# High-Dimensional Similarity Search and Sketching: Algorithms and Hardness 

by<br>Ilya Razenshteyn<br>B.S., Lomonosov Moscow State University (2012)<br>S.M., Massachusetts Institute of Technology (2014)<br>Submitted to the Department of Electrical Engineering and Computer Science<br>in partial fulfillment of the requirements for the degree of<br>Doctor of Philosophy<br>at the<br>MASSACHUSETTS INSTITUTE OF TECHNOLOGY<br>September 2017<br>© Massachusetts Institute of Technology 2017. All rights reserved.

Author
Department of Electrical Engineering and Computer Science August 31, 2017

Certified by
Piotr Indyk
Professor of Electrical Engineering and Computer Science Thesis Supervisor

Accepted by
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science Chair, Department Committee on Graduate Students

# High-Dimensional Similarity Search and Sketching: Algorithms and Hardness 

by

Ilya Razenshteyn

Submitted to the Department of Electrical Engineering and Computer Science on August 31, 2017, in partial fulfillment of the requirements for the degree of

Doctor of Philosophy


#### Abstract

We study two fundamental problems that involve massive high-dimensional datasets: approximate near neighbor search (ANN) and sketching. We obtain a number of new results including: - An algorithm for the ANN problem over the $\ell_{1}$ and $\ell_{2}$ distances that, for the first time, improves upon the Locality-Sensitive Hashing (LSH) framework. The key new insight is to use random space partitions that depend on the dataset. - An implementation of the core component of the above algorithm, which is released as FALCONN: a new C++ library for high-dimensional similarity search. - An efficient algorithm for the ANN problem over any distance that can be expressed as a symmetric norm. - For norms, we establish the equivalence between the existence of short and accurate sketches and good embeddings into $\ell_{p}$ spaces for $0<p \leq 2$. We use this equivalence to show the first sketching lower bound for the Earth Mover's Distance (EMD).


Thesis Supervisor: Piotr Indyk
Title: Professor of Electrical Engineering and Computer Science

## Acknowledgments

Throughout my life I have been extremely lucky to be surrounded with incredible people: mentors, teachers, colleagues, classmates, friends and family. Without your support and investment in me I would not be where I am now.

Let me start with my Ph. D. advisor Piotr Indyk. Being Piotr's student for the past five years has been both an honor and a pleasure. In addition to being a brilliant researcher, Piotr has all the other qualities that a great advisor needs: a warm and enthusiastic personality, a great sense of care for all of his students, extremely high ethical standards when it comes to collaborations and publishing papers, and a perfect balance between being hands-on and giving enough freedom. I would also like to point out Piotr's talent of finding "just right" problems for each particular student (with an amazing precision!), and a strong much needed support during my academic job search. Last but not least, as a well-established professor, Piotr has built a developed ecosystem of current and former students and postdocs, with many of whom I had a privilege to work with...
... which brings me to the second person I would like to thank: Alex Andoni, a former student of Piotr. Alex has effectively been my informal advisor. Both scientifically: for instance, all the papers this thesis is based on are co-authored with Alex, and more generally: he has given me lots of academic and life advice. Alex's cheerful personality made our five-year long (extremely fruitful) collaboration very enjoyable, even during harder times when not all the problems were being solved to our satisfaction. In fact, I met Alex for the first time even before Piotr: back in Summer 2011 at Microsoft Research, where I was an intern. Our short conversation and Alex's (unintended?) pitch for MIT and Piotr were decisive later when I needed to choose a school and an advisor for my Ph. D.

My fascination with Theory started during my undergraduate studies at Moscow State University, and the two people responsible for this are Maxim Babenko and Sasha Shen, who were kind enough to advise a clueless and naive but eager undergrad. Both Maxim and Sasha served as great role models both in academic and non-academic
matters. Among many other things, I would like to thank Maxim for providing me with a part-time job at Yandex, when I needed money, and Sasha for teaching me (by many examples) how to give good talks and lectures and write papers well. Uncharacteristically for Russian researchers, both Maxim and Sasha encouraged me to go abroad for a graduate school, which in retrospect was essentially the only viable option.

Let me turn to my summer internship mentors: Andrew Goldberg, Renato Werneck, Vahab Mirrokni, Silvio Lattanzi, Alex Andoni, David Woodruff, and Konstantin Makarychev. Between 2010 and 2016 I spent only one summer outside of a research lab, and all of the remaining summers were formative for me as a researcher. I would like to thank Andrew and Renato for strongly influencing my research taste ${ }^{1}$, and Andrew for all of the support throughout the years. I thank Vahab and Silvio for expanding my research horizons. I interned with Alex and Konstantin at Microsoft Research, and both of these summers were perhaps the most productive time in my life so far. Finally, I thank David for being a very dedicated mentor.

Summer internships are not as much fun without fellow interns. I would like to thank Andreea, Avi, Cyrus, Daniel, Dimitris, Fan, Gopi, Irina, Jana, Konstantina, Melanie, Mika, Mohsen, Nina, Thomas, Tomas, Li-Yang, Rad, Reza, Sadra, Sahil, Tom, Tselil, and Zhao for great time we spent together. On a loosely related note, I would like to thank Clement, Dimitris and Erik for a great social atmosphere in a place, which I am not allowed to name here.

I would like to thank all my co-authors and collaborators, past and present: Thomas Ahle, Zeyuan Allen-Zhu, Alex Andoni, Misha Andreev, Maxim Babenko, Arturs Backurs, Daniel Delling, Daniel Fleischman, Rati Gelashvili, Andrew Goldberg, Alexey Gusakov, Piotr Indyk, Michael Kapralov, Ignat Kolesnichenko, Robi Krauthgamer, Thijs Laarhoven, Silvio Lattanzi, Stefano Leonardi, Sepideh Mahabadi, Konstantin Makarychev, Yury Makarychev, Vahab Mirrokni, Huy Nguyen, Sasho Nikolov, Rasmus Pagh, Eric Price, Ruslan Savchenko, Ludwig Schmidt, Negev Shekel Nosatzki, Sasha Shen, Francesco Silvestri, Zhao Song, Ameya Velingker, Tal Wagner, Erik Waingarten,

[^0]Renato Werneck, and David Woodruff. I have learnt a lot from each of you, and all the collaborations were not only productive, but also very enjoyable.

MIT and especially the Theory group of CSAIL have become my second home for the past five years. I would like to thank fellow graduate students who made this place particularly cozy for me: Adrian, Arturs, Badih, Fereshte, G, Govind, Henry, Madalina, Nicholas, Nishanth, Rati, Sepideh, Tal, and Zeyuan. No matter what format our interactions took-going for hikes, walks, runs or drinks, organizing parties at conferences, participating in ICPC contests or writing papers together-it was all great fun. I thank Henry for hosting me during the EECS visit day back in 2012 and showing me what a wonderful place MIT is.

I would like to thank my friends back from Russia: Alexey, Boris, Ilya, Maxim, Polina, Sasha and Sasha. You are all over the world, and I may not be meeting you that often, but whenever we meet, I have this feeling that we never really parted.

I thank my sister Eleonora and my parents Lyudmila and Petr for believing in me Without your love, support and passion for education and science, nothing would be possible.

My wife Oksana has been the main person in my life and the best friend for the past eight years. Thank you for everything!

The person who introduced the world of computer science and programming to me is Vladimir Denisovich Lelyukh, my school informatics teacher. It all started when my parents brought me, a ten-year old, to his class almost seventeen years ago. Vladimir Denisovich was an amazing teacher, a charismatic mentor, and a good friend to all of his many students. He passed away four years ago, but I still can't fully believe it happened. I dedicate this thesis to him.

## Contents

1 Introduction ..... 19
1.1 Introduction ..... 19
1.1.1 Overview of the problems ..... 20
1.1.2 Main contributions: a bird's eye view ..... 22
1.2 Near neighbor search (NNS) ..... 23
1.2.1 Approximate near neighbor search (ANN) ..... 25
1.2.2 Locality-sensitive hashing (LSH) ..... 27
1.2.3 Data-dependent LSH ..... 29
1.2.4 Time-space trade-offs for ANN ..... 31
1.2.5 FALCONN: practical and optimal LSH for unit sphere ..... 33
1.2.6 ANN for general symmetric norms ..... 35
1.3 Characterization of sketchable distances ..... 36
1.4 Preliminaries and notation ..... 38
1.4.1 Metric and normed spaces ..... 38
1.4.2 Examples of metrics and norms ..... 40
1.4.3 Approximate near neighbor search (ANN) ..... 41
1.4.4 Locality-sensitive hashing (LSH) ..... 42
1.4.5 Gaussian distribution ..... 43
2 New ANN algorithms for $\ell_{1}$ and $\ell_{2}$ ..... 45
2.1 Problem definition ..... 45
2.1.1 Parameters and guarantees ..... 46
2.2 Prior work ..... 47
2.2.1 Locality-sensitive hashing (LSH) ..... 47
2.2.2 Time-space trade-offs ..... 48
2.3 The main result ..... 48
2.3.1 Comparison with the prior work ..... 49
2.4 Overview of the algorithm ..... 51
2.5 Simplification of the problem ..... 55
2.6 Random ANN instances ..... 56
2.7 Gaussians and tail bounds ..... 57
2.8 Data-independent partitions ..... 61
2.8.1 Results ..... 62
2.8.2 Description ..... 63
2.8.3 Analysis ..... 65
2.8.4 Setting parameters: idealized version ..... 67
2.8.5 Setting parameters: final version ..... 72
2.9 Data-dependent partitions ..... 74
2.9.1 Overview ..... 74
2.9.2 Description ..... 76
2.9.3 Setting parameters ..... 79
2.9.4 Analysis ..... 82
2.9.5 Fast preprocessing ..... 93
2.9.6 Handling insertions and deletions ..... 95
3 FALCONN: practical and optimal LSH for unit sphere ..... 97
3.1 Introduction ..... 97
3.1.1 Our results ..... 98
3.1.2 Related work ..... 102
3.2 Preliminaries ..... 102
3.3 Cross-polytope LSH ..... 103
3.3.1 Making the cross-polytope LSH practical ..... 104
3.4 Lower bound ..... 106
3.5 Multiprobe LSH for the cross-polytope LSH ..... 108
3.6 Experiments ..... 110
3.7 Appendix: Gaussian measure of a planar set ..... 112
3.8 Appendix: Proof of Theorem 3.3.1 ..... 114
3.8.1 Idealized proof ..... 115
3.8.2 Real proof ..... 116
3.9 Appendix: Proof of Theorem 3.4.1 ..... 117
3.10 Appendix: Further description of experiments ..... 120
4 ANN algorithms for general symmetric norms ..... 125
4.1 Introduction ..... 125
4.1.1 The main result ..... 127
4.1.2 Why symmetric norms? ..... 127
4.1.3 Prior work: ANN for norms beyond $\ell_{1}$ and $\ell_{2}$ ..... 129
4.1.4 Overview of the proof of Theorem 4.1.3 ..... 130
4.1.5 Optimality of Theorem 4.1.3 ..... 131
4.1.6 Lower bounds for general norms ..... 131
4.1.7 Other related work: dealing with general norms ..... 132
4.2 Preliminaries ..... 133
4.2.1 Norms and products ..... 133
4.2.2 ANN for $\ell_{\infty}$ and $\ell_{\infty}$-products ..... 134
4.3 ANN for Orlicz and top- $k$ norms ..... 135
4.4 Embedding into product spaces ..... 139
4.4.1 Proof of Lemma 4.4.14: bounding the net size ..... 145
4.5 Proof of the main theorem ..... 149
4.6 Lower bounds ..... 151
4.7 Appendix: Bounding space in Theorem 4.2.9 ..... 153
4.8 Appendix: $\widetilde{O}(\log d)$-ANN for symmetric norms ..... 155
4.8.1 The $\log ^{\Omega(1)} d$-approximation is necessary ..... 158
4.9 Appendix: Lower bound for arbitrary metrics via expander graphs ..... 158
5 Sketching and embedding are equivalent for norms ..... 163
5.1 Introduction ..... 163
5.1.1 The embedding approach ..... 165
5.1.2 Our results ..... 166
5.1.3 Applications ..... 168
5.1.4 Other related work ..... 170
5.1.5 Proof overview ..... 170
5.2 Preliminaries on functional analysis ..... 172
5.3 Preliminaries on communication complexity ..... 174
5.3.1 Information complexity: private-coins vs. public-coins ..... 176
5.3.2 From countable to finite number of coins ..... 177
5.4 From sketches to uniform embeddings ..... 178
5.4.1 Sketching implies the absence of Poincaré inequalities ..... 179
5.4.2 The absence of Poincaré inequalities implies threshold maps ..... 181
5.4.3 Threshold maps imply uniform embeddings ..... 187
5.4.4 Putting it all together ..... 192
5.5 Quantitative bounds ..... 193
5.6 Embedding into $\ell_{1}$ via sum-products ..... 199
5.7 Appendix: EMD reduction ..... 200
6 Hardness results for the ANN problem ..... 203
6.1 Introduction ..... 203
6.1.1 Our results ..... 204
6.2 Preliminaries ..... 208
6.2.1 Random Hamming instances ..... 208
6.2.2 Graphical neighbor search and robust expansion ..... 209
6.2.3 Locally-decodable codes (LDC) ..... 209
6.3 One-probe data structures ..... 210
6.3.1 Robust expansion of the Hamming space ..... 210
6.3.2 One-probe data structures ..... 213
6.4 List-of-points data structures ..... 214
6.5 Two-probe data structures ..... 219
6.5.1 Deterministic data structures ..... 221
6.5.2 Making low-contention data structures ..... 222
6.5.3 Datasets which shatter ..... 223
6.5.4 Corrupting some cell contents of shattered points ..... 226
6.5.5 Decreasing the word size ..... 229
6.5.6 Connection to locally-decodable codes ..... 230

## List of Figures

1-1 The bag of words representation ..... 20
1-2 Near neighbor search ..... 21
1-3 Sketching ..... 22
1-4 The definition of the ANN problem ..... 26
1-5 Distribution of distances for word embeddings ..... 27
1-6 Time-space trade-offs ..... 33
1-7 Density of $N(0,1)$. ..... 43
2-1 Various specializations of the main result ..... 50
2-2 Illustration of the main result ..... 51
2-3 Comparison of the main result with the state of the art ..... 52
2-4 Plot for the data-independent bound ..... 54
2-5 Illustration of the definition of $\alpha(s)$ and $\beta(s)$. ..... 58
2-6 Illustration of the definition of $A$. ..... 58
2-7 Two cases for $\Delta_{A}$ ..... 60
2-8 Illustration of the definition of $\mu_{A}(r)$. ..... 62
2-9 Pseudocode for data-independent partitions ..... 65
2-10 Covering a spherical cap of radius $(\sqrt{2}-\varepsilon) R$ on a ball of radius $R$. ..... 75
2-11 The definition of Project ..... 76
2-12 The definition of $\varepsilon^{\prime}>0$. ..... 80
2-13 Pseudocode of the data-dependent data-structure ..... 83
3-1 Comparison of the new LSH family and Hyperplane LSH ..... 100
3-2 Two plots ..... 106
3-3 Gaussian measure of planar sets ..... 123
3-4 Distance to the nearest neighbor for the datasets. ..... 124

## List of Tables

3.1 Average running times ..... 112
3.2 Comparison of vanilla LSH and Multiprobe LSH ..... 122
3.3 Average running times for random data ..... 122

## Chapter 1

## Introduction

### 1.1 Introduction

In this thesis, we show new algorithms and hardness results for fundamental problems that involve massive high-dimensional data processing. Such datasets are pervasive in many applications. Let us list just a few immediate examples.

- Feature vectors A popular approach in data analysis is to map a dataset into a feature vector space, and reduce a given task to a geometric computational problem. The simplest example of this sort is the bag of words representation: one maps a text document into a vector of word counts (see Figure 1-1). In many applications, both the number of data items and features is huge. For instance, Twitter [1] receives around 200 billion tweets per year. Thus, if we wanted to analyze them using the bag of words approach, we would need to process 200B vectors with hundreds of thousands of dimensions, since the number of coordinates is equal to the number of distinct words.
- Signals Typically, signals of various nature (sound, images, etc.) can be modeled by real or complex vectors after appropriate sampling/discretization. As in the previous example, one often needs to deal with extremely high-dimensional vectors. For instance, if we sample audio at 44.1 kHz , one minute corresponds to a vector with more than two million dimensions.

```
everything
excluding
in
including
moderation
out
```

Figure 1-1: The bag of words representation of the sentence "Everything in moderation, including moderation."

- Data streams When processing a long stream of data items that one can not even store, thinking of the stream as a sequence of updates to a high-dimensional vector is often useful. For instance, if one needs to process a stream of network packets and detect potential DDoS attacks, one can imagine a vector indexed by all possible $2^{32}$ IP addresses ${ }^{1}$. Each packet is interpreted as an update to this vector ("increase the corresponding coordinate by 1 "), and the goal is to keep track of its largest coordinates, which correspond to high-traffic source/destination addresses.

These and other applications motivate the field of high-dimensional computational geometry, which this thesis contributes to. It should be contrasted with the classical, "low-dimensional" computational geometry, where most of the efficient algorithms have running time exponential in the dimension, thus making their applicability to the high-dimensional regime limited.

### 1.1.1 Overview of the problems

This thesis consists of two parts. Each part is devoted to a separate problem, which we briefly and informally introduce below.

- Near neighbor search (see Figure 1-2) Given a dataset of $n$ points (vectors) in a $d$-dimensional space $\mathbb{R}^{d}$, we would like to preprocess it and build a data structure for answering nearest neighbor queries. Each query is a d-dimensional point, and the goal is to return the data point closest to it. In a simpler version of

[^1]

Figure 1-2: Near neighbor search
the problem, we allow answers to be approximate. The parameters we are trying to optimize are: space used by a data structure, query time, and preprocessing time.

A naïve solution is the linear scan: omit the preprocessing stage altogether, and, whenever given a query, try all the data points and choose the closest one. However, this is too slow for large datasets, and one typically uses the preprocessing stage to compute and store auxiliary information which allows to speed up the query procedure.

Perhaps the most "obvious" application of near neighbor search is similarity search over various kinds of data, which corresponds to nearest neighbor queries in the feature space [162].

- Sketching (succinct summarization) (see Figure 1-3) Sketching involves (lossy) compression of a massive object into a short summary which can be used as a proxy in further computations. Sketching has been used for similarity search [85, 108, 171, 172], fast algorithms for numerical linear algebra [178, 154], processing data streams [136, 5], and speeding up other algorithms [22, 99].

In this thesis, we study sketching of high-dimensional vectors with the goal of estimating a given distance function between two vectors from their (succinct) summaries.


Figure 1-3: Sketching high-dimensional vectors; here $s \ll d$

Superficially, the common theme among these problems is working with highdimensional vectors/datasets. However, the relation is more fundamental. First of all, the problems are closely related on a technical level: sketching can be used for near neighbor search. Second, both of these problems build on a coherent toolbox of mathematical/algorithmic primitives, which can be called "efficient representations of data": randomized hashing, dimension reduction, metric embeddings, etc. Hence, this thesis can be seen as a step towards building a unified theory of such efficient representations.

### 1.1.2 Main contributions: a bird's eye view

Here we briefly describe the main contributions of the thesis. A more detailed description can be found in the further sections of the introduction.

Near neighbor search We present a number of new algorithms for the near neighbor search problem. Most notably, our algorithms provide the first improvement over Locality-Sensitive Hashing (LSH), a popular framework in theory and applications. We implement a core component of one of the new algorithms and release it as a part of FALCONN: a new open-source C++ library for similarity search over high-dimensional data. In addition to algorithms, we show several impossibility results, which we hope will guide further algorithmic research on the problem.

This part is based on the following papers:

- Alexandr Andoni, Piotr Indyk, Thijs Laarhoven, Ilya Razenshteyn, Ludwig Schmidt, Practical and Optimal LSH for Angular Distance, appeared at NIPS

2015 [17],

- Alexandr Andoni, Thijs Laarhoven, Ilya Razenshteyn, Erik Waingarten, Optimal Hashing-Based Time-Space Trade-offs for Approximate Near Neighbors, appeared at SODA 2017 [24], invited to the special issue of ACM Transactions on Algorithms; this result heavily builds on:
- Alexandr Andoni, Piotr Indyk, Huy Nguyen, Ilya Razenshteyn, Beyond Locality-Sensitive Hashing, appeared at SODA 2014 [18],
- Alexandr Andoni, Ilya Razenshteyn, Optimal Data-Dependent Hashing for Approximate Near Neighbors, appeared at STOC 2015 [26],

Alexandr Andoni, Huy Nguyen, Aleksandar Nikolov, Ilya Razenshteyn, Erik Waingarten, Approximate Near Neighbors for General Symmetric Norms, appeared at STOC 2017 [25].

Sketching Various distance functions (metrics) behave very differently when it comes to sketching: for some distances there are extremely efficient compression techniques (such as the celebrated Johnson-Lindentrauss lemma), for other distances essentially nothing is possible. For a large class of distances, we characterize completely when efficient sketching is possible. As a byproduct of this characterization, we obtain lower bounds on sketches for several distances, including the Earth Mover's Distance (EMD). This part of the thesis is based on the paper:

- Alexandr Andoni, Robert Krauthgamer, Ilya Razenshteyn, Sketching and Embedding are Equivalent for Norms, appeared at STOC 2015 [23], accepted to the special issue of SIAM Journal on Computing.


### 1.2 Near neighbor search (NNS)

The Near Neighbor Search problem (NNS) is defined as follows. Given a dataset of $n$ points, preprocess it to prepare for answering queries of the following kind: given
a query point, find a data point within distance $r$ from the query provided that it exists. Here $r$ is a parameter that is known from the beginning and is the same for all queries.

One can define the NNS problem with respect to any distance function. However, we will mostly consider the case when the dataset and queries lie in a $d$-dimensional vector space $\mathbb{R}^{d}$ equipped with the $\ell_{1}$ or $\ell_{2}$ distances ${ }^{2}$. A useful special case of the $\ell_{1}$ scenario is when the dataset and queries lie in the hypercube $\{0,1\}^{d}$. This corresponds to the Hamming distance ${ }^{3}$. An important special case of the $\ell_{2}$ setting is when data points and queries lie on a unit sphere $S^{d-1} \subset \mathbb{R}^{d}$ : this is equivalent to the near neighbor search with respect to the cosine similarity.

The NNS problem has many applications. Perhaps the most "obvious" one is similarity search for various types of data: text documents, images, audio files, proteins, etc. A typical approach is to take a dataset and map it into $\mathbb{R}^{d}$ by computing a certain feature representation [162]. This step typically requires significant domain expertise and often utilizes techniques from machine learning and/or optimization. But after the feature vectors are computed, the similarity search directly corresponds to executing the NNS queries in the feature space. A sample of other applications includes faster algorithms for cryptanalysis [109], optimization [67, 79, 8], and training neural networks [165], to name a few.

The simplest NNS algorithm is the linear scan: omit the preprocessing stage, and whenever given a query, try all the data points and choose the closest one. However, linear scan ends up being too slow for large datasets and we will be looking for algorithms that have query time sublinear in the number of points $n$. In order for this to be possible, we are going to utilize the preprocessing stage to compute some auxiliary information that will be useful during the query stage.

Three most important parameters of an NNS data structure are:

- Space the data structure occupies. Storing the dataset itself requires $O(d n)$ space,

[^2]but auxiliary information computed during the preprocessing takes additional space to store;

- Time it takes to answer a query. Ideally, processing a query should take time $O(d)$, but usually the query time depends on $n$. In any case, the query time must be strongly sublinear in $n$ (say, at most $n^{0.99}$ );
- Time it takes to build a data structure. Clearly, it cannot be smaller than the amount of space used, but typically the preprocessing time can be made to be near-linear in the space.

The NNS problem exhibits the so-called "curse of dimensionality". Namely, all the known data structures with query time strongly sublinear in $n$ require space $n^{\Omega(d)}$ (see, e.g., $[59,125]$ ), which is super-polynomial as soon as $d$ is super-constant. Moreover, modulo a plausible hardness assumption, there is no data structure for NNS with polynomial in $n$ and $d$ preprocessing time, and query time that is polynomial in $d$ and strongly sublinear in $n[175]^{4}$. To overcome the curse of dimensionality, we will relax the NNS problem somewhat and then observe that the techniques we develop for the relaxed version can be used to solve the exact NNS problem on many real-world datasets.

### 1.2.1 Approximate near neighbor search (ANN)

The relaxed version is called the Approximate Near Neighbor Search problem (ANN) and is defined similarly to the NNS. However, whenever we receive a query with a data point within distance $r$ from it, we are now allowed to return any data point within distance $c r$ from the query (see Figure 1-4). The new parameter $c>1$ is the approximation we are willing to tolerate. Allowing answers to be approximate overcomes the curse of dimensionality and enables very efficient data structures with polynomial dependence on the dimension $d$, space polynomial in $n$ and query time sublinear in $n$ (see [77] for a survey of the state of affairs back in 2012).

[^3]

Figure 1-4: Illustration for the definition of the ANN problem. The red point is a query. We are guaranteed that the smaller ball is non-empty, but are satisfied with any point from the larger ball.

Thus, for the ANN problem we will be studying the trade-off between approximation $c$, space, query time and preprocessing time. We would hope to achieve for every $c>1$ :

- Space polynomial (ideally, near-linear) in $n$ and $d$;
- Query time polynomial (ideally, near-linear) in $d$ and strongly sublinear in $n$;
- Preprocessing time near-linear in the space used.

How good are approximate answers in practice? It turns out that for many real-world datasets, approximation is not really necessary and one can quickly find exact near neighbors. The reason is that oftentimes the following "gap phenomenon" holds: most of the data points are much further from the query than the nearest neighbor ${ }^{5}$ (see Figure 1-5 for an example). Under this gap assumption, most of the ANN algorithms (including the ones described in this thesis) can be shown to return exact nearest neighbors in time sublinear in $n$.

For the ANN problem over $\ell_{1}$ and $\ell_{2}$ distances, we can assume that ${ }^{6} d=\widetilde{O}(\log n)$. Indeed, for the $\ell_{2}$ distance we can apply a random projection onto $\widetilde{O}(\log n)$ dimensions: by the Johnson-Lindenstrauss lemma [93, 65], this procedure incurs distortion of

[^4]distances by a multiplicative factor $1+o(1)$. For the $\ell_{1}$ distance the situation is a bit more delicate, since the dimension reduction as above is known to be impossible [46]. However, since we care only about a single distance scale (" $r$ vs. $c r$ "), there are known procedures for the dimension reduction for this relaxed setting [108]. From now on, we will be implicitly assuming that $d=\widetilde{O}(\log n)$ unless stated otherwise.


Figure 1-5: A distribution of distances in a dataset of GloVe word embeddings [148]; a random data point used as a query.

### 1.2.2 Locality-sensitive hashing (LSH)

One of the main techniques for the high-dimensional ANN problem is locality-sensitive hashing (LSH) [85, 77]. The idea is to design and use a random partition $\mathcal{P}$ of the ambient space $\mathbb{R}^{d}$ with the following properties ${ }^{7}$. For every $x, y \in \mathbb{R}^{d}$ the following conditions are satisfied:

- If the distance between $x$ and $y$ is at most $r$, then

$$
\begin{equation*}
\operatorname{Pr}_{\mathcal{P}}[x \text { and } y \text { end up in the same part of } \mathcal{P}] \geq p_{1} \tag{1.1}
\end{equation*}
$$

[^5]- If the distance between $x$ and $y$ is more than $c r$, then

$$
\begin{equation*}
\operatorname{Pr}_{\mathcal{P}}[x \text { and } y \text { end up in the same part of } \mathcal{P}] \leq p_{2} \tag{1.2}
\end{equation*}
$$

Here $p_{1}$ and $p_{2}$ are two parameters that characterize the quality of an LSH partition. Of course, we must have $p_{1}>p_{2}$, and the bigger the gap between $p_{1}$ and $p_{2}$ is, the higher quality a partition is of.

More precisely, as has been shown in [85], the existence of an "efficient" ${ }^{8}$ partition $\mathcal{P}$ with parameters $p_{1}$ and $p_{2}$ implies a data structure for ANN with preprocessing time and space $n^{1+\rho+o(1)} / p_{1}$ and query time $n^{\rho+o(1)} / p_{1}$, where

$$
\rho=\frac{\log \left(1 / p_{1}\right)}{\log \left(1 / p_{2}\right)}
$$

Thus, the bigger the gap between $p_{1}$ and $p_{2}$ is, the smaller the value of $\rho$ is, and the better bounds on the data structure parameters one is getting. Typically, $\rho=\rho(c)$ is equal to 1 for $c=1$ and $\rho(c) \rightarrow 0$ as $c \rightarrow \infty$.

Having the above reduction in mind, we can try to optimize the value of $\rho$ for an LSH partition. Of course, the optimal value crucially depends on the underlying distance metric, and, as it turns out, the best partitions for the $\ell_{1}$ and $\ell_{2}$ distances are known. Namely, for $\ell_{1}$ the original LSH paper [85] gives $\rho \leq \frac{1}{c}+o(1)$, while for the $\ell_{2}$ distance a longer line of work $[85,66,13]$ leads to the value $\rho \leq \frac{1}{c^{2}}+o(1)$. Translating these bounds into query times, we can get, for instance, query time $n^{1 / 2+o(1)}$ for $\ell_{1}$ and $n^{1 / 4+o(1)}$ for $\ell_{2}$ for approximation $c=2$. Moreover, the above values of $\rho$ are optimal for both $\ell_{1}$ and $\ell_{2}[129,140]$. Thus, we understand LSH almost completely. However:

- A data structure for ANN does not have to fit the LSH framework. Can we get an ANN data structure that bypasses the limitations of LSH?
- For many important metrics beyond $\ell_{1}$ and $\ell_{2}$, LSH with non-trivial guarantees does not exist; a good example is the $\ell_{\infty}$ distance ${ }^{9}$ [81]. How tractable the ANN

[^6]problem is for such metrics?

- The optimal LSH for $\ell_{2}$ from [13], which gives the exponent $\rho \leq \frac{1}{c^{2}}+o(1)$, is highly impractical. Is there a practical analog of it?

In this thesis we make significant progress on all of the above questions.

### 1.2.3 Data-dependent LSH

In $[18,26]$ we show new data structures for high-dimensional ANN over $\ell_{1}$ and $\ell_{2}$ that overcome the above LSH barrier. The new space and time bounds are still of the form $n^{1+\rho+o(1)}$ and $n^{\rho+o(1)}$, but the new values of $\rho$ are:

$$
\begin{equation*}
\rho \leq \frac{1}{2 c-1}+o(1) \tag{1.3}
\end{equation*}
$$

for the $\ell_{1}$ distance, and

$$
\begin{equation*}
\rho \leq \frac{1}{2 c^{2}-1}+o(1) \tag{1.4}
\end{equation*}
$$

for $\ell_{2}$. For instance, for approximation $c=2$, the new data structures give query time $n^{1 / 3+o(1)}$ for $\ell_{1}$ and $n^{1 / 7+o(1)}$ for $\ell_{2}$. At the same time, we now require less memory than the best possible LSH-based approaches.

The main idea behind these results is to allow (random) partitions to depend on the dataset. It turns out that we do not need the full power of conditions (1.1) and (1.2) for every pair of points $x, y \in \mathbb{R}^{d}$ from the ambient $d$-dimensional space that the universal LSH constructions provide. Instead, we can try to design a random partition that works well for a given dataset of $n$ points we know in advance.

At a high level, the new algorithms consist of two parts. First, we design and analyze Voronoi LSH: a new LSH family that yields improved exponents (1.3) and (1.4) under the assumption that a dataset is random ${ }^{10}$. In Voronoi LSH, one samples a number of random landmark points and hashes a point to the closest landmark. Second, to handle the general case, we show how to reduce a worst-case dataset to several pieces

[^7]that look random enough so that Voronoi LSH works well on them. This is done by removing dense clusters of points iteratively and handling them separately. The overall algorithm can be seen as a random collection of $n^{\rho+o(1)}$ decision trees; during the query stage, we query all of them.

Let us make two remarks. First, data-dependent hashing (partitioning) has been ubiquitous in practice (see, e.g., [171, 172]), but, to the best of our knowledge, all the known constructions implicitly or explicitly exploit some nice structure hidden in a dataset. In contrast to these, our algorithms give a provable improvement upon LSH for worst-case datasets. Second, by being data-dependent, we seemingly sacrifice one important property of LSH: it supports adding and removing data points for free. However, we can utilize a generic data structure technique of dynamization [142] and support insertions/deletions of data points in time $n^{\rho+o(1)}$.

In a follow-up work [27] we show optimality of the above bounds (1.3) and (1.4) for any data structure based on data-dependent LSH. We need to be careful which partitions we allow and which we do not. After all, for every dataset its Voronoi diagram gives an "unreasonably high-quality" partition, but it is not very useful algorithmically, since point location for the Voronoi diagram is equivalent to an exact NNS query. Ideally, we would hope to prove a lower bound against partitions that are algorithmically efficient. However, this would require showing data structure lower bounds that are way beyond reach of the currently known techniques. Instead, we require partitions to have strongly sublinear in $n$ description complexity. Indeed, it is the case for all the known algorithms, including [18, 26], and overcoming this barrier is a notorious open problem ${ }^{11}$. The hard distribution used to prove the lower bound is simply a random instance as introduced above.

However, in this thesis we do not describe the above new results, since after publishing the results [18, 26, 27], we significantly extended the data-dependent LSH framework. We introduce this extension next.

[^8]
### 1.2.4 Time-space trade-offs for ANN

Data-independent and data-dependent LSH both provide a very specific relation between space and query time: one gets space $n^{1+\rho+o(1)}$ and query time $n^{\rho+o(1)}$, where $\rho$ is a function of the approximation $c$. However, one can ask if it is possible to trade space for query time and vice versa. For example, in practice space is often a scarcer resource than time, and one may require the space used by a data structure to be $n^{1+o(1)}$. Can we still get sublinear query time in this space-constrained regime?

Researchers have studied various time-space trade-offs for ANN [85, 108, 80, 144, 13, 97]. The culmination of the prior work has been the paper [97] by Kapralov, which shows how to smoothly interpolate between: the space-constrained regime (space $n^{1+o(1)}$, sublinear query time), the balanced regime (space $n^{1+\rho+o(1)}$, query time $n^{\rho+o(1)}$ as for the LSH framework), for which it essentially matches the best data-independent LSH constructions (from [85] and [13]), and the fast queries regime (query time $n^{o(1)}$, polynomial space).

In Chapter 2, which is based on the paper [24], we describe how to extend the data-dependent LSH framework and show how to get an improved time-space trade-off for ANN which:

- Interpolates between the near-linear space regime, the balanced regime, and the fast queries regime;
- In the balanced regime (space $n^{1+\rho+o(1)}$, query time $\left.n^{\rho+o(1)}\right)$, matches the best data-dependent LSH data structures (achieving (1.3) and (1.4) for $\ell_{1}$ and $\ell_{2}$, respectively);
- In all the other regimes, significantly improves upon all the previous work, including the result from [97].

More precisely, we show how to get space $n^{1+\rho_{u}+o(1)}$ and query time $n^{\rho_{q}+o(1)}$ for every $\rho_{u}, \rho_{q} \geq 0$ such that:

$$
\begin{equation*}
c \sqrt{\rho_{q}}+(c-1) \sqrt{\rho_{u}}=\sqrt{2 c-1} \tag{1.5}
\end{equation*}
$$

for $\ell_{1}$ and

$$
\begin{equation*}
c^{2} \sqrt{\rho_{q}}+\left(c^{2}-1\right) \sqrt{\rho_{u}}=\sqrt{2 c^{2}-1} \tag{1.6}
\end{equation*}
$$

for $\ell_{2}$.
For example, for approximation $c=2$ and the $\ell_{1}$ distance, we can interpolate between:

- Space $n^{1+o(1)}$, query time $n^{3 / 4+o(1)}$;
- Space $n^{4 / 3+o(1)}$, query time $n^{1 / 3+o(1)}$;
- Space $n^{4+o(1)}$, query time $n^{o(1)}$.

Similarly, for the $\ell_{2}$ distance, we can get:

- Space $n^{1+o(1)}$, query time $n^{7 / 16+o(1)}=n^{0.4375+o(1)}$;
- Space $n^{8 / 7+o(1)}$, query time $n^{1 / 7+o(1)}$;
- Space $n^{16 / 9+o(1)}=n^{1.77 \ldots}$, query time $n^{o(1)}$.
(See Figure 1-6 for the illustration).
In Chapter 6 (also based on [24]), we show that the time-space trade-offs (1.5) and (1.6) are optimal in two settings:
- When a data structure fits a certain general "hashing framework", which all the known data structures fit;
- If a data structure is arbitrary, but is allowed to inspect at most two memory locations during a query stage;

In both cases we use random instances mentioned in Section 1.2.3 as a hard distribution. The latter lower bound shows and uses an interesting connection between ANN data structures and locally-decodable codes (LDC) (see, e.g., [179] for a survey). Namely, we show that if there had been a too-good-to-be-true ANN data structure that touches $t$ memory locations, we would have gotten a too-good-to-be-true LDC that performs $t$ queries. Such a reduction allows us to utilize a strong lower-bound


Figure 1-6: Illustration of the new time-space trade-off for the $\ell_{2}$ distance and approximation $c=2$. We can interpolate between fast queries $\left(\rho_{q}=0, \rho_{s}=7 / 9=\right.$ $0.77 \ldots$ ), the balanced regime ( $\rho_{q}=\rho_{s}=1 / 7=0.14 \ldots$ ), and the near-linear space regime ( $\rho_{q}=7 / 16=0.43 \ldots, \rho_{s}=0$ ).
for two-query LDCs [101], which gives the desired space lower bound for the ANN problem. However, this approach breaks down completely for $t \geq 3$, since there are strong LDC constructions that make three or more queries. In light of this transition, one can not help but ask if LDCs might be useful for obtaining upper bounds that would bypass the abovementioned lower bounds for the "hashing framework", where the bounds (1.5) and (1.6) are tight.

### 1.2.5 FALCONN: practical and optimal LSH for unit sphere

The next natural question is how practical the new algorithms are. After all, datadependent partitioning is very popular in practice, so it would be natural to try to eliminate the gap between the theory and the practice completely. Unfortunately, in full generality the new algorithms are impractical, especially the reduction from worst-case datasets to "random looking" pieces ${ }^{12}$.

However, at the core of the above algorithms there is a simple construction of

[^9]Voronoi LSH, which is a data-independent LSH partition that gives an improved exponent $\rho$ for random instances. We notice that in order for Voronoi LSH to work well, a dataset does not have to be completely random. Rather, it is enough that the dataset lies on a sphere and is not too concentrated in any small region of it. As it turns out, the latter is a reasonable assumption for many real-world datasets, thus it makes sense to investigate if Voronoi LSH is practical.

Voronoi LSH can also be seen as a (theoretical) improvement upon the celebrated and widely-used Hyperplane LSH [55]. Moreover, Voronoi LSH can be shown to give guarantees similar to the optimal LSH construction for $\ell_{2}$ from [13] for worst-case data in $\mathbb{R}^{d}$. Thus, the question of practicality of Voronoi LSH becomes especially pressing.

But, as it turns out, even Voronoi LSH by itself is not practical due to prohibitively high hashing time. In Chapter 3, which is based on the paper [17], we show how to modify it to make it practical, while essentially preserving the theoretical guarantees.

