Iterations Even in the Plane. *Discrete & Computational Geometry*, 45(4):596–616. [28](#)