The new construction can be seen as a combination of three ingredients:

- Cross-Polytope LSH from [168], where it was proposed but not analyzed. We show that the Cross-Polytope LSH, which can be seen as a structured version of the Voronoi LSH, gives the same theoretical guarantees as the vanilla Voronoi LSH.
- Fast pseudo-random rotations based on the Fast Hadamard Transform (FHT) from [6]. Unlike the first ingredient, this step is heuristic: we conjecture that the new LSH scheme achieves the same theoretical guarantees as the pure CrossPolytope LSH or Voronoi LSH, but we currently do not know how to show it. However, there have been follow-up works [100, 43], where this issue has been (partially) addressed.
- If a dataset is extremely high-dimensional and very sparse (e.g., bag of words for text documents), then it pays off to first perform feature hashing [173] down to the intermediate dimension and then apply the above fast variant of the Cross-Polytope LSH.

In practice, the space bound $n^{1+\rho+o(1)}$ is usually too high. This is the reason why we need to adapt the technique of Multiprobe LSH [118] to the Cross-Polytope LSH, which reduces the space usage at a cost of somewhat increased query time. Multiprobe LSH can be seen as a practical analog of the space-constrained algorithm from Section 1.2.4.

We implement the resulting algorithm and release it as a part of FALCONN: a new $\mathrm{C}++$ library with Python bindings for similarity search over high-dimensional data [152]. The implementation requires quite a bit of careful engineering: in particular, we extensively use SIMD to speed up various operations, and we spend lots of effort to make sure that all the necessary data structures occupy as little space as possible. We release our implementation of the Fast Hadamard Transform separately [153]: as it turns out, it is faster than what one can get out of FFTW [75], and we hope it might be useful for other implementations.

We perform a set of experiments on various kinds of data and conclude that FALCONN indeed significantly outperforms Hyperplane LSH as the theory predicts, and is competitive with the state of the art heuristics for high-dimensional similarity search.

### 1.2.6 ANN for general symmetric norms

In light of the above results, the ANN problem for $\ell_{1}$ and $\ell_{2}$ distances has been studied pretty thoroughly. However, our understanding of ANN for more general classes of distances is still very rudimentary.

Since for distances other than $\ell_{1}$ and $\ell_{2}$ there are no known dimension reduction results, we cannot assume that $d=\widetilde{O}(\log n)$ anymore. However, it is convenient to assume that $d=n^{o(1)}$, which means that we are fine with any polynomial dependence on the dimension $d$.

For instance, for the $\ell_{\infty}$ distance, the best known algorithm from [81] achieves for every $\varepsilon>0$ space $n^{1+\varepsilon}$, query time $n^{o(1)}$ and approximation $O_{\varepsilon}(\log \log d)$. This is qualitatively different from $\ell_{1} / \ell_{2}$, where in this regime we would be getting approximation $O_{\varepsilon}(1)$. Somewhat unexpectedly, the $O_{\varepsilon}(\log \log d)$ approximation turns out to
be tight for certain restricted models of computation [11, 98].
In Chapter 4, which is based on the paper [25], we present a new general algorithm that solves ANN for any symmetric norm ${ }^{13}$ (that is, a norm invariant under permutations of coordinates and sign flips), which achieves space $n^{1+o(1)}$, query time $n^{o(1)}$ and approximation $(\log \log n)^{O(1)}$.

The approach that does not quite work is to observe that any norm embeds into $\ell_{\infty}$ with distortion $1+\varepsilon$ for every $\varepsilon>0^{14}$ [176], and then use the data structure for the $\ell_{\infty}$ distance from [81]. The problem with this reduction is that the required dimension of $\ell_{\infty}$ can be as high as $2^{\Omega(d)}$ even if the original norm is as simple as $\ell_{2}$ [32].

The alternative approach is to design an embedding into a space more complicated than $\ell_{\infty}$, but which is nevertheless tractable. This would allow us to save on the dimension making it only polynomial in $d$. Indeed, we show that any symmetric norm embeds into a $d^{O(1)}$-dimensional normed space that is a direct sum of certain sufficiently simple spaces. Then, we invoke a general result from $[82,10]$ that gives an ANN data structure for such direct sums.

Finally, we show that for general, not necessarily symmetric norms a similar approach would not work: any universal normed space with dimension $d^{O(1)}$ that could potentially "host" all the $d$-dimensional normed spaces must incur distortion $d^{1 / 2-o(1)}$.

### 1.3 Characterization of sketchable distances

Sketching is a general notion of compression of a massive object while preserving its useful properties. It is typically used either to reduce the amount of space necessary to store a dataset (in particular, for streaming algorithms [136, 5] or similarity search [171, 172]) or to speed-up various computations [22, 99, 178].

[^10]${ }^{14}$ That is, there is a map $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d^{\prime}}$ such that for every vector $x \in \mathbb{R}^{d}$, one has $\|f(x)\|_{\infty} \in$ $(1 \pm \varepsilon) \cdot\|x\|$.

In this thesis, we are concerned with sketching distance functions. Given a metric ${ }^{15}$ on $\mathbb{R}^{d}$, can we compress vectors so that we can accurately estimate the distance between two vectors given only their compressed versions (sketches)? It turns out that the answer to this question depends crucially on a metric at hand.

For the $\ell_{2}$ distance and for every $\varepsilon>0$, one can compress vectors into $O\left(1 / \varepsilon^{2}\right)$ bits in a randomized manner such that from these sketches we can distinguish $\ell_{2}$ distance at most 1 vs. more than $1+\varepsilon$ with high probability. This result follows from the Johnson-Lindenstrauss Lemma [93, 65] together with a discretization trick from [66]. Note that the sketch size is completely independent of the original dimension $d$ and is determined solely by a desired level of accuracy $\varepsilon>0$. The main idea is to project $\mathbb{R}^{d}$ on a random subspace of dimension $O\left(1 / \varepsilon^{2}\right)$; with high probability, this preserves the Euclidean distance between a given pair of points up to a multiplicative factor of $(1 \pm \varepsilon)$.

For the $\ell_{p}$ distances for $0<p<2$ (including $\ell_{1}$ ), the random projections approach has been generalized by Indyk [84], who showed how to achieve the same bound (sketch size $O\left(1 / \varepsilon^{2}\right)$ for distinguishing distances at most 1 vs. more than $1+\varepsilon$ ). For every $0<p \leq 2$, the bound $O\left(1 / \varepsilon^{2}\right)$ on the sketch size is, in fact, tight [177].

On the other hand, for $\ell_{p}$ distances with $p>2$, the situation is not as nice. It has been shown $[9,34,87]$ that in this case the sketch size must be $\widetilde{\Theta}\left(d^{1-2 / p}\right)$ for any constant $\varepsilon$. The lower bound is shown using communication complexity. In particular, the result for $\ell_{\infty}$ follows from the celebrated linear lower bound for the disjointness problem [96]. Thus, even the simplest case of $\ell_{p}$ distances is quite delicate and mysterious, which justifies the further study of the problem.

The next observation is that a metric admits efficient sketches if it embeds into an $\ell_{p}$ space with $0<p \leq 2$ with a small distortion $D$. That is, there exists a map $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d^{\prime}}$ such that for every $x_{1}, x_{2} \in \mathbb{R}^{d}$ it is true that $\left\|f\left(x_{1}\right)-f\left(x_{2}\right)\right\|_{p}$ is within a multiplicative factor of $D$ from the distance between $x_{1}$ and $x_{2}$ with respect to the

[^11]- $d(x, y)=0$ iff $x=y ;$
- $d(x, z) \leq d(x, y)+d(y, z)$.
metric. If there exists such an embedding, we can sketch our metric with approximation $O(D)$ and constant sketch size by applying the embedding and then using sketches for $\ell_{p}$ from [84].

In Chapter 5, which is based on the paper [23], we show that these are essentially the only possibilities if the metric is a norm. More precisely, we show that if a norm $\|$. admits sketches of size $s$ bits and approximation $D$, then one can embed $\|$.$\| into \ell_{1-\varepsilon}$ with distortion $O(s D / \varepsilon)$ for every $\varepsilon>0$. This provides a complete characterization of "efficiently sketchable" norms. The proof of this result uses a combination of information-theoretic and analytic tools.

If we take the result in the contrapositive, we conclude that for a norm nonembeddability into $\ell_{p}$ spaces for $0<p<2$ implies a lower bound on general sketches. This is very useful, since there are many known tools to show non-embeddability statements, and our result allows to lift them automatically to lower bounds for sketches.

As two concrete applications, we show first sketching lower bounds for the Earth Mover's Distance (EMD) [155] using a non-embeddability result from [134], and for the trace norm (a.k. a. the nuclear norm) using a result from [149].

### 1.4 Preliminaries and notation

We denote by $\mathbb{R}_{+}$the set of non-negative real numbers. Let us denote $\langle\cdot, \cdot\rangle$ the standard dot product on $\mathbb{R}^{d}$. For a positive integer $n$, we define $[n]$ to be the set $\{1,2, \ldots, n\}$. Whenever we use the notation $o(\cdot), O(\cdot), \omega(\cdot)$ or $\Omega(\cdot)$, we write all the parameters the implicit constant depends on as subscripts.

### 1.4.1 Metric and normed spaces

A metric space is a pair $M=(X, d)$, where $X$ is an arbitrary set and $d: X \times X \rightarrow \mathbb{R}_{+}$ is a distance function that satisfies the following properties:

- For every $x, y \in X, d(x, y)=0$ iff $x=y$;
- For every $x, y \in X, d(x, y)=d(y, x)$;
- For every $x, y, z \in X, d(x, z) \leq d(x, y)+d(y, z)$.

We denote $B_{M}(x, R)$ a closed ball in a metric space $M$ centered in $x$ and of radius $R$. More formally, $B_{M}(x, R)=\left\{x^{\prime} \in X \mid d_{X}\left(x, x^{\prime}\right) \leq R\right\}$. Oftentimes, we will omit the subscript if the metric space is clear from the context.

An embedding of a metric space $\left(X, d_{X}\right)$ into a metric space $\left(Y, d_{Y}\right)$ with distortion $D \geq 1$ is a function $f: X \rightarrow Y$ such that there exists a constant $C>0$ such that for every $x_{1}, x_{2} \in X$ one has:

$$
C \cdot d_{X}\left(x_{1}, x_{2}\right) \leq d_{Y}\left(f\left(x_{1}\right), f\left(x_{2}\right)\right) \leq D C \cdot d_{X}\left(x_{1}, x_{2}\right)
$$

A normed space is a pair $\left(\mathbb{R}^{d},\|\cdot\|\right)$, where the function $\|\cdot\|: \mathbb{R}^{d} \rightarrow \mathbb{R}_{+}$is called a norm and satisfies the following properties:

- For every $x \in \mathbb{R}^{d},\|x\|=0$ iff $x=0$;
- For every $\alpha \in \mathbb{R}$ and $x \in \mathbb{R}^{d},\|\alpha x\|=|\alpha| \cdot\|x\|$;
- For every $x, y \in \mathbb{R}^{d},\|x+y\| \leq\|x\|+\|y\|$.

It is immediate to check that any normed space defines a metric on $\mathbb{R}^{d}$ via the formula $d(x, y)=\|x-y\|$.

If we consider embeddings of a normed space into another normed space, we may want to restrict ourselves to linear embeddings, for which the mapping is given by a linear operator $A: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d^{\prime}}$.

If $X=\left(\mathbb{R}^{d},\|\cdot\|\right)$ is a normed space, we can define the dual space $X^{*}=\left(\mathbb{R}^{d},\|\cdot\|^{*}\right)$ as follows:

$$
\|y\|^{*}=\sup _{\|x\| \leq 1}|\langle x, y\rangle| .
$$

It is immediate to check that $\|\cdot\|^{*}$ is indeed a norm. Moreover, one has $\left(X^{*}\right)^{*}=X$. This follows from a finite-dimensional version of the Hahn-Banach theorem [156].

For a normed space $X$, we denote $B_{X}$ the unit ball $B_{X}(0,1)$, which is a centrallysymmetric convex body. Finally, for a ball $B(x, R)$ in a normed space, we denote by $\partial B(x, R)$ its boundary.

### 1.4.2 Examples of metrics and norms

For every $p \geq 1$ and $d \geq 1$ we can define the $\ell_{p}$ norm on $\mathbb{R}^{d}$ as follows:

$$
\|x\|_{p}=\left(\sum_{i=1}^{d}\left|x_{i}\right|^{p}\right)^{1 / p}
$$

We can also define the $\ell_{\infty}$ norm as follows:

$$
\|x\|_{\infty}=\max _{1 \leq i \leq d}\left|x_{i}\right| .
$$

The proof of the fact that the above functions are indeed norms can be found in any introductory textbook in real or functional analysis.

For $0<p<1$, the function $\|\cdot\|_{p}$ is not a norm, since it might be the case that $\|x+y\|_{p}>\|x\|_{p}+\|y\|_{p}$. However, one has:

$$
\|x+y\|_{p} \leq 2^{\frac{1}{p}-1} \cdot\left(\|x\|_{p}+\|y\|_{p}\right) .
$$

Such functions are called quasi-norms. We denote $\ell_{p}^{d}$ the space $\mathbb{R}^{d}$ equipped with the $\ell_{p}$ (quasi-)norm.

The dual to the $\ell_{p}^{d}$ space for $p \geq 1$ is the $\ell_{q}^{d}$ space, where $q \geq 1$ is such that

$$
\frac{1}{p}+\frac{1}{q}=1
$$

For $1 \leq k \leq d$, we define the top- $k$ norm of a vector to be the sum of $k$ largest absolute values of the coordinates. The top-1 norm coincides with the $\ell_{\infty}$ norm, while the top- $d$ norm coincides with the $\ell_{1}$ norm.

For $1 \leq p \leq \infty$, a Schatten- $p$ norm of a square matrix $A \in \mathbb{R}^{n \times n}$ is defined as the
$\ell_{p}$ norm of the vector of the singular values of $A$. The Schatten- 1 norm is also called the trace or the nuclear norm. The Schatten-2 norm coincides with the Frobenius norm, and the Schatten- $\infty$ norm is also called the spectral or operator norm of a matrix.

Let $f:[n]^{2} \rightarrow \mathbb{R}$ be a function such that $\sum_{i, j} f((i, j))=0$ (such functions form an $\left(n^{2}-1\right)$-dimensional vector space). Then the Lipschitz norm of $f$ is defined as:

$$
\|f\|_{\text {Lip }}=\max _{\left(i_{1}, j_{1}\right) \neq\left(i_{2}, j_{2}\right)} \frac{\left|f\left(\left(i_{1}, j_{1}\right)\right)-f\left(\left(i_{2}, j_{2}\right)\right)\right|}{\left|i_{1}-i_{2}\right|+\left|j_{1}-j_{2}\right|} .
$$

The Earth Mover's Distance (EMD) norm is defined to be the dual of the Lipschitz norm (see [134] for more information).

For metric spaces $M_{1}=\left(X_{1}, d_{1}\right), M_{2}=\left(X_{2}, d_{2}\right), \ldots, M_{n}=\left(X_{n}, d_{n}\right)$ and every $p \geq 1$, one can define an $\ell_{p}$-direct sum of $M_{1}, \ldots, M_{n}$ as follows. It is a metric space denoted by $M_{1} \bigoplus_{\ell_{p}} M_{2} \bigoplus_{\ell_{p}} \ldots \bigoplus_{\ell_{p}} M_{n}$ whose ground set is $M_{1} \times M_{2} \times \ldots \times M_{n}$ and the distance function is defined as follows:

$$
d\left(\left(x_{1}, x_{2}, \ldots, x_{n}\right),\left(y_{1}, y_{2}, \ldots, y_{n}\right)\right)=\left\|\left(d_{1}\left(x_{1}, y_{1}\right), d_{2}\left(x_{2}, y_{2}\right), \ldots, d_{n}\left(x_{n}, y_{n}\right)\right)\right\|_{p}
$$

The $\ell_{p}$-direct sum of normed spaces is also a normed space. Confusingly enough, $\ell_{p}$-direct sums are sometimes called $\ell_{p}$-direct products or $\ell_{p}$-products.

### 1.4.3 Approximate near neighbor search (ANN)

Here we introduce one of the two central problems we study in this thesis formally.
Definition 1.4.1 (The $(c, r)$-ANN problem). Given an $n$-point dataset $P$ lying in a metric space $M=\left(X, d_{X}\right)$, the goal is to preprocess it to answer the following queries. Given a query point $q \in X$ such that there exists a data point within distance $r$ from $q$, return a data point within distance $c r$ from $q$. All distances are computed in $M$ using the function $d_{X}(\cdot, \cdot)$.

Both approximation $c$ and the distance scale $r$ are known from the beginning rather than being a part of a query. Usually the ANN problem is studied for the case
when $M$ is a normed $d$-dimensional space, where $d=n^{o(1)}$. We assume that the norm of a vector $x \in \mathbb{R}^{d}$ can be computed in time $d^{O(1)}=n^{o(1)}$.

### 1.4.4 Locality-sensitive hashing (LSH)

Here we introduce locality-sensitive hashing, one of the canonical tools to solve the ANN problem. Next comes the central definition.

Definition 1.4.2. We say that a hash family $\mathcal{H}$ on a metric space $M=\left(X, d_{X}\right)$ is $\left(r_{1}, r_{2}, p_{1}, p_{2}\right)$-sensitive, if for every $p, q \in X$ one has $\operatorname{Pr}_{h \sim \mathcal{H}}[h(x)=h(y)] \geq p_{1}$ if $d_{X}(x, y) \leq r_{1}$, and $\underset{h \sim \mathcal{H}}{\operatorname{Pr}}[h(x)=h(y)] \leq p_{2}$ if $d_{X}(x, y)>r_{2}$,

It is known [85, 77] that an efficient ( $r, c r, p_{1}, p_{2}$ )-sensitive hash family on a $n^{o(1)}$-dimensional normed space implies a data structure for $(c, r)$-ANN with space $n^{1+\rho+o(1)} / p_{1}$ and query time $n^{\rho+o(1)} / p_{1}$, where $\rho=\frac{\log \left(1 / p_{1}\right)}{\log \left(1 / p_{2}\right)}$.

Note that sometimes it is more convenient to talk about random partitions of the metric space $M$. Clearly, random partitions and hash families are equivalent in this context.

Let us briefly sketch this reduction, see [77] for the further details. The reduction depends on two integer parameters: a number of tables $L$ and a number of hash functions in a single table $K$. In each of the $L$ tables, we sample $K$ hash functions $h_{1}, h_{2}, \ldots, h_{K}$ from the LSH family, and hash the dataset using the composite hash function $p \mapsto\left(h_{1}(p), h_{2}(p), \ldots, h_{K}(p)\right)$. Now, to answer a query $q$, we iterate over all the $L$ tables, and in each table we retrieve ${ }^{16}$ the data points that hash to the same bucket as $q$; as soon as we find a point within distance $c r$ from the query, we stop immediately.

Let us show how to set the parameters $K$ and $L$. First, we choose $K$ to be the smallest integer such that $p_{2}^{K}$ is at most $1 / n$. Then, for a fixed query and a given table, there will be at most $n \cdot p_{2}^{K} \leq 1$ far points that fall into the same bucket with the query. Thus, the expected query time is $n^{o(1)} \cdot L$ given that we can evaluate hash functions from the LSH family efficiently. It remains to choose $L$. Since we are a data

[^12]

Figure 1-7: Density of $N(0,1)$.
point within distance $r$ from the query, the probability of success for a single table is at least $p_{1}^{K}=p_{1}^{\left\lceil\frac{\log n}{\log \left(1 / p_{2}\right)}\right\rceil} \geq n^{-\rho} \cdot p_{1}$. Thus, to get a high probability of success, one needs to choose $L=O\left(n^{\rho} / p_{1}\right)$.

### 1.4.5 Gaussian distribution

Our algorithms crucially use the Gaussian distribution $N\left(a, \sigma^{2}\right)$, which is defined as the distribution over the reals $\mathbb{R}$ with density:

$$
f(x)=\frac{1}{\sqrt{2 \pi \sigma^{2}}} \cdot e^{-\frac{(x-a)^{2}}{2 \sigma^{2}}}
$$

(see Figure 1-7). The mean and the variance of $N\left(a, \sigma^{2}\right)$ are $a$ and $\sigma^{2}$, respectively. We denote by $N\left(a, \sigma^{2}\right)^{d}$ a distribution over $\mathbb{R}^{d}$ where each coordinate is an i.i.d. sample from $N\left(a, \sigma^{2}\right)$.

A property of the Gaussian distribution most useful for us is its spherical symmetry. Namely, the density of $N(0,1)^{d}$ in point $v \in \mathbb{R}^{d}$ depends only on the norm $\|v\|_{2}$. In particular, if a vector distributed according to $N(0,1)^{d}$ is multiplied by a fixed orthogonal matrix, the resulting vector is also distributed according to $N(0,1)^{d}$. Spherical symmetry implies 2-stability: if $v$ is distributed according to $N(0,1)^{d}$ and $u \in \mathbb{R}^{d}$ is an arbitrary vector, then $\langle u, v\rangle$ is distributed according to $N\left(0,\|u\|_{2}^{2}\right)$. Indeed,

$$
\langle u, v\rangle=\langle A u, A v\rangle
$$

for any fixed orthogonal matrix $A$. We can choose $A$ so that $A u=\left(\|u\|_{2}, 0, \ldots, 0\right)$, and remembering that, by the spherical symmetry, $A v$ has the same distribution as $v$, we get the desired 2-stability property.

## Chapter 2

## New ANN algorithms for $\ell_{1}$ and $\ell_{2}$

### 2.1 Problem definition

In this chapter, we will show new algorithms for the Approximate Nearest Neighbor search problem for the $\ell_{p}$ distances for $1 \leq p \leq 2$. Recall that the $\ell_{p}$ distance between two vectors $x, y \in \mathbb{R}^{d}$ is defined as follows:

$$
\|x-y\|_{p}= \begin{cases}\left(\sum_{i=1}^{d}\left|x_{i}-y_{i}\right|^{p}\right)^{1 / p} & 1 \leq p<\infty \\ \max _{1 \leq i \leq d}\left|x_{i}-y_{i}\right| & p=\infty\end{cases}
$$

For applications, the two most important cases are $\ell_{1}$ (a.k.a. Manhattan distance) and $\ell_{2}$ (a.k.a. Euclidean distance).

First, let us define the actual problem we will be solving in this chapter.

Definition 2.1.1 (The $(c, r)$-ANN problem). Given an $n$-point dataset $P$ in $\mathbb{R}^{d}$ with $d=n^{o(1)}$, the goal is to preprocess it to answer the following queries. Given a query point $q \in \mathbb{R}^{d}$ such that there exists a data point within $\ell_{p}$ distance $r$ from $q$, return a data point within $\ell_{p}$ distance $c r$ from $q$.

Both approximation $c$ and the distance scale $r$ are known from the beginning rather than being a part of a query.

As we pointed out before, the two most important $\ell_{p}$ distances, both for theory
and for applications, are $\ell_{1}$ and $\ell_{2}$. There are two further important special cases:

- The $(c, r)$-ANN problem for the $\ell_{1}$ distance, when the dataset and queries are promised to lie in the hypercube $\{0,1\}^{d}$. The hypercube equipped with the $\ell_{1}$ distance is often called the Hamming space. Sometimes, it will be convenient for us to work with $\{-1,1\}^{d}$ instead of $\{0,1\}^{d}$.
- The $(c, r)$-ANN problem for the $\ell_{2}$ distance, when the dataset and queries are promised to lie on the unit sphere $S^{d-1}=\left\{x \in \mathbb{R}^{d}:\|x\|_{2}=1\right\} \subset \mathbb{R}^{d}$.

For the $(c, r)$-ANN problem over the whole $\mathbb{R}^{d}$, we can assume w.l.o.g. that $r=1$ by rescaling the dataset and queries. However, this is not the case for the Hamming space and the unit sphere settings. As we will see later in this chapter, for both of these settings, the $(c, r)$-ANN problem becomes easier as $r$ increases ${ }^{1}$.

### 2.1.1 Parameters and guarantees

Data structures for the ANN problem have several characteristics we would like to optimize over:

- Space the data structure occupies;
- Query time: time it takes to answer a single query;
- Preprocessing time: time it takes to build the data structure;
- Insertion/deletion time: in some scenarios we might allow the dataset to change, and we would like to minimize time it takes to add/remove a single point.

As it turns out, the two most fundamental parameters are space and query time: we will be mostly focusing on those, briefly addressing the remaining two parameters towards the end of the chapter. Finally, all the data structures we present are randomized. We guarantee that we return a correct answer for a given query with probability, say, 0.9. This probability is taken over both preprocessing and query

[^13]stages. Such a guarantee is similar to, say, what universal hashing [51] gives for the dictionary problem.

### 2.2 Prior work

In this section we briefly survey the state of the art for the ANN algorithms over the $\ell_{1}$ and $\ell_{2}$ distances prior to the work from this thesis.

### 2.2.1 Locality-sensitive hashing (LSH)

A classic technique for ANN over $\ell_{1}$ and $\ell_{2}$ is Locality-Sensitive Hashing (LSH), introduced in 1998 [85, 77]. The main idea is to use random space partitions of $\mathbb{R}^{d}$, for which a pair of close points (at a distance at most $r$ ) is more likely to belong to the same part than a pair of far points (at a distance more than $c r$ ). Given such a partition, the data structure splits the dataset $P$ according to the partition, and, given a query, checks all the data points which belong to the same part as the query. In order to return a near neighbor with high probability, one needs to maintain several partitions and to check all of them during the query stage. The LSH approach yields data structures with space $n^{1+\rho+o(1)}$ and query time $n^{\rho+o(1)}$, where $\rho$ is the key quantity measuring the quality of a space partition for a particular distance function and approximation $c \geq 1$. Usually, $\rho=1$ for $c=1$ and $\rho \rightarrow 0$ as $c \rightarrow \infty$.

More specifically, the value of $\rho$ is equal to $\frac{\log \left(1 / p_{1}\right)}{\log \left(1 / p_{2}\right)}$, where:

- $p_{1}$ is the lower bound on the probability that two points within distance at most $r$ end up in the same part of a random partition;
- $p_{2}$ is the upper bound on a similar probability for pairs within distance more than cr .

For a sketch of the reduction see Section 1.4.4.
Since the introduction of LSH in [85], subsequent research established optimal values of the LSH exponent $\rho$ for several metrics of interest, including $\ell_{1}$ and $\ell_{2}$. For the $\ell_{1}$ distance, the optimal value is $\rho=\frac{1}{c} \pm o(1)$ [ $\left.85,129,140\right]$. For the $\ell_{2}$ case, the
best possible exponent is $\rho=\frac{1}{c^{2}} \pm o(1)[85,66,13,129,140]$. Since the LSH approach has been really successful both in theory and in practice, it has become a pressing open problem to break through this barrier and obtain algorithms for the ANN problem that are better than the best possible LSH-based algorithms.

### 2.2.2 Time-space trade-offs

Since the early results on LSH, a natural question has been whether one can obtain query time vs. space trade-offs for a fixed approximation $c$. Indeed, data structures with polynomial space and polylogarithmic query time were introduced [85, 108] simultaneously with LSH. In practice, the most important regime is that of near-linear space, since space is usually a harder constraint than time: see, e.g., [118]. This regime has been studied since [80], with subsequent improvements in [144, 13, 97].

For the $\ell_{2}$ distance, the work of Kapralov [97] shows a smooth time-space trade-off that interpolates between all the three abovementioned regimes: near-linear space, balanced LSH regime, and the regime of the low query time. In the LSH regime, it almost matches the best LSH construction by [13], while in other regimes it improves upon the prior work significantly.

### 2.3 The main result

The following is the main result of the present chapter.

Theorem 2.3.1. For every $1 \leq p \leq 2, c>1$ and $\rho_{s}, \rho_{q} \geq 0$ such that:

$$
c^{p} \cdot \sqrt{\rho_{q}}+\left(c^{p}-1\right) \cdot \sqrt{\rho_{s}} \geq \sqrt{2 c^{p}-1}
$$

there exists a data structure for the $c-A N N$ problem over the $\ell_{p}$ distance and n-point datasets in $\mathbb{R}^{d}$ with $d=n^{o(1)}$ that has the following parameters:

- space $n^{1+\rho_{s}+o(1)}$,
- query time $n^{\rho_{q}+o(1)}$,
- preprocessing time $n^{1+\rho_{s}+o(1)}$,
- insertion/deletion time $n^{\rho_{s}+o(1)}$.

In Figure 2-1, we illustrate the main result by instantiating it for various settings. See Figure 2-2 for the plot showing the trade-off between $\rho_{s}$ and $\rho_{q}$ we achieve for $c=2$ and the $\ell_{2}$ distance.

### 2.3.1 Comparison with the prior work

Balanced regime. For the balanced regime (space $n^{1+\rho}$, query time $n^{\rho}$ ), the new result is the first to bypass the LSH barrier. For the $\ell_{1}$ distance, the best possible LSH construction from [85] obtains the exponent $\rho=\frac{1}{c}+o(1)$, while the new exponent is $\frac{1}{2 c-1}+o(1)$. Similarly, for the $\ell_{2}$ distance the best possible LSH from [13] gives $\rho=\frac{1}{c^{2}}+o(1)$, while the new bound is $\frac{1}{2 c^{2}-1}+o(1)$. Thus, in the balanced regime we obtain a polynomial improvement both in space and in query time.

The main insight that allows us to overcome the LSH lower bounds [129, 140] is to allow random partitions to depend on the dataset. The definition of LSH requires certain conditions (close pairs collide with probability at least $p_{1}$, far pairs collide with probability at most $p_{2}$ ) to hold for every pair of points from the ambient space $\mathbb{R}^{d}$. However, for the ANN application, it would be enough if these conditions held for pairs, where one of the points is one of the $n$ dataset points, which we know in advance. That is exactly the route we take. Data-dependent partitions for similarity search and related problems are ubiquitous in practice [171, 172], but the main contribution of this work is to establish their utility for getting improved worst-case bounds. Informally speaking, we show that every dataset has some structure to exploit.

Other regimes. In other regimes we also significantly improve upon the state of the art. In particular, for the near-linear space regime, the new algorithm is the first which achieves sub-linear in $n$ query time for every approximation $c>1$. The previous best result from [97] gets sub-linear time only for $c>\sqrt{3}$.

| Distance | Approximation | Regime | Space | Query time |
| :---: | :---: | :---: | :---: | :---: |
| $\ell_{1}$ | General $c>1$ | Low space | $n^{1+o(1)}$ | $n^{\frac{2}{c}-\frac{1}{c^{2}}+o(1)}$ |
|  |  | Balanced | $n^{1+\frac{1}{2 c-1}+o(1)}$ | $n^{\frac{1}{2 c-1}+o(1)}$ |
|  |  | Fast queries | $n^{\frac{c^{2}}{(c-1)^{2}}+o(1)}$ | $n^{o(1)}$ |
|  | $c=1+\varepsilon, \varepsilon \rightarrow 0$ | Low space | $n^{1+o(1)}$ | $n^{1-\varepsilon^{2}+O\left(\varepsilon^{3}\right)}$ |
|  |  | Balanced | $n^{2-2 \varepsilon+O\left(\varepsilon^{2}\right)}$ | $n^{1-2 \varepsilon+O\left(\varepsilon^{2}\right)}$ |
|  |  | Fast queries | $n^{\frac{1}{\varepsilon^{2}}+O\left(\frac{1}{\varepsilon}\right)}$ | $n^{o(1)}$ |
|  | $c=2$ | Low space | $n^{1+o(1)}$ | $n^{\frac{3}{4}+o(1)}=n^{0.75+o(1)}$ |
|  |  | Balanced | $n^{\frac{4}{3}+o(1)} \approx n^{1.34}$ | $n^{\frac{1}{3}+o(1)} \approx n^{0.34}$ |
|  |  | Fast queries | $n^{4+o(1)}$ | $n^{o(1)}$ |
|  | $c \rightarrow \infty$ | Low space | $n^{1+o(1)}$ | $n^{\frac{2}{c}+O\left(\frac{1}{c^{2}}\right)}$ |
|  |  | Balanced | $n^{1+\frac{1}{2 c}+O\left(\frac{1}{c^{2}}\right)}$ | $n^{\frac{1}{2 c}+O\left(\frac{1}{c^{2}}\right)}$ |
|  |  | Fast queries | $n^{1+\frac{2}{c}+O\left(\frac{1}{c^{2}}\right)}$ | $n^{o(1)}$ |
| $\ell_{2}$ | General $c>1$ | Low space | $n^{1+o(1)}$ | $n{ }^{\frac{2}{c^{2}}-\frac{1}{c^{4}}+o(1)}$ |
|  |  | Balanced | $n^{1+\frac{1}{2 c^{2}-1}+o(1)}$ | $n^{\frac{1}{2 c^{2}-1}+o(1)}$ |
|  |  | Fast queries | $n^{\frac{c^{4}}{\left(c^{2}-1\right)^{2}}+o(1)}$ | $n^{o(1)}$ |
|  | $c=1+\varepsilon, \varepsilon \rightarrow 0$ | Low space | $n^{1+o(1)}$ | $n^{1-4 \varepsilon^{2}+O\left(\varepsilon^{3}\right)}$ |
|  |  | Balanced | $n^{2-4 \varepsilon+O\left(\varepsilon^{2}\right)}$ | $n^{1-4 \varepsilon+O\left(\varepsilon^{2}\right)}$ |
|  |  | Fast queries | $n^{\frac{1}{4 \varepsilon^{2}}+O\left(\frac{1}{\varepsilon}\right)}$ | $n^{o(1)}$ |
|  | $c=2$ | Low space | $n^{1+o(1)}$ | $n^{\frac{7}{16}+o(1)} \approx n^{0.44}$ |
|  |  | Balanced | $n^{\frac{8}{7}+o(1)} \approx n^{1.15}$ | $n^{\frac{1}{7}+o(1)} \approx n^{0.15}$ |
|  |  | Fast queries | $n^{\frac{16}{9}+o(1)} \approx n^{1.78}$ | $n^{o(1)}$ |
|  | $c \rightarrow \infty$ | Low space | $n^{1+o(1)}$ | $n n^{\frac{2}{c^{2}}+O\left(\frac{1}{c^{4}}\right)}$ |
|  |  | Balanced | $n^{1+\frac{1}{2 c^{2}}+O\left(\frac{1}{c^{4}}\right)}$ | $n^{\frac{1}{2 c^{2}}+O\left(\frac{1}{c^{4}}\right)}$ |
|  |  | Fast queries | $n^{1+\frac{2}{c^{2}}+O\left(\frac{1}{c^{4}}\right)}$ | $n^{o(1)}$ |

Figure 2-1: Specializations of the main result (Theorem 2.3.1) for various settings. This includes: $\ell_{1}$ and $\ell_{2}$ distances; general approximation $c, c$ close to $1, c=2$, large $c$; the regime of space $n^{1+o(1)}$, the balanced regime (space $n^{1+\rho}$, query time $n^{\rho}$ ) and the regime of fast query time $n^{o(1)}$.


Figure 2-2: Illustration of the main result (Theorem 2.3.1) for the $\ell_{2}$ distance and approximation $c=2$. The time-space trade-off interpolates between fast queries $\left(\rho_{q}=0, \rho_{s}=7 / 9=0.77 \ldots\right)$, the balanced regime $\left(\rho_{q}=\rho_{s}=1 / 7=0.14 \ldots\right)$, and the near-linear space regime ( $\rho_{q}=7 / 16=0.43 \ldots, \rho_{s}=0$ ).

We compare the new algorithm with the prior state of the art ([13] in the balanced regime, and [97] in other regimes) for the $\ell_{2}$ distance and $c=2$ on Figure 2-3.

### 2.4 Overview of the algorithm

We now describe the proof of Theorem 2.3.1 at a high level. First of all, in Section 2.5, using by now standard tools, we reduce the general $(c, r)$-ANN problem for the $\ell_{p}$ distances to the $(c-o(1), r)$-ANN problem for the $\ell_{2}$ distance on the unit sphere $S^{d-1} \subset \mathbb{R}^{d}$, where $d=\log ^{1+o(1)} n$ and $r=\frac{1}{\log \log n}$. For this case it is enough to achieve the following trade-off:

$$
\begin{equation*}
c^{2} \cdot \sqrt{\rho_{q}}+\left(c^{2}-1\right) \cdot \sqrt{\rho_{s}} \geq \sqrt{2 c^{2}-1} \tag{2.1}
\end{equation*}
$$

The new algorithm for the unit sphere case consists of two major stages. In the first stage (Section 2.8), we give an algorithm for random instances (introduced formally in Section 2.6). To generate such an instance, we sample a dataset uniformly at


Figure 2-3: Comparison of the main result (Theorem 2.3.1) with the prior state of the art for the $\ell_{2}$ distance and $c=2$. In the balanced regime, the best prior bound follows from the work of [13], while in all the other regimes, the best bound has been shown in [97]. For $\rho_{q}=0$, the algorithm from [97] gives $\rho_{s}=4$ (this point does not fit on the plot).
random on a unit sphere $S^{d-1} \subset \mathbb{R}^{d}$ and plant a query at random within distance $\sqrt{2} / c$ from a randomly chosen data point. In the second stage, we show the claimed result for the worst-case instances by combining the algorithm from the first stage with data-dependent partitioning.

Data-independent partitions. To handle random instances, we use a certain data-independent partitioning process, which we briefly introduce below. It can be seen as a modification of spherical Locality-Sensitive Filtering from [38], and is related to a data structure from [146]. While this data-independent approach can be extended to worst case instances, it gives a bound worse than the desired (2.1).

We now describe the partitioning process which produces a decision tree to solve an ANN instance on the unit sphere $S^{d-1}$. We take our initial dataset $P \subset S^{d-1}$ and sample $T$ i.i.d. standard Gaussian $d$-dimensional vectors $z_{1}, z_{2}, \ldots, z_{T} \in \mathbb{R}^{d}$. The sets
$P_{i} \subseteq P$ (not necessarily disjoint) are defined for each $z_{i}$ as follows:

$$
P_{i}=\left\{p \in P \mid\left\langle z_{i}, p\right\rangle \geq \eta_{s}\right\} .
$$

We then recurse and repeat the above procedure for each non-empty $P_{i}$. We stop the recursion once we reach depth $K$. The above procedure generates a tree of depth $K$ and degree at most $T$, where each leaf explicitly stores the corresponding subset of the dataset. To answer a query $q \in S^{d-1}$, we start at the root and descend into (potentially multiple) $P_{i}$ 's for which $\left\langle z_{i}, q\right\rangle \geq \eta_{q}$. When we eventually reach the $K$-th level, we iterate through all the points stored in the leaf searching for a near neighbor.

The parameters $T, K, \eta_{s}$ and $\eta_{q}$ depend on the distance threshold $r$, the approximation factor $c$, as well as the desired space and query time exponents $\rho_{s}$ and $\rho_{q}$. The special case of $\eta_{s}=\eta_{q}$ corresponds to the balanced regime $\rho_{s}=\rho_{q} ; \eta_{s}<\eta_{q}$ corresponds to the "fast queries" regime $\rho_{q}<\rho_{s}$ (the query procedure is more selective); and $\eta_{s}>\eta_{q}$ corresponds to the "low memory" regime $\rho_{s}<\rho_{q}$. The analysis of this algorithm relies on bounds on the Gaussian area of certain two-dimensional sets (see Section 2.7).

This algorithm has two important consequences. First, we obtain the desired trade-off (2.1) for random instances by setting $r=\frac{\sqrt{2}}{c}$. Second, we obtain a worse trade-off trade-off for $r=\frac{1}{\log \log n}$. Namely, we get:

$$
\begin{equation*}
\left(c^{2}+1\right) \sqrt{\rho_{q}}+\left(c^{2}-1\right) \cdot \sqrt{\rho_{s}} \geq 2 c \tag{2.2}
\end{equation*}
$$

Even though it is worse than the desired bound from $(2.1)^{2}$, it is already non-trivial. In particular, (2.2) is already better than all the prior work on time-space trade-offs for ANN.

Data-dependent partitions. We then improve upon (2.2) for worst-case instances and obtain the final result, Theorem 2.3.1. For this we employ data-dependent partitioning.

[^14]

Figure 2-4: Illustration of the data-independent bound (2.2) for $c=2$. In the balanced regime, it matches the best LSH from [13], while in other regimes it improves upon the work of [97]. The final bound is polynomially better than the data-independent one.

If pairwise distances between data points are distributed roughly like for a random instance, we could apply the data-independent procedure. In absence of such a guarantee, we manipulate the dataset in order to reduce it to a random-looking case. Namely, we search for low-diameter clusters that contain many data points. We extract these clusters, and we enclose each of them in a ball of radius non-trivially smaller than one, and we recurse on each cluster. For the remaining points, which do not lie in any cluster, we perform one step of the data-independent algorithm: we sample $T$ Gaussian vectors, form $T$ subsets of the dataset, and recurse on each subset. Overall, we make progress in two ways: for the clusters, we make them a bit more isotropic after re-centering and shrinking, which, after several re-centerings, makes the instance amenable to the data-independent algorithm, and for the remainder of the points, we can show that the absence of dense clusters makes the data-independent algorithm work for a single level of the tree (though, when recursing into $P_{i}$ 's, dense clusters may re-appear, which we will need to extract again).

While the above intuition is very simple and, in hindsight, natural, the actual execution requires a good amount of work. For example, we need to formalize "low-
diameter", "lots of points", "more isotropic", etc.
Finally, getting fast preprocessing time $n^{1+\rho_{s}+o(1)}$ and fast insertion/deletion times $n^{\rho_{s}+o(1)}$ requires additional work. For the former, we show how to quickly find dense cluster using sampling (see Section 2.9.5). For the latter we employ a standard dynamization technique [142] (see Section 2.9.6).

### 2.5 Simplification of the problem

It can be shown that in order to prove Theorem 2.3.1, it is enough to establish the following special case of it.

Theorem 2.5.1. For every $c>1$ and $\rho_{s}, \rho_{q} \geq 0$ such that:

$$
\begin{equation*}
c^{2} \cdot \sqrt{\rho_{q}}+\left(c^{2}-1\right) \cdot \sqrt{\rho_{s}} \geq \sqrt{2 c^{2}-1} \tag{2.3}
\end{equation*}
$$

there exists a data structure for the ( $c, r$ )-ANN problem over the $\ell_{2}$ distance with $n$-point datasets and queries lying on the unit sphere $S^{d-1} \subset \mathbb{R}^{d}$ with $d=\log ^{1+o(1)} n$ and $r=\frac{1}{\sqrt{\log \log n}}$ that has the following parameters:

- space $n^{1+\rho_{s}+o(1)}$,
- query time $n^{\rho_{q}+o(1)}$,
- preprocessing time $n^{1+\rho_{s}+o(1)}$,
- insertion/deletion times $n^{\rho_{s}+o(1)}$.

In short, it is enough to handle the $(c, r)$-ANN problem over the $\ell_{2}$ distance on the unit sphere in $\mathbb{R}^{d}$ for $d=\log ^{1+o(1)} n$ and $r=\frac{1}{\log \log n}=o(1)$.

The proof of the fact that Theorem 2.5.1 implies Theorem 2.3.1 is by now standard if a bit tedious. Let us list the necessary steps and provide pointers to the literature.

- We reduce the $(c, r)$-ANN problems over the $\ell_{p}$ distance to the $\left(c^{p / 2}, 1\right)$-ANN problem over the $\ell_{2}$ distance. This reduction is described in detail in [137].
- We reduce the dimension to $d^{\prime}=\log ^{1+o(1)} n$ by using the dimension reduction lemma of Johnson and Lindenstrauss [93, 65]. This step introduces multiplicative distortion $1 \pm o(1)$ for pairwise distances, which is acceptable for us.
- Next, we reduce the diameter of the dataset to $O\left((\log \log n)^{1 / 4}\right)$. This can be done by partitioning the dataset using LSH family from [66] and querying the part, where the query belongs. We need to repeat this procedure $n^{O\left(\frac{1}{(\log \log n)^{1 / 4}}\right)}=n^{o(1)}$ times to get high probability of success.
- Finally, we reduce the problem to the unit sphere case with $r=\frac{1}{\sqrt{\log \log n}}$. This reduction can be found in [169].

Furthermore, from now on we will mostly focus on bounding the space and query time, postponing the preprocessing and insertion/deletion times until Sections 2.9.5 and 2.9.6, respectively.

### 2.6 Random ANN instances

Let us look at the complexity of $(c, r)$-ANN on the unit sphere $S^{d-1}$ when $c$ is fixed. As we saw in Section 2.5, the case $r=o(1)$ is the hardest, since it implies a $c$-ANN algorithm for the whole $\mathbb{R}^{d}$ with essentially the same guarantees. At the other extreme is the case $r \geq \frac{2}{c}$. Since the diameter of the sphere is 2 , if $c r \geq 2$, any data point is an answer for any valid query. Hence, in this case, the $(c, r)$-ANN problem is trivial.

As it turns out, there exists an intermediate regime: $r=\frac{\sqrt{2}}{c}$, which corresponds to the following random instances.

- A dataset $P \subset S^{d-1}$ is given by $n$ unit vectors, where each vector is drawn independently and uniformly at random from $S^{d-1}$. We assume that $d=\omega(\log n)$.
- A query $q \in S^{d-1}$ is drawn by first choosing a dataset point $p \in P$ uniformly at random, and then choosing $q$ uniformly at random from all points in $S^{d-1}$ within distance $\frac{\sqrt{2}}{c}$ from $p$.
- The goal of the data structure is to preprocess $P$ in order to recover the data point $p$ from the query point $q$ quickly.

Any $\left(c-o(1), \frac{\sqrt{2}}{c}\right)$-ANN data structure for the unit sphere must be able to find $p$ given $q$ with high probability. Indeed, the distance between $p$ and $q$ is at most $\frac{\sqrt{2}}{c}$, while all the distances from $q$ to other data points are at least $\sqrt{2}-o(1)$ with high probability, since $d=\omega(\log n)^{3}$.

As it turns out, handling these random instances are crucial for getting improved ANN algorithms for the general case. Namely, the proof of Theorem 2.5.1 proceeds in two big steps:

- We show how to get the desired bounds (2.1) on space and query time for $(c-o(1), r)$-ANN with $r=\frac{\sqrt{2}}{c}$, and, as discussed above, this setting is enough to handle random instances;
- Then, we will see how to reduce the setting of Theorem 2.5.1, $r=o(1)$, to the above case of $r=\frac{\sqrt{2}}{c}$ by finding subsets of the dataset that "look random".

Section 2.8 is devoted to the first part, while (a much more involved) Section 2.9 describes the second part.

### 2.7 Gaussians and tail bounds

Here we develop a bit the theory of the Gaussian distribution introduced in Section 1.4.5.

For $0 \leq s \leq 2$, let $\alpha(s)=1-\frac{s^{2}}{2}$ be the cosine of the angle between two points on a unit Euclidean sphere $S^{d-1}$ with distance $s$ between them, and $\beta(s)=\sqrt{1-\alpha^{2}(s)}$ be the sine of the same angle (see Figure 2-5).

We introduce two functions that will be useful later. First, for $\rho>0$, let

$$
F(\rho)=\operatorname{Pr}_{z \sim N(0,1)^{d}}[\langle z, u\rangle \geq \rho],
$$

[^15]

Figure 2-5: Illustration of the definition of $\alpha(s)$ and $\beta(s)$.


Figure 2-6: Illustration of the definition of $A$.
where $u \in S^{d-1}$ is an arbitrary point on the unit sphere. Note that $F(\rho)$ does not depend on the specific choice of $u$ due to the spherical symmetry of $N(0,1)^{d}$.

Second, for $0 \leq s \leq 2$ and $\rho, \sigma>0$, let

$$
G(s, \rho, \sigma)=\operatorname{Pr}_{z \sim N(0,1)^{d}}[\langle z, u\rangle \geq \rho \text { and }\langle z, v\rangle \geq \sigma]
$$

where $u, v \in S^{d-1}$ are arbitrary points from the unit sphere with $\|u-v\|_{2}=s$. As with $F$, the value of $G(s, \rho, \sigma)$ does not depend on the specific choice of $u$ and $v$; it only depends on the distance $s=\|u-v\|_{2}$ between them. Clearly, $G(s, \rho, \sigma)$ is non-increasing in $s$, for fixed $\rho$ and $\sigma$.

We state two useful bounds on $F(\cdot)$ and $G(\cdot, \cdot, \cdot)$. The first is a standard tail bound for the Gaussian distribution and the second can be seen as a certain two-dimensional generalization of the first.

Lemma 2.7.1 ([74]). For $\rho \rightarrow \infty$,

$$
F(\rho)=e^{-\left(1 \pm O\left(\frac{\log \rho}{\rho^{2}}\right)\right) \cdot \frac{\rho^{2}}{2}}
$$

Proof. By the 2-stability,

$$
F(\rho)=\operatorname{Pr}_{z \sim N(0,1)^{d}}[\langle z, u\rangle \geq \rho]=\operatorname{Pr}_{x \sim N(0,1)}[x \geq \rho]=\frac{1}{\sqrt{2 \pi}} \int_{\rho}^{\infty} e^{-x^{2} / 2} d x
$$

We can bound the latter integral as follows:

$$
\begin{aligned}
\left(\frac{1}{\rho}-\frac{1}{\rho^{3}}\right) \cdot \frac{e^{-\rho^{2} / 2}}{\sqrt{2 \pi}} & =\frac{1}{\sqrt{2 \pi}} \int_{\rho}^{\infty}\left(1-\frac{3}{x^{4}}\right) e^{-x^{2} / 2} d x \\
& \leq \frac{1}{\sqrt{2 \pi}} \int_{\rho}^{\infty} e^{-x^{2} / 2} d x \\
& \leq \frac{1}{\sqrt{2 \pi}} \int_{\rho}^{\infty}\left(1+\frac{1}{x^{2}}\right) e^{-x^{2} / 2} d x=\frac{1}{\rho} \cdot \frac{e^{-\rho^{2} / 2}}{\sqrt{2 \pi}}
\end{aligned}
$$

which implies the required statement.

Lemma 2.7.2. If $\rho, \sigma \rightarrow \infty$, then, for every $0<s<2$, one has:

$$
G(s, \rho, \sigma)=e^{-\left(1 \pm O\left(\frac{\log \Delta}{\Delta^{2}}\right)\right) \cdot \frac{\Delta^{2}}{2}}
$$

where:

$$
\Delta^{2}= \begin{cases}\max \left\{\rho^{2}, \sigma^{2}\right\}, & \text { if } \frac{\min \{\rho, \sigma\}}{\max \{\rho, \sigma\}}<\alpha(s) \leq 1  \tag{2.4}\\ \frac{\rho^{2}+\sigma^{2}-2 \alpha(s) \cdot \rho \sigma}{\beta^{2}(s)}, & \text { otherwise }\end{cases}
$$

Proof. By the spherical symmetry,

$$
\begin{aligned}
G(s, \rho, \sigma) & =\operatorname{Pr}_{z \sim N(0,1)^{d}}[\langle z, u\rangle \geq \rho \text { and }\langle z, v\rangle \geq \sigma] \\
& =\operatorname{Pr}_{x, y \sim N(0,1)}[x \geq \rho \text { and } \alpha(s) \cdot x+\beta(s) \cdot y \geq \sigma] .
\end{aligned}
$$


(a) The closest point lies "on a side".

Figure 2-7: Two cases for $\Delta_{A}$ correspond to the two cases in the statement of Lemma 2.7.2.

The latter quantity is the Gaussian measure $\mathcal{G}(A)$ of the two-dimensional set

$$
A=\left\{(x, y) \in \mathbb{R}^{2} \mid x \geq \rho \text { and } \alpha(s) \cdot x+\beta(s) \cdot y \geq \sigma\right\}
$$

(see Figure 2-6). Let us denote $\Delta_{A}$ the distance from the origin to $A$. It is not hard to check that $\Delta_{A}$ is equal to $\Delta$ from the statement of the lemma. The two cases in (2.4) correspond to Figure 2-7a and Figure 2-7b, respectively. Thus, it remains to show that if $\rho, \sigma \rightarrow \infty$, then

$$
\mathcal{G}(A)=e^{-\left(1 \pm O\left(\frac{\log \Delta_{A}}{\Delta_{A}^{2}}\right)\right) \cdot \frac{\Delta_{A}^{2}}{2}}
$$

For $r>0$, let us denote $0 \leq \mu_{A}(r) \leq 1$ the normalized measure of the intersection $A \cap \partial B(0, r):$

$$
\mu_{A}(r)=\frac{\mu(A \cap \partial B(0, r))}{2 \pi r}
$$

where $\mu$ is the standard one-dimension Lebesgue measure (see Figure 2-8). Then, integrating in the polar coordinates, we get:

$$
\begin{equation*}
\mathcal{G}(A)=\int_{\Delta_{A}}^{\infty} \mu_{A}(r) \cdot r e^{-r^{2} / 2} d r \tag{2.5}
\end{equation*}
$$

We can upper bound (2.5) as follows:

$$
\int_{\Delta_{A}}^{\infty} \mu_{A}(r) \cdot r e^{-r^{2} / 2} d r \leq \int_{\Delta_{A}}^{\infty} r e^{-r^{2} / 2} d r=e^{-\Delta_{A}^{2} / 2}
$$

On the lower bound side, since $\mu_{A}(r)$ is non-decreasing as $r$ increases, we get for every $\varepsilon>0$ :
$\int_{\Delta_{A}}^{\infty} \mu_{A}(r) \cdot r e^{-r^{2} / 2} d r \geq \mu_{A}\left((1+\varepsilon) \Delta_{A}\right) \int_{(1+\varepsilon) \Delta_{A}}^{\infty} r e^{-r^{2} / 2} d r=\mu_{A}\left((1+\varepsilon) \Delta_{A}\right) e^{-(1+\varepsilon)^{2} \Delta_{A}^{2} / 2}$.

Thus, to get the desired bound, it is enough to show that as $\varepsilon$ tends to zero, we have:

$$
\mu_{A}\left((1+\varepsilon) \Delta_{A}\right) \geq\left(\frac{\varepsilon}{\Delta_{A}}\right)^{O(1)}
$$

This is a pretty straightforward computation. If we are in the situation of Figure 2-7b, then one can show that:

$$
\begin{aligned}
\mu\left((1+\varepsilon) \Delta_{A}\right) & =\frac{1}{2 \pi} \cdot\left(\arccos \left(\frac{\rho}{(1+\varepsilon) \Delta_{A}}\right)-\arccos \left(\frac{\rho}{\Delta_{A}}\right)\right. \\
& \left.+\arccos \left(\frac{\sigma}{(1+\varepsilon) \Delta_{A}}\right)-\arccos \left(\frac{\sigma}{\Delta_{A}}\right)\right) \\
& \geq \Omega\left(\frac{\varepsilon}{\Delta_{A}^{2}}\right)
\end{aligned}
$$

as required. Otherwise, the situation is like on Figure 2-7a, and we have $\mu\left((1+\varepsilon) \Delta_{A}\right)=$ $\Omega\left(\varepsilon^{1 / 2}\right)$, again, as required.

### 2.8 Data-independent partitions

Here we fulfill the first step of the program outlined in Section 2.4. We show a simple algorithm based on data-independent partitioning that gives the trade-off (2.2) for worst-case instances of $(c, r)$-ANN on the unit sphere and a better trade-off (2.1) for random instances (as defined in Section 2.6). Later we will show how to achieve


Figure 2-8: Illustration of the definition of $\mu_{A}(r)$.
the trade-off (2.1) for worst-case instances, which will imply Theorem 2.5.1 and, by Section 2.5, the final result (Theorem 2.3.1). See Figure 2-4 for the comparison of the trade-offs for $c=2$.

### 2.8.1 Results

Below is the main theorem we prove in Section 2.8.

Theorem 2.8.1. Let $\varepsilon_{0}>0$ be a (small) fixed positive constant. For every $c>1$, $\frac{1}{\log \log n} \leq r \leq \sqrt{2}-\varepsilon_{0}, \rho_{q} \geq 0$ and $\rho_{s} \geq 0$ such that $c r \leq 2-\varepsilon_{0}$ and

$$
\begin{equation*}
(1-\alpha(r) \alpha(c r)) \cdot \sqrt{\rho_{q}}+(\alpha(r)-\alpha(c r)) \cdot \sqrt{\rho_{s}} \geq \beta(r) \beta(c r), \tag{2.6}
\end{equation*}
$$

there exists a data structure for $(c, r)-A N N$ on a unit sphere $S^{d-1} \subset \mathbb{R}^{d}$, where $d=n^{o(1)}$, with space $n^{1+\rho_{s}+o(1)}$ and query time $n^{\rho_{q}+o(1)}$.

Recall that $\alpha(\cdot)$ and $\beta(\cdot)$ are the functions defined in the beginning of Section 2.7. Let us see what this theorem gives for $r=o(1)$ and $r=\frac{\sqrt{2}}{c}$.

Corollary 2.8.2. If $\frac{1}{\log \log n} \leq r=o(1)$, then Theorem 2.8.1 holds as soon as:

$$
\begin{equation*}
\left(c^{2}+1\right) \cdot \sqrt{\rho_{q}}+\left(c^{2}-1\right) \cdot \sqrt{\rho_{s}} \geq 2 c \tag{2.7}
\end{equation*}
$$

Proof. We use the following estimates that are easy to check:

$$
\begin{aligned}
1-\alpha(r) \alpha(c r) & =\frac{\left(c^{2}+1\right) r^{2}}{2}+O_{c}\left(r^{4}\right) \\
\alpha(r)-\alpha(c r) & =\frac{\left(c^{2}-1\right) r^{2}}{2} \\
\beta(r) \beta(c r) & =c r^{2}+O_{c}\left(r^{3}\right)
\end{aligned}
$$

Plugging these estimates into (2.6), we obtain the result.

Thus, using the algorithm from this section, one can immediately obtain a version of Theorem 2.5.1 with the trade-off (2.7) instead of the required (2.1).

Corollary 2.8.3. If $r=\frac{\sqrt{2}}{c}$, then Theorem 2.8.1 holds as soon as:

$$
\begin{equation*}
c^{2} \cdot \sqrt{\rho_{q}}+\left(c^{2}-1\right) \cdot \sqrt{\rho_{s}} \geq \sqrt{2 c^{2}-1} \tag{2.8}
\end{equation*}
$$

Proof. Since $c>1$, one has $r \leq \sqrt{2}-\Omega(1)$, so we can apply Theorem 2.8.1. We use the following identities: $\alpha(r)=1-\frac{1}{c^{2}}, \alpha(c r)=\alpha(\sqrt{2})=0, \beta(r)=\sqrt{\frac{2}{c^{2}}-\frac{1}{c^{4}}}$, $\beta(c r)=\beta(\sqrt{2})=1$, which, as one can check, imply the required bound.

This corollary implies that the algorithm from this section applied to random instances gives the trade-off (2.8), which we would like to get in full generality to conclude the proof of Theorem 2.5.1. The remainder of Section 2.8 is devoted to proving Theorem 2.8.1.

### 2.8.2 Description

Fix $K$ and $T$ to be positive integers, we determine their exact value later. Our data structure is a rooted tree where each node corresponds to a spherical cap. The tree consists of $K+1$ levels of nodes where each node has out-degree at most $T$. We will index the levels by $0,1, \ldots, K$, where the 0 -th level consists of the root denoted by $v_{0}$, and each node up to the $(K-1)$-th level has at most $T$ children. Therefore, there are at most $T^{K}$ nodes at the $K$-th level.

For every node $v$ in the tree, let $\mathcal{L}_{v}$ be the set of nodes on the path from $v$ to the root $v_{0}$ excluding the root (but including $v$ ). Each node $v$, except for the root, stores a random Gaussian vector $z_{v} \sim N(0,1)^{d}$. For each node $v$, we define the following subset of the dataset $P_{v} \subseteq P$ :

$$
P_{v}=\left\{p \in P \mid \forall v^{\prime} \in \mathcal{L}_{v}\left\langle z_{v^{\prime}}, p\right\rangle \geq \eta_{s}\right\},
$$

where $\eta_{s}>0$ is a parameter to be chosen later.

At the root node $v_{0}, P_{v_{0}}=P$, since $\mathcal{L}_{v_{0}}=\emptyset$. Intuitively, each set $P_{v}$ corresponds to a subset of the dataset lying in the intersection of spherical caps centered around $z_{v^{\prime}}$ for all $v^{\prime} \in \mathcal{L}_{v}$. Every leaf $\ell$ at the level $K$ stores the subset $P_{\ell}$ explicitly.

We build the tree recursively. For a given node $v$ in levels $0, \ldots, K-1$, we first sample $T$ i.i.d. Gaussian vectors $g_{1}, g_{2}, \ldots, g_{T} \sim N(0,1)^{d}$. Then, for every $i$ such that $\left\{p \in P_{v} \mid\left\langle g_{i}, p\right\rangle \geq \eta_{s}\right\}$ is non-empty, we create a new child $v^{\prime}$ with $z_{v^{\prime}}=g_{i}$ and recursively process $v^{\prime}$. At the $K$-th level, each node $v$ stores $P_{v}$ as a list of points.

In order to answer a query $q \in S^{d-1}$, we start from the root $v_{0}$ and descend down the tree. We consider every child $v$ of the root for which $\left\langle z_{v}, q\right\rangle \geq \eta_{q}$, where $\eta_{q}>0$ is another parameter to be chosen later ${ }^{4}$. After identifying all such children, we proceed down the children recursively. If we reach leaf $\ell$ at level $K$, we scan through all the points in $P_{\ell}$ and compute their distance to the query $q$. If a point lies at a distance at most $c r$ from the query, we return it and stop immediately.

We provide pseudocode for the data structure above in Figure 2-9. The procedure $\operatorname{Build}(P, 0, \perp)$ builds the data structure for dataset $P$ and returns the root of the tree, $v_{0}$. The procedure $\operatorname{QUERy}\left(q, v_{0}\right)$ queries the data structure with root $v_{0}$ at point $q$.

[^16]```
function \(\operatorname{Build}\left(P^{\prime}, l, z\right)\)
    create a tree node \(v\)
    store \(l\) as \(v . l\)
    store \(z\) as \(v . z\)
    if \(l=K\) then
        store \(P^{\prime}\) as \(v . P\)
    else
        for \(i \leftarrow 1 \ldots T\) do
        sample a Gaussian vector \(z^{\prime} \sim N(0,1)^{d}\)
        \(P^{\prime \prime} \leftarrow\left\{p \in P^{\prime} \mid\left\langle z^{\prime}, p\right\rangle \geq \eta_{s}\right\}\)
        if \(P^{\prime \prime} \neq \emptyset\) then
            add \(\operatorname{Build}\left(P^{\prime \prime}, l+1, z^{\prime}\right)\) as a child of \(v\)
    return \(v\)
function \(\operatorname{QUERY}(q, v)\)
    if \(v . l=K\) then
        for \(p \in v . P\) do
        if \(\|p-q\| \leq c r\) then
            return \(p\)
    else
        for \(v^{\prime}: v^{\prime}\) is a child of \(v\) do
            if \(\left\langle v^{\prime} . z, q\right\rangle \geq \eta_{q}\) then
            \(p \leftarrow \operatorname{QUERy}\left(q, v^{\prime}\right)\)
            if \(p \neq \perp\) then
                return \(p\)
    return \(\perp\)
```

Figure 2-9: Pseudocode for data-independent partitions

### 2.8.3 Analysis

In this section we analyze the above data structure. During the analysis, we will use functions $F(\cdot)$ and $G(\cdot, \cdot, \cdot)$ defined in Section 2.7.

Probability of success. We first analyze the probability of success of the data structure. We assume that a query $q$ has some $p \in P$ where $\|p-q\|_{2} \leq r$. The data structure succeeds when $\operatorname{QUERY}\left(q, v_{0}\right)$ returns some point $p^{\prime} \in P$ with $\left\|q-p^{\prime}\right\|_{2} \leq c r$.

Lemma 2.8.4. If

$$
T \geq \frac{3}{G\left(r, \eta_{s}, \eta_{q}\right)},
$$

then with probability at least $0.9, \operatorname{QUERY}\left(q, v_{0}\right)$ finds some point within distance cr

## from $q$.

Proof. We prove the lemma by induction on the depth of the tree. Let $q \in S^{d-1}$ be a query point and $p \in P$ its near neighbor. Suppose we are within the recursive call $\operatorname{Query}(q, v)$ for some node $v$ in the tree. Suppose we have not yet failed, that is, $p \in P_{v}$. We would like to prove that-if the condition of the lemma is met-the probability that this call returns some point within distance $c r$ is at least 0.9.

When $v$ is a node in the last level $K$, the algorithm enumerates $P_{v}$ and, since we assume $p \in P_{v}$, some good point will be discovered (though not necessarily $p$ itself). Therefore, this case is trivial. Now suppose that $v$ is not from the $K$-th level. Using the inductive assumption, suppose that the statement of the lemma is true for all $T$ potential children of $v$, i.e., if $p \in P_{v^{\prime}}$, then with probability 0.9 , $\operatorname{QUERY}\left(q, v^{\prime}\right)$ returns some point within distance $c r$ from $q$. Then,

$$
\begin{aligned}
\operatorname{Pr}[\text { failure }] & \leq \prod_{i=1}^{T}\left(1-\underset{g \sim N(0,1)^{d}}{\operatorname{Pr}}\left[\langle g, p\rangle \geq \eta_{s} \text { and }\langle g, q\rangle \geq \eta_{q}\right] \cdot 0.9\right) \\
& \leq\left(1-G\left(r, \eta_{s}, \eta_{q}\right) \cdot 0.9\right)^{T} \leq e^{-G\left(r, \eta_{s}, \eta_{q}\right) \cdot 0.9 \cdot T} \leq 0.1,
\end{aligned}
$$

where the first step follows from the inductive assumption and independence between the children of $v$ during the preprocessing phase. The second step follows by monotonicity of $G(s, \rho, \sigma)$ in $s$, and the fourth is due to the assumption of the lemma.

Space. We now analyze the space consumption of the data structure.
Lemma 2.8.5. The expected space consumption of the data structure is at most

$$
n^{1+o(1)} \cdot K \cdot\left(T F\left(\eta_{s}\right)\right)^{K}
$$

Proof. We compute the expected total size of the sets $P_{\ell}$ for leaves $\ell$ at $K$-th level. There are at most $T^{K}$ such nodes, and for a fixed point $p \in P$ and a fixed leaf $\ell$ the probability that $p \in P_{\ell}$ is equal to $F\left(\eta_{s}\right)^{K}$. Thus, the expected total size is at most $n \cdot\left(T F\left(\eta_{s}\right)\right)^{K}$. Since we only store a node $v$ if $P_{v}$ is non-empty, the number of nodes
stored is at most $K+1$ times the number of points stored at the leaves. The Gaussian vectors stored at each node require space $O(d)$, which is at most $n^{o(1)}$.

Query time. Finally, we analyze the query time.

Lemma 2.8.6. If $T F\left(\eta_{q}\right) \geq 3$, then the expected query time is at most

$$
\begin{equation*}
n^{o(1)} \cdot T \cdot\left(T F\left(\eta_{q}\right)\right)^{K}+n^{1+o(1)} \cdot\left(T G\left(c r, \eta_{s}, \eta_{q}\right)\right)^{K} . \tag{2.9}
\end{equation*}
$$

Proof. First, we compute the expected query time spent going down the tree, without scanning the leaves. The expected number of nodes the query procedure reaches is:

$$
1+T F\left(\eta_{q}\right)+\left(T F\left(\eta_{q}\right)\right)^{2}+\ldots+\left(T F\left(\eta_{q}\right)\right)^{K}=O(1) \cdot\left(T F\left(\eta_{q}\right)\right)^{K}
$$

since we assume that $T F\left(\eta_{q}\right) \geq 3$. In each of node, we spend time $n^{o(1)} \cdot T$. The product of the two expressions gives the first term in (2.9).

The expected time spent scanning points in the leaves is at most $n^{o(1)}$ times the number of points scanned at the leaves reached. The number of points scanned is always at most one more than the number of far points, i.e., lying a distance greater than $c r$ from $q$, that reached the same leaf. There are at most $n-1$ far points and $T^{K}$ leaves. For each far point $p^{\prime}$ and each leaf $\ell$ the probability that both $p^{\prime}$ and $q$ end up in $P_{\ell}$ is at most $G\left(c r, \eta_{s}, \eta_{q}\right)^{K}$. For each such pair, we spend time at most $n^{o(1)}$ processing the corresponding $p^{\prime}$. This gives the second term in (2.9).

### 2.8.4 Setting parameters: idealized version

We end the section by describing how to set parameters $T, K, \eta_{s}$ and $\eta_{q}$ to prove Theorem 2.8.1. We first show how we would do it under the following simplifying assumptions:

- Lemma 2.7.1 holds in the form $F(\rho)=e^{-\rho^{2} / 2}$;
- Lemma 2.7.2 holds as $G(s, \rho, \sigma)=e^{-\Delta^{2} / 2}$;
- There is no additional factor of $T$ in the first term of (2.9): that is, the query time is $n^{o(1)} \cdot\left(T F\left(\eta_{q}\right)\right)^{K}+n^{1+o(1)} \cdot\left(T G\left(c r, \eta_{s}, \eta_{q}\right)\right)^{K}$.

After we will be done under these assumptions, we will argue that in reality we can achieve similar guarantees by essentially setting parameters to their "ideal" values. For the idealized setting, the value of $K$ will not be too important, we will only need to have $K=o(\log n)$. For $\tau_{q}, \tau_{s}>0$ (which we fix later), we choose $\eta_{q}, \eta_{s}$ such that: $F\left(\eta_{q}\right)^{K}=n^{-\tau_{q}}$ and $F\left(\eta_{s}\right)^{K}=n^{-\tau_{s}}$. By the idealized version of Lemma 2.7.1 $\left(F(\rho)=e^{-\rho^{2} / 2}\right)$, this means the following:

$$
\begin{align*}
& \eta_{s}=\sqrt{\frac{2 \tau_{s} \ln n}{K}}  \tag{2.10}\\
& \eta_{q}=\sqrt{\frac{2 \tau_{q} \ln n}{K}} \tag{2.11}
\end{align*}
$$

We will set $\tau_{s}, \tau_{q}>0$ such that:

$$
\begin{equation*}
\alpha(c r) \leq \alpha(r) \leq \frac{\min \left\{\eta_{q}, \eta_{s}\right\}}{\max \left\{\eta_{q}, \eta_{s}\right\}}=\frac{\min \left\{\sqrt{\tau_{s}}, \sqrt{\tau_{q}}\right\}}{\max \left\{\sqrt{\tau_{s}}, \sqrt{\tau_{q}}\right\}} \tag{2.12}
\end{equation*}
$$

(the last step follows from (2.10) and (2.11)), so the idealized version of Lemma 2.7.2 $\left(G(s, \rho, \sigma)=e^{-\frac{\Delta^{2}}{2}}\right)$ combined with (2.10) and (2.11) gives:

$$
\begin{gather*}
G\left(r, \eta_{s}, \eta_{q}\right)^{K}=e^{-\frac{\eta_{s}^{2}+\eta_{q}^{2}-2 \alpha(r) \cdot \eta_{s} \eta_{q}}{2 \beta^{2}(r)} \cdot K}=n^{-\frac{\tau_{s}+\tau_{q}-2 \alpha(r) \cdot \sqrt{\tau_{s} \tau_{q}}}{\beta^{2}(r)}}  \tag{2.13}\\
G\left(c r, \eta_{s}, \eta_{q}\right)^{K}=e^{-\frac{\eta_{s}^{2}+\eta_{q}^{2}-2 \alpha(c r) \cdot \eta_{s} \eta_{q}}{2 \beta^{2}(c r)} \cdot K}=n^{-\frac{\tau_{s}+\tau_{q}-2 \alpha(c r) \cdot \sqrt{\tau_{s} \tau_{q}}}{\beta^{2}(c r)}} \tag{2.14}
\end{gather*}
$$

Now let us check that we set the parameters such that the terms in (2.9) are approximately the same. Namely, we aim at satisfying the following:

$$
\begin{equation*}
F\left(\eta_{q}\right)^{K}=n \cdot G\left(c r, \eta_{s}, \eta_{q}\right)^{K} \tag{2.15}
\end{equation*}
$$

which by $(2.14)$ and the relation between $\eta_{q}$ and $\tau_{q}$ means that:

$$
n^{-\tau_{q}}=n^{1-\frac{\tau_{s}+\tau_{q}-2 \alpha(c r) \cdot \sqrt{\tau s} \tau_{q}}{\beta^{2}(c r)}}
$$

or, looking at the exponents:

$$
-\tau_{q}=1-\frac{\tau_{s}+\tau_{q}-2 \alpha(c r) \cdot \sqrt{\tau_{s} \tau_{q}}}{\beta^{2}(c r)} .
$$

Remembering that $1-\beta^{2}(c r)=\alpha^{2}(c r)$, we get that, in order to balance the terms in the idealized version of (2.9) we need to have:

$$
\begin{equation*}
\left|\sqrt{\tau_{s}}-\alpha(c r) \sqrt{\tau_{q}}\right|=\beta(c r) \tag{2.16}
\end{equation*}
$$

We set $T=\frac{3}{G\left(r, \eta_{s}, \eta_{q}\right)}$. By Lemma 2.8.4 it means that the probability of success of the data structure is at least 0.9 . By $(2.13)$ and $K=o(\log n)$, we get:

$$
\begin{equation*}
T^{K}=n^{\frac{\tau_{s}+\tau_{q}-2 \alpha(r) \cdot \sqrt{\tau_{s} \tau_{q}}}{\beta^{2}(r)}+o(1)} . \tag{2.17}
\end{equation*}
$$

Now let us look at space and query time. By Lemma 2.8.5, the space can be upper bounded as follows:

$$
\begin{aligned}
n^{1+o(1)} \cdot\left(T F\left(\eta_{s}\right)\right)^{K} & =n^{1+\frac{\tau_{s}+\tau_{q}-2 \alpha(r) \cdot \sqrt{\tau_{s} \tau_{q}}}{\beta^{2}(r)}-\tau_{s}+o(1)} \\
& =n^{1+\frac{\left(\sqrt{\tau_{q}}-\alpha(r) \cdot \sqrt{\tau s}\right)^{2}}{\beta^{2}(r)}+o(1)}
\end{aligned}
$$

where the first step follows from (2.13), (2.17) and the relation between $\eta_{s}$ and $\tau_{s}$ and the second step follows from $1-\beta^{2}(r)=\alpha^{2}(r)$. Thus, we get the following identity for the space exponent:

$$
\begin{equation*}
\sqrt{\rho_{s}}=\frac{\left|\sqrt{\tau_{q}}-\alpha(r) \sqrt{\tau_{s}}\right|}{\beta(r)} . \tag{2.18}
\end{equation*}
$$

Finally, let us look at the query time. By the assumed idealized version of the Lemma 2.8.6 (no factor $T$ in the first term of (2.9)) and (2.15), we get the following upper bound on the query time:

$$
\begin{aligned}
n^{o(1)} \cdot\left(T F\left(\eta_{q}\right)\right)^{K} & =n^{\frac{\tau_{s}+\tau_{q}-2 \alpha(r) \cdot \sqrt{\tau_{s} \tau_{q}}}{\beta^{2}(r)}-\tau_{q}+o(1)} \\
& =n^{1+\frac{\left(\sqrt{\tau_{s}}-\alpha(r) \cdot \sqrt{\tau}\right)^{2}}{\beta^{2}(r)}+o(1)},
\end{aligned}
$$

where the first step follows from (2.13), (2.17) and the relation between $\eta_{q}$ and $\tau_{q}$ and the second step follows from $1-\beta^{2}(r)=\alpha^{2}(r)$. Thus, we get the following identity for the query time exponent:

$$
\begin{equation*}
\sqrt{\rho_{q}}=\frac{\left|\sqrt{\tau_{s}}-\alpha(r) \sqrt{\tau_{q}}\right|}{\beta(r)} . \tag{2.19}
\end{equation*}
$$

Now by varying positive $\tau_{s}, \tau_{q}>0$ such that (2.16) and (2.12) hold, we can vary the exponents $\rho_{s}, \rho_{q}$ according to (2.18) and (2.19), respectively. The equation (2.12) restricts us to the range:

$$
\begin{equation*}
\frac{\sqrt{\tau_{s}}}{\sqrt{\tau_{q}}} \in\left[\alpha(r) ; \frac{1}{\alpha(r)}\right] . \tag{2.20}
\end{equation*}
$$

Note that $\alpha(r)>0$, since $r \leq \sqrt{2}-\Omega(1)$. But for this range, (2.16), (2.18), and (2.19) turn into:

$$
\begin{aligned}
& \sqrt{\tau_{s}}=\alpha(c r) \cdot \sqrt{\tau_{q}}+\beta(c r), \\
& \sqrt{\rho_{s}}=\frac{\sqrt{\tau_{q}}-\alpha(r) \sqrt{\tau_{s}}}{\beta(r)}, \\
& \sqrt{\rho_{q}}=\frac{\sqrt{\tau_{s}}-\alpha(r) \sqrt{\tau_{q}}}{\beta(r)},
\end{aligned}
$$

respectively. Thus, if we move $\sqrt{\rho_{q}}$ between 0 and $\frac{\beta(r) \beta(c r)}{1-\alpha(r) \alpha(c r)}>0$, we get:

$$
\begin{align*}
& \sqrt{\tau_{s}}=\frac{\alpha(r) \beta(c r)-\alpha(c r) \beta(r) \sqrt{\rho_{q}}}{\alpha(r)-\alpha(c r)},  \tag{2.21}\\
& \sqrt{\tau_{q}}=\frac{\beta(c r)-\beta(r) \sqrt{\rho_{q}}}{\alpha(r)-\alpha(c r)},  \tag{2.22}\\
& \sqrt{\rho_{s}}=\frac{\beta(r) \beta(c r)-(1-\alpha(r) \alpha(c r)) \sqrt{\rho_{q}}}{\alpha(r)-\alpha(c r)} . \tag{2.23}
\end{align*}
$$

We need to show that for this range of $\sqrt{\rho_{q}}$ the constraint (2.20) holds and, moreover, the right-hand sides of (2.21) and (2.22) are positive. Let us start with the former. We have:

$$
\begin{equation*}
\frac{\sqrt{\tau_{s}}}{\sqrt{\tau_{q}}}=\frac{\alpha(r) \beta(c r)-\alpha(c r) \beta(r) \sqrt{\rho_{q}}}{\beta(c r)-\beta(r) \sqrt{\rho_{q}}} . \tag{2.24}
\end{equation*}
$$

If we treat $\sqrt{\rho_{q}}$ as a formal variable, then the derivative of (2.24) with respect to it
is positive. At the boundary points, the ratio is equal to $\alpha(r)$ and $\frac{1}{\alpha(r)}$. Thus, (2.20) holds. Now let us check that the right-hand sides of (2.21) and (2.22) are indeed positive. They are linear functions of $\sqrt{\rho_{q}}$, so it is enough to check non-negativity for $\sqrt{\rho_{q}} \in\left\{0, \frac{\beta(r) \beta(c r)}{1-\alpha(r) \alpha(c r)}\right\}$. For $\sqrt{\rho_{q}}=0$, we have:

$$
\begin{aligned}
\sqrt{\tau_{s}} & =\frac{\alpha(r) \beta(c r)}{\alpha(r)-\alpha(c r)} \\
\sqrt{\tau_{q}} & =\frac{\beta(c r)}{\alpha(r)-\alpha(c r)}
\end{aligned}
$$

For $\sqrt{\rho_{q}}=\frac{\beta(r) \beta(c r)}{1-\alpha(r) \alpha(c r)}$, we have:

$$
\begin{aligned}
\sqrt{\tau_{s}} & =\frac{\beta(c r)}{1-\alpha(r) \alpha(c r)} \\
\sqrt{\tau_{q}} & =\frac{\alpha(r) \beta(c r)}{1-\alpha(r) \alpha(c r)}
\end{aligned}
$$

All four values are positive, since $\alpha(r)>0$ due to $r \leq \sqrt{2}-\Omega(1)$, and $\beta(c r)>0$, since $0<c r<2$. The smallest of these values is:

$$
\frac{\alpha(r) \beta(c r)}{1-\alpha(r) \alpha(c r)},
$$

since $\alpha(r)-\alpha(c r) \leq 1-\alpha(r) \alpha(c r)$ and $\alpha(r) \leq 1$. The largest is:

$$
\frac{\beta(c r)}{\alpha(r)-\alpha(c r)} .
$$

Finally, (2.23) gives the desired trade-off (2.6). When $\sqrt{\rho_{q}}$ changes between 0 and $\frac{\beta(r) \beta(c r)}{1-\alpha(r) \alpha(c r)}, \sqrt{\rho_{s}}$ changes from $\frac{\beta(r) \beta(c r)}{\alpha(r)-\alpha(c r)}$ to 0 (and remains non-negative).

### 2.8.5 Setting parameters: final version

Here we get rid of the (unrealistic) assumptions we made before. Namely, we need to make sure that the identities (2.13) and (2.14) (approximately) hold, and that $T=n^{o(1)}$ if we set all the parameters appropriately. We start with fixing $K=\sqrt{\ln n}$. Suppose that the query exponent we would like to achieve is $0 \leq \sqrt{\rho_{q}} \leq \frac{\beta(r) \beta(c r)}{1-\alpha(r) \alpha(c r)}$. Let us choose $\sqrt{\tau_{s}}$ and $\sqrt{\tau_{q}}$ according to the formulas (2.21) and (2.22), respectively. From Section 2.8.4, we know that:

$$
\begin{equation*}
\frac{\alpha(r) \beta(c r)}{1-\alpha(r) \alpha(c r)} \leq \sqrt{\tau_{s}}, \sqrt{\tau_{q}} \leq \frac{\beta(c r)}{\alpha(r)-\alpha(c r)} \tag{2.25}
\end{equation*}
$$

The lower bound in (2.25) is at least $\Omega_{c}\left(\varepsilon_{0}^{3 / 2}\right)$, since $r \leq \sqrt{2}-\varepsilon_{0}$ and cr $\leq 2-\varepsilon_{0}$, while the upper bound is at most $O_{c}(\log \log n)$, since $r \geq \frac{1}{\log \log n}$. Thus, we have:

$$
\begin{equation*}
\Omega_{c}\left(\varepsilon_{0}^{3}\right) \leq \tau_{s}, \tau_{q} \leq O_{c}\left(\log ^{2} \log n\right) \tag{2.26}
\end{equation*}
$$

Now we need to choose $\eta_{s}, \eta_{q}>0$. Let us choose them as in the idealized scenario:

$$
\begin{align*}
& \eta_{s}=\sqrt{\frac{2 \tau_{s} \ln n}{K}}=\sqrt{2 \tau_{s}} \cdot \ln ^{1 / 4} n  \tag{2.27}\\
& \eta_{q}=\sqrt{\frac{2 \tau_{q} \ln n}{K}}=\sqrt{2 \tau_{q}} \cdot \ln ^{1 / 4} n \tag{2.28}
\end{align*}
$$

Now, we have by Lemma 2.7.1:

$$
\begin{align*}
F\left(\eta_{s}\right)^{K} & =n^{-\tau_{s}} \cdot e^{ \pm O_{c}(K \cdot \log \log n)} \\
& \in n^{-\tau_{s}} \cdot e^{ \pm O_{c}(\sqrt{\log n} \cdot \log \log n)}  \tag{2.29}\\
F\left(\eta_{q}\right)^{K} & =n^{-\tau_{q}} \cdot e^{ \pm O_{c}(K \cdot \log \log n)} \\
& \in n^{-\tau_{q}} \cdot e^{ \pm O_{c}(\sqrt{\log n} \cdot \log \log n)} \tag{2.30}
\end{align*}
$$

since $\tau_{s}, \tau_{q}=O_{c}\left(\log ^{2} \log n\right)$ by (2.26). Finally, as in the idealized scenario, we have:

$$
\frac{\min \left\{\eta_{s}, \eta_{q}\right\}}{\max \left\{\eta_{s}, \eta_{q}\right\}}=\frac{\min \left\{\sqrt{\tau_{s}}, \sqrt{\tau_{q}}\right\}}{\max \left\{\sqrt{\tau_{s}}, \sqrt{\tau_{q}}\right\}} \geq \alpha(r) \geq \alpha(c r)
$$

thus we can use Lemma 2.7.2 the same way as in Section 2.8.4. Namely, if we denote:

$$
\begin{aligned}
\Delta_{r}^{2} & =\frac{\eta_{s}^{2}+\eta_{q}^{2}-2 \alpha(r) \cdot \eta_{s} \eta_{q}}{\beta^{2}(r)} \\
\Delta_{c r}^{2} & =\frac{\eta_{s}^{2}+\eta_{q}^{2}-2 \alpha(c r) \cdot \eta_{s} \eta_{q}}{\beta^{2}(c r)}
\end{aligned}
$$

we get:

$$
\begin{aligned}
G\left(r, \eta_{s}, \eta_{q}\right) & =e^{-\frac{\Delta_{r}^{2}}{2} \pm O\left(\log \Delta_{r}\right)} \\
G\left(c r, \eta_{s}, \eta_{q}\right) & =e^{-\frac{\Delta_{c r}^{2}}{2} \pm O\left(\log \Delta_{c r}\right)}
\end{aligned}
$$

Since $\Delta_{r}^{2}, \Delta_{c r}^{2} \leq O_{c}\left(\sqrt{\log n} \cdot \log ^{4} \log n\right)$, we have $\log \Delta_{r}, \log \Delta_{c r} \leq O_{c}(\log \log n)$. Thus,

$$
\begin{align*}
G\left(r, \eta_{s}, \eta_{q}\right)^{K} & =n^{-\frac{\tau_{s}+\tau_{q}-2 \alpha(r) \cdot \sqrt{\tau_{s} \tau q}}{\beta^{2}(r)}} \cdot e^{ \pm O_{c}(\sqrt{\log n \cdot} \cdot \log \log n)}  \tag{2.31}\\
G\left(c r, \eta_{s}, \eta_{q}\right)^{K} & =n^{-\frac{\tau_{s}+\tau_{q}-2 \alpha(c r) \cdot \sqrt{\tau_{s} \tau q}}{\beta^{2}(c r)}} \cdot e^{ \pm O_{c}(\sqrt{\log n} \cdot \log \log n)} \tag{2.32}
\end{align*}
$$

Thus, using (2.29), (2.30), (2.31) and (2.32), we can bound the query and the space exponents exactly the same way as in the idealized setting. The only thing remains to be checked is that $T=n^{o(1)}$. But this is immediate:

$$
T=\frac{3}{G\left(r, \eta_{s}, \eta_{q}\right)}=O(1) \cdot e^{\frac{\Delta_{r}^{2}}{2} \pm O\left(\log \Delta_{r}\right)} \leq e^{O_{c}\left(\sqrt{\log n \cdot} \cdot \log ^{4} \log n\right)} \leq n^{o(1)}
$$

Finally, to prepare for the next section, let us state the relation between parameters we can use as a black box later.

Lemma 2.8.7. If $c>1$ and $\frac{1}{\log \log n} \leq r \leq \sqrt{2}-\Omega(1)$, then for every $\rho_{s}, \rho_{q} \geq 0$ such that (2.6) holds there exist $\eta_{s}, \eta_{q}>0$ such that:

- $F\left(\eta_{s}\right) / G\left(r, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{s}+o(1)\right) / K} ;$
- $F\left(\eta_{q}\right) / G\left(r, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{q}+o(1)\right) / K}$;
- $G\left(c r, \eta_{s}, \eta_{q}\right) / G\left(r, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{q}-1+o(1)\right) / K} ;$
- $1 / G\left(r, \eta_{s}, \eta_{q}\right) \leq n^{o(1)}$,
where $K=\sqrt{\ln n}$.


### 2.9 Data-dependent partitions

In this section we prove Theorem 2.5.1, which we restate below:

Theorem 2.9.1. For every $c>1, \rho_{q} \geq 0$ and $\rho_{s} \geq 0$ such that

$$
\begin{equation*}
c^{2} \sqrt{\rho_{q}}+\left(c^{2}-1\right) \sqrt{\rho_{s}} \geq \sqrt{2 c^{2}-1} \tag{2.33}
\end{equation*}
$$

there exists a data structure for $(c, r)-A N N$ for the unit sphere $S^{d-1}$ with $r=\frac{1}{\sqrt{\log \log n}}$ and $d=\log ^{1+o(1)} n$ with space $n^{1+\rho_{s}+o(1)}$ and query time $n^{\rho_{q}+o(1)}$.

### 2.9.1 Overview

We start with a high-level overview. Consider a dataset $P_{0}$ of $n$ points. We partition $P_{0}$ into various components: $s$ dense components, denoted by $C_{1}, C_{2}$, $\ldots, C_{s}$, and one pseudo-random component, denoted by $\widetilde{P}$. The partition is designed to satisfy the following properties. Each dense component $C_{i}$ satisfies $\left|C_{i}\right| \geq$ $\tau \cdot\left|P_{0} \backslash\left(C_{1} \cup C_{2} \cup \ldots \cup C_{i-1}\right)\right|$ and can be covered by a spherical cap of radius $\sqrt{2}-\varepsilon$ (see Figure 2-10). Here $\tau, \varepsilon>0$ are small quantities to be chosen later. One should think of $C_{i}$ 's as clusters consisting of many points which are closer than random points would be. The pseudo-random component $\widetilde{P}$ consists of the remaining points without any dense clusters inside.

We process each $C_{i}$ and $\widetilde{P}$ separately. We enclose every dense component $C_{i}$ in a smaller ball $E_{i}$ of radius $1-\Omega\left(\varepsilon^{2}\right)$ (see Figure 2-10). For simplicity, we first ignore the fact that $C_{i}$ does not necessarily lie on the boundary $\partial E_{i}$. Once we enclose each dense cluster with a smaller ball, we recurse on each resulting spherical instance of radius $1-\Omega\left(\varepsilon^{2}\right)$. We treat the pseudo-random component $\widetilde{P}$ similarly to the


Figure 2-10: Covering a spherical cap of radius $(\sqrt{2}-\varepsilon) R$ on a ball of radius $R$.
random instance described in Section 2.6. Namely, we sample $T$ Gaussian vectors $z_{1}, z_{2}, \ldots, z_{T} \sim N(0,1)^{d}$, and form $T$ subsets of $\widetilde{P}$ :

$$
\widetilde{P}_{i}=\left\{p \in \widetilde{P} \mid\left\langle z_{i}, p\right\rangle \geq \eta_{s} R\right\}
$$

where $\eta_{s}>0$ is a parameter to be chosen later (for each pseudo-random remainder separately). Then, we recurse on each $\widetilde{P}_{i}$. Note that after we recurse, new dense clusters may appear in some $\widetilde{P}_{i}$ since it becomes easier to satisfy the minimum size constraint.

During the query procedure, we recursively query each $C_{i}{ }^{5}$ with the query point $q$. For the pseudo-random component $\widetilde{P}$, we identify all $i$ 's such that $\left\langle z_{i}, q\right\rangle \geq \eta_{q} R$, and query all the corresponding children recursively. Here $\eta_{q}>0$ is a new parameter that need to be chosen carefully (for each pseudo-random remainder separately).

Our algorithm makes progress in two ways. For dense clusters, we reduce the radius of the enclosing sphere by a factor of $\left(1-\Omega\left(\varepsilon^{2}\right)\right)$. Since initially $r=\frac{1}{\sqrt{\log \log n}}$, in $O\left(\log \log \log n / \varepsilon^{2}\right)$ iterations we would arrive to the regime, where Corollary 2.8.3 begins to apply. For the pseudo-random component $\widetilde{P}$, most points will lie within a distance at least $\sqrt{2}-\varepsilon$ from each other. In particular, a typical inter-point distance is approximately $\sqrt{2}$, exactly like for the random case. For this reason, we call $\widetilde{P}$ pseudo-random. In this setting, the data structure from Section 2.8 performs well, at least for one step. However, after we recurse, we will need to extract dense clusters

[^17]

Figure 2-11: The definition of Project
again.
We now address the issue deferred in the above high-level description: a dense component $C_{i}$ does not generally lie on $\partial E_{i}$, but rather can occupy the interior of $E_{i}$. In this case, we partition $E_{i}$ into thin annuli of carefully chosen width and treat each annulus as a sphere. This discretization adds to the complexity of the analysis, but does not seem to be too fundamental. However, now it becomes harder to argue about clusters making progress, since the value of $r$ may decrease. As a result, we need quite a delicate potential function argument to reason about it. In particular, we need to make sure that $r$ does not drop below $\frac{1}{\log \log n}$ in order for the techniques from Section 2.8 to be applicable.

### 2.9.2 Description

We are now ready to describe the data structure formally. It depends on the (small positive) parameters $\tau, \varepsilon, \varepsilon^{\prime}$ and $\delta$, an integer parameter $K \sim \sqrt{\ln n}$ and a large parameter $\gamma_{\text {stop }}>1$. We also need to choose parameters $T, \eta_{s}>0, \eta_{q}>0$ for each pseudo-random remainder separately. Figure 2-13 provides pseudocode for the algorithm.

Preprocessing. Our preprocessing algorithm consists of the following functions:

- ProcessSphere $\left(P, r_{1}, r_{2}, o, R, l\right)$ builds the data structure for a dataset $P$ lying on a sphere $\partial B(o, R)$, assuming we need to solve ANN with distance
thresholds $r_{1}$ and $r_{2}$. Moreover, we are guaranteed that queries will lie on the sphere $\partial B(o, R)$ as well. The parameter $l$ counts the number of non-cluster nodes in the recursion stack we have encountered so far. Similarly to the dataindependent algorithm from Section 2.8, we stop as soon as we encounter $K$ of them. We start the preprocessing with calling ProcessSphere $\left(P_{0}, r, c r, 0,1\right.$, $0)$.
- ProcessBall $\left(P, r_{1}, r_{2}, o, R, l\right)$ builds the data structure for a dataset $P$ lying inside the ball $B(o, R)$, assuming we need to solve ANN with distance thresholds $r_{1}$ and $r_{2}$. Unlike ProcessSphere, here queries can be arbitrary. The parameter $l$ has the same meaning as for ProcessSphere.
- $\operatorname{Project}\left(R_{1}, R_{2}, r\right)$ is an auxiliary function used when we discretize a ball into annuli. Suppose we have two spheres $S_{1}$ and $S_{2}$ with a common center and radii $R_{1}$ and $R_{2}$, respectively. Suppose there are points $p_{1} \in S_{1}$ and $p_{2} \in S_{2}$ with $\left\|p_{1}-p_{2}\right\|=r . \operatorname{Project}\left(R_{1}, R_{2}, r\right)$ returns the distance between $p_{1}$ and the point $\widetilde{p_{2}}$ that lies on $S_{1}$ and is the closest to $p_{2}$ (see Figure 2-11). This is implemented by a formula.

We now elaborate on the above descriptions of ProcessSphere and ProcessBALL.

ProcessSphere. We consider three base cases.

1. If $l=K$, we stop and store $P$ explicitly. This corresponds to having reached a leaf in the algorithm from Section 2.8.
2. If $r_{2} \geq 2 R$, then we may only store one point, since any point in $P$ is a valid answer to any query made on a sphere of radius $R$ containing $P$.
3. If $r_{2}^{2} / r_{1}^{2} \geq \gamma_{\text {stop }}$, then the instance can be handled by using an LSH-based data structure from [66]. We will set $\gamma_{\text {stop }}$ to be super-constant, thus the data structure occupies space $|P| \cdot n^{o(1)}$, and allows query time $n^{o(1)}$.

If none of the above cases apply, we proceed by removing the dense components and then handling the pseudo-random remainder. The dense components $C_{i}$ are clusters of at least $\tau\left|P \backslash\left(C_{1} \cup C_{2} \cup \ldots \cup C_{i-1}\right)\right|$ points lying in a ball of radius $(\sqrt{2}-\varepsilon) R$ with its center on $\partial B(o, R)$. These balls can be enclosed by smaller balls of radius $\widetilde{R} \leq\left(1-\Omega\left(\varepsilon^{2}\right)\right) R$ (with unrestricted centers). For each of these smaller balls, we invoke ProcessBall with the same $l$.

After we extract the dense components, we consider two cases. If $r_{2} \geq\left(\sqrt{2}+\varepsilon^{\prime}\right) \cdot R$ (where $\varepsilon^{\prime}=\Theta(\varepsilon)$ ), we are essentially done. Indeed, in this case for every query either an answer lies in one of the dense components, or $\Omega(1)$-fraction of the remainder can serve as a valid answer. We can check the latter (during the query stage) using uniform subsampling.

Finally, if $r_{2}<\left(\sqrt{2}+\varepsilon^{\prime}\right) \cdot R$, then we build a single level of the tree from Section 2.8 for the remaining pseudo-random remainder. We pick the appropriate $\eta_{s}, \eta_{q}>0$ and $T$ and recurse on each part with ProcessSphere with $l$ increased by 1.

ProcessBall. Similarly to ProcessSphere, if $r_{1}+2 R \leq r_{2}$, then any point from $B(o, R)$ is a valid answer to any query in $B\left(o, R+r_{2}\right)$.

If we are not in the trivial setting above, we reduce the ball to the spherical case via a discretization of the ball $B(o, R)$ into thin annuli of width $\delta r_{1}$. First, we round all distances from points and queries to $o$ to a multiple of $\delta r_{1}$. This rounding can change the distance between any pair of points by at most $2 \delta r_{1}$ by the triangle inequality.

Then, we handle each non-empty annulus separately. In particular, for a fixed annulus at distance $\delta i r_{1}$ from $o$, a possible query can lie at most a distance $\delta j r_{1}$ from $o$, where $\delta r_{1}|i-j| \leq r_{1}+2 \delta r_{1}$. For each such case, we recursively build a data structure with ProcessSphere. However, when projecting points, the distance thresholds of $r_{1}$ and $r_{2}$ change, and this change is computed using Project.

Overall, the preprocessing creates a rooted tree. The root corresponds to the first ProcessSphere call, and subsequent nodes correspond to procedures ProcessSphere and ProcessBall. We refer to the tree nodes correspondingly, using the
labels in the description of the query algorithm from below.

Query algorithm. Consider a query point $q \in \mathbb{R}^{d}$. We run the query on the tree, starting with the root node, and applying the following algorithms depending on the label of the nodes:

- For a ProcessSphere node, we first handle the three base cases in a straightforward way. If $l=K$, then we just enumerate the whole $P$ until we find a point within distance $r_{2}$. If $r_{2} \geq 2 R$, we simply return an arbitrary data point. If $r_{2}^{2} / r_{1}^{2} \geq \gamma_{\text {stop }}$, then we query the data structure from [66] we build for this case and stop.

If none of the base cases apply, we start with querying the data structures corresponding to the clusters recursively.

Then, if $r_{2} \geq\left(\sqrt{2}+\varepsilon^{\prime}\right) R$, we check $O(1)$ uniformly random points from $P$ against the query point $q$. If $r_{2}<\left(\sqrt{2}+\varepsilon^{\prime}\right) R$, we locate $q$ in the spherical caps (with threshold $\eta_{q}$, like in Section 2.8), and query data structures we built for the corresponding subsets of $P$. If one of these calls finds a near neighbor, we return.

- In ProcessBall, we first consider the base case, where we just return the stored point if it is close enough. In general, we check whether $\|q-o\|_{2} \leq R+r_{1}$. If not, we return with no neighbor, since each dataset point lies within a ball of radius $R$ from $o$, but the query point is at least $R+r_{1}$ away from $o$. If $\|q-o\|_{2} \leq R+r_{1}$, we round $q$ so the distance from $o$ to $q$ is a multiple of $\delta r_{1}$ and enumerate all possible distances from $o$ to the potential near neighbor we are looking for. For each possible distance, we query the corresponding ProcessSphere children after projecting $q$ on the sphere with a tentative near neighbor using Project.


### 2.9.3 Setting parameters

We complete the description of the data structure by setting the remaining parameters. Recall that the dimension is $d=\log ^{1+o(1)} n$ and $K \sim \sqrt{\ln n}$. We set $\varepsilon, \varepsilon^{\prime}, \delta, \tau, \gamma_{\text {stop }}$ as


Figure 2-12: The definition of $\varepsilon^{\prime}>0$.
follows:

- $\varepsilon=\frac{1}{(\log \log \log n)^{0.01}}$;
- $\varepsilon^{\prime}=\Theta(\varepsilon)$ according to Figure 2-12;
- $\delta=\frac{1}{\log n}$;
- $\tau=\frac{1}{\exp \left(\log ^{2 / 3} n\right)}$;
- $\gamma_{\text {stop }}=(\log \log \log n)^{0.01}$.

Now we specify how to set $\eta_{s}, \eta_{q}>0$ and $T$ for each pseudo-random remainder. The idea will be to try to replicate the parameter settings of Section 2.8.4 corresponding to the random instance. Since we remove all the dense components, we can use $(\sqrt{2}-\varepsilon) R$ as an "effective $r_{2}$." In particular, we let

$$
T=\frac{3}{G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right)}
$$

in order to achieve a high probability of success. Then we let $\eta_{s}$ and $\eta_{q}$ such that

- $F\left(\eta_{s}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{s}+o_{c}(1)\right) / K} ;$
- $F\left(\eta_{q}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{q}+o_{c}(1)\right) / K} ;$
- $G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{q}-1+o_{c}(1)\right) / K} ;$
- $1 / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{o_{c}(1)}$.
which correspond to the parameter settings achieving the trade-off of Section 2.8.3. It is not immediate why we can accomplish this, we will show this later in Lemma 2.9.7 (namely, we will need to show that $\frac{1}{\log \log n} \leq \frac{r_{1}}{R} \leq\left(1+o_{c}(1)\right) \cdot \frac{\sqrt{2}}{c}$, and then we simply invoke Lemma 2.8.7).

A crucial relation between the parameters is that $\tau$ should be much smaller than $G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right)$. This implies that the "effective $r_{2}$ " is equal to $(\sqrt{2}-\varepsilon) R$, at least for the sake of a single step of the random partition.

We collect some basic facts from the data structure which will be useful for the analysis. These facts follow trivially from the pseudocode in Figure 2-13.

- The children to ProcessSphere may contain at most $O\left(\frac{\log n}{\tau}\right)$ many calls to ProcessBall, corresponding to cluster nodes, and $T$ calls to ProcessSphere. Each ProcessBall call of ProcessSphere handles a disjoint subset of the dataset. Points can be replicated in the pseudo-random remainder, when a point lies in the intersection of two or more caps.
- A ProcessBall node has many children, all of which are ProcessSphere which do not increment $l$. Each of these children corresponds to a call for a specific annulus of width $\delta r_{1}$ as well as a possible distance for a query. For each annulus, there are at most $\frac{2}{\delta}+4=O\left(\frac{1}{\delta}\right)$ "notable distances": after rounding by $\delta r_{1}$, a valid query can be at most $r_{1}+2 \delta r_{1}$ away from a particular annulus in both directions, thus, each point gets duplicated at most $O\left(\frac{1}{\delta}\right)=O(\log n)$ many times.
- For each possible point $p \in P$, we may consider the subtree of nodes which process that particular point. We make the distinction between two kinds of calls to ProcessSphere: calls where $p$ lies in a dense cluster, and calls where $p$ lies in a pseudo-random remainder. If $p$ lies in a dense cluster, $l$ is not incremented; if $p$ lies in the pseudo-random remainder, $l$ is incremented. The point $p$ may be processed by various rounds of calls to ProcessBall and ProcessSphere without incrementing $l$; however, there will be a moment when $p$ is not in a dense cluster and will be part of the pseudo-random remainder. In that setting,
$p$ will be processed by a call to ProcessSphere which increments $l$.


### 2.9.4 Analysis

In this section we will analyze the above data structure. We start with a somewhat technical part.

Lemma 2.9.2. Suppose we are within a call to ProcessSphere and none of the three base cases took place. Then:

- $\frac{r_{2}}{r_{1}} \geq\left(1-o_{c}(1)\right) \cdot c$;
- $\frac{r_{1}}{R} \geq \frac{1}{\log \log n} ;$
- The number of calls to ProcessBall in the recursion stack is at most

$$
O_{c}\left((\log \log \log n)^{2}\right)
$$

The values of $r_{1}, r_{2}$ and $R$ can only change when we call ProcessBall (from which ProcessSphere is called): namely, $R$ changes when we call ProcessBall, and then all three parameters change during a subsequent ProcessSphere call.

Let us outline the proof of Lemma 2.9.2. Suppose that we are within a call to ProcessSphere such that none of the three base cases took place. Consider the sequence of the triples $\left(r_{1}^{(i)}, r_{2}^{(i)}, R^{(i)}\right)$ as they evolved from the root of the tree to the current node (the total number of the triples is equal to the number of the ProcessBall calls in the recursion stack). For this sequence, we will consider a new "idealized" sequence $\left(\widetilde{r_{1}}{ }^{(i)}, \widetilde{r_{2}}{ }^{(i)}, R^{(i)}\right)$ defined as follows. We rerun the same branch of the recursion, but in ProcessBall we do not add $\delta r_{1}$ to $r_{1}$ and do not subtract $\delta r_{1}$ from $r_{2}$ when computing new thresholds (but we use the same values of $i$ and $j$ ). Note that at some point the idealized sequence may stop being well-defined, but the actual sequence still is: for instance, when in the idealized sequence $r_{1}<\delta|j-i| r_{1} \leq(1+2 \delta) r_{1}$. Now the proof of Lemma 2.9.2 consists of two steps. First, we show it for the idealized sequence by keeping track of the following quantities: $\gamma=\frac{r_{2}^{2}}{r_{1}^{2}}$ and $\xi=\frac{r_{2}^{2}}{R^{2}}$. Then, we

```
function ProcessSphere \(\left(P, r_{1}, r_{2}, o, R, l\right)\)
    if \(l=K\) then
        store \(P\) explicitly
        return
    if \(r_{2} \geq 2 R\) then
        store any point from \(P\)
        return
    if \(r_{2}^{2} / r_{1}^{2} \geq \gamma_{\text {stop }}\) then
        build a data structure for \(P\) from [66]
        return
    \(\widehat{R} \leftarrow(\sqrt{2}-\varepsilon) R\)
    while \(\exists x \in \partial B(o, R):|B(x, \widehat{R}) \cap P| \geq \tau|P|\) do
        \(B(\widetilde{o}, \widetilde{R}) \leftarrow\) the SEB for \(P \cap B(x, \widehat{R})\)
        \(\operatorname{ProcessBall}\left(P \cap B(x, \widehat{R}), r_{1}, r_{2}, \widetilde{o}, \widetilde{R}, l\right)\)
        \(P \leftarrow P \backslash B(x, \widehat{R})\)
    if \(r_{2} \geq\left(\sqrt{2}+\varepsilon^{\prime}\right) R\) then
        return \(\triangleright\) We will compare queries against a (small) uniform subsample of \(P\).
    choose \(\eta_{s}\) and \(\eta_{q}\) such that:
- \(F\left(\eta_{s}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{s}+o(1)\right) / K} ;\)
- \(F\left(\eta_{q}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{q}+o(1)\right) / K} ;\)
- \(G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{q}-1+o(1)\right) / K}\);
- \(1 / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{o(1)}\).
\(T \leftarrow 3 / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right)\)
for \(i \leftarrow 1 \ldots T\) do
sample \(z \sim N(0,1)^{d}\)
\(P^{\prime} \leftarrow\left\{p \in P \mid\langle z, p\rangle \geq \eta_{s} R\right\}\)
if \(P^{\prime} \neq \emptyset\) then
\(\operatorname{ProcessSphere}\left(P^{\prime}, r_{1}, r_{2}, o, R, l+1\right)\)
function \(\operatorname{ProcessBalL}\left(P, r_{1}, r_{2}, o, R, l\right)\)
if \(r_{1}+2 R \leq r_{2}\) then
store any point from \(P\)
return
\(P \leftarrow\left\{\left.o+\delta r_{1}\left\lceil\frac{\|p-o\|}{\delta r_{1}}\right\rceil \cdot \frac{p-o}{\|p-o\|} \right\rvert\, p \in P\right\}\)
```



```
\(\widetilde{P} \leftarrow\left\{p \in P:\|p-o\|=\delta i r_{1}\right\}\)
if \(\widetilde{P} \neq \emptyset\) then
for \(j \leftarrow 1 \ldots\left\lceil\frac{R+r_{1}+2 \delta r_{1}}{\delta r_{1}}\right\rceil\) do if \(\delta|i-j| \leq r_{1}+2 \delta r_{1}\) then
\(\widetilde{r_{1}} \leftarrow \operatorname{Project}\left(\delta i r_{1}, \delta j r_{1}, r_{1}+2 \delta r_{1}\right)\)
\(\widetilde{r_{2}} \leftarrow \operatorname{Project}\left(\delta i r_{1}, \delta j r_{1}, r_{2}-2 \delta r_{1}\right)\)
\(\operatorname{ProcessSphere}\left(\widetilde{P}, \widetilde{r_{1}}, \widetilde{r_{2}}, o, \delta i r_{1}, l\right)\)
function \(\operatorname{Project}\left(R_{1}, R_{2}, r\right)\)
return \(\sqrt{R_{1}\left(r^{2}-\left(R_{1}-R_{2}\right)^{2}\right) / R_{2}}\)
```

Figure 2-13: Pseudocode of the data structure (SEB stands for smallest enclosing ball)
show that the actual sequence can not diverge too much from the idealized one (or stop being well-defined) unless one of the base cases in ProcessSphere is triggered for the actual sequence.

Let us start with the idealized sequence. First of all, the initial values of $\gamma$ and $\xi$ are $c^{2}$ and $\frac{c^{2}}{\log \log n}$, respectively. Second, the largest $\gamma$ and $\xi$ can be are $\gamma_{\text {stop }}$ and 4, respectively, since otherwise base cases in PROCESSSPHERE are triggered. Now supposed we call ProcessBall followed by ProcessSphere for some values $i$ and $j$. Suppose that $\gamma$ and $\xi$ evolved into $\widetilde{\gamma}$ and $\widetilde{\xi}$, respectively. We have the following claims which immediately follow from the formula for PRoJECT.

## Claim 2.9.3.

$$
\widetilde{\gamma}=\frac{r_{2}^{2}-\delta^{2}(i-j)^{2} r_{1}^{2}}{r_{1}^{2}-\delta^{2}(i-j)^{2} r_{1}^{2}} .
$$

## Claim 2.9.4.

$$
\widetilde{\xi}=\frac{r_{2}^{2}-\delta^{2}(i-j)^{2} r_{1}^{2}}{\delta^{2} i j r_{1}^{2}}
$$

We observe that $\gamma$ never decreases. On the other hand, $\xi$ can decrease but not drastically.

## Claim 2.9.5.

$$
\frac{\widetilde{\xi}}{\xi} \geq \Omega_{c}(1)
$$

Proof. We have: $\delta i r_{1} \leq\left(1-\Omega\left(\varepsilon^{2}\right)\right) \cdot R+\delta r_{1} \leq O(R)$, since $r_{1} \leq O(R)$ and $\delta=O(1)$. Besides that, $\delta j r_{1} \leq \delta i r_{1}+r_{1} \leq O(R)$. Thus, by Claim 2.9.4, since $\gamma$ is non-decreasing and starts with $c^{2}$, and since $\delta^{2}(i-j)^{2} \leq 1$,

$$
\frac{\tilde{\xi}}{\xi}=\frac{\left(r_{2}^{2}-\delta^{2}(i-j)^{2} r_{1}^{2}\right) / \delta^{2} i j r_{1}^{2}}{r_{2}^{2} / R^{2}} \geq \frac{\left(r_{2}^{2}-r_{1}^{2}\right) / O\left(R^{2}\right)}{r_{2}^{2} / R^{2}} \geq \Omega\left(c^{2}-1\right) \geq \Omega_{c}(1)
$$

The following claim is the key to the overall analysis for the idealized case.

Claim 2.9.6. If $\frac{\tilde{\xi}}{\xi} \leq 1+\Theta\left(\varepsilon^{2}\right)$, then $\frac{\tilde{\gamma}}{\gamma} \geq 1+\Omega_{c}\left(\varepsilon^{4}\right)$.

Proof. By Claim 2.9.4 one has:

$$
\frac{\widetilde{\xi}}{\xi}=\frac{\left(r_{2}^{2}-\delta^{2}(i-j)^{2} r_{1}^{2}\right) R^{2}}{\delta^{2} i j r_{1}^{2} r_{2}^{2}}
$$

We have $\delta i r_{1} \leq\left(1-\Omega\left(\varepsilon^{2}\right)+\delta\right) R$ and

$$
\delta j r_{1} \leq \delta i r_{1}+\delta|i-j| r_{1} \leq\left(1-\Omega\left(\varepsilon^{2}\right)+\delta\right) R+\delta|i-j| r_{1}
$$

Thus,

$$
\frac{\widetilde{\xi}}{\xi} \geq \frac{1-\frac{\delta^{2}(i-j)^{2} r_{1}^{2}}{r_{2}^{2}}}{\left(1-\Omega\left(\varepsilon^{2}\right)+\delta\right)\left(1-\Omega\left(\varepsilon^{2}\right)+\delta+\frac{\delta|i-j| r_{1}}{R}\right)}
$$

Since $\delta \ll \varepsilon^{2}, r_{2} \geq r_{1}$, and $R \geq \Omega\left(r_{1}\right)$, we have that $\frac{\tilde{\xi}}{\xi} \leq 1+\Theta\left(\varepsilon^{2}\right)$ implies

$$
\delta|i-j|=\Omega_{c}\left(\varepsilon^{2}\right) .
$$

But this implies by Claim 2.9.3:

$$
\frac{\widetilde{\gamma}}{\gamma}=\frac{1-\delta^{2}(i-j)^{2} \cdot \frac{r_{1}^{2}}{r_{2}^{2}}}{1-\delta^{2}(i-j)^{2}} \geq 1+\Omega_{c}\left(\varepsilon^{4}\right)
$$

since $\frac{r_{1}}{r_{2}} \leq \frac{1}{c}<1$.

Thus, the number of times $\xi$ increases by at most $1+\Theta\left(\varepsilon^{2}\right)$ is at most the number of times $\gamma$ increases by at least $1+\Omega_{c}\left(\varepsilon^{4}\right)$. Since $\gamma$ is non-decreasing, starts at $c^{2}$ and grows until $\gamma_{\text {stop }}$, the number of such steps is at most $O_{c}\left(\frac{\log \gamma_{\text {stop }}}{\varepsilon^{4}}\right)$. Since $\xi$ decreases by at most $O_{c}(1)$ each step, in the worst case it can go down to:

$$
\frac{c^{2}}{\log \log n} \cdot 2^{-O_{c}\left(\frac{\log \gamma_{\text {stop }}}{\varepsilon^{4}}\right)} \geq \frac{1}{\log \log n} \cdot 2^{-O_{c}\left((\log \log \log n)^{0.05}\right)} \gg \frac{\gamma_{\text {stop }}}{\log ^{2} \log n}
$$

Thus,

$$
\frac{r_{1}^{2}}{R^{2}} \geq \frac{\xi}{\gamma_{\text {stop }}} \geq \frac{1}{\log ^{2} \log n}
$$

Since during the remaining steps, $\xi$ increases by a factor at least $1+\Omega\left(\varepsilon^{2}\right)$, and can
be at most 4 , the number of the remaining iterations is at most:

$$
O_{c}\left(\frac{1}{\varepsilon^{2}} \cdot\left(\frac{1}{\varepsilon^{4}}+\log \gamma_{\text {stop }}+\log \log \log n\right)\right) \ll O_{c}\left((\log \log \log n)^{2}\right) .
$$

Thus, we proved Lemma 2.9.2 for the idealized sequence. Let us now argue about the actual sequence. In order to do this, we will consider both sequences side by side. Namely, let $\left(r_{1}, r_{2}, R\right)$ be an element of the idealized sequence, and let $\left(\widetilde{r_{1}}, \widetilde{r_{2}}, R\right)$ be the corresponding element of the real sequence. Let us assume that we have already shown that:

- $\widetilde{r_{1}} \in\left(1 \pm \tau_{1}\right) \cdot r_{1}$, where $\tau_{1}=o(1)$;
- $\widetilde{r_{2}} \in\left(1 \pm \tau_{2}\right) \cdot r_{2}$, where $\tau_{2}=o(1)$;
- In particular, $\frac{\widetilde{r_{2}}}{\widetilde{r_{1}}} \geq(1-o(1)) \cdot c$.

Let us see how these conditions evolve during one step. Let us show that the new gap between $\widetilde{r_{1}}$ and $r_{1}$ is $\left(1 \pm \tau_{2}\right)(1 \pm O(\sqrt{\delta}))$ unless a base case within PROCESSSPHERE is triggered. One of the two undesired things can happen:

- First, it can be the case that $\delta|i-j|>1$. In this case, the next step for the idealized sequence is not defined.
- One has

$$
\begin{equation*}
\frac{\operatorname{ProJECT}\left(\delta i r_{1}, \delta j r_{1}, r_{1}\right)^{2}}{\operatorname{PROJECT}\left(\delta i \widetilde{r_{1}}, \delta j \widetilde{r_{1}}, \widetilde{r_{1}}+2 \delta \widetilde{r_{1}}\right)^{2}} \notin(1 \pm \Theta(\sqrt{\delta})) \cdot \frac{r_{1}^{2}}{\widetilde{r_{1}^{2}}} \tag{2.34}
\end{equation*}
$$

In this case the idealized and the real $r_{1}$ 's diverge too much.
We will show that in both of these cases, for the real sequence, during the next iteration a base case will be triggered. First, we will show that in both cases, $\delta|i-j| \geq(1-O(\sqrt{\delta}))$. For the first case, this is immediate. Let us handle the case of (2.34) now:
$\frac{\operatorname{Project}\left(\delta i r_{1}, \delta j r_{1}, r_{1}\right)^{2}}{\operatorname{PROJECT}\left(\delta i \widetilde{r_{1}}, \delta j \widetilde{r_{1}}, \widetilde{r_{1}}+2 \delta \widetilde{r_{1}}\right)^{2}}=\frac{r_{1}^{2}-\delta^{2}(i-j)^{2} r_{1}^{2}}{{\widetilde{r_{1}}}^{2}(1+2 \delta)^{2}-\delta^{2}(i-j)^{2} \widetilde{r_{1}}} \notin(1 \pm \Theta(\sqrt{\delta})) \cdot \frac{r_{1}^{2}}{\widetilde{r_{1}}}$
iff

$$
\frac{1-\delta^{2}(i-j)^{2}}{(1+2 \delta)^{2}-\delta^{2}(i-j)^{2}} \notin(1 \pm \Theta(\sqrt{\delta}))
$$

iff

$$
\delta^{2}(i-j)^{2} \geq 1-O(\sqrt{\delta})
$$

iff

$$
\delta|i-j| \geq 1-O(\sqrt{\delta})
$$

But if $\delta|i-j| \geq 1-O(\sqrt{\delta})$, then at the next iteration the new value of $\gamma$ is:

$$
\begin{aligned}
\frac{\operatorname{PROJECT}\left(\delta i \widetilde{r_{1}}, \delta j \widetilde{r_{1}}, \widetilde{r_{2}}-2 \delta \widetilde{r_{1}}\right)^{2}}{\operatorname{PROJECT}\left(\delta i \widetilde{r_{1}}, \delta j \widetilde{r_{1}}, \widetilde{r_{1}}+2 \delta \widetilde{r_{1}}\right)^{2}} & =\frac{\left(\widetilde{r_{2}}-2 \delta \widetilde{r_{1}}\right)^{2}-\delta^{2}(i-j)^{2}{\widetilde{r_{1}}}^{2}}{\left(\widetilde{r_{1}}+2 \delta \widetilde{r_{1}}\right)^{2}-\delta^{2}(i-j)^{2}{\widetilde{r_{1}}}^{2}} \\
& =\Omega_{c}\left(\frac{1}{\sqrt{\delta}}\right)=\Omega_{c}(\sqrt{\log n}) \gg \gamma_{\text {stop }}
\end{aligned}
$$

where we use that $\frac{{\widetilde{r_{2}}}^{2}}{\widetilde{r}_{1}^{2}} \geq(1-o(1)) c$.
Overall, the gap between $r_{1}$ and $\widetilde{r_{1}}$ grows by at most a factor of $(1 \pm O(\sqrt{\delta}))$.
Now let us handle the relation between $r_{2}$ and $\widetilde{r_{2}}$.

$$
\frac{\operatorname{PROJECT}\left(\delta i r_{1}, \delta j r_{1}, r_{2}\right)^{2}}{\operatorname{PROJECT}\left(\delta i \widetilde{r_{1}}, \delta j \widetilde{r_{1}}, \widetilde{r_{2}}-2 \delta \widetilde{r_{1}}\right)^{2}}=\frac{r_{2}^{2}-\delta^{2}(i-j)^{2} r_{1}^{2}}{\left(\widetilde{r_{2}}-2 \delta \widetilde{r_{1}}\right)^{2}-\delta^{2}(i-j)^{2}{\widetilde{r_{1}}}^{2}} \in 1 \pm O_{c}\left(\tau_{2}+\delta\right)
$$

Thus, after all the $O_{c}(\log \log \log n)$ iterations we will have:

$$
\begin{equation*}
\frac{r_{1}}{\widetilde{r_{1}}}=(1+O(\sqrt{\delta}))^{O_{c}(\log \log \log n)}=1+o_{c}(1) \tag{2.35}
\end{equation*}
$$

and

$$
\begin{equation*}
\frac{r_{2}}{\widetilde{r_{2}}}=1+\delta \cdot 2^{O_{c}(\log \log \log n)}=1+o_{c}(1) \tag{2.36}
\end{equation*}
$$

Thus,

$$
\frac{\widetilde{r_{2}}}{\widetilde{r_{1}}} \geq\left(1-o_{c}(1)\right) \cdot c
$$

as required. Other desired statements follow from their counterparts for the idealized sequence, (2.35) and (2.36).

Lemma 2.9.7. During the algorithm we are always able to choose $\eta_{s}$ and $\eta_{q}$ such
that:

- $F\left(\eta_{s}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{s}+o_{c}(1)\right) / K} ;$
- $F\left(\eta_{q}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{q}+o_{c}(1)\right) / K} ;$
- $G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right) / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right) \leq n^{\left(\rho_{q}-1+o_{c}(1)\right) / K}$;
- $1 / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right)=n^{o_{c}(1)}$.

Proof. If we need to choose $\eta_{s}$ and $\eta_{q}$, then $r_{2} \leq\left(\sqrt{2}+\varepsilon^{\prime}\right) R$. Since by Lemma 2.9.2, $\frac{r_{2}}{r_{1}} \geq$ $\left(1-o_{c}(1)\right) c$, we have $r_{1} \leq \frac{\left(1+o_{c}(1)\right)\left(\sqrt{2}+\varepsilon^{\prime}\right) R}{c}$ (in particular, it is at most $\left.\left(\sqrt{2}-\Omega_{c}(1)\right) R\right)$. Besides, by Lemma 2.9.2, we have $r_{1} \geq \frac{R}{\log \log n}$. This allows us to choose the required $\eta_{s}, \eta_{q}$ by using Lemma 2.8.7, since $\varepsilon, \varepsilon^{\prime}=o(1)$.

Lemma 2.9.8. The probability of success of the data structure is at least 0.9.

Proof. There are three places the preprocessing and query procedures are not deterministic.

First, when we build a data structure from [66] (in case $r_{2}^{2} / r_{1}^{2} \geq \gamma_{\text {stop }}$ ). In this case the probability of success is at least 0.9 by construction.

Second, when in ProcessSphere we have $r_{2} \geq\left(\sqrt{2}+\varepsilon^{\prime}\right) R$, and, after querying the dense components, we check the query against a uniform subsample of the remainder. Let us show that if a near neighbor is in the reminder, we succeed with high probability. Indeed, if in this case we have more than $\tau$-fraction of the remainder further than $\left(\sqrt{2}+\varepsilon^{\prime}\right) R$ from the query, it means that there is a dense component (see Figure 212), which contradicts the definition of a remainder. Thus, a $(1-\tau)$-fraction of the remainder can serve as a valid answer to the query, which means that uniform sampling finds it with probability at least $1-\tau>0.9$.

Third, when we perform a single step of the algorithm from Section 2.8. In this case, we can repeat the proof of Lemma 2.8.4 and conclude that we succeed with probability at least 0.9 , since $T=3 / G\left(r_{1} / R, \eta_{s}, \eta_{q}\right)$.

Lemma 2.9.9. The total space the data structure occupies is at most $n^{1+\rho_{s}+o(1)}$ in expectation.

Proof. We will prove that the total number of explicitly stored points (when $l=K$ ) is at most $n^{1+\rho_{s}+o(1)}$. We will count the contribution from each point separately, and use linearity of expectation to sum up the contributions. In particular, for a point $p \in P_{0}$, we want to count the number of lists where $p$ appears in the data structure. Each root to leaf path of the tree has at most $K$ calls to ProcessSphere which increment $l$, and at most $O_{c}\left((\log \log \log n)^{2}\right)$ calls to ProcessBall, and thus $O_{c}\left((\log \log \log n)^{2}\right)$ calls to ProcessSphere which do not increment $l$. Thus, once we count the number of lists, we may multiply by $K+O_{c}\left((\log \log \log n)^{2}\right)=n^{o(1)}$ to count the size of the whole tree.

For each point, we will consider the subtree of the data structure where the point was processed. In particular, we may consider the tree corresponding to calls to ProcessSphere and ProcessBall which process $p$. As discussed briefly in Section 2.9.3, we distinguish between calls to ProcessSphere which contain $p$ in a dense cluster, and calls to ProcessSphere which contain $p$ in the pseudo-random remainder. We increment $l$ only when $p$ lies in the pseudo-random remainder.

Claim 2.9.10. It suffices to consider the data structure where each node is a function call to ProcessSphere which increments l, i.e., when p lies in the pseudo-random remainder, since the total amount of duplication of points corresponding to other nodes is $n^{o(1)}$.

We will account for the duplication of points in calls to ProcessBall and calls to ProcessSphere which do not increment $l$. Consider the first node $v$ in a path from the root which does not increment $l$, this corresponds to a call to ProcessSphere which had $p$ in some dense cluster. Consider the subtree consisting of descendants of $v$ where the leaves correspond to the first occurrence of ProcessSphere which increments $l$. We claim that every internal node of the tree corresponds to alternating calls to ProcessBall and ProcessSphere which do not increment $l$. Note that calls to ProcessSphere which do not increment $l$ never replicate $p$. Each call to ProcessBall replicates $p$ in $b:=O(1 / \delta)=O(\log n)$ many times. We may consider contracting the tree and at edge, multiplying by the number of times we encounter

## ProcessBall.

Note that $p$ lies in a dense cluster if and only if it does not lie in the pseudo-random remainder. Thus, our contracted tree looks like a tree of $K$ levels, each corresponding to a call to ProcessSphere which contained $p$ in the pseudo-random remainder.

The number of children of some nodes may be different; however, the number of times ProcessBall is called in each branch of computation is

$$
U:=O_{c}\left((\log \log \log n)^{2}\right)
$$

the total amount of duplication of points due to ProcessBALL is at most $b^{U}=n^{o(1)}$. Now, the subtree of nodes processing $p$ contains $K$ levels with each $T$ children, exactly like the data structure for Section 2.8.

Claim 2.9.11. A node $v$ corresponding to ProcessSphere $\left(P, r_{1}, r_{2}, o, R, l\right)$ has, in expectation, $p$ appearing in $n^{\left((K-l) \rho_{s}+o_{c}(1)\right) / K}$ many lists in the subtree of $v$.

The proof is an induction over the value of $l$ in a particular node. For our base case, consider some node $v$ corresponding to a function call of ProcessSphere which is a leaf, so $l=K$, in this case, each point is only stored at most once, so the claim holds.

Suppose for the inductive assumption the claim holds for some $l$, then for a particular node at level $l-1$, consider the point when $p$ was part of the pseudo-random remainder. In this case, $p$ is duplicated in

$$
T \cdot F\left(\eta_{s}\right)=\frac{3 \cdot F\left(\eta_{s}\right)}{G\left(r_{1} / R, \eta_{s}, \eta_{q}\right)} \leq n^{\left(\rho_{s}+o_{c}(1)\right) / K}
$$

many children, and in each child, the point appears $n^{\left((K-l) \rho_{s}+o(1)\right) / K}$ many times. Therefore, in a node $v, p$ appears in $n^{\left((K-l+1) \rho_{s}+o(1)\right) / K}$ many list in its subtree. Letting $l=0$ for the root gives the desired outcome.

Lemma 2.9.12. The expected query time is at most $n^{\rho_{q}+o(1)}$.
Proof. We need to bound the expected number of nodes we traverse as well as the number of points we enumerate for nodes with $l=K$.

We first bound the number of nodes we traverse. Let $A(u, l)$ be an upper bound on the expected number of visited nodes when we start in a ProcessSphere node such that there are $u$ ProcessBall nodes in the stack and $l$ non-cluster nodes. By Lemma 2.9.2,

$$
u \leq U:=O_{c}\left((\log \log \log n)^{2}\right)
$$

and from the description of the algorithm, we have $l \leq K$. We will prove $A(0,0) \leq$ $n^{\rho_{q}+o_{c}(1)}$, which corresponds to the expected number of nodes we touch starting from the root.

We claim

$$
\begin{equation*}
A(u, l) \leq \exp \left(\log ^{2 / 3+o(1)} n\right) \cdot A(u+1, l)+n^{\left(\rho_{q}+o_{c}(1)\right) / K} \cdot A(u, l+1) \tag{2.37}
\end{equation*}
$$

There are at most $O(\log n / \tau)=\exp \left(\log ^{2 / 3+o(1)} n\right)$ cluster nodes, and in each node, we recurse on $O(1 / \delta)=\exp \left(\log ^{o(1)} n\right)$ possible annuli with calls to ProcessSphere nodes where $u$ increased by 1 and $l$ remains the same. On the other hand, there are

$$
T \cdot F\left(\eta_{q}\right)=\frac{3 \cdot F\left(\eta_{q}\right)}{G\left(r_{1} / R, \eta_{s}, \eta_{q}\right)} \leq n^{\left(\rho_{q}+o_{c}(1)\right) / K}
$$

caps, where the query falls, in expectation. Each calls ProcessSphere where $u$ remains the same and $l$ increased by 1 .

Solving (2.37):

$$
A(0,0) \leq\binom{ U+K}{K} \exp \left(U \cdot \log ^{2 / 3+o(1)} n\right) \cdot n^{\rho_{q}+o_{c}(1)} \leq n^{\rho_{q}+o_{c}(1)}
$$

We now give an upper bound on the number of points the query algorithm will test at level $K$. Let $B(u, l)$ be an upper bound on the expected fraction of the dataset in the current node that the query algorithm will eventually test at level $K$ (where we count multiplicities). The variables $u$ and $l$ have the same meaning as discussed above.

We claim

$$
B(u, l) \leq O\left(\frac{\log n}{\tau}\right) \cdot B(u+1, l)+n^{\left(\rho_{q}-1+o_{c}(1)\right) / K} \cdot B(u, l+1)
$$

The first term comes from recursing down dense clusters. The second term, which corresponds to random spherical caps, is a bit more subtle. We have $T$ spherical caps. For each cap, where the query belongs to, we have some number of close points (closer than $(\sqrt{2}-\varepsilon) R$ to the query) and some number of far points. Thus, the expected number on the fraction of the tested points in this case is at most:

$$
\begin{aligned}
& T \cdot\left(F\left(\eta_{q}\right) \cdot \tau+G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right)\right) \cdot B(u, l+1) \\
& \leq \frac{3 \cdot\left(F\left(\eta_{q}\right) \cdot \tau+G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right)\right)}{G\left(r_{1} / R, \eta_{s}, \eta_{q}\right)} \cdot B(u, l+1) \\
& \leq \frac{4 \cdot G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right)}{G\left(r_{1} / R, \eta_{s}, \eta_{q}\right)} \cdot B(u, l+1) \\
& \leq n^{\left(\rho_{q}-1+o_{c}(1)\right) / K} \cdot B(u, l+1),
\end{aligned}
$$

where the second step follows from the fact that

$$
\left.G\left(\sqrt{2}-\varepsilon, \eta_{s}, \eta_{q}\right)=2^{-O_{c}(\sqrt{\log n} \cdot \log O(1)} \log n\right)
$$

and that $\tau=2^{-\Theta\left(\log ^{2 / 3} n\right)}$. Unraveling the recursion, we note that $u \leq U$ and $l \leq K \sim \sqrt{\ln n}$. Additionally, we have that $B(u, K) \leq 1$, since we do not store duplicates in the last level. Therefore,

$$
B(0,0) \leq\binom{ U+K}{U} O\left(\frac{\log n}{\tau}\right)^{U} \cdot\left(n^{\left(\rho_{q}-1+o(1)\right) / K}\right)^{K}=n^{\rho_{q}-1+o_{c}(1)}
$$

### 2.9.5 Fast preprocessing

A priori, it is not even clear how to implement the preprocessing in polynomial time, let alone near-linear in the space. We will first show how to get preprocessing time to $n^{2+\rho_{s}+o_{c}(1)}$ and then reduce it to $n^{1+\rho_{s}+o_{c}(1)}$.

## Near-quadratic time

To get preprocessing time $n^{2+\rho_{s}+o_{c}(1)}$ we need to observe that during the clustering step in ProcessSphere we may look only for balls with centers being points from $P$. We build upon the following lemma.

Proposition 2.9.13 (van der Corput Lemma). For any $v^{*}, v_{1}, v_{2}, \ldots, v_{n} \in S^{d-1}$ one has

$$
\sum_{i, j}\left\langle v_{i}, v_{j}\right\rangle \geq\left|\sum_{i}\left\langle v^{*}, v_{i}\right\rangle\right|^{2}
$$

Proof. We have

$$
\left|\sum_{i}\left\langle v^{*}, v_{i}\right\rangle\right|^{2}=\left|\left\langle v^{*}, \sum_{i} v_{i}\right\rangle\right|^{2} \leq\left\|v^{*}\right\|^{2} \cdot\left\|\sum_{i} v_{i}\right\|^{2}=\left\|\sum_{i} v_{i}\right\|^{2}=\sum_{i, j}\left\langle v_{i}, v_{j}\right\rangle,
$$

where the second step is an application of the Cauchy-Schwartz inequality.

The following claim is the main estimate we use to analyze the variant of ProcessSphere, where we are looking only for clusters with centers in data points. Informally, we prove that if a non-trivially small spherical cap covers $n$ points, then there is a non-trivially small cap centered in one of the points that covers a substantial fraction of points.

Claim 2.9.14. Fix $\varepsilon>0$. Suppose that $U \subset S^{d-1}$ with $|U|=n$. Suppose that there exists $u^{*} \in S^{d-1}$ such that $\left\|u^{*}-u\right\| \leq \sqrt{2}-\varepsilon$ for every $u \in U$. Then, there exists $u_{0} \in U$ such that

$$
\left|\left\{u \in U:\left\|u-u_{0}\right\| \leq \sqrt{2}-\Omega\left(\varepsilon^{2}\right)\right\}\right| \geq \Omega\left(\varepsilon^{2} n\right) .
$$

Proof. First, observe that $\left\|u^{*}-u\right\| \leq \sqrt{2}-\varepsilon$ iff $\left\langle u^{*}, u\right\rangle \geq \Omega(\varepsilon)$. By van der Corput Lemma (Proposition 2.9.13),

$$
\sum_{u, v \in U}\langle u, v\rangle \geq\left|\sum_{u \in U}\left\langle u^{*}, u\right\rangle\right|^{2} \geq \Omega\left(\varepsilon^{2} n^{2}\right)
$$

Thus, there exists $u_{0} \in U$ such that

$$
\sum_{u \in U}\left\langle u_{0}, u\right\rangle \geq \Omega\left(\varepsilon^{2} n\right) .
$$

This implies

$$
\left|\left\{u \in U \mid\left\langle u_{0}, u\right\rangle \geq \Omega\left(\varepsilon^{2}\right)\right\}\right| \geq \Omega\left(\varepsilon^{2} n\right)
$$

which is equivalent to

$$
\left|\left\{u \in U \mid\left\|u-u_{0}\right\| \leq \sqrt{2}-\Omega\left(\varepsilon^{2}\right)\right\}\right| \geq \Omega\left(\varepsilon^{2} n\right) .
$$

It means that if in ProcessSphere we search for clusters of radius $\sqrt{2}-\Omega\left(\varepsilon^{2}\right)$ centered in data points that cover at least $\Theta\left(\varepsilon^{2} \tau\right)$-fraction of the remaining data points, then after we remove all of them, we are sure that there are no clusters of radius $\sqrt{2}-\varepsilon$ with arbitrary centers that cover at least $\tau$-fraction of the remaining points. It is immediate to check that we can adjust the analysis of the data structure accordingly to accommodate this change.

It is easy to see that by reducing the time of each clustering step to near-quadratic, we reduce the total preprocessing time to $n^{2+\rho_{s}+o_{c}(1)}$. This follows from the proof of the space bound (intuitively, each point participates in $n^{o_{c}(1)}$ instances of the clustering subroutine).

## Near-linear time

To get $n^{1+\rho_{s}+o_{c}(1)}$ preprocessing time, we just subsample the dataset before finding each dense component. Indeed, since we care about the clusters with at least $\varepsilon^{2} \tau=n^{-o_{c}(1)}$ fraction of the remaining points, we can sample $n^{o_{c}(1)}$ points from the dataset and find dense components for the sample. Then, using the fact that the VC-dimension for balls in $\mathbb{R}^{d}$ is $O(d)$, we can argue that this sample is accurate enough with probability at least $1-n^{-10 \rho_{s}}$. Thus, with high probability, all the clustering steps are correct.

### 2.9.6 Handling insertions and deletions

The paper [142] shows how to convert a static data structure into a dynamic one if the following two conditions hold:

- The problem we are solving is decomposable. This means that if we partition a dataset into two parts and answer a query for both of the parts, we can compute the answer for the whole dataset efficiently. This condition clearly holds for the ANN problem.
- The preprocessing time of the static data structure must be small.

Then, [142] shows how to obtain a dynamic data structure whose query time is the static query time times $\log n$, and insertion/deletion time is the preprocessing time times $\frac{\log n}{n}$. This gives query time $n^{\rho_{q}+o_{c}(1)}$ and insertion/deletion time $n^{\rho_{s}+o(1)}$.

## Chapter 3

## FALCONN: practical and optimal LSH for unit sphere

### 3.1 Introduction

In Chapter 2 we developed a $(c, r)$-ANN data structure over the $\ell_{p}$ distance with strong guarantees (Theorem 2.3.1). A natural question is how practical the resulting data structure is. More specifically, we will focus on the case of the unit sphere $S^{d-1} \subset \mathbb{R}^{d}$, thus asking for a practical counterpart of Theorem 2.5.1. Indeed, in theory Theorem 2.5.1 implies Theorem 2.3.1, while in practice, the unit sphere case corresponds to similarity search with respect to the cosine similarity, which is widely used in applications such as comparing image feature vectors [91], speaker representations [158], and tf-idf data sets [166].

Unfortunately, the data structure we developed to prove Theorem 2.5.1 (for an overview, see Section 2.4) is far from being practical, the biggest problem being the "worst case to random" reduction from Section 2.91. But can we at least make the data structure based on data-independent partitions (Theorem 2.8.1 or, more specifically, Figure 2-9) practical? This is the problem we address in the present chapter.

As it turns out, as described in Section 2.8, even the simple data-independent approach is impractical. One immediate problem is the following: in order to get good

[^18]space and query exponents $\rho_{s}, \rho_{q}$, we must allow tree nodes to have many ${ }^{2}$ children, and during the query stage, we have to enumerate all of them.

Let us focus on the balanced regime (space $n^{1+\rho+o(1)}$ and query time $n^{\rho+o(1)}$ ). In this regime, the data-independent approach from Section 2.8 gives the bound:

$$
\sqrt{\rho}=\frac{\beta(r) \beta(c r)}{(1+\alpha(r))(1-\alpha(c r))},
$$

where functions $\alpha(\cdot)$ and $\beta(\cdot)$ are defined in the beginning of Section 2.7, or, expanding $\alpha(\cdot)$ and $\beta(\cdot)$,

$$
\begin{equation*}
\rho=\frac{4-c^{2} r^{2}}{4-r^{2}} \cdot \frac{1}{c^{2}} \tag{3.1}
\end{equation*}
$$

Thus, we can ask: can one achieve the exponent (3.1) with a practical data structure?

### 3.1.1 Our results

In this chapter we obtain a practical data structure that achieves the exponent (3.1) for the $(c, r)$-ANN on the unit sphere. In particular, when $r=o(1)$, we get $\rho=\frac{1}{c^{2}}+o_{c}(1)$ matching the best possible data-independent LSH for $\ell_{2}$ from [13]. Matching the guarantee from [13] by a practical algorithm has been an open problem since 2006 and here we essentially resolve it.

## New practical LSH family for the unit sphere

The new data structure fits the Locality-Sensitive Hashing (LSH) framework outlined in Section 1.2.2 and Section 2.2.1. We describe the relation between LSH and the ANN problem in more detail in Section 1.4.4, but for now we just recall the definition of LSH partitions. We say that a random partition $\mathcal{P}$ of the unit sphere $S^{d-1}$ is $\left(r, c r, p_{1}, p_{2}\right)$-sensitive, if for every $x, y \in S^{d-1}$ one has:

- if $\|x-y\|_{2} \leq r$, then $\operatorname{Pr}_{\mathcal{P}}[x$ and $y$ fall into the same part of $\mathcal{P}] \geq p_{1}$;
- if $\|x-y\|_{2}>c r$, then $\underset{\mathcal{P}}{\operatorname{Pr}}[x$ and $y$ fall into the same part of $\mathcal{P}] \leq p_{2}$.

[^19]It is known [85, 77] that an efficient ( $r, c r, p_{1}, p_{2}$ )-sensitive LSH family implies a data structure for $(c, r)$-ANN with space $n^{1+\rho+o(1)} / p_{1}$ and query time $n^{\rho+o(1)} / p_{1}$, where:

$$
\rho=\frac{\log \left(1 / p_{1}\right)}{\log \left(1 / p_{2}\right)} .
$$

In light of this reduction, it is enough to construct an efficient $\left(c, c r, p_{1}, p_{2}\right)$-sensitive LSH family for the unit sphere, for which

$$
\begin{equation*}
\frac{\log \left(1 / p_{1}\right)}{\log \left(1 / p_{2}\right)} \leq \frac{4-c^{2} r^{2}}{4-r^{2}} \cdot \frac{1}{c^{2}}+o_{c, r}(1) \tag{3.2}
\end{equation*}
$$

This is essentially what we do in this chapter. Let us compare the new LSH family with the existing families.

- Voronoi LSH from [18, 26] also achieves (3.2), but it is impractical for more or less the same reason as the data-independent approach from Section 2.8. Thus, our result can be seen as a practical counterpart of the Voronoi LSH.
- Hyperplane LSH from [55] gives the exponent worse than (3.2). See Figure 3-1 for the comparison for the case $c=2$. Nevertheless, Hyperplane LSH is very simple and is widely used (see, e.g., [118, 166]). The new LSH family described in this chapter significantly outperforms Hyperplane LSH in practice.
- Cross-polytope LSH introduced in [168] (and then appeared in [72]) is one of the main building blocks of the new family. To sample a cross-polytope LSH partition, we consider a randomly rotated set $\left\{ \pm e_{i}\right\}_{i=1}^{d}$, where $e_{i}$ are the standard basis vectors in $\mathbb{R}^{d}$; then we use their Voronoi diagram as a partitions. Prior to this work, this LSH construction has not been analyzed, but in this chapter we analyze it rigorously and show that it achieves (3.2). However, by itself, the Cross-polytope LSH is still impractical, and we need to improve it using additional algorithmic ideas.

Thus, the new LSH family can be seen as the Cross-polytope LSH "on steroids".


Figure 3-1: The dependence of the exponent $\rho$ on $r$ for the Hyperplane LSH and the LSH family developed in this chapter for $c=2$. The dots on the left correspond to $r=0$, while the ones on the right correspond to $r=\frac{\sqrt{2}}{2}$ (the regime of random instances).

## Fine-grained lower bound for LSH for unit sphere

Let us note that the bound on the exponent (3.2) is known to be optimal for LSH over the sphere in the case $r=\frac{\sqrt{2}}{c} \pm o(1)$, which corresponds to the random case [27]. But in this chapter we obtain a somewhat finer lower bound.

To highlight the difficulty of obtaining optimal and practical LSH schemes, we prove the first non-asymptotic lower bound on the trade-off between the collision probabilities $p_{1}$ and $p_{2}$ for $r=\frac{\sqrt{2}}{c} \pm o(1)$. So far, the optimal LSH upper bound $\rho=\frac{1}{2 c^{2}-1}$ (Voronoi LSH from [18, 26] and the Cross-polytope LSH from this chapter) attain this bound only in the limit, as $p_{1}, p_{2} \rightarrow 0$. Very small $p_{1}$ and $p_{2}$ are undesirable since the hash evaluation time is often proportional to $1 / p_{2}$. Our lower bound proves this is unavoidable: if we require $p_{2}$ to be large, $\rho$ has to be suboptimal.

This result has two important implications for designing practical hash functions. First, it shows that the trade-offs achieved by the cross-polytope LSH and the Voronoi LSH from [18, 26] are essentially optimal. Second, the lower bound guides design of future LSH functions: if one is to significantly improve upon the cross-polytope LSH,
one has to design a hash function that is computed more efficiently than by explicitly enumerating its range (see Section 3.4 for a more detailed discussion).

## Multiprobe scheme for the cross-polytope LSH

The space complexity of an LSH data structure is sub-quadratic, but even this is often too large (i.e., strongly super-linear in the number of points), and several methods have been proposed to address this issue. Empirically, the most efficient scheme is multiprobe LSH [118], which leads to a significantly reduced memory footprint for the hyperplane LSH. In order to make the cross-polytope LSH competitive in practice with the multiprobe hyperplane LSH, we propose a novel multiprobe scheme for the cross-polytope LSH.

This part of the algorithm can be seen as a heuristic analogue of the time-space trade-off from Chapter 2 (rather, the half of the trade-off, which corresponds to $\left.\rho_{s} \leq \rho_{q}\right)$.

## Experimental evaluation

We complement these contributions with an experimental evaluation on both real and synthetic data (SIFT vectors, tf-idf data, and a random point set). In order to make the cross-polytope LSH practical, we combine it with fast pseudo-random rotations [6] via the Fast Hadamard Transform, and feature hashing [173] to exploit sparsity of data. Our results show that for data sets with around $10^{5}$ to $10^{8}$ points, our multiprobe variant of the cross-polytope LSH is up to $10 \times$ faster than an efficient implementation of the hyperplane LSH, and up to $700 \times$ faster than a linear scan. To the best of our knowledge, our combination of techniques provides the first "exponent-optimal" algorithm that empirically improves over the hyperplane LSH in terms of query time for an exact nearest neighbor search.

We also conducted preliminary comparison with other practical ANN algorithms and concluded that the algorithm from this chapter is competitive with the state of the art. However, we will report on these experiments in future work.

Finally, we release the implementation as a new $\mathrm{C}++$ library with Python bindings called FALCONN [153].

### 3.1.2 Related work

The cross-polytope LSH functions were originally proposed in [168]. However, the analysis in that paper was mostly experimental. Specifically, the probabilities $p_{1}$ and $p_{2}$ of the proposed LSH functions were estimated empirically using the Monte Carlo method. Similar hash functions were later proposed in [72]. The latter paper also uses DFT to speed-up the random matrix-vector matrix multiplication operation. Both of the aforementioned papers consider only the single-probe algorithm.

There are several works that show lower bounds on the quality of LSH hash functions [129, 70, 140, 27]. However, those papers provide only a lower bound on the $\rho$ parameter for asymptotic values of $p_{1}$ and $p_{2}$, as opposed to an actual tradeoff between these two quantities. In this paper we provide such a trade-off, with implications as outlined in the introduction.

Finally, let us point out that the algorithm developed in this chapter has strong parallels with the algorithm from Chapter 2 as well as with Product Quantization [91]. We will elaborate on these similarities in the future work.

### 3.2 Preliminaries

In this chapter we use $\|$.$\| to denote the Euclidean (a.k.a. \ell_{2}$ ) norm on $\mathbb{R}^{d}$. The Gaussian distribution with mean zero and variance of one is denoted by $N(0,1)$. Let $\mu$ be a normalized Haar measure on $S^{d-1}$ (that is, $\mu\left(S^{d-1}\right)=1$ ). Note that $\mu$ it corresponds to the uniform distribution over $S^{d-1}$. We also let $u \sim S^{d-1}$ be a point sampled from $S^{d-1}$ uniformly at random. For $\eta \in \mathbb{R}$ we denote

$$
\Phi_{c}(\eta)=\operatorname{Pr}_{X \sim N(0,1)}[X \geq \eta]=\frac{1}{\sqrt{2 \pi}} \int_{\eta}^{\infty} e^{-t^{2} / 2} d t
$$

This chapter relies heavily on the definition of Locality-Sensitive Hashing (LSH)
and the relation between LSH and the ANN problem. See Section 1.4.4 or [77] for the introduction.

### 3.3 Cross-polytope LSH

In this section, we describe the cross-polytope LSH, analyze it, and show how to make it practical. First, we recall the definition of the cross-polytope LSH [168]: Consider the following hash family $\mathcal{H}$ for points on a unit sphere $S^{d-1} \subset \mathbb{R}^{d}$. Let $A \in \mathbb{R}^{d \times d}$ be a random matrix with i.i.d. Gaussian entries ("a random rotation"). To hash a point $x \in S^{d-1}$, we compute $y=A x /\|A x\| \in S^{d-1}$ and then find the point closest to $y$ from $\left\{ \pm e_{i}\right\}_{1 \leq i \leq d}$, where $e_{i}$ is the $i$-th standard basis vector of $\mathbb{R}^{d}$. We use the closest neighbor as a hash of $x$.

The following theorem bounds the collision probability for two points under the above family $\mathcal{H}$.

Theorem 3.3.1. Suppose that $p, q \in S^{d-1}$ are such that $\|p-q\|=\tau$, where $0<\tau<2$. Then,

$$
\ln \frac{1}{\underset{h \sim \mathcal{H}}{\operatorname{Pr}}[h(p)=h(q)]}=\frac{\tau^{2}}{4-\tau^{2}} \cdot \ln d+O_{\tau}(\ln \ln d)
$$

Before we show how to prove this theorem, we briefly describe its implications. Theorem 3.3.1 shows that the cross-polytope LSH achieves almost the same bounds on the collision probabilities as the (theoretically) optimal LSH for the sphere from [26] (see Section "Spherical LSH" there). In particular, substituting the bounds from Theorem 3.3.1 for the cross-polytope LSH into the standard reduction from Near Neighbor Search to LSH (see [77] or Section 1.4.4), we obtain the following data structure with sub-quadratic space and sublinear query time for Near Neighbor Search on a sphere.

Corollary 3.3.2. The $(c, r)-A N N$ on a unit sphere $S^{d-1}$ can be solved in space $O\left(n^{1+\rho}+d n\right)$ and query time $O\left(d \cdot n^{\rho}\right)$, where $\rho=\frac{1}{c^{2}} \cdot \frac{4-c^{2} r^{2}}{4-r^{2}}+o(1)$.

We now outline the proof of Theorem 3.3.1. For the full proof, see Section 3.8.

Due to the spherical symmetry of Gaussians, we can assume that $p=e_{1}$ and $q=\alpha e_{1}+\beta e_{2}$, where $\alpha, \beta$ are such that $\alpha^{2}+\beta^{2}=1$ and $(\alpha-1)^{2}+\beta^{2}=\tau^{2}$. Then, we expand the collision probability:

$$
\begin{align*}
& \operatorname{Pr}_{h \sim \mathcal{H}}[h(p)=h(q)] \\
& =2 d \cdot \underset{h \sim \mathcal{H}}{\operatorname{Pr}}\left[h(p)=h(q)=e_{1}\right] \\
& =2 d \cdot \underset{u, v \sim N(0,1)^{d}}{\operatorname{Pr}}\left[\forall i\left|u_{i}\right| \leq u_{1} \text { and }\left|\alpha u_{i}+\beta v_{i}\right| \leq \alpha u_{1}+\beta v_{1}\right] \\
& =2 d \cdot \underset{X_{1}, Y_{1}}{\mathrm{E}}\left[\operatorname{Pr}_{X_{2}, Y_{2}}^{\operatorname{Pr}}\left[\left|X_{2}\right| \leq X_{1} \text { and }\left|\alpha X_{2}+\beta Y_{2}\right| \leq \alpha X_{1}+\beta Y_{1}\right]^{d-1}\right], \tag{3.3}
\end{align*}
$$

where $X_{1}, Y_{1}, X_{2}, Y_{2} \sim N(0,1)$. Indeed, the first step is due to the spherical symmetry of the hash family, the second step follows from the above discussion about replacing a random orthogonal matrix with a Gaussian one and that one can assume w.l.o.g. that $p=e_{1}$ and $q=\alpha e_{1}+\beta e_{2}$; the last step is due to the independence of the entries of $u$ and $v$.

Thus, proving Theorem 3.3.1 reduces to estimating the right-hand side of (3.3). Note that the probability $\operatorname{Pr}\left[\left|X_{2}\right| \leq X_{1}\right.$ and $\left.\left|\alpha X_{2}+\beta Y_{2}\right| \leq \alpha X_{1}+\beta Y_{1}\right]$ is equal to the Gaussian area of the planar set $S_{X_{1}, Y_{1}}$ shown in Figure 3-2a. The latter is heuristically equal to $1-e^{-\Delta^{2} / 2}$, where $\Delta$ is the distance from the origin to the complement of $S_{X_{1}, Y_{1}}$, which is easy to compute (see Section 3.7 for the precise statement of this argument). Using this estimate, we compute (3.3) by taking the outer expectation.

### 3.3.1 Making the cross-polytope LSH practical

As described above, the cross-polytope LSH is not quite practical. The main bottleneck is sampling, storing, and applying a random rotation. In particular, to multiply a random Gaussian matrix with a vector, we need time proportional to $d^{2}$, which is infeasible for large $d$.

Pseudo-random rotations. To rectify this issue, we instead use pseudo-random rotations. Instead of multiplying an input vector $x$ by a random Gaussian matrix,
we apply the following linear transformation: $x \mapsto H D_{3} H D_{2} H D_{1} x$, where $H$ is the Hadamard transform, and $D_{i}$ for $i \in\{1,2,3\}$ is a random diagonal $\pm 1$-matrix. Clearly, this is an orthogonal transformation, which one can store in space $O(d)$ and evaluate in time $O(d \log d)$ using the Fast Hadamard Transform. This is similar to pseudorandom rotations used in the context of LSH [64], dimensionality reduction [6], or compressed sensing [7]. While we are currently not aware how to prove rigorously that such pseudo-random rotations perform as well as the fully random ones, empirical evaluations show that three applications of $H D_{i}$ are exactly equivalent to applying a true random rotation (when $d$ tends to infinity). We note that only two applications of $H D_{i}$ are not sufficient.

Feature hashing. While we can apply a pseudo-random rotation in time $O(d \log d)$, even this can be too slow. E.g., consider an input vector $x$ that is sparse: the number of non-zero entries of $x$ is $s$ much smaller than $d$. In this case, we can evaluate the hyperplane LSH from [55] in time $O(s)$, while computing the cross-polytope LSH (even with pseudo-random rotations) still takes time $O(d \log d)$. To speed-up the crosspolytope LSH for sparse vectors, we apply feature hashing [173]: before performing a pseudo-random rotation, we reduce the dimension from $d$ to $d^{\prime} \ll d$ by applying a linear map $x \mapsto S x$, where $S$ is a random sparse $d^{\prime} \times d$ matrix, whose columns have one non-zero $\pm 1$ entry sampled uniformly. This way, the evaluation time becomes $O\left(s+d^{\prime} \log d^{\prime}\right) .{ }^{3}$
"Partial" cross-polytope LSH. In the above discussion, we defined the crosspolytope LSH as a hash family that returns the closest neighbor among $\left\{ \pm e_{i}\right\}_{1 \leq i \leq d}$ as a hash (after a (pseudo-)random rotation). In principle, we do not have to consider all $d$ basis vectors when computing the closest neighbor. By restricting the hash to $d^{\prime} \leq d$ basis vectors instead, Theorem 3.3.1 still holds for the new hash family (with $d$ replaced by $d^{\prime}$ ) since the analysis is essentially dimension-free. This slight

[^20]Figure 3-2

generalization of the cross-polytope LSH turns out to be useful for experiments (see Section 3.6). Note that the case $d^{\prime}=1$ corresponds to the hyperplane LSH.

### 3.4 Lower bound

Let $\mathcal{H}$ be a hash family on $S^{d-1}$. For $0<r_{1}<r_{2}<2$ we would like to understand the trade-off between $p_{1}$ and $p_{2}$, where $p_{1}$ is the smallest probability of collision under $\mathcal{H}$ for points at distance at most $r_{1}$ and $p_{2}$ is the largest probability of collision for points at distance at least $r_{2}$. We focus on the case $r_{2} \approx \sqrt{2}$ because setting $r_{2}$ to $\sqrt{2}-o(1)$ (as $d$ tends to infinity) allows us to replace $p_{2}$ with the following quantity that is somewhat easier to handle:

$$
p_{2}^{*}=\operatorname{Pr}_{\substack{h \sim \mathcal{H} \\ u, v \sim S^{d-1}}}[h(u)=h(v)] .
$$

This quantity is at most $p_{2}+o(1)$, since the distance between two random points on a unit sphere $S^{d-1}$ is tightly concentrated around $\sqrt{2}$. So for a hash family $\mathcal{H}$ on a unit sphere $S^{d-1}$, we would like to understand the upper bound on $p_{1}$ in terms of $p_{2}^{*}$ and $0<r_{1}<\sqrt{2}$.

For $0 \leq \tau \leq \sqrt{2}$ and $\eta \in \mathbb{R}$, we define $\Lambda(\tau, \eta)$ to be:

$$
\operatorname{Pr}_{X, Y \sim N(0,1)}\left[X \geq \eta \text { and }\left(1-\frac{\tau^{2}}{2}\right) \cdot X+\sqrt{\tau^{2}-\frac{\tau^{4}}{4}} \cdot Y \geq \eta\right] / \operatorname{Pr}_{X \sim N(0,1)}[X \geq \eta] .
$$

We are now ready to formulate the main result of this section.

Theorem 3.4.1. Let $\mathcal{H}$ be a hash family on $S^{d-1}$ such that every function in $\mathcal{H}$ partitions the sphere into at most $T$ parts of measure at most $1 / 2$. Then we have $p_{1} \leq \Lambda\left(r_{1}, \eta\right)+o(1)$, where $\eta \in \mathbb{R}$ is such that $\Phi_{c}(\eta)=p_{2}^{*}$ and o(1) is a quantity that depends on $T$ and $r_{1}$ and tends to 0 as d tends to infinity.

The idea of the proof is first to reason about one part of the partition using the isoperimetric inequality from [73], and then to apply a certain averaging argument by proving concavity of a function related to $\Lambda$ using a delicate analytic argument. For the full proof, see Section 3.9.

We note that the above requirement of all parts induced by $\mathcal{H}$ having measure at most $1 / 2$ is only a technicality. We conjecture that Theorem 3.4.1 holds without this restriction. In any case, as we will see below, in the interesting range of parameters this restriction is essentially irrelevant.

One can observe that if every hash function in $\mathcal{H}$ partitions the sphere into at most $T$ parts, then $p_{2}^{*} \geq \frac{1}{T}$ (indeed, $p_{2}^{*}$ is precisely the average sum of squares of measures of the parts). This observation, combined with Theorem 3.4.1, leads to the following interesting consequence. Specifically, we can numerically estimate $\Lambda$ in order to give a lower bound on $\rho=\frac{\log \left(1 / p_{1}\right)}{\log \left(1 / p_{2}\right)}$ for any hash family $\mathcal{H}$ in which every function induces at most $T$ parts of measure at most $1 / 2$. See Figure 3-2b, where we plot this lower bound for $r_{1}=\sqrt{2} / 2,{ }^{4}$ together with an upper bound that is given by the cross-polytope $\mathrm{LSH}^{5}$ (for which we use numerical estimates for (3.3)). We can make several conclusions from this plot. First, the cross-polytope LSH gives an almost optimal trade-off between $\rho$ and $T$. Given that the evaluation time for the

[^21]cross-polytope LSH is $O(T \log T)$ (if one uses pseudo-random rotations), we conclude that in order to improve upon the cross-polytope LSH substantially in practice, one should design an LSH family with $\rho$ being close to optimal and evaluation time that is sublinear in $T$. We note that none of the known LSH families for a sphere has been shown to have this property. This direction looks especially interesting since the convergence of $\rho$ to the optimal value (as $T$ tends to infinity) is extremely slow (for instance, according to Figure $3-2 \mathrm{~b}$, for $r_{1}=\sqrt{2} / 2$ and $r_{2} \approx \sqrt{2}$ we need more than $10^{5}$ parts to achieve $\rho \leq 0.2$, whereas the optimal $\rho$ is $1 / 7 \approx 0.143$ ).

### 3.5 Multiprobe LSH for the cross-polytope LSH

We now describe our multiprobe scheme for the cross-polytope LSH, which is a method for reducing the number of independent hash tables in an LSH data structure. Given a query point $q$, a "standard" LSH data structure considers only a single cell in each of the $L$ hash tables (the cell is given by the hash value $h_{i}(q)$ for $i \in[L]$ ). In multiprobe LSH, we consider candidates from multiple cells in each table [118]. The rationale is the following: points $p$ that are close to $q$ but fail to collide with $q$ under hash function $h_{i}$ are still likely to hash to a value that is close to $h_{i}(q)$. By probing multiple hash locations close to $h_{i}(q)$ in the same table, multiprobe LSH achieves a given probability of success with a smaller number of hash tables than "standard" LSH. Multiprobe LSH has been shown to perform well in practice $[118,164]$.

The main ingredient in multiprobe LSH is a probing scheme for generating and ranking possible modifications of the hash value $h_{i}(q)$. The probing scheme should be computationally efficient and ensure that more likely hash locations are probed first. For a single cross-polytope hash, the order of alternative hash values is straightforward: let $x$ be the (pseudo-)randomly rotated version of query point $q$. Recall that the "main" hash value is $h_{i}(q)=\operatorname{argmax}_{j \in[d]}\left|x_{j}\right| .{ }^{6}$ Then it is easy to see that the second highest probability of collision is achieved for the hash value corresponding to the

[^22]coordinate with the second largest absolute value, etc. Therefore, we consider the indices $i \in[d]$ sorted by their absolute value as our probing sequence or "ranking" for a single cross-polytope.

The remaining question is how to combine multiple cross-polytope rankings when we have more than one hash function. As in the analysis of the cross-polytope LSH (see Section 3.3, we consider two points $q=e_{1}$ and $p=\alpha e_{1}+\beta e_{2}$ at distance $R$. Let $A^{(i)}$ be the i.i.d. Gaussian matrix of hash function $h_{i}$, and let $x^{(i)}=A^{(i)} e_{1}$ be the randomly rotated version of point $q$. Given $x^{(i)}$, we are interested in the probability of $p$ hashing to a certain combination of the individual cross-polytope rankings. More formally, let $r_{v_{i}}^{(i)}$ be the index of the $v_{i}$-th largest element of $\left|x^{(i)}\right|$, where $v \in[d]^{k}$ specifies the alternative probing location. Then we would like to compute

$$
\begin{aligned}
A_{A^{(1)}, \ldots, A^{(k)}} \operatorname{Pr}_{i} & {\left[h_{i}(p)=r_{v_{i}}^{(i)} \text { for all } i \in[k] \mid A^{(i)} q=x^{(i)}\right] } \\
& =\prod_{i=1}^{k} \operatorname{Pr}_{A^{(i)}}\left[\operatorname{argmax}_{j \in[d]}\left|\left(\alpha \cdot A^{(i)} e_{1}+\beta \cdot A^{(i)} e_{2}\right)_{j}\right|=r_{v_{i}}^{(i)} \mid A^{(i)} e_{1}=x^{(i)}\right] .
\end{aligned}
$$

If we knew this probability for all $v \in[d]^{k}$, we could sort the probing locations by their probability. We now show how to approximate this probability efficiently for a single value of $i$ (and hence drop the superscripts to simplify notation). WLOG, we permute the rows of $A$ so that $r_{v}=v$ and get

$$
\begin{aligned}
& \operatorname{Pr}_{A}\left[\operatorname{argmax}_{j \in[d]}\left|\left(\alpha x+\beta \cdot A e_{2}\right)_{j}\right|=v \mid A e_{1}=x\right] \\
& =\operatorname{Pr}_{y \sim N\left(0, I_{d}\right)}\left[\operatorname{argmax}_{j \in[d]}\left|\left(x+\frac{\beta}{\alpha} \cdot y\right)_{j}\right|=v\right] .
\end{aligned}
$$

The RHS is the Gaussian measure of the set $S=\left\{\left.y \in \mathbb{R}^{d}\left|\operatorname{argmax}_{j \in[d]}\right|\left(x+\frac{\beta}{\alpha} y\right)_{j} \right\rvert\,=v\right\}$. Similar to the analysis of the cross-polytope LSH, we approximate the measure of $S$ by its distance to the origin. Then the probability of probing location $v$ is proportional to $\exp \left(-\left\|y_{x, v}\right\|^{2}\right)$, where $y_{x, v}$ is the shortest vector $y$ such that $\operatorname{argmax}_{j}|x+y|_{j}=v$. Note that the factor $\beta / \alpha$ becomes a proportionality constant, and hence the probing scheme does not require to know the distance $R$. For computational performance and
simplicity, we make a further approximation and use $y_{x, v}=\left(\max _{i}\left|x_{i}\right|-\left|x_{v}\right|\right) \cdot e_{v}$, i.e., we only consider modifying a single coordinate to reach the set $S$.

Once we have estimated the probabilities for each $v_{i} \in[d]$, we incrementally construct the probing sequence using a binary heap, similar to the approach in [118]. For a probing sequence of length $m$, the resulting algorithm has running time $O(L \cdot d \log d+m \log m)$. In our experiments, we found that the $O(L \cdot d \log d)$ time taken to sort the probing candidates $v_{i}$ dominated the running time of the hash function evaluation. In order to circumvent this issue, we use an incremental sorting approach that only sorts the relevant parts of each cross-polytope and gives a running time of $O(L \cdot d+m \log m)$.

### 3.6 Experiments

We now show that the cross-polytope LSH, combined with our multiprobe extension, leads to an algorithm that is also efficient in practice and improves over the hyperplane LSH on several data sets. The focus of our experiments is the query time for an exact nearest neighbor search. Since hyperplane LSH has been compared to other nearest-neighbor algorithms before [158], we limit our attention to the relative speed-up compared with hyperplane hashing.

We evaluate the two hashing schemes on three types of data sets. We use a synthetic data set of randomly generated points because this allows us to vary a single problem parameter while keeping the remaining parameters constant. We also investigate the performance of our algorithm on real data: two tf-idf data sets [113] and a set of SIFT feature vectors [91]. We have chosen these data sets in order to illustrate when the cross-polytope LSH gives large improvements over the hyperplane LSH, and when the improvements are more modest. See Section 3.10 for a more detailed description of the data sets and our experimental setup (implementation details, CPU, etc.).

In all experiments, we set the algorithm parameters so that the empirical probability of successfully finding the exact nearest neighbor is at least 0.9 . Moreover, we set the
number of LSH tables $L$ so that the amount of additional memory occupied by the LSH data structure is comparable to the amount of memory necessary for storing the data set. We believe that this is the most interesting regime because significant memory overheads are often impossible for large data sets. In order to determine the parameters that are not fixed by the above constraints, we perform a grid search over the remaining parameter space and report the best combination of parameters. For the cross-polytope hash, we consider "partial" cross-polytopes in the last of the $k$ hash functions in order to get a smooth trade-off between the various parameters (see Section 3.3.1).

Multiprobe experiments. In order to demonstrate that the multiprobe scheme is critical for making the cross-polytope LSH competitive with hyperplane hashing, we compare the performance of a "standard" cross-polytope LSH data structure with our multiprobe variant on an instance of the random data set ( $n=2^{20}, d=128$ ). As can be seen in Table 3.2 (Section 3.10), the multiprobe variant is about $13 \times$ faster in our memory-constrained setting $(L=10)$. Note that in all of the following experiments, the speed-up of the multiprobe cross-polytope LSH compared to the multiprobe hyperplane LSH is less than $11 \times$. Hence without our multiprobe addition, the crosspolytope LSH would be slower than the hyperplane LSH, for which a multiprobe scheme is already known [118].

Experiments on random data. Next, we show that the better time complexity of the cross-polytope LSH already applies for moderate values of $n$. In particular, we compare the cross-polytope LSH, combined with fast rotations (Section 3.3.1) and our multiprobe scheme, to a multi-probe hyperplane LSH on random data. We keep the dimension $d=128$ and the distance to the nearest neighbor $R=\sqrt{2} / 2$ fixed, and vary the size of the data set from $2^{20}$ to $2^{28}$. The number of hash tables $L$ is set to 10. For $2^{20}$ points, the cross-polytope LSH is already $3.5 \times$ faster than the hyperplane LSH, and for $n=2^{28}$ the speedup is $10.3 \times$ (see Table 3.3 in Section 3.10). Compared to a linear scan, the speed-up achieved by the cross-polytope LSH ranges from $76 \times$

| Data set | Method | Query <br> time (ms) | Speed-up <br> vs HP | Best $k$ | Number of <br> candidates | Hashing <br> time $(\mathrm{ms})$ | Distances <br> time $(\mathrm{ms})$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| NYT | HP | 120 ms |  | 19 | 57,200 | 16 | 96 |
| NYT | CP | 35 ms | $\mathbf{3 . 4 \times}$ | $2(64)$ | 17,900 | 3.0 | 30 |
| pubmed | HP | 857 ms |  | 20 | $1,480,000$ | 36 | 762 |
| pubmed | CP | 213 ms | $\mathbf{4 . 0 \times}$ | $2(512)$ | 304,000 | 18 | 168 |
| SIFT | HP | 3.7 ms |  | 30 | 18,628 | 0.2 | 3.0 |
| SIFT | CP | 3.1 ms | $\mathbf{1 . 2 \times}$ | $6(1)$ | 13,000 | 0.6 | 2.2 |

Table 3.1: Average running times for a single nearest neighbor query with the hyperplane (HP) and cross-polytope (CP) algorithms on three real data sets. The cross-polytope LSH is faster than the hyperplane LSH on all data sets, with significant speed-ups for the two tf-idf data sets NYT and pubmed. For the cross-polytope LSH, the entries for $k$ include both the number of individual hash functions per table and (in parenthesis) the dimension of the last of the $k$ cross-polytopes.
for $n=2^{20}$ to about $700 \times$ for $n=2^{28}$.

Experiments on real data. On the SIFT data set ( $n=10^{6}$ and $d=128$ ), the cross-polytope LSH achieves a modest speed-up of $1.2 \times$ compared to the hyperplane LSH (see Table 3.1). On the other hand, the speed-up is is $3-4 \times$ on the two tf-idf data sets, which is a significant improvement considering the relatively small size of the NYT data set $(n \approx 300,000)$. One important difference between the data sets is that the typical distance to the nearest neighbor is smaller in the SIFT data set, which can make the nearest neighbor problem easier (see Section 3.10). Since the tf-idf data sets are very high-dimensional but sparse ( $d \approx 100,000$ ), we use the feature hashing approach described in Section 3.3.1 in order to reduce the hashing time of the cross-polytope LSH (the standard hyperplane LSH already runs in time proportional to the sparsity of a vector). We use 512 and 2048 as feature hashing dimensions for NYT and pubmed, respectively.

### 3.7 Appendix: Gaussian measure of a planar set

In this Section we formalize the intuition that the standard Gaussian measure of a closed subset $A \subseteq \mathbb{R}^{2}$ behaves like $e^{-\Delta_{A}^{2} / 2}$, where $\Delta_{A}$ is the distance from the origin
to $A$, unless $A$ is quite special.
For a closed subset $A \subseteq \mathbb{R}^{2}$ and $r>0$ denote $0 \leq \mu_{A}(r) \leq 1$ the normalized measure of the intersection $A \cap r S^{1}$ ( $A$ with the circle centered in the origin and of radius $r$ ):

$$
\mu_{A}(r):=\frac{\mu\left(A \cap r S^{1}\right)}{2 \pi r} ;
$$

here $\mu$ is the standard one-dimensional Lebesgue measure (see Figure 3-3a). Denote $\Delta_{A}:=\inf \left\{r>0: \mu_{A}(r)>0\right\}$ the (essential) distance from the origin to $A$. Let $\mathcal{G}(A)$ be the standard Gaussian measure of $A$.

Lemma 3.7.1. Suppose that $A \subseteq \mathbb{R}^{2}$ is a closed set such that $\mu_{A}(r)$ is non-decreasing. Then,

$$
\sup _{r>0}\left(\mu_{A}(r) \cdot e^{-r^{2} / 2}\right) \leq \mathcal{G}(A) \leq e^{-\Delta_{A}^{2} / 2}
$$

Proof. For the upper bound, we note that

$$
\mathcal{G}(A)=\int_{0}^{\infty} \mu_{A}(r) \cdot r e^{-r^{2} / 2} d r \leq \int_{\Delta_{A}}^{\infty} r e^{-r^{2} / 2} d r=e^{-\Delta_{A}^{2} / 2}
$$

For the lower bound, we similarly have, for every $r^{*}>0$,

$$
\mathcal{G}(A)=\int_{0}^{\infty} \mu_{A}(r) \cdot r e^{-r^{2} / 2} d r \geq \mu_{A}\left(r^{*}\right) \cdot \int_{r^{*}}^{\infty} r e^{-r^{2} / 2} d r=\mu_{A}\left(r^{*}\right) e^{-\left(r^{*}\right)^{2} / 2}
$$

where we use that $\mu_{A}\left(r^{*}\right)$ is non-decreasing.
Now we derive two corollaries of Lemma 3.7.1.
Lemma 3.7.2. Let $K \subseteq \mathbb{R}^{2}$ be the complement of an open convex subset of the plane that is symmetric around the origin. Then, for every $0<\varepsilon<1 / 3$,

$$
\Omega\left(\varepsilon^{1 / 2} \cdot e^{-(1+\varepsilon) \cdot \Delta_{K}^{2} / 2}\right) \leq \mathcal{G}(K) \leq e^{-\Delta_{K}^{2} / 2}
$$

Proof. This follows from Lemma 3.7.1: indeed, due to the convexity of the complement of $K, \mu_{K}(r)$ is non-decreasing. It is easy to check that

$$
\mu_{K}\left((1+\varepsilon) \Delta_{K}\right)=\Omega\left(\varepsilon^{1 / 2}\right)
$$

again, due to the convexity (see Figure 3-3b). Thus, the required bounds follow.
Lemma 3.7.3. Let $K \subseteq \mathbb{R}^{2}$ be an intersection of two closed half-planes such that:

- $K$ does not contain a line;
- the "corner" of $K$ is the closest point of $K$ to the origin;
- the angle between half-planes equals to $0<\alpha<\pi$.

Then, for every $0<\varepsilon<1 / 2$,

$$
\Omega_{\alpha}\left(\varepsilon \cdot e^{-(1+\varepsilon) \cdot \Delta_{K}^{2}}\right) \leq \mathcal{G}(K) \leq e^{-\Delta_{K}^{2} / 2}
$$

Proof. This, again, follows from Lemma 3.7.1. The second condition implies that $\mu_{K}(r)$ is non-decreasing, and an easy computation shows that

$$
\mu_{K}\left((1+\varepsilon) \Delta_{K}\right) \geq \Omega_{\alpha}(\varepsilon)
$$

(see Figure 3-3c).

### 3.8 Appendix: Proof of Theorem 3.3.1

In this section we complete the proof of Theorem 3.3.1, following the outline from Section 3.3. Our starting point is the collision probability bound from Eqn. (3.3).

For $u, v \in \mathbb{R}$ with $u \geq 0$ and $\alpha u+\beta v \geq 0$ define,

$$
\sigma(u, v)=\operatorname{Pr}_{X_{2}, Y_{2} \sim N(0,1)}\left[\left|X_{2}\right| \leq u \text { and }\left|\alpha X_{2}+\beta Y_{2}\right| \leq \alpha u+\beta v\right] .
$$

Then, the right-hand side of (3.3) is equal to

$$
2 d \cdot \underset{X_{1}, Y_{1} \sim N(0,1)}{\mathrm{E}}\left[\sigma\left(X_{1}, Y_{1}\right)^{d-1}\right] .
$$

Let us define

$$
\Delta(u, v)=\min \{u, \alpha u+\beta v\} .
$$

Lemma 3.8.1. For every $0<\varepsilon<1 / 3$,

$$
1-e^{-\Delta(u, v)^{2} / 2} \leq \sigma(u, v) \leq 1-\Omega\left(\varepsilon^{1 / 2} \cdot e^{-(1+\varepsilon) \Delta(u, v)^{2} / 2}\right)
$$

Proof. This is a combination of Lemma 3.7.2 together with the following obvious observation: the distance from the origin to the set $\{(x, y):|x| \geq u$ or $|\alpha x+\beta y| \geq$ $\alpha u+\beta v\}$ is equal to $\Delta(u, v)$ (see Figure 3-2a).

Lemma 3.8.2. For every $t \geq 0$ and $0<\varepsilon<1 / 3$,

$$
\Omega_{\tau}\left(\varepsilon \cdot e^{-(1+\varepsilon) \cdot \frac{4}{4-\tau^{2}} \cdot \frac{t^{2}}{2}}\right) \leq \operatorname{Pr}_{X_{1}, Y_{1} \sim N(0,1)}\left[\Delta\left(X_{1}, Y_{1}\right) \geq t\right] \leq e^{-\frac{4}{4-\tau^{2}} \cdot \frac{t^{2}}{2}} .
$$

Proof. Similar to the previous lemma, this is a consequence of Lemma 3.7.3 together with the fact that the squared distance from the origin to the set $\{(x, y): x \geq$ $t$ and $\alpha x+\beta y \geq t\}$ is equal to $\frac{4}{4-\tau^{2}} \cdot t^{2}$.

### 3.8.1 Idealized proof

Let us expand (3.3) further, assuming that the "idealized" versions of Lemma 3.8.1 and Lemma 3.8.2 hold. Namely, we assume that

$$
\begin{equation*}
\sigma(u, v)=1-e^{-\Delta(u, v)^{2} / 2} \tag{3.4}
\end{equation*}
$$

and

$$
\begin{equation*}
\operatorname{Pr}_{X_{1}, Y_{1} \sim N(0,1)}\left[\Delta\left(X_{1}, Y_{1}\right) \geq t\right]=e^{-\frac{4}{4-\tau^{2}} \cdot \frac{t^{2}}{2}} . \tag{3.5}
\end{equation*}
$$

In the next section we redo the computations using the precise bounds for $\sigma(u, v)$ and $\operatorname{Pr}\left[\Delta\left(X_{1}, Y_{1}\right) \geq t\right]$.

Expanding Eqn. (3.3), we have

$$
\begin{align*}
\underset{X_{1}, Y_{1} \sim N(0,1)}{\mathrm{E}}\left[\sigma\left(X_{1}, Y_{1}\right)^{d-1}\right] & =\int_{0}^{1} \operatorname{Pr}_{X_{1}, Y_{1} \sim N(0,1)}\left[\sigma\left(X_{1}, Y_{1}\right) \geq t^{\frac{1}{d-1}}\right] d t \\
& =\int_{0}^{1} \operatorname{Pr}_{X_{1}, Y_{1} \sim N(0,1)}\left[e^{-\Delta\left(X_{1}, Y_{1}\right)^{2} / 2} \leq 1-t^{\frac{1}{d-1}}\right] d t \\
& =\int_{0}^{1}\left(1-t^{\frac{1}{d-1}}\right)^{\frac{4}{4-\tau^{2}}} d t \\
& =(d-1) \cdot \int_{0}^{1}(1-u)^{\frac{4}{4-\tau^{2}}} u^{d-2} d t \\
& =(d-1) \cdot B\left(\frac{8-\tau^{2}}{4-\tau^{2}} ; d-1\right) \\
& =\Theta_{\tau}(1) \cdot d^{-\frac{4}{4-\tau^{2}}}, \tag{3.6}
\end{align*}
$$

where:

- the first step is a standard expansion of an expectation;
- the second step is due to (3.4);
- the third step is due to (3.5);
- the fourth step is a change of variables;
- the fifth step is a definition of the Beta function;
- the sixth step is due to the Stirling approximation.

Overall, substituting (3.6) into (3.3), we get:

$$
\ln \frac{1}{\operatorname{Pr}_{h \sim \mathcal{H}}[h(p)=h(q)]}=\frac{\tau^{2}}{4-\tau^{2}} \cdot \ln d \pm O_{\tau}(1) .
$$

### 3.8.2 Real proof

We now perform calculations using the bounds (involving $\varepsilon$ ) from Lemma 3.8.1 and Lemma 3.8.2. We set $\varepsilon=1 / d$ and obtain the following asymptotic statements:

$$
\sigma(u, v)=1-d^{ \pm O(1)} \cdot e^{-\left(1 \pm d^{-\Omega(1)}\right) \cdot \Delta(u, v)^{2} / 2}
$$

and

$$
\operatorname{Pr}_{X, Y \sim N(0,1)}[\Delta(X, Y) \geq t]=d^{ \pm O(1)} \cdot e^{-\left(1 \pm d^{-\Omega(1)}\right) \cdot \frac{4}{4-\tau^{2}} \cdot \frac{t^{2}}{2}}
$$

Then, we can repeat the "idealized" proof (see Eqn. (3.6)) verbatim with the new estimates and obtain the final form of Theorem 3.3.1:

$$
\ln \frac{1}{\underset{h \sim \mathcal{H}}{\operatorname{Pr}}[h(p)=h(q)]}=\frac{\tau^{2}}{4-\tau^{2}} \cdot \ln d \pm O_{\tau}(\ln \ln d)
$$

Note the difference in the low order term between idealized and the real version. As we argue in Section 3.4, the latter $O_{\tau}(\ln \ln d)$ is, in fact, tight.

### 3.9 Appendix: Proof of Theorem 3.4.1

Lemma 3.9.1. Let $A \subset S^{d-1}$ be a measurable subset of a sphere with $\mu(A)=\mu_{0} \leq 1 / 2$. Then, for $0<\tau<\sqrt{2}$, one has

$$
\begin{align*}
& \operatorname{Pr}_{u, v \sim S^{d-1}}[v \in A \mid u \in A,\|u-v\| \leq \tau] \\
& =\frac{\operatorname{Pr}_{X, Y \sim N(0,1)}[X \geq \eta \text { and } \alpha X+\beta Y \geq \eta]+o(1)}{\operatorname{Pr}_{X \sim N(0,1)}[X \geq \eta]+o(1)}, \tag{3.7}
\end{align*}
$$

where:

- $\alpha=1-\frac{\tau^{2}}{2}$;
- $\beta=\sqrt{\tau^{2}-\frac{\tau^{4}}{4}} ;$
- $\eta \in \mathbb{R}$ is such that $\operatorname{Pr}_{X \sim N(0,1)}[X \geq \eta]=\mu_{0}$.

In particular, if $\mu_{0}=\Omega(1)$, then

$$
\operatorname{Pr}_{u, v \sim S^{d-1}}[v \in A \mid u \in A,\|u-v\| \leq \tau]=\Lambda\left(\tau, \Phi_{c}^{-1}\left(\mu_{0}\right)\right)+o(1)
$$

Proof. First, the left-hand side of (3.7) is maximized by a spherical cap of measure $\mu_{0}$. This follows from Theorem 5 of [73]. So, from now on we assume that $A$ is a spherical
cap.
Second, one has

$$
\begin{aligned}
& \operatorname{Pr}_{u, v \sim S^{d-1}}[v \in A \mid u \in A,\|u-v\| \leq \tau] \\
& =\operatorname{Pr}_{u, v \sim S^{d-1}}[v \in A \mid u \in A,\|u-v\|=\tau \pm o(1)]+o(1) \\
& =\frac{\operatorname{Pr}_{u \sim S^{d-1}}\left[u_{1} \geq \widetilde{\eta} \text { and }(\alpha \pm o(1)) u_{1}+(\beta \pm o(1)) u_{2} \geq \widetilde{\eta}\right]}{\operatorname{Pr}_{u \sim S^{d-1}}\left[u_{1} \geq \widetilde{\eta}\right]}+o(1) \\
& =\frac{\operatorname{Pr}_{X, Y \sim N(0,1)}[X \geq \eta \text { and } \alpha X+\beta Y \geq \eta]+o(1)}{\operatorname{Pr}_{X \sim N(0,1)}[X \geq \eta]+o(1)}
\end{aligned}
$$

where $\widetilde{\eta}$ is such that $\underset{u \sim S^{d-1}}{\operatorname{Pr}}\left[u_{1} \geq \widetilde{\eta}\right]=\mu_{0}$ and:

- the first step is due to the concentration of measure on the sphere;
- the second step is expansion of the conditional probability;
- the third step is due to the fact that a $O(1)$-dimensional projection of the uniform measure on a sphere of radius $\sqrt{d}$ in $\mathbb{R}^{d}$ converges in total variation to a standard Gaussian measure [68].

Lemma 3.9.2. For every $0<\tau<\sqrt{2}$, the function $\mu \mapsto \Lambda\left(\tau, \Phi_{c}^{-1}(\mu)\right)$ is concave for $0<\mu<1 / 2$.

Proof. Abusing notation, for this proof we denote $\Lambda(\eta)=\Lambda(\tau, \eta)$ and

$$
I(\eta)=\operatorname{Pr}_{X, Y \sim N(0,1)}[X \geq \eta \text { and } \alpha X+\beta Y \geq \eta]
$$

(that is, $\left.\Lambda(\eta)=I(\eta) / \Phi_{c}(\eta)\right)$. One has $\Phi_{c}^{\prime}(\eta)=-\frac{e^{-\eta^{2} / 2}}{\sqrt{2 \pi}}$ and

$$
I^{\prime}(\eta)=-\sqrt{\frac{2}{\pi}} \cdot e^{-\eta^{2} / 2} \cdot \Phi_{c}\left(\frac{(1-\alpha) \eta}{\beta}\right) .
$$

Combining, we get

$$
\Lambda^{\prime}(\eta)=\frac{e^{-\eta^{2} / 2}}{\sqrt{2 \pi}} \cdot \frac{I(\eta)-2 \Phi_{c}(\eta) \Phi_{c}\left(\frac{(1-\alpha) \eta}{\beta}\right)}{\Phi_{c}(\eta)^{2}}
$$

and

$$
\frac{d \Lambda\left(\Phi_{c}^{-1}(\mu)\right)}{d \mu}=\frac{2 \Phi_{c}\left(\eta^{*}\right) \Phi_{c}\left(\frac{(1-\alpha) \eta^{*}}{\beta}\right)-I\left(\eta^{*}\right)}{\Phi_{c}\left(\eta^{*}\right)^{2}}=: \Pi\left(\eta^{*}\right)
$$

where $\eta^{*}=\eta^{*}(\mu)=\Phi_{c}^{-1}(\mu)$. It is sufficient to show that $\Pi\left(\eta^{*}\right)$ is non-decreasing in $\eta^{*}$ for $\eta^{*} \geq 0$.

We have

$$
\begin{array}{r}
\Pi^{\prime}(\eta)=\sqrt{\frac{2}{\pi}} \cdot \frac{e^{-\eta^{2} / 2}}{\Phi_{c}(\eta)^{3}}\left(2 \cdot \Phi_{c}(\eta) \Phi_{c}\left(\frac{(1-\alpha) \eta}{\beta}\right)-I(\eta)-\frac{1-\alpha}{\beta} \cdot e^{\frac{\alpha(1-\alpha)}{\beta^{2}} \cdot \eta^{2}} \Phi_{c}(\eta)^{2}\right) \\
=: \sqrt{\frac{2}{\pi}} \cdot \frac{e^{-\eta^{2} / 2}}{\Phi_{c}(\eta)^{3}} \cdot \Omega(\eta)
\end{array}
$$

We need to show that $\Omega(\eta) \geq 0$ for $\eta \geq 0$. We will do this by showing that $\Omega^{\prime}(\eta) \leq 0$ for $\eta \geq 0$ and that $\lim _{\eta \rightarrow \infty} \Omega(\eta)=0$. The latter is obvious, so let us show the former.

$$
\Omega^{\prime}(\eta)=-\frac{2 \alpha(1-\alpha)^{2}}{\beta^{3}} \cdot e^{\frac{\alpha(1-\alpha)}{\beta^{2}} \eta^{2}} \cdot \Phi_{c}(\eta)^{2} \cdot \eta \leq 0
$$

for $\eta \geq 0$.

Now we are ready to prove Theorem 3.4.1. Let us first assume that all the parts have measure $\Omega(1)$. Later we will show that this assumption can be removed. W.l.o.g.
we assume that functions from the family have subsets integers as a range. We have,

$$
\begin{aligned}
p_{1} & \leq \operatorname{Pr}_{\substack{u, v \sim S^{d-1} \\
h \sim \mathcal{H}}}[h(u)=h(v) \mid\|u-v\| \leq \tau] \\
& =\underset{h \sim \mathcal{H}}{\mathrm{E}}\left[\sum_{i} \mu\left(h^{-1}(i)\right) \operatorname{Pr}\left[v \in h^{-1}(i) \mid u \in h^{-1}(i),\|u-v\| \leq \tau\right]\right] \\
& \leq \underset{h \sim \mathcal{H}}{\mathrm{E}}\left[\sum_{i} \mu\left(h^{-1}(i)\right) \Lambda\left(\tau, \Phi_{c}^{-1}\left(\mu\left(h^{-1}(i)\right)\right)\right)\right]+o(1) \\
& \leq \Lambda\left(\tau, \Phi_{c}^{-1}\left(\underset{h \sim \mathcal{H}}{\mathrm{E}}\left[\sum_{i} \mu\left(h^{-1}(i)\right)^{2}\right]\right)\right)+o(1) \\
& \leq \Lambda\left(\tau, \Phi_{c}^{-1}\left(p^{*}(\mathcal{H})\right)\right)+o(1),
\end{aligned}
$$

where:

- the first step is by the definition of $p_{1}$;
- the third step is due to the condition $\mu\left(h^{-1}(i)\right)=\Omega(1)$ and Lemma 3.9.1;
- the fourth step is due to Lemma 3.9.2 and the assumption $\mu\left(h^{-1}(i)\right) \leq 1 / 2$;
- the final step is due to the definition of $p^{*}(\mathcal{H})$.

To get rid of the assumption that a measure of every part is $\Omega(1)$ observe that all parts with measure at most $\varepsilon$ contribute to the expectation at most $\varepsilon \cdot T$, since there are at most $T$ pieces in total. Note that if $\varepsilon=o(1)$, then $\varepsilon \cdot T=o(1)$, since we assume $T$ being fixed.

### 3.10 Appendix: Further description of experiments

In order to compare meaningful running time numbers, we have written fast $\mathrm{C}++$ implementations of both the cross-polytope LSH and the hyperplane LSH. This enables a fair comparison since both implementations have been optimized by us to the same degree. In particular, hyperplane hashing can be implemented efficiently using a matrix-vector multiplication sub-routine for which we use the Eigen library (Eigen is
also used for all other linear algebra operations). For the fast pseudo-random rotation in the cross-polytope LSH, we have written a SIMD-optimized version of the Fast Hadamard Transform (FHT). We compiled our code with g++4.9 and the -O3 flag. All experiments except those in Table 3.3 ran on an Intel Core i5-2500 CPU (3.3-3.7 $\mathrm{GHz}, 6 \mathrm{MB}$ cache) with 8 GB of RAM. Since 8 GB of RAM was too small for the larger values of $n$, we ran the experiments in Table 3.3 on a machine with an Intel Xeon E5-2690 v2 CPU (3.0 GHz, 25 MB cache) and 512 GB of RAM.

In our experiments, we evaluate the performance of the cross-polytope LSH on the following data sets. Figure 3-4 shows the distribution of distances to the nearest neighbor for the four data sets.
random For the random data sets, we generate a set of $n$ points uniformly at random on the unit sphere. In order to generate a query, we pick a random point $q^{\prime}$ from the data set and generate a point at distance $R$ from $q^{\prime}$ on the unit sphere. In our experiments, we vary the dimension of the point set between 128 and 1,024 . Experiments with the random data set are useful because we can study the impact of various parameters (e.g., the dimension $d$ or the number of points $n$ ) while keeping the remaining parameters constant.
pubmed / NYT The pubmed and NYT data sets contain bag-of-words representations of medical paper abstracts and newspaper articles, respectively [113]. We convert this representation into standard tf-idf feature vectors with dimensionality about 100,000 . The number of points in the pubmed data set is about 8 million, for NYT it is 300,000 . Before setting up the LSH data structures, we set 1000 data points aside as query vectors. When selecting query vectors, we limit our attention to points for which the inner product with the nearest neighbor is between 0.3 and 0.8 . We believe that this is the most interesting range since near-duplicates (inner product close to 1) can be identified more efficiently with other methods, and points without a close nearest neighbor (inner product less than 0.3 ) often do not have a semantically meaningful match.

SIFT We use the standard data set of one million SIFT feature vectors from [91],
which also contains a set of 10,000 query vectors. The SIFT feature vectors have dimension 128 and (approximately) live on a sphere. We normalize the feature vectors to unit length but keep the original nearest neighbor assignments-this is possible because only a very small fraction of nearest neighbors changes through normalization. We include this data set as an example where the speed-up of the cross-polytope LSH is more modest.

| Method | $k$ | Last CP <br> dimension | Extra <br> probes | Query <br> time (ms) | Number of <br> candidates | CP hashing <br> time (ms) | Distances <br> time (ms) |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Single-probe | 1 | 128 | 0 | $\mathbf{6 . 7}$ | 39800 | 0.01 | 6.3 |
| Multiprobe | 3 | 16 | 896 | $\mathbf{0 . 5 1}$ | 867 | 0.22 | 0.16 |

Table 3.2: Comparison of "standard" LSH using the cross-polytope (CP) hash vs. our multiprobe variant ( $L=10$ in both cases). On a random data set with $n=2^{20}, d=128$, and $R=\sqrt{2} / 2$, the single-probe scheme requires $13 \times$ more time per query. Due to the larger value of $k$, the multiprobe variant performs fewer distance computations, which leads to a better trade-off between the hash computation time and the time spent on computing distances to candidates from the hash tables.

| Data set size $n$ | $2^{20}$ | $2^{22}$ | $2^{24}$ | $2^{26}$ | $2^{28}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| HP query time $(\mathrm{ms})$ | 2.6 | 7.4 | 25 | 63 | 185 |
| CP query time $(\mathrm{ms})$ | 0.75 | 1.4 | 3.1 | 8.8 | 18 |
| Speed-up | $\mathbf{3 . 5 \times}$ | $\mathbf{5 . 3 \times}$ | $\mathbf{8 . 1 \times}$ | $\mathbf{7 . 2 \times}$ | $\mathbf{1 0 . 3 \times}$ |
| $k$ for CP | $3(16)$ | $3(64)$ | $3(128)$ | $4(2)$ | $4(64)$ |

Table 3.3: Average running times for a single nearest neighbor query with the hyperplane (HP) and cross-polytope (CP) algorithms on a random data set with $d=128$ and $R=\sqrt{2} / 2$. The cross-polytope LSH is up to $10 \times$ faster than the hyperplane LSH. The last row of the table indicates the optimal choice of $k$ for the cross-polytope LSH and (in parenthesis) the dimension of the last of the $k$ cross-polytopes; all other crosspolytopes have full dimension 128. Note that the speed-up ratio is not monotonically increasing because the cross-polytope LSH performs better for values of $n$ where the optimal setting of $k$ uses a last cross-polytope with high dimension.

Figure 3-3: Gaussian measure of planar sets

(a) Defintion of $\mu_{A}(r)$

(b) For Lemma 3.7.2

(c) For Lemma 3.7.3


Figure 3-4: Distance to the nearest neighbor for the four data sets used in our experiments. The SIFT data set has the closest nearest neighbors.

## Chapter 4

## ANN algorithms for general symmetric norms

### 4.1 Introduction

In this chapter we will study the ANN problem for a general class of metrics. The best-studied metrics are the Hamming $\left(\ell_{1}\right)$ and the Euclidean $\left(\ell_{2}\right)$ distances. There are good reasons for this: $\ell_{1}$ and $\ell_{2}$ are very common in applications and admit very efficient algorithms based on hashing, in particular, Locality-Sensitive Hashing (LSH) [85, 13] and its data-dependent versions [18, 26]. Hashing-based algorithms for ANN over $\ell_{1} / \ell_{2}$ have now been the subject of a two-decade-long line of work, leading to a very good understanding of algorithms and their limitations. All such algorithms for $c$-approximate ANN obtain space $n^{1+\rho_{u}+o(1)}$ and query time $n^{\rho_{q}+o(1)}$ for some exponents $\rho_{u}$ and $\rho_{q}<1$ dependent on $c$; e.g., the most recent algorithm from Chapter 2 gives tight time-space trade-offs for every approximation factor $c>1$. ${ }^{1}$ We point the reader the introduction and Chapter 2, which summarize the state of affairs of the high-dimensional ANN over $\ell_{1} / \ell_{2}$. A practical perspective is presented in the surveys [171, 172].

Beyond $\ell_{1}$ and $\ell_{2}$, the landscape of ANN is much more mysterious, despite having

[^23]received significant attention. In 1998, [81] showed an efficient data structure for $\ell_{\infty}$ for $c=O(\log \log d)$ approximation. There are a few extensions of this result to other metrics, some of which proceed via embedding a metric into $\ell_{\infty}$ (see Section 4.1.3). However, we are still very far from having a general recipe for ANN data structures for general metrics with a non-trivial approximation; this is in stark contrast with the success of the low-dimensional regime. This state of affairs motivates the following broad question.

Problem 4.1.1. For a given approximation $c>1$, which metric spaces allow efficient ANN algorithms?

An algorithm for general metrics is highly desirable both in theory and in practice. From the theoretical perspective, we are interested in a common theory of ANN algorithms for a wide class of distances. Such a theory would yield data structures (or impossibility results) for a variety of important distance measures for which we still do not know efficient ANN algorithms (e.g., matrix norms, the Earth Mover's Distance (EMD), the edit distance, etc.). Perhaps even more tantalizing is understanding what exactly makes some distances harder than others, and how to quantify that hardness. From the practical perspective, it is also desirable to have a generic algorithm: one that either uses the underlying distance measure as a black box, or provides a "knob" to easily specialize to any desired distance. In practice, one must oftentimes tune the distance to the specifics of the application, and hence algorithms that allow such tuning without major re-implementations are preferred.

In this paper, we focus on the following important case of Problem 4.1.1.
Problem 4.1.2. Solve Problem 4.1.1 for high-dimensional normed spaces.
(See Section 1.4.1 for the definition of a norm).
Norms are important for two reasons. First, most metric spaces arising in applications are actually norms (e.g., the Earth-Mover Distance [134]). Second, norms are geometrically nicer than general metrics, so there is hope for a coherent theory (e.g., for the problems of sketching and streaming norms, see the generic results of Chapter 5 and [42]). Using embeddings into $\ell_{2}$ [92, 32], one can solve ANN for any norm with
approximation $O(\sqrt{d / \varepsilon})$, space $n^{1+\varepsilon}$, and query time $n^{\varepsilon}$, where $0<\varepsilon<1 / 2$ is a constant; however, no better results are known in general.

### 4.1.1 The main result

In this paper we nearly settle Problem 4.1.2 for symmetric norms, i.e., norms that are invariant under all permutations and changes of signs of the coordinates of a vector. We show the following general result:

Theorem 4.1.3. For every $n, d=n^{o(1)}$, and every $d$-dimensional symmetric norm $\|\cdot\|$, there exists a data structure for ANN over $\|\cdot\|$ for n-point datasets with approximation $(\log \log n)^{O(1)}$ space $n^{1+o(1)}$, and query time $n^{o(1)}$.

We note that the techniques behind Theorem 4.1.3 cannot be extended to general norms; see details in Section 4.1.6.

### 4.1.2 Why symmetric norms?

The class of symmetric norms is, in some sense, a sweet spot. On the one hand, symmetric norms are mathematically nice and, as we show, allow for a clean characterization that leads to an efficient ANN data structure (see the proof overview from Section 4.1.4). On the other hand, symmetric norms vastly generalize $\ell_{p}$ distances and enable many new interesting examples, some of which arise in applications. We first consider the following two examples of symmetric norms, which are crucial for the subsequent discussion.

The first important example is the top- $k$ norm: the sum of $k$ largest absolute values of the coordinates of a vector; $k=1$ corresponds to $\ell_{\infty}$, while $k=d$ corresponds to $\ell_{1}$. Another rich set of examples is that of Orlicz norms: for any non-zero convex function $G: \mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$such that $G(0)=0$, we define the unit ball of a norm $\|\cdot\|_{G}$ to be:

$$
\left\{x \in \mathbb{R}^{d} \mid \sum_{i=1}^{d} G\left(\left|x_{i}\right|\right) \leq 1\right\}
$$

Clearly, for $1 \leq p<\infty$ the $\ell_{p}$ norm is Orlicz via $G(t)=t^{p}$.

In statistics and machine learning, Orlicz norms are known as $M$-estimators (for the case of convex losses) [60]. A specific example is the Huber loss. Even though non-convex losses do not correspond to norms, our algorithm still can handle them (see Section 4.3).

Other examples of symmetric norms used in applications include:

- $k$-support norm [29] used for the sparse regression problem; its unit ball is the convex hull of $\left\{x \mid x\right.$ is $k$-sparse, $\left.\|x\|_{2} \leq 1\right\}$,
- box- $\Theta$ norm [123] (again, used for sparse regression), defined for $0<a<b \leq c$ and $\Theta=\left\{\theta \in[a, b]^{d} \mid\|\theta\|_{1} \leq c\right\}$ as $\|x\|=\min _{\theta \in \Theta}\left(\sum_{i=1}^{d} \frac{x_{i}^{2}}{\theta_{i}}\right)^{1 / 2}$, and its dual;
- K-functional [69] used to show tight tail bounds, defined for $t>0$ as $\|x\|=$ $\min \left\{\left\|x_{1}\right\|_{1}+t \cdot\left\|x_{2}\right\|_{2} \mid x_{1}+x_{2}=x\right\}$,
- $\|\cdot\|_{1,2, s}$ norms [105] used for dimension reduction, defined as $\|x\|=\left(\sum_{i}\left\|x_{S_{i}}\right\|_{1}^{2}\right)^{1 / 2}$, where $S_{1}$ is the set of $s$ largest absolute values of coordinates of $x, S_{2}$ is the set of next $s$ largest coordinates, etc.

Finally, we show two simple ways to construct many interesting examples of symmetric norms. Let $0=a_{0} \leq a_{1} \leq a_{2} \leq \ldots \leq a_{d}$ be a non-decreasing sub-additive ${ }^{2}$ sequence. We can define two norms associated with it [40]: a minimal norm is defined as

$$
\|x\|=\max _{1 \leq k \leq d} a_{k} \cdot(\text { average of the largest } k \text { absolute values of the coordinates of } x),
$$

and a maximal norm is equal to

$$
\|x\|=\sum_{k=1}^{d}\left(a_{k}-a_{k-1}\right) \cdot(k \text {-th largest absolute value of a coordinate of } x) .
$$

The minimal norm is the smallest norm such that for every $k$ one has:

$$
\|(\underbrace{1,1, \ldots, 1}_{k}, 0,0, \ldots, 0)\|=a_{k} .
$$

[^24]Similarly, the maximal norm is the largest such norm. Minimal norms will provide hard examples of symmetric norms that preclude some simple(r) approaches to ANN (see Section 4.8.1). We also note that the dual (with respect to the standard dot product) of any symmetric norm is symmetric as well.

### 4.1.3 Prior work: ANN for norms beyond $\ell_{1}$ and $\ell_{2}$

For norms beyond $\ell_{1}$ and $\ell_{2}$, the cornerstone result in ANN is a data structure for $\ell_{\infty}$ due to Indyk [81]. For every $\varepsilon>0$, the data structure achieves space $n^{1+\varepsilon}$, query time $n^{o(1)}$, and approximation $O_{\varepsilon}(\log \log d)$. This is a doubly-exponential improvement over embeddings of $\ell_{\infty}$ into $\ell_{1} / \ell_{2}$ which require distortion $\Omega(\sqrt{d})$.

It is well-known [176] that any $d$-dimensional normed space embeds into $\ell_{\infty}$ with distortion $(1+\varepsilon)$, which raises the question: can we combine this embedding with the result from [81] to solve ANN for any norm? It turns out that the answer is negative: accommodating a norm of interest may require embedding into a very high-dimensional $\ell_{\infty}$. In the worst case, we need $2^{O_{\varepsilon}(d)}$ dimensions, and this bound is known to be tight [32], even for spaces as simple as $\ell_{2}$. Even though this approach would give a non-trivial approximation of $O\left(\log \log 2^{O(d)}\right)=O(\log d)$, the resulting data structure has query time which is exponential in $d$; thus, this approach is interesting only for the low-dimensional regime $d=o(\log n)$.

The result of [81] has been extended as follows. In [82, 83, 16, 10] it was shown how to build data structures for ANN over arbitrary $\ell_{p}$-products of metrics given that there exists an ANN data structure for every factor. Recall that the $\ell_{p}$-product of metric spaces $M_{1}, M_{2}, \ldots, M_{k}$ is a metric space with the ground set $M_{1} \times M_{2} \times \ldots \times M_{k}$ and the following distance function:

$$
d\left(\left(x_{1}, x_{2}, \ldots, x_{k}\right),\left(y_{1}, y_{2}, \ldots, y_{k}\right)\right)=\left\|\left(d_{M_{1}}\left(x_{1}, y_{1}\right), d_{M_{2}}\left(x_{2}, y_{2}\right), \ldots, d_{M_{k}}\left(x_{k}, y_{k}\right)\right)\right\|_{p}
$$

In a nutshell, if we can build efficient ANN data structures for every $M_{i}$ with approximation $c$, there exist an efficient data structure for ANN over the product space with approximation $O(c \cdot \log \log n)$. Note that the above also implies ANN for the standard
$\ell_{p}$, though for this case a better approximation $O(\log \log d)$ is possible via randomized embeddings into $\ell_{\infty}[10]$. For small values of $p$, one can also get approximation $c=2^{O(p)}[133,36]$ for ANN for the $\ell_{p}$ distance using different techniques.

### 4.1.4 Overview of the proof of Theorem 4.1.3

We prove Theorem 4.1.3 in three steps.

- First, we build a data structure for $d$-dimensional top- $k$ norms. We proceed by constructing a randomized embedding into $d$-dimensional $\ell_{\infty}$ with constant distortion, and then invoke the data structure for ANN over $\ell_{\infty}$ from [81].

Our embedding is a refinement of the technique of max-p-stable distributions used in [10] to embed $\ell_{p}$ into $\ell_{\infty}$. Surprisingly, the technique turns out to be very general, and can handle top- $k$ norms as well an arbitrary Orlicz norm.

While this technique can handle even arbitrary symmetric norms (see Section 4.8), there exist symmetric norms, for which this approach leads to merely a $\log ^{\Omega(1)} d$ approximation, which is exponentially worse than the bound we are aiming at (see Section 4.8.1).

- To bypass the above limitation and obtain the desired $(\log \log n)^{O(1)}$-approximation, we show the following structural result: any $d$-dimensional symmetric norm allows a constant-distortion (deterministic) embedding into a low-dimensional iterated product of top- $k$ norms. More specifically, the host space $Y$ is an $\ell_{\infty}$-product of $d^{O(1)}$ copies of the $\ell_{1}$-product of $X_{1}, X_{2}, \ldots, X_{d}$, where $X_{k}$ is $\mathbb{R}^{d}$ equipped with the top- $k$ norm. The dimension of $Y$ is $d^{O(1)}$ which is significantly better than the bound $2^{\Omega(d)}$ necessary to embed symmetric norms (even $\ell_{2}$ ) into $\ell_{\infty}$. It is exactly this improvement over the naïve approach that allows us to handle any dimension $d=n^{o(1)}$ as opposed to the trivial $o(\log n)$.
- Finally, we use known results $[82,10]$, which allow us to construct a data structure for ANN over a product space if we have ANN data structures for the individual factors. Each such step incurs an additional $\log \log n$ factor in the resulting approximation.

Since we have built a data structure for top- $k$ norms, and can embed a symmetric norm into an iterated product of top- $k$ norms, we are done!

Embeddings into iterated product spaces have been successfully used before for constructing data structures for ANN over Fréchet distance [82], edit distance [83], and Ulam distance [16]. Theorem 4.1.3 gives yet another confirmation of the power of the technique.

### 4.1.5 Optimality of Theorem 4.1.3

There remains one aspects of Theorem 4.1.3 that can potentially be improved: the approximation factor $(\log \log n)^{O(1)}$.

One bottleneck for our algorithm is the ANN data structure for $\ell_{\infty}$ from [81], which gives $O(\log \log d)$ approximation. This bound is known to be tight [11, 98] for certain models of computation (in particular, for decision trees, which captures the result of [81]). Thus, going beyond approximation $\Omega(\log \log d)$ in Theorem 4.1.3 might be hard; however, it remains entirely possible to improve the approximation from $(\log \log n)^{O(1)}$ to $O(\log \log d)$, which we leave as an open question.

### 4.1.6 Lower bounds for general norms

The second step of the proof of Theorem 4.1.3 (see Section 4.1.4) shows how to embed any $d$-dimensional symmetric norm into a universal normed space of dimension $d^{O(1)}$ with a constant distortion. In contrast, we show that for general norms a similar universal construction is impossible. More formally, for a fixed $0<\varepsilon<1 / 3$, suppose $U$ is a normed space such that for every $d$-dimensional normed space $X$ there exists a randomized linear embedding of $X$ into $U$ with distortion $O\left(d^{1 / 2-\varepsilon}\right)$. Then, $U$ must have dimension at least $\exp \left(d^{\Omega_{\varepsilon}(1)}\right)$. By John's theorem [92], $d$-dimensional $\ell_{2}$ is a universal space for distortion $\sqrt{d}$, so our lower bound is tight up to sub-polynomial factors. See Section 4.6 for details.

To take this a step further, it would be highly desirable to prove stronger hardness results for ANN over general norms. One approach would be to show that such a
norm $X$ has high robust expansion, which is a property used to deduce ANN lower bounds [146, 24]. There exist metrics $M$ that have high robust expansion, such as the shortest path metric of a spectral expander (see Section 4.9). To obtain a hard norm, it suffices to embed such an $N$-point metric $M$ into a $\log ^{O(1)} N$-dimensional norm with a constant distortion. The result of [120] shows that there exist $N$-point metrics $M$ which cannot be embedded into any norm of dimension $N^{o(1)}$. However, these metrics are not expanders, and for expanders such a dimension reduction procedure might be possible. ${ }^{3}$

### 4.1.7 Other related work: dealing with general norms

The recent result of [42] completely characterizes the streaming complexity of any symmetric norm. Even though many symmetric norms (including $\ell_{\infty}$ ) are hard in the streaming model, the state of affairs with ANN is arguably much nicer. In particular, our results imply that all symmetric norms have highly efficient ANN data structures. We also point out that streaming algorithms for the special case of Orlicz norms have been studied earlier [45].

Another related work is [23], which shows that for norms, the existence of good sketches is equivalent to uniform embeddability into $\ell_{2}$. Sketches are known to imply efficient ANN data structures, but since many symmetric norms do not embed into $\ell_{2}$ uniformly, we conclude that ANN is provably easier than sketching for a large class of norms.

Finally, we also mention the work of [2], who study ANN under the class of highdimensional distances which are Bregman divergences. These results are somewhat disjoint since the Bregman divergences are not norms.

[^25]
### 4.2 Preliminaries

### 4.2.1 Norms and products

For any subset $A \subseteq \mathbb{R}$, we let $\chi_{A}: \mathbb{R} \rightarrow\{0,1\}$ be the indicator function of $A$. For a vector $x \in \mathbb{R}^{d}$ we define $|x|=\left(\left|x_{1}\right|,\left|x_{2}\right|, \ldots,\left|x_{d}\right|\right)$ to be the vector of the absolute values of the coordinates of $x$.

Definition 4.2.1. For any vector $x \in \mathbb{R}^{d}$, we let $x^{*}=P|x|$ be the vector obtained by applying the permutation matrix $P$ to $|x|$ so coordinates of $x^{*}$ are sorted in non-increasing absolute value.

Definition 4.2.2 (Symmetric norm). A norm $\|\cdot\|_{X}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is symmetric if for every $x \in \mathbb{R}^{d},\|x\|_{X}=\||x|\|_{X}=\left\|x^{*}\right\|_{X}$.

See the introduction for examples of symmetric norms. We note once again that the dual norm of a symmetric norm is also symmetric.

A natural way to combine norms is via product spaces, which we will heavily exploit in this paper.

Definition 4.2.3 (Product space). Let $1 \leq p \leq \infty$. Let $\left(X_{1}, d_{X_{1}}\right),\left(X_{2}, d_{X_{2}}\right), \ldots$, $\left(X_{k}, d_{X_{k}}\right)$ be metric spaces. We define the $\ell_{p}$-product space, denoted $\bigoplus_{\ell_{p}} X_{i}$, to be a metric space whose ground set is $X_{1} \times X_{2} \times \ldots \times X_{k}$, and the distance function is defined as follows: the distance between $\left(x_{1}, x_{2}, \ldots, x_{k}\right)$ and $\left(x_{1}^{\prime}, x_{2}^{\prime}, \ldots, x_{k}^{\prime}\right)$ is defined as the $\ell_{p}$ norm of the vector $\left(d_{X_{1}}\left(x_{1}, x_{1}^{\prime}\right), d_{X_{2}}\left(x_{1}, x_{2}^{\prime}\right), \ldots, d_{X_{k}}\left(x_{k}, x_{k}^{\prime}\right)\right)$.

Next we define the top- $k$ norm:
Definition 4.2.4. For any $k \in[d]$, the top- $k$ norm, $\|\cdot\|_{T(k)}: \mathbb{R}^{d} \rightarrow \mathbb{R}$, is the sum of the absolute values of the top $k$ coordinates. In other words,

$$
\|x\|_{T(k)}=\sum_{i=1}^{k}\left|x_{i}^{*}\right|
$$

where $x^{*}$ is the vector obtained in Definition 4.2.1.

Definition 4.2.5. Given vectors $x, y \in \mathbb{R}^{d}$, we say $x$ weakly majorizes $y$ if for all $k \in[d]$,

$$
\sum_{i=1}^{k}\left|x_{i}^{*}\right| \geq \sum_{i=1}^{k}\left|y_{i}^{*}\right| .
$$

Lemma 4.2.6 (Theorem B. 2 in [119]). If $x, y \in \mathbb{R}^{d}$ where $x$ weakly majorizes $y$, then for any symmetric norm $\|\cdot\|_{X}$,

$$
\|x\|_{X} \geq\|y\|_{X}
$$

Definition 4.2.7. For $i \in[d]$, let $\xi^{(i)} \in \mathbb{R}^{d}$ be the vector

$$
\xi^{(i)}=(\underbrace{1, \ldots, 1}_{i}, \underbrace{0, \ldots, 0}_{d-i})
$$

consisting of exactly $i$ 1's, and $d-i 0$ 's.

### 4.2.2 ANN for $\ell_{\infty}$ and $\ell_{\infty}$-products

We will crucially use the following two powerful results of Indyk. The first result is for the standard $d$-dimensional $\ell_{\infty}$ space.

Theorem 4.2.8 ([81, Theorem 1]). For any $\varepsilon \in(0,1 / 2)$, there exists a data structure for ANN for n-points datasets in the $\ell_{\infty}^{d}$ space with approximation $O\left(\frac{\log \log d}{\varepsilon}\right)$, space $O\left(d \cdot n^{1+\varepsilon}\right)$, and query time $O(d \cdot \log n)$.

The second is a generalization of the above theorem, which applies to an $\ell_{\infty^{-}}$ product of $k$ metrics $X_{1}, \ldots X_{k}$, and achieves approximation $O(\log \log n)$. It only needs black-box ANN schemes for each metric $X_{i}$.

Theorem 4.2.9 ([82, Theorem 1]). Let $X_{1}, X_{2}, \ldots, X_{k}$ be metric space, and let $c>1$ be a real number. Suppose that for every $1 \leq i \leq k$ and every $n$ there exists a data structure for $A N N$ for $n$-point datasets from $X_{i}$ with approximation $c$, space $S(n) \geq n$, query time $Q(n)$, and probability of success 0.99 . Then, for every $\varepsilon>0$, there exists ANN under $\bigoplus_{\ell_{\infty}}^{k} \mathcal{M}$ with:

- $O\left(\varepsilon^{-1} \log \log n\right)$ approximation,
- $O(Q(n) \log n+d k \log n)$ query time, where $d$ is the time to compute distances in each $X_{i}$, and
- $S(n) \cdot O\left(k n^{\varepsilon}\right)$ space/preprocessing.

Strictly speaking, we need to impose a technical condition on the ANN for each $X_{i}$ - that it reports the point with the smallest priority - which is satisfied in all our scenarios; see [82, Section 2] for details. Also, the original statement of [82] gave a somewhat worse space bound. The better space results simply from a better analysis of the algorithm, as was observed in [16]; we include a proof in Section 4.7.

### 4.3 ANN for Orlicz and top- $k$ norms

Before showing a data structure for general symmetric norms, we give an algorithm for general Orlicz norms. We then show how to apply these ideas to top- $k$ norms. This restricted setting has a simple analysis and illustrates one of the main techniques used in the rest of the paper. A similar approach was used in prior work to construct randomized embeddings of $\ell_{p}$ norms into $\ell_{\infty}$, and solve the ANN search problem; here we show that these techniques are in fact applicable in much greater generality.

Lemma 4.3.1. Let $\|\cdot\|_{G}$ be an Orlicz norm. For every $D, \alpha>1$ and every $\mu \in(0,1 / 2)$ there exists a randomized linear map $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d}$ such that for every $x \in \mathbb{R}^{d}$ :

- if $\|x\|_{G} \leq 1$, then $\operatorname{Pr}_{f}\left[\|f(x)\|_{\infty} \leq 1\right] \geq \mu$;
- if $\|x\|_{G}>\alpha D$, then $\operatorname{Pr}_{f}\left[\|f(x)\|_{\infty}>D\right] \geq 1-\mu^{\alpha}$.

Proof. Let the distribution $\mathcal{D}$ over $\mathbb{R}_{+}$have the following CDF $F: \mathbb{R}_{+} \rightarrow[0,1]$ :

$$
F(t)=\operatorname{Pr}_{u \sim \mathcal{D}}[u \leq t]=1-\mu^{G(t)}
$$

Consider the following randomized linear map $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d}$ :

$$
\left(x_{1}, x_{2}, \ldots, x_{d}\right) \stackrel{f}{\mapsto}\left(\frac{x_{1}}{u_{1}}, \frac{x_{2}}{u_{2}}, \ldots, \frac{x_{d}}{u_{d}}\right)
$$

where $u_{1}, \ldots, u_{d} \sim \mathcal{D}$ are i.i.d. samples from $\mathcal{D}$. Suppose that $\|x\|_{G} \leq 1$. Then, $\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right) \leq 1$. This, in turn, implies:

$$
\operatorname{Pr}_{f}\left[\|f(x)\|_{\infty} \leq 1\right]=\prod_{i=1}^{d} \operatorname{Pr}_{u_{i} \sim \mathcal{D}}\left[\left|\frac{x_{i}}{u_{i}}\right| \leq 1\right]=\prod_{i=1}^{d} \mu^{G\left(\left|x_{i}\right|\right)}=\mu^{\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right)} \geq \mu .
$$

Now suppose that $\|x\|_{G}>\alpha D$. This, together with the convexity of $G(\cdot)$, implies:

$$
\sum_{i=1}^{d} G\left(\frac{\left|x_{i}\right|}{D}\right) \geq(1-\alpha) G(0)+\alpha \cdot \sum_{i=1}^{d} G\left(\frac{\left|x_{i}\right|}{\alpha D}\right) \geq \alpha
$$

Thus, we have:

$$
\operatorname{Pr}_{f}\left[\|f(x)\|_{\infty} \leq D\right]=\prod_{i=1}^{d} \operatorname{Pr}_{u_{i} \sim \mathcal{D}}\left[\left|\frac{x_{i}}{u_{i}}\right| \leq D\right]=\prod_{i=1}^{d} \mu^{G\left(\left|x_{i}\right| / D\right)}=\mu^{\sum_{i=1}^{d} G\left(\left|x_{i}\right| / D\right)} \leq \mu^{\alpha}
$$

Theorem 4.3.2. For every d-dimensional Orlicz norm $\|\cdot\|_{G}$ and every $\varepsilon \in(0,1 / 2)$, there exists a data structure for ANN over $\|\cdot\|_{G}$, which achieves approximation $O\left(\frac{\log \log d}{\varepsilon^{2}}\right)$ using space $O\left(d n^{1+\varepsilon}\right)$ and query time $O\left(d n^{\varepsilon}\right)$.

Proof. Let $P \subset \mathbb{R}^{d}$ be a dataset of $n$ points. Consider the data structure which does the following:

1. For all $1 \leq i \leq n^{\varepsilon}$, we independently apply the randomized linear map $f$ from Lemma 4.3.1 with parameters $\mu=n^{-\varepsilon}, D=O\left(\frac{\log \log d}{\varepsilon}\right)$, and $\alpha=\frac{2}{\varepsilon}$. We define

$$
P_{i}=\left\{f_{i}(x) \mid x \in P\right\}
$$

to be the image of the dataset under $f_{i}$, where $f_{i}$ is the $i$-th independent copy of $f$.
2. For each $1 \leq i \leq n^{\varepsilon}$, we use Theorem 4.2 .8 to build a data structure for ANN over $\ell_{\infty}$ with approximation $D$ for dataset $P_{i}$. We refer to the $i$-th data structure as $T_{i}$.

Each $T_{i}$ occupies space $O\left(d n^{1+\varepsilon}\right)$ and achieves approximation $D$ with query time $O(d \log n)$. To answer a query $q \in \mathbb{R}^{d}$, we query $T_{i}$ with $f_{i}(q)$ for each $i \in\left[n^{\varepsilon}\right]$. Let $x_{i}$ be the point returned by $T_{i}$, and let $p_{i} \in P$ be the pre-image of $x_{i}$ under $f_{i}$, so that $f_{i}\left(p_{i}\right)=x_{i}$. If for some $T_{i}$, the point returned satisfies $\left\|p_{i}-q\right\|_{G} \leq \alpha D$, then we return $p_{i}$.

- If there exists some $p \in P$ with $\|p-q\|_{G} \leq 1$, then by Lemma 4.3.1, with probability $1-\left(1-n^{-\varepsilon}\right)^{n^{\varepsilon}} \geq \frac{3}{5}$, some $f_{i}$ has $\left\|f_{i}(p-q)\right\|_{\infty} \leq 1$. Since $f_{i}$ is linear, $\left\|f_{i}(p)-f_{i}(q)\right\|_{\infty} \leq 1$ as well.
- Let $i \in\left[n^{\varepsilon}\right]$ be an index where some $p \in P$ with $\|p-q\|_{G} \leq 1$ has $\| f_{i}(p)-$ $f_{i}(q) \|_{\infty} \leq 1$. Every other $p^{\prime} \in P$ with $\left\|p^{\prime}-q\right\|_{G} \geq \alpha D$ satisfies

$$
\operatorname{Pr}\left[\left\|f_{i}\left(p^{\prime}\right)-f_{i}(q)\right\|_{\infty} \leq D\right] \leq \frac{1}{n^{2}}
$$

A union bound over at most $n$ points with distance greater than $\alpha D$ to $q$ shows that except with probability at most $\frac{1}{n}, T_{i}$ returns some $p_{i} \in P$ with $\left\|p_{i}-q\right\|_{G} \leq \alpha D$. Thus, the total probability of success of the data structure is at least $\frac{3}{5}-\frac{1}{n}$.

The total query time is $O\left(d n^{\varepsilon} \cdot \log n\right)$ and the total space used is $O\left(d n^{1+2 \varepsilon}\right)$. This data structure achieves approximation $\alpha D=O\left(\frac{\log \log d}{\varepsilon^{2}}\right)$. Decreasing $\varepsilon$ by a constant factor, we get the desired guarantees.

Remark. The construction of the randomized embedding in Lemma 4.3.1 and the data structure from Theorem 4.3.2 work in a somewhat more general setting, rather than just for Orlicz norms. For a fixed norm $\|\cdot\|$, we can build a randomized map $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d}$ with the guarantees of Lemma 4.3 .1 if there exists a non-decreasing $G: \mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$where $G(0)=0, G(t) \rightarrow \infty$ as $t \rightarrow \infty$, and for every $x \in \mathbb{R}^{d}:$

- if $\|x\| \leq 1$, then $\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right) \leq 1$, and
- if $\|x\| \geq \alpha D$, then $\sum_{i=1}^{d} G\left(\frac{\left|x_{i}\right|}{D}\right) \geq \alpha$.

The data structure itself just requires the existence of a randomized linear map satisfying the conditions of Lemma 4.3.1.

We now describe how to obtain a data structure for ANN for any top- $k$ norm.

Lemma 4.3.3. Fix any $k \in[d]$. For every $D, \alpha>1$ and every $\mu \in(0,1 / 2)$, there exists a randomized linear map $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d}$ such that for every $x \in \mathbb{R}^{d}$ :

- if $\|x\|_{T(k)} \leq 1$, then $\operatorname{Pr}_{f}\left[\|f(x)\|_{\infty} \leq 1\right] \geq \mu$;
- if $\|x\|_{T(k)}>\alpha D$, then $\operatorname{Pr}_{f}\left[\|f(x)\|_{\infty}>D\right] \geq 1-\mu^{\alpha-1}$.

Proof. We define $G: \mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$where for every $x \in \mathbb{R}^{d}$,

$$
G(t)=t \cdot \chi_{\left[\frac{1}{k}, \infty\right)}(t)
$$

If $\|x\|_{T(k)} \leq 1$, there are at most $k$ coordinates where $\left|x_{i}\right| \geq \frac{1}{k}$. Therefore,

$$
\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right) \leq\|x\|_{T(k)} \leq 1
$$

If $\|x\|_{T(k)} \geq \alpha D$, then $\sum_{i=1}^{k}\left|x_{i}^{*}\right| \geq \alpha D$. Therefore, $\sum_{i=1}^{d} G\left(\frac{\left|x_{i}^{*}\right|}{D}\right) \geq \sum_{i=1}^{k} G\left(\frac{\left|x_{i}^{*}\right|}{D}\right) \geq$ $\alpha-1$. The proof now follows in the same way as Lemma 4.3.1.

Lemma 4.3.3 gives us a data structure for any top- $k$ norm with approximation $O(\log \log d)$ applying Theorem 4.3.2.

One could imagine using a similar argument to design an algorithm for general symmetric norms. This idea indeed works and yields an algorithm with approximation $\widetilde{O}(\log d)$ for a general symmetric norm (see Section 4.8 for a detailed analysis of this approach). However, we show this strategy cannot achieve an approximation better than $\Omega(\sqrt{\log d})$ (see the end of the same Section 4.8).

### 4.4 Embedding into product spaces

In this section, we construct an embedding of general symmetric norms into product spaces of top- $k$ norms. To state the main result of this section, we need the following definition.

Definition 4.4.1. For any $c_{1}, \ldots, c_{d} \geq 0$, let $\bigoplus_{\ell_{1}}^{d} T^{(c)} \subset \mathbb{R}^{d^{2}}$ be the space given by the seminorm $\|\cdot\|_{T, 1}^{(c)} \cdot \mathbb{R}^{d^{2}} \rightarrow \mathbb{R}$ where for $x=\left(x_{1}, \ldots, x_{d}\right) \in \mathbb{R}^{d^{2}}$ and $x_{1}, \ldots, x_{d} \in \mathbb{R}^{d}$ :

$$
\|x\|_{T, 1}^{(c)}=\sum_{k=1}^{d} c_{k}\left\|x_{k}\right\|_{T(k)} .
$$

Theorem 4.4.2 (Embedding into a product space). For any constant $\delta \in(0,1 / 2)$, any symmetric norm $\|\cdot\|_{X}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ can be embedded linearly with distortion $1+\delta$ into $\bigoplus_{\ell_{\infty}}^{t} \bigoplus_{\ell_{1}}^{d} T^{(c)}$ where $t=d^{O\left(\log (1 / \delta) \delta^{-1}\right)}$. In particular, there exists $c \in \mathbb{R}_{+}^{t \times d}$ such that for every $x \in \mathbb{R}^{d}$,

$$
\begin{equation*}
(1-\delta)\|x\|_{X} \leq \max _{i \in[t]}\left(\sum_{k=1}^{d} c_{i, k}\|x\|_{T(k)}\right) \leq(1+\delta)\|x\|_{X} \tag{4.1}
\end{equation*}
$$

The vectors in $\bigoplus_{\ell_{\infty}}^{t} \bigoplus_{\ell_{1}}^{d} T^{(c)} \subset \mathbb{R}^{t d^{2}}$ can be broken up into $t d$ blocks of $d$ coordinates each. The embedding referenced above will simply map $x \in \mathbb{R}^{d}$ into $\mathbb{R}^{t d^{2}}$ by making each of the $t d$ many blocks equal to a copy of $x$. The non-trivial part of the above theorem is setting the constants $c_{i, k}$ for $i \in[t]$ and $k \in[d]$ so (4.1) holds. Before going on to give the proof of Theorem 4.4.2, we establish some definitions and propositions which will be used in the proof. For the subsequent sections, let $\beta \in(1,2)$ be considered a constant close to 1 .

Definition 4.4.3 (Levels and Level Vectors). For any fixed vector $x \in \mathbb{R}^{d}$ and any $k \in \mathbb{Z}$, we define level $k$ with respect to $x$ as $B_{k}(x)=\left\{i \in[d]\left|\beta^{-k-1}<\left|x_{i}\right| \leq \beta^{-k}\right\}\right.$. Additionally, we let $b_{k}(x)=\left|B_{k}(x)\right|$ be the size of level $k$ with respect to $x$. The level
vector of $x, V(x) \in \mathbb{R}^{d}$ is given by

$$
V(x)=(\underbrace{\beta^{k}, \ldots, \beta^{k}}_{b_{-k}(x) \text { times }}, \underbrace{\beta^{k-1}, \ldots, \beta^{k-1}}_{b_{-k+1}(x) \text { times }}, \ldots, \underbrace{\beta^{-k}, \ldots, \beta^{-k}}_{b_{k}(x) \text { times }}, 0, \ldots 0)
$$

where $k$ is some integer such that all non-zero coordinates lie in some level between $-k$ and $k$. We say the $i$-th level vector $V_{i}(x) \in \mathbb{R}^{d}$ is given by

$$
V_{i}(x)=(\underbrace{\beta^{-i}, \ldots, \beta^{-i}}_{b_{i}(x) \text { times }}, 0, \ldots, 0)
$$

The notation used for level vectors appears in [42]; however, we refer to level $k$ as the coordinates of $x$ lying in $\left(\beta^{-k-1}, \beta^{-k}\right]$; whereas [42] refers to level $k$ as the coordinates of $x$ lying in $\left[\beta^{k-1}, \beta^{k}\right)$.

Definition 4.4.4. Fix some $\tau>0$. For any vector $x \in \mathbb{R}^{d}$, let $C(x) \in \mathbb{R}^{d}$ be the vector where each $i \in[d]$ sets

$$
C(x)_{i}=\left\{\begin{array}{cc}
x_{i} & \left|x_{i}\right| \geq \tau \\
0 & \left|x_{i}\right|<\tau
\end{array}\right.
$$

Proposition 4.4.5 (Proposition 3.4 in [42]). Let $\|\cdot\|_{X}$ be any symmetric norm and $x \in \mathbb{R}^{d}$ be any vector. Then

$$
\frac{1}{\beta}\|V(x)\|_{X} \leq\|x\|_{X} \leq\|V(x)\|_{X}
$$

Proposition 4.4.6. Let $\|\cdot\|_{X}$ be any symmetric norm. For any vector $x \in \mathbb{R}^{d}$,

$$
\|x\|_{X}-\tau d \leq\|C(x)\|_{X} \leq\|x\|_{X} .
$$

Proof. Note that $x$ weakly majorizes $C(x)$, so $\|C(x)\|_{X} \leq\|x\|_{X}$. For the other direction, let $v=x-C(x)$. Then $v$ has all coordinates with absolute value at most $\tau$,
so $\tau d \xi^{(1)}$ weakly majorizes $v$. Therefore,

$$
\|x\|_{X} \leq\|C(x)\|_{X}+\|v\|_{X} \leq\|C(x)\|_{X}+\tau d .
$$

Intuitively, the above two propositions say that up to multiplicative loss $\beta$ and additive loss $\tau d$ in the norm of the vector, we may assume that all coordinates are exactly $\beta^{j}$ for $j \geq \log _{\beta}(\tau)$. Thus, if $x \in \mathbb{R}^{d}$, then

$$
\|x\|_{X}-\tau d \leq\|V(C(x))\|_{X} \leq \beta\|x\|_{X}
$$

If additionally, we let $\tau=\frac{\beta}{d^{2}}$, so when $\|x\|_{X} \leq 1$ there are at most $2 \log _{\beta} d$ non-empty levels in $V(C(x))$.

Definition 4.4.7 (Rounded counts vector). Fix any level vector $x \in \mathbb{R}^{d}$. The rounded counts vector of $x, W(x) \in \mathbb{R}^{d}$ is given by $y$ where the $y \in \mathbb{R}^{d}$ is constructed using the following procedure:

1: Initialize $y=(0, \ldots, 0) \in \mathbb{R}^{d}$ and $c=d$.
2: for $k=-\infty, \ldots, 2 \log _{\beta}(d)-1$ do
3: $\quad$ if $b_{k}(x) \geq 0$ then
4: $\quad$ Let $j \in \mathbb{Z}_{+}$be the integer where $\beta^{j-1}<b_{k}(x) \leq \beta^{j}$.
5: $\quad$ if $c \geq\left\lfloor\beta^{j}\right\rfloor$ then
6: $\quad$ Set the first $\left\lfloor\beta^{j}\right\rfloor$ zero-coordinates of $y$ with $\beta^{-k}$. Update $c \leftarrow c-\left\lfloor\beta^{-k}\right\rfloor$.
7: Return $y$

Intuitively, $W(x)$ represents the level vector of $x$ where we ignore coordinates smaller than $\frac{\beta}{d^{2}}$, and additionally, we round the counts of coordinates to powers of $\beta$.

Lemma 4.4.8. For every vector $x \in \mathbb{R}^{d}$ and any symmetric norm $\|\cdot\|_{X}$,

$$
\|x\|_{X}-\tau d \leq\|W(V(C(x)))\|_{X} \leq \beta^{2}\|x\|_{X}
$$

Proof. The bound $\|x\|_{X}-\tau d \leq\|W(V(C(x)))\|_{X}$ follows by combining Proposition 4.4.5 and Proposition 4.4.6, as well as the monotonicity of norms. The bound $\|W(V(C(x)))\|_{X} \leq \beta^{2}\|x\|_{X}$ follows from Proposition 4.4.5, Proposition 4.4.6, as well as Lemma 3.5 from [42].

In order to simplify notation, we let $R: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d}$ given by $R(x)=W(V(C(x)))$.
Definition 4.4.9. Let the set $\mathcal{L} \subset \mathbb{R}_{+}^{d}$ be given by

$$
\mathcal{L}=\left\{y \in \mathbb{R}_{+}^{d} \mid y_{1} \geq \ldots y_{d} \geq 0\right\} .
$$

Additionally, for an arbitrary symmetric norm $\|\cdot\|_{X}$ with dual norm $\|\cdot\|_{X^{*}}$, we let the set $\mathcal{R} \subset \mathcal{L}$ be given by

$$
\mathcal{R}=\left\{R(y) \in \mathbb{R}_{+}^{d} \mid y \in \mathcal{L} \cap B_{X^{*}}\right\} .
$$

Definition 4.4.10. Fix a vector $y \in \mathcal{L} \backslash\{0\}$ ( $y$ has non-negative, non-increasing coordinates). Let the maximal seminorm with respect to $y,\|\cdot\|_{y}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ be the seminorm where for every $x \in \mathbb{R}^{d}$,

$$
\|x\|_{y}=\langle | x^{*}|, y\rangle .
$$

We first show there exists some setting of $c \in \mathbb{R}^{d}$ such that we may compute $\|x\|_{y}$ as $\oplus_{\ell_{1}}^{d} T^{(c)}$.

Lemma 4.4.11. For every vector $y \in \mathcal{L} \backslash\{0\}$, there exists $c_{1}, \ldots, c_{d} \geq 0$ where for all $x \in \mathbb{R}^{d}$,

$$
\|x\|_{y}=\|x\|_{T, 1}^{(c)} .
$$

Proof. For $k \in[d]$, we let $c_{k}=y_{k}-y_{k+1}$, where $y_{d+1}=0$.

$$
\langle | x^{*}|, y\rangle=\sum_{i=1}^{d}\left|x_{i}^{*}\right| y_{i}=\sum_{i=1}^{d}\left|x_{i}^{*}\right|\left(\sum_{k=i}^{d} c_{k}\right)=\sum_{k=1}^{d} c_{k}\left(\sum_{i=1}^{k}\left|x_{i}^{*}\right|\right)=\sum_{k=1}^{d} c_{k}\|x\|_{T(k)}
$$

Given Lemma 4.4.11, it suffices to show that for an arbitrary symmetric norm $\|\cdot\|_{X}$, we may compute $\|x\|_{X}$ (with some distortion) as a maximum over many maximal seminorms. In the following lemma, we show that taking the maximum over maximal norms from $\mathcal{R}$ suffices, but gives sub-optimal parameters. We then improve the parameters to prove Theorem 4.4.2.

Lemma 4.4.12. Let $\|\cdot\|_{X}$ be an arbitrary symmetric norm and let $\|\cdot\|_{X^{*}}$ be its dual norm. Then for any $\|x\|_{X} \leq 1$,

$$
\|x\|_{X}-\tau d \leq \max _{y \in \mathcal{R}}\|x\|_{y} \leq \beta^{2}\|x\|_{X}
$$

Proof. Without loss of generality, we rescale the norm so that $\left\|e_{1}\right\|_{X^{*}}=1$, where $e_{1}$ is the first standard basis vector. Consider any $x \in \mathbb{R}^{d}$ with $\|x\|_{X} \leq 1$. Then since $\|\cdot\|_{X}$ is symmetric, we may assume without loss of generality that all coordinates of $x$ are non-negative and in non-increasing order. Thus for each $y \in \mathcal{L} \cap\{0\}$, we have $\|x\|_{y}=\langle x, y\rangle$.

The lower bound simply follows from the fact that $R(z)$, other than coordinates less than $\tau$, is monotonically above $z$, and all coordinates in $x$ are non-negative. More specifically,

$$
\|x\|_{X}=\sup _{z \in \mathcal{R} \cap B_{X^{*}}}\langle x, z\rangle \leq \sup _{z \in \mathcal{L} \cap B_{X^{*}}}\langle x, R(z)\rangle+\tau d=\max _{y \in \mathcal{R}}\langle x, y\rangle+\tau d,
$$

where $\tau d$ comes from the fact that because $\|x\|_{X} \leq 1$, every coordinate of $x$ is at most 1. On the other hand, we have

$$
\max _{y \in \mathcal{R}}\langle x, y\rangle=\beta^{2} \max _{y \in \mathcal{R}}\left\langle x, \frac{y}{\beta^{2}}\right\rangle \leq \beta^{2} \sup _{z \in B_{X^{*}}}\langle x, z\rangle=\beta^{2}\|x\|_{X},
$$

where we used the fact that $\left\|\frac{y}{\beta^{2}}\right\|_{X^{*}} \leq 1$ by Lemma 4.4.8.
Given Lemma 4.4.12, it follows that we may linearly embed $X$ into $\oplus_{\ell_{\infty}}^{t} \oplus_{\ell_{1}}^{d} T^{(c)}$ where $t=|\mathcal{R}|$, with distortion $\frac{\beta^{2}}{1-\tau d} \leq \beta^{3}$ (where we used the fact $\tau=\frac{\beta}{d}$ and that $1+\beta / d \leq \beta$ for a large enough $d$ ). The embedding follows by copying the vector $x$
into the $t$ spaces $\oplus_{\ell_{1}}^{d} T^{(c)}$ corresponding to each vector $y \in \mathcal{R}$ given in Lemma 4.4.11. The one caveat is that this embedding requires $t$ copies of $\oplus_{\ell_{1}}^{d} T^{(c)}$, and $t$ is as large as $\left(\log _{\beta} d+1\right)^{2 \log _{\beta} d}=d^{O(\log \log d)}$. This is because there are at most $2 \log _{\beta} d$ many levels, and each contains has number of coordinates being some value in $\left\{\beta^{i}\right\}_{i=0}^{\log _{\beta} d}$. Thus, our algorithm becomes inefficient once $d \geq 2^{\omega\left(\frac{\log n}{\log \log n}\right)}$.

In order to avoid this problem, we will make the embedding more efficient by showing that we do not need all of $\mathcal{R}$, but rather a fine net of $\mathcal{R}$. In addition, our net will be of polynomial size in the dimension, which gives an efficient algorithm for all $\omega(\log n) \leq d \leq n^{o(1)}$. We first show that it suffices to consider fine nets of $\mathcal{R}$, and then build a fine net of $\mathcal{R}$ of size poly $(d)$.

Lemma 4.4.13. Fix an $\gamma \in(0,1 / 2)$. Let $\|\cdot\|_{X}$ be an arbitrary symmetric norm and $\|\cdot\|_{X^{*}}$ be its dual norm. If $N$ is a $\gamma$-net of $\mathcal{R}$ with respect to distance given by $\|\cdot\|_{X^{*}}$, then

$$
(1-\gamma-\tau d)\|x\|_{X} \leq \max _{y \in N}\|x\|_{y} \leq\left(\beta^{2}+\gamma\right)\|x\|_{X}
$$

Proof. Since the embedding we build is linear, it suffices to show that every vector $x \in \mathbb{R}^{d}$ with $\|x\|_{X}=1$ has

$$
1-\gamma-\tau d \leq \max _{y \in N}\|x\|_{y} \leq\left(\beta^{2}+\gamma\right)
$$

Consider a fixed vector $x \in \mathbb{R}^{d}$ with $\|x\|_{X}=1$. Additionally, we may assume the coordinates of $x$ are non-negative and in non-increasing order. We simply follow the computation:

$$
\begin{aligned}
\|x\|_{X} & =\left\|\left|x^{*}\right|\right\|_{X} \leq \max _{y \in \mathcal{R}}\langle | x^{*}|, y\rangle+\tau d \\
& =\max _{y \in N}\left(\langle | x^{*}|, y\rangle+\langle | x^{*}|, v\rangle\right)+\tau d \leq \max _{y \in N}\langle | x^{*}|, y\rangle+\gamma\|x\|_{X}+\tau d
\end{aligned}
$$

where we used Lemma 4.4.12 and the fact that $\|v\|_{X^{*}} \leq \gamma$ in a $\gamma$-net of $\mathcal{R}$ with respect to the distance given by $\|\cdot\|_{X^{*}}$. On the other hand,

$$
\max _{y \in N}\|x\|_{y}=\max _{y \in \mathcal{R}}\left(\langle | x^{*}|, y\rangle+\langle | x^{*}|, v\rangle\right) \leq \max _{y \in \mathcal{R}}\|x\|_{y}+\gamma\|x\|_{X} \leq\left(\beta^{2}+\gamma\right)\|x\|_{X}
$$

where again, we used Lemma 4.4.12 and the fact that $\|v\|_{X^{*}} \leq \gamma$.

Finally, we conclude the theorem by providing a $\gamma$-net for $\mathcal{R}$ of size $d^{O\left(\log (1 / \gamma) \gamma^{-1}\right)}$.

Lemma 4.4.14. Fix any symmetric space $X$ with dual $X^{*}$. There exists an $8(\beta-1)$ net of size $d^{O(\log (1 /(\beta-1)) / \log \beta)}$ for $\mathcal{R}$ with respect to distances given by $\|\cdot\|_{X^{*}}$.

We defer the proof of Lemma 4.4.14 to the next section. The proof of Theorem 4.4.2 follows by combining Lemma 4.4.11, Lemma 4.4.13, and Lemma 4.4.14. In particular, given a $\frac{\beta-1}{8}$-net of $\mathcal{R}$, we get an embedding with distortion at most $\left(\beta^{2}+8(\beta-\right.$ $1))\left(1+(8(\beta-1)+\tau d)^{2}\right)$ from Lemma 4.4.13. We let $\tau=\frac{\beta}{d^{2}}$ and $\beta=\sqrt{1+\delta / 100}$ to get the desired linear embedding with distortion $1+\delta$. We now proceed to proving Lemma 4.4.14, which gives the desired upper bound on the size of the net.

### 4.4.1 Proof of Lemma 4.4.14: bounding the net size

We now give an upper bound on the size of a fine net of $\mathcal{R}$. We proceed by constructing a further simplification of $R(x)$. Intuitively we show that one can ignore the higher levels if there are fewer coordinates in the higher levels than some lower level.

Lemma 4.4.15. Let $\|\cdot\|_{X}$ be a symmetric norm. Consider any nonnegative vector $x \in \mathbb{R}_{+}^{d}$ as well as two indices $u, v \in[d]$. Let $y \in \mathbb{R}_{+}^{d}$ be the vector with:

$$
y_{k}=\left\{\begin{array}{cc}
x_{k} & k \in[d] \backslash\{u, v\} \\
x_{u}+x_{v} & k=u \\
0 & k=v
\end{array} .\right.
$$

Then $\|y\|_{X} \geq\|x\|_{X}$.

Proof. Consider the vector $z \in B_{X^{*}}$ where $\langle x, z\rangle=\|x\|_{X}$. Now, we let $z^{\prime} \in \mathbb{R}^{d}$ be given by

$$
z_{k}^{\prime}=\left\{\begin{array}{cc}
z_{k} & k \in[d] \backslash\{u, v\} \\
\max \left\{z_{u}, z_{v}\right\} & k=u \\
\min \left\{z_{u}, z_{v}\right\} & k=v
\end{array}\right.
$$

Note that $z^{\prime}$ is a permutation of $z$, so $z^{\prime} \in B_{X^{*}}$. Now,

$$
\left\langle y, z^{\prime}\right\rangle=\left(x_{u}+x_{v}\right) \max \left\{z_{u}, z_{v}\right\}+\sum_{k \in[d] \backslash\{u, v\}} x_{k} z_{k} \geq \sum_{k \in[d]} x_{k} z_{k}=\langle x, z\rangle=\|x\|_{X} .
$$

Definition 4.4.16. Consider a vector $x \in \mathcal{R}$. We define the simplified rounded vector $S(x)$ as the vector returned by the following procedure.

1: Initialize $z=x$
2: for $k=0,1, \ldots, 2 \log _{\beta}(d)-1$ do
3: $\quad$ if $b_{k}(z) \leq \max _{j<k+3 \log _{\beta}(\beta-1)} b_{j}(z)$ then
4: $\quad$ Set all coordinates of $z$ of value $\beta^{-k}$ to 0 i.e. set $b_{k}(z)=0$.
5: Sort the coordinates of $z$ in non-increasing order and return $z$.
Next we show that the simplified rounded vector is close to the rounded counts vector.

Lemma 4.4.17. Let $\|\cdot\|_{X}$ be a symmetric norm and let $x \in \mathcal{R}$. Then $\|S(x)-x\|_{X} \leq$ $2(\beta-1)\|x\|_{X}$.

Proof. Consider some $k \in\left[2 \log _{\beta} d-1\right]$ and let $C_{k}(x) \subset[d]$ be set of coordinates where $x$ is at level $k$ and does not equal $S(x)=z$, i.e.,

$$
C_{k}(x)=\left\{i \in[d] \mid x_{i}=\beta^{-k} \text { and } x_{i} \neq z_{i}\right\} .
$$

Additionally, for each $k \in\left[2 \log _{\beta} d-1\right]$, let $T_{k} \subset[d]$ be the coordinates at level $k$ in $x$ which trigger line 3 of $S(x)$, and thus become 0 's in $z$ (we do not need to consider the case $k=0$ since line 3 never triggers, in fact, we do not need to consider $k \in\left[-3 \log _{\beta}(\beta-1)\right]$ either $)$. In other words,

$$
T_{k}(x)=\left\{i \in[d] \mid x_{i}=\beta^{-k} \text { and at iteration } k \text { of } S(x), b_{k}(z) \leq \max _{j<k+3 \log _{\beta}(\beta-1)} b_{j}(z)\right\}
$$

Note that $T_{1}(x), \ldots, T_{2 \log _{\beta} d-1}(x)$ are all disjoint, and $\left|C_{k}(x)\right| \leq \sum_{j \in[k]}\left|T_{j}(x)\right|$, since whenever we zero out coordinates in levels less than or equal to $k, S(x)$ will shift
the coordinates when we sort, causing $x_{i} \neq z_{i}$. Thus, we may consider an injection $s_{k}: C_{k}(x) \rightarrow \bigcup_{j \in[k]} T_{j}(x)$, which charges coordinates in $C_{k}(x)$ to coordinates which were zeroed out in line 3 of $S(x)$.

Additionally, for each $j \in\left[2 \log _{\beta} d-1\right]$ where $T_{j}(x) \neq \emptyset$, we let $q_{j}$ be the integer between 0 and $j+3 \log _{\beta}(\beta-1)$ which triggered line 3 of $S(x)$ at $k=j$. More specifically, $0 \leq q_{j} \leq j+3 \log _{\beta}(\beta-1)$ is the smallest integer for which $b_{j}(x) \leq b_{q_{j}}(x)$.

Finally, for each $j \in\left[2 \log _{\beta} d-1\right]$ where $T_{j}(x) \neq \emptyset$, we let $g_{j}: T_{j}(x) \rightarrow B_{q_{j}}(x)$ (recall that $B_{q_{j}}(x) \subset[d]$ are the indices of $x$ at level $q_{j}$ ) be an arbitrary injection. Such an injection exists because $b_{j}(x) \leq b_{q_{j}}(x)$. We may consider the mapping $F: \bigcup_{k \in\left[2 \log _{\beta} d-1\right]} C_{k}(x) \rightarrow[d]$ where

$$
F(i)=g_{j}\left(s_{k}(i)\right) \quad \text { where } k \text { and } j \text { are such that } i \in C_{k}(x) \text { and } s_{k}(i) \in T_{j}(x) .
$$

Let $y$ be the vector where we "aggregate" coordinates of $\bigcup_{k \in\left[2 \log _{\beta}(d)-1\right]} C_{k}(x)$ of $x$ according to the map $F$ according to Lemma 4.4.15. In particular, we define $y \in \mathbb{R}^{d}$ where for $i \in[d]$, we let

$$
y_{i}=\sum_{i^{\prime} \in F^{-1}(i)} x_{i^{\prime}}
$$

Note that for each $i \in[d], 0 \leq(x-z)_{i} \leq x_{i}$, and $\bigcup_{k \in\left[2 \log _{\beta}(d)-1\right]} C_{k}(x) \subset[d]$ are the non-zero coordinates of $x-z$. Thus, from Lemma 4.4.15, we conclude that $\|x-z\|_{X} \leq\|y\|_{X}$. We now turn to upper-bounding $\|y\|_{X}$.

Fix some $i \in[d]$ where $x_{i}=\beta^{-j}$. Then

$$
y_{i}=\sum_{k>j-3 \log _{\beta}(\beta-1)}\left(\sum_{k^{\prime} \geq k} x_{s_{k^{\prime}}^{-1}\left(g_{k}^{-1}(i)\right)}\right),
$$

where we interpret $x_{s_{k^{\prime}}^{-1}\left(g_{k}^{-1}(i)\right)}$ as 0 when $g_{k}^{-1}(i)=\emptyset$, or $s_{k^{\prime}}^{-1}\left(g_{k}^{-1}(i)\right)=\emptyset$. Note that
whenever $x_{s_{k^{\prime}}^{-1}\left(g_{k}^{-1}(i)\right)} \neq 0, x_{s_{k^{\prime}}^{-1}\left(g_{k}^{-1}(i)\right)}=\beta^{-k^{\prime}}$. Thus,

$$
\begin{aligned}
y_{i} & \leq \sum_{k>j-3 \log _{\beta}(\beta-1)}\left(\sum_{k^{\prime} \geq k} \beta^{-k^{\prime}}\right) \\
& \leq \sum_{k>j-3 \log _{\beta}(\beta-1)} \frac{\beta^{1-k}}{\beta-1} \leq \frac{\beta^{1-j+3 \log _{\beta}(\beta-1)}}{(\beta-1)^{2}} \leq \beta(\beta-1) \cdot \beta^{-j} .
\end{aligned}
$$

Recall that $x_{i}=\beta^{-j}$, so $\beta(\beta-1) x$ weakly majorizes $y$, and thus

$$
\|x-z\|_{X} \leq\|y\|_{X} \leq \beta(\beta-1) \cdot\|x\|_{X}
$$

Hence, when $\beta \leq 2$, we have $\|x-z\|_{X} \leq 2(\beta-1)\|x\|_{X}$.

Proof of Lemma 4.4.14. We now prove the theorem by showing that the set

$$
N=\left\{S(x) \in \mathbb{R}^{d} \mid x \in \mathcal{R}\right\}
$$

is a $\gamma$-net of $\mathcal{R}$, and we give an upper bound on the size. By Lemma 4.4.17, $\| x-$ $S(x)\left\|_{X} \leq 2(\beta-1)\right\| x \|_{X} \leq 8(\beta-1)$. So it suffices to give an upper bound on the size of $N$.

We bound from above the number of net points by an encoding argument. Let $z=S(x)$ and let

$$
t_{k}=\frac{b_{k}(z)}{\max _{j<k+3 \log _{\beta}(\beta-1)} b_{j}(z)}
$$

Let $k^{*} \in\left\{0, \ldots, 2 \log _{\beta} d-1\right\}$ be the smallest level $k$ with non-zero $b_{k}(z)$. For all $j>k^{*}-3 \log _{\beta}(\beta-1)$, we either have $t_{j}(z)=0$ or $t_{j}(z) \geq 1$. Additionally, $z$ has $d$ coordinates, so

$$
\prod_{j=k^{*}-3 \log _{\beta}(\beta-1)}^{2 \log _{\beta} d-1} \max \left(t_{j}, 1\right) \leq d^{-3 \log _{\beta}(\beta-1)},
$$

since terms of the product "cancel" except for at most $-3 \log _{\beta}(\beta-1)$, which are each at most $d$.

We will encode $z \in N$ in three steps. In the first step, we use $2 \log _{\beta} d-1$ bits
in order to encode whether $b_{k}(z)=0$ or $b_{k}(z)>0$. In the second step, we then encode the values $b_{k^{*}+j}(z)$ for $j \in\left\{0, \ldots, 3 \log _{\beta}(1 /(\beta-1))\right\}$. Finally, we go through $j>k^{*}+3 \log _{\beta}\left(\frac{1}{\beta-1}\right)$, and encode $t_{i}$ using a prefix-free code, where writing $t_{i}$ uses at most $O\left(\log \max \left(t_{i}, 1\right)\right)$ many bits. Thus, in total, the number of bits we use is

$$
\begin{aligned}
& O\left(\log _{\beta} d+\log d \log _{\beta}\left(\frac{1}{\beta-1}\right)+\sum_{j=k^{*}-3 \log _{\beta}(\beta-1)}^{2 \log _{\beta} d-1} \log \max \left(t_{j}, 1\right)\right) \\
& \quad=O\left(\frac{\log d \cdot \log \left(\frac{1}{\beta-1}\right)}{\log \beta}+\log \left(\prod_{j=k^{*}-3 \log _{\beta}(\beta-1)}^{2 \log _{\beta} d-1} \max \left(t_{j}, 1\right)\right)\right) \\
& \quad=O\left(\frac{\log d \cdot \log \left(\frac{1}{\beta-1}\right)}{\log \beta}\right) .
\end{aligned}
$$

Thus, we obtain $N$ is a $8(\beta-1)$-net, and the size of $N$ is $d^{O(\log (1 /(\beta-1)) / \log \beta)}$.

### 4.5 Proof of the main theorem

We now prove our main result, Theorem 4.1.3. The algorithm here achieves approximation

$$
O\left(\frac{\log ^{2} \log n \cdot \log \log d}{\varepsilon^{5}}\right)
$$

We proceed by giving an algorithm for $\bigoplus_{\ell_{\infty}}^{t} \bigoplus_{\ell_{1}}^{d} T^{(c)}$ using Theorem 4.2.8, Theorem 5.1.2 from [10], and Theorem 4.2.9.

Lemma 4.5.1. Fix some $c_{1}, \ldots, c_{d} \geq 0$. Let $\bigoplus_{\ell_{\infty}} T^{(c)}$ be the space with

$$
\|\cdot\|_{T, \infty}^{(c)}: \mathbb{R}^{d^{2}} \rightarrow \mathbb{R}
$$

seminorm where for every $x=\left(x_{1}, \ldots, x_{d}\right) \in \mathbb{R}^{d^{2}}$,

$$
\|x\|_{T, \infty}^{(c)}=\max _{k \in[d]} c_{k}\left\|x_{k}\right\|_{T(k)}
$$

For every $\varepsilon \in(0,1 / 2)$, there exists a data structure for ANN over $\|\cdot\|_{T, \infty}^{(c)}$ which
achieves approximation $O\left(\frac{\log \log n \cdot \log \log d}{\varepsilon^{3}}\right)$ using space $O\left(d^{2} \cdot n^{1+\varepsilon}\right)$ and query time $O\left(d^{2} \cdot n^{\varepsilon}\right)$.

Proof. Given the randomized embedding from Lemma 4.3.3, we can build a data structure for $c_{k}\|\cdot\|_{T(k)}$ achieving approximation $O\left(\frac{\log \log d}{\varepsilon^{2}}\right)$ using space $O\left(d^{2} n^{1+\varepsilon / 2}\right)$ and query time $O\left(d^{2} n^{\varepsilon / 2}\right)$. This data structure works in the same way as in the proof of Theorem 4.3.2. We handle the constant $c_{k}$ by rescaling the norm, and since the embeddings are linear, it does not affect the correctness of the data structure. Then we apply Theorem 4.2.8.

Lemma 4.5.2. Fix some $c_{1}, \ldots, c_{d} \geq 0$. Let $\bigoplus_{\ell_{1}} T^{(c)}$ be the space with $\|\cdot\|_{T, 1}^{(c)}: \mathbb{R}^{d^{2}} \rightarrow$ $\mathbb{R}$ seminorm where $x=\left(x_{1}, \ldots, x_{m}\right) \in \mathbb{R}^{d^{2}}$,

$$
\|x\|_{T, 1}^{(c)}=\sum_{k=1}^{d} c_{k}\left\|x_{k}\right\|_{T(k)} .
$$

For every $\varepsilon \in(0,1 / 2)$, there exists a data structure for ANN over $\|\cdot\|_{T, 1}^{(c)}$ which achieves approximation $O\left(\frac{\log \log n \cdot \log \log d}{\varepsilon^{4}}\right)$ using space $O\left(d^{2} \cdot n^{1+\varepsilon}\right)$ and query time $O\left(d^{2} \cdot n^{\varepsilon}\right)$.

Proof. The proof follows from Theorem 5.1.2 in [10] and Lemma 4.5.1.

Finally, we are combine the above results to get an improved algorithm for general symmetric norms.

Theorem 4.5.3. For every d-dimensional symmetric norm $\|\cdot\|_{X}$ and every $\varepsilon \in$ $(0,1 / 2)$, there exists a data structure for ANN over $\|\cdot\|_{X}$ which achieves approximation $O\left(\frac{\log ^{2} \log n \log \log d}{\varepsilon^{5}}\right)$ using space $d^{O(1)} \cdot O\left(n^{1+\varepsilon}\right)$ and query time $d^{O(1)} \cdot O\left(n^{\varepsilon}\right)$.

Proof. Given Theorem 4.4.2, we embed $\|\cdot\|_{X}$ into $\bigoplus_{\ell_{\infty}} \bigoplus_{\ell_{1}} T^{(c)}$ with approximation $\left(1 \pm \frac{1}{10}\right)$. The result from Lemma 4.5.2 allows we to apply Theorem 4.2.9 to obtain the desired data structure.

Theorem 4.5.3 implies our main result Theorem 4.1.3 stated in the introduction.

### 4.6 Lower bounds

In this section, we show that our techniques do not extend to general norms. In particular, we show there does not exist a universal norm $U$ for which any norm embeds (possibly randomized) with constant distortion, unless the blow-up in dimension is exponential. Hence the result from below applies to cases of $U=\ell_{\infty}$ as well as an (low-dimensional) product spaces.

Theorem 4.6.1. For any $\varepsilon>0$, let $U$ be a $d^{\prime}$-dimensional normed space such that for any d-dimensional normed space $X$, there exists a distribution $\mathcal{D}$ supported on linear embeddings $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d^{\prime}}$ where for every $x \in \mathbb{R}^{d}$,

$$
\|x\|_{X} \leq\|f(x)\|_{U} \leq D\|x\|_{U}
$$

holds with probability at least $\frac{2}{3}$ over the draw of $f \sim \mathcal{D}$, for $D=O\left(d^{1 / 2-\varepsilon}\right)$. Then $d^{\prime}=\exp \left(\Omega\left(d^{2 \varepsilon}\right)\right)$.

We will prove the above theorem by showing that if there exists a universal normed space $U$ satisfying the conditions of Theorem 4.6.1 above, then two parties, call them Alice and Bob, can use the embeddings to solve the communication problem Index with only a few bits. Let $U$ be a proposed $d^{\prime}$-dimensional normed space satisfying the conditions of Theorem 4.6.1. By the John's theorem [32], we may apply a linear transform so that:

$$
B_{\ell_{2}} \subset B_{U} \subset \sqrt{d^{\prime}} B_{\ell_{2}}
$$

Lemma 4.6.2. For any $\varepsilon>0$, there exists a set of $\exp \left(\Omega\left(d^{2 \varepsilon}\right)\right)$ many points on the unit sphere $S^{d-1}$ such that pairwise inner-products are at most $\frac{1}{d^{1 / 2-\varepsilon}}$. In fact, these points may consist of points whose coordinates are $\pm \frac{1}{\sqrt{d}}$.

Proof. Consider picking two random points $x, y \in S^{d-1}$ where each entry is $\pm \frac{1}{\sqrt{d}}$. Then by Bernstein's inequality,

$$
\operatorname{Pr}_{x, y}\left[|\langle x, y\rangle| \geq \frac{1}{d^{1 / 2-\varepsilon}}\right] \leq 2 \exp \left(-\Omega\left(d^{2 \varepsilon}\right)\right)
$$

We may pick $\exp \left(\Omega\left(d^{2 \varepsilon}\right)\right)$ random points and union bound over the probability that some pair has large inner product.

Fix $\varepsilon>0$ and $C=d^{1 / 2-\varepsilon}$, and let $P$ be set a set of unit vectors with pairwise inner-product at most $\frac{1}{C}$ of size $\exp \left(\Omega\left(d^{2 \varepsilon}\right)\right)$. For each $a \in\{0,1\}^{P}$ consider the following norm:

$$
\|x\|_{a}=C \cdot \max _{y \in P: a_{y}=1}|\langle x, y\rangle| .
$$

Assume there exists a randomized linear embedding $f: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d^{\prime}}$ with the following guarantees:

- For every $x \in \mathbb{R}^{d},\|x\|_{a} \leq\|f(x)\|_{U} \leq D\|x\|_{a}$ with probability at least $\frac{2}{3}$.

Note the embedding $f$ can be described by $M$, a $d^{\prime} \times d$ matrix of real numbers. Additionally, we consider rounding each entry of $M$ by to the nearest integer multiple of $\frac{1}{\text { poly(d) }}$ to obtain $M^{\prime}$. For each $x \in S^{d-1},\left\|\left(M-M^{\prime}\right) x\right\|_{U} \leq\left\|\left(M-M^{\prime}\right) x\right\|_{2} \leq \frac{1}{\operatorname{poly}(d)}$. Thus, we may assume each entry of $M$ is an integer multiple of $\frac{1}{\operatorname{poly}(d)}$, and lose $\left(1 \pm \frac{1}{\operatorname{poly}(d)}\right)$ factor in the distortion of the embedding for vectors in $B_{2}$.

We now show that the existence of the randomized embedding implies a one-way randomized protocol for the communication problem Index. We first describe the problem. In an instance of Index:

- Alice receives a string $a \in\{0,1\}^{n}$.
- Bob receives an index $i \in[n]$.
- Alice communicates with Bob so that he can output $a_{i}$.

Theorem 4.6.3 ([107]). The randomized one-way communication complexity of INDEX is $\Omega(n)$.

We give a protocol for Index:

1. Suppose Alice has input $a \in\{0,1\}^{P}$. She will generate the norm $\|\cdot\|_{a}$ described above. Note that $f \sim \mathcal{D}$ has that for each $x \in \mathbb{R}^{d}$, the embedding preserves the norm of $x$ up to $D$ with probability $\frac{2}{3}$. In particular, if Bob's input is $i \in|P|$,
corresponding to point $y$, then an embedding $f \in \mathcal{D}$, which we represent as a $d^{\prime} \times d$ matrix $M$, satisfies:

$$
\|y\|_{a} \leq\|M y\|_{U} \leq D\|y\|_{a}
$$

with probability $\frac{2}{3}$. In particular, with probability $\frac{2}{3}$ :

- If $a_{i}=0$, then $\|y\|_{a} \leq 1$, which implies $\|M y\|_{U} \leq D$.
- If $a_{i}=1$, then $\|y\|_{a} \geq C$, which implies $\|M y\|_{U} \geq C$.

Alice computes the set $P_{c} \subset P$ of vectors which satisfy the above property (i.e. the embedding $M$ preserves increases the norm by at most a factor $D$ ).
2. Alice finds a subset $B \subset P_{c}$ of linearly independent vectors such that every $x \in P_{c}$ we have $x \in \operatorname{span}(B)$. Note that $|B| \leq d$ and for all $x \in B,\|M x\|_{2} \leq$ $\sqrt{d^{\prime}}\|M x\|_{U} \leq C \cdot D \cdot \sqrt{d^{\prime}}$. Therefore, each $M x \in \mathbb{R}^{d^{\prime}}$ can be written with $\tilde{O}\left(d^{\prime}\right)$ bits. So Alice sends the set $B$, as well as $M x$ for each $x \in B$ using $\tilde{O}\left(d d^{\prime}\right)$ bits.
3. In order for Bob to decode $a_{i}$, he first checks whether $y \in \operatorname{span}(B)$, and if not, he guesses. If $y \in \operatorname{span}(B)$, which happens with probability $\frac{2}{3}$, then Bob writes

$$
y=\sum_{b_{i} \in B} c_{i} b_{i}
$$

and $M y=\sum_{b_{i} \in B} c_{i} M b_{i}$. If $\|M y\|_{U} \leq D$, then $a_{i}=0$ and if $\|M y\|_{U} \geq C$ then $a_{i}=1$. Thus, if $D<\frac{C}{2}$, Bob can recover $a_{i}$ with probability $\frac{2}{3}$.

Alice communicates $\tilde{O}\left(d d^{\prime}\right)$ bits, and Bob is able to recover $a_{i}$ with probability $\frac{2}{3}$. By Theorem 4.6.3, $d d^{\prime} \geq \tilde{\Omega}(|P|)$, which in turn implies $d^{\prime} \geq \exp \left(\Omega\left(d^{2 \varepsilon}\right)\right)$.

### 4.7 Appendix: Bounding space in Theorem 4.2.9

Here we justify the space bound of the algorithm from Theorem 4.2.9 (from [82]). We note that the improved bound was also claimed in [16], albeit without a proof.

First of all, as suggested at the end of Section 3 of [82], one modifies the algorithm to obtain space of the form of $n^{1+\varepsilon}$, at the expense of increasing the approximation to $O\left(\varepsilon^{-1} \log \log n\right)$. This is done by replacing the conditions in Case 2 and 3 by respectively:

$$
\left[\frac{\left|B(s, R(s)+c+1) \cap S_{\mid i}\right|}{\left|S_{\mid i}\right|}\right]^{1+\varepsilon}<\frac{\left|B(s, R(s)) \cap S_{\mid i}\right|}{\left|S_{\mid i}\right|}
$$

and

$$
\left[\frac{\left|S_{\mid i}-B\left(p, R^{\prime}\right)\right|}{|S|}\right]^{1+\varepsilon}<\frac{\left|S_{\mid i}-B\left(s, R^{\prime}+2\right)\right|}{|S|}
$$

With the remaining algorithm being precisely the same, our only task here is to argue the space bound. First of all we bound the sum of the number of points stored in all the leaves. For a tree with $m$ nodes, let $L(m)$ be an upper bound on this count. We would like to prove that $L(m) \leq m^{1+\varepsilon}$. As in [82], we only need to focus on cases 2 and 3 of the construction, as case 1 does not replicate the points. We will consider the case 2 (case 3 is exactly similar).

Let $m_{j}=\left|S_{j}\right|$ and $m_{j}^{\prime}=\left|S_{i} \cap \cup_{s \in S_{j}} B(s, c+1)\right|$, whereas $|S|=m$. By construction, we have that $\sum m_{j}=m$ and $m_{j} / m>\left(m_{j}^{\prime} / m\right)^{1+\varepsilon}$ for all $j$.

By induction, assume $L\left(m_{j}^{\prime}\right) \leq\left(m_{j}^{\prime}\right)^{1+\varepsilon}$ for all children. Then, we have that:

$$
L(m) \leq \sum_{j} L\left(m_{j}^{\prime}\right) \leq \sum_{j}\left(m_{j}^{\prime}\right)^{1+\varepsilon}<m^{\varepsilon} \sum_{j} m_{j}=m^{1+\varepsilon}
$$

We now argue the total space is $O\left(S(n) \cdot k \log n \cdot n^{\epsilon}\right)$. Since the depth of the tree is $O(k \log n)$, we have that the total number of points stored in the ANN data structures is $O(k \log n \cdot C(n))=O\left(k \log n \cdot n^{1+\epsilon}\right)$. Since each ANN is on at most $n$ points, we have that, for each occurrence of a point in the ANN data structure, we have an additional factor of $S(n) / n .{ }^{4}$ Hence the total space occupied by all the ANN data structures is $O\left(S(n) / n \cdot k \log n \cdot n^{1+\epsilon}\right)$. Using a smaller $\varepsilon$ (to hide the $\log n$ factor), we obtain the stated space bound of $O\left(S(n) \cdot k \cdot n^{\epsilon}\right)$.

[^26]
### 4.8 Appendix: $\widetilde{O}(\log d)$-ANN for symmetric norms

We provide a simple ANN algorithm for general symmetric norms with approximation $O(\log d \log \log d)$ using near-linear space and sub-linear query time. The algorithm will leverage the results in the previous section by relating general symmetric norms to Orlicz norms. Recall the definition of level vectors in Definition 4.4.3.

Definition 4.8.1. Let $\|\cdot\|_{X}$ be any symmetric norm. Let $L_{k}>0$ be the minimum number of coordinates needed at level $k$ to have norm at least 1 . In other words,

$$
L_{k}=\min \left\{j \in[d] \mid\left\|\beta^{-i} \xi^{(j)}\right\|_{X}>1\right\} .
$$

At a high level, we will relate the norm of a vector $x \in \mathbb{R}^{d}$ to the norm of its level vectors $V_{k}(x)$. The definition above gives a way to measure the contribution of level $k$ to the norm. For example, if $x \in \mathbb{R}^{d}$ has norm $\|x\|_{X} \geq D$, and there are only $2 \log _{\beta} d$ non-zero levels with respect to $x$, then some level vector $\left\|V_{k}(x)\right\|_{X} \geq \frac{D}{2 \log _{\beta} d}$. This implies $b_{k}=\Omega\left(\frac{D L_{k}}{\log _{\beta} d}\right)$, since we may divide $V_{k}(x)$ into a sum of vectors with $L_{k}$ coordinates at level $k$.

On the other hand, if $x \in \mathbb{R}^{d}$ has $\|x\|_{X} \leq 1$, then $b_{k}<L_{k}$ for each $k$. Since we consider only $2 \log _{\beta} d$ relevant levels, for $\|x\|_{S} \leq 1$,

$$
\sum_{k=0}^{2 \log _{\beta} d-1} \frac{b_{k}}{L_{k}} \leq 2 \log _{\beta} d
$$

Additionally, $\sum_{k=0}^{2 \log _{\beta} d-1}\left(b_{k} / L_{k}\right)$ can be decomposed as an additive contribution of coordinates. In particular, coordinate $x_{i}$ contributes $1 / L_{k}$ if $i \in B_{k}$. Therefore, we can hope to approximate the symmetric norm by an Orlicz norms and apply the arguments from Lemma 4.3.1.

The lemma below formalizes the ideas discussed above.
Lemma 4.8.2. Let $\|\cdot\|_{X}$ be any symmetric norm. For any $D, \alpha>1$, there exists a non-decreasing function $G: \mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$with $G(0)=0$ and $G(t) \rightarrow \infty$ as $t \rightarrow \infty$, where every vector $x \in \mathbb{R}^{d}$ satisfies the following:

- If $\|x\|_{X} \leq 1$, then $\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right) \leq 2 \log _{\beta} d$.
- If $\|x\|_{X}>\alpha D \cdot 7 \log _{\beta} d$, then $\sum_{i=1}^{d} G\left(\frac{\left|x_{i}\right|}{D}\right) \geq \alpha \cdot 2 \log _{\beta} d$.

Proof. For $i \geq 0$, let $A_{i}=\left(\beta^{-i-1}, \beta^{-i}\right]$. The function $G: \mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$is defined as

$$
\begin{equation*}
G(t)=\sum_{i=0}^{2 \log _{\beta} d-1} \frac{\chi_{A_{i}}(t)}{L_{i}}+\alpha \cdot 2 \log _{\beta} d \cdot t \cdot \chi_{(1, \infty)}(t) \tag{4.2}
\end{equation*}
$$

Note that $G(0)=0$ and $G(t) \rightarrow \infty$ as $t \rightarrow \infty$.
Recall the norm satisfies, $\left\|\xi^{(1)}\right\|_{X}=1$, so if $\|x\|_{X} \leq 1$, then $\left|x_{i}\right| \leq 1$ for all $i \in[d]$. This means $\chi_{(1, \infty)}\left(\left|x_{i}\right|\right)=0$ so the second term of the RHS of (4.2) is zero. Therefore,

$$
\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right)=\sum_{i=1}^{d} \sum_{k=0}^{2 \log _{\beta} d-1} \frac{\chi_{A_{k}}\left(\left|x_{i}\right|\right)}{L_{k}}=\sum_{k=0}^{2 \log _{\beta} d-1} \frac{b_{k}}{L_{k}}
$$

where $b_{k}$ is defined with respect to $x$. Since, $b_{k}<L_{k}$ for all $0 \leq k<2 \log _{\beta} d$,

$$
\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right) \leq 2 \log _{\beta} d
$$

If $x \in \mathbb{R}^{d}$ where $\|x\|_{X}>\alpha D \cdot 7 \log _{\beta} d$, then the vector $\left\|\frac{x}{D}\right\|_{X}>\alpha \cdot 7 \log _{\beta} d$. So it suffices to prove that for any vector $x \in \mathbb{R}^{d}$ with $\|x\|_{X}>\alpha \cdot 7 \log _{\beta} d$,

$$
\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right) \geq \alpha \cdot 2 \log _{\beta} d
$$

Additionally, for any vector $x \in \mathbb{R}^{d}$, we may consider the vector $C(x) \in \mathbb{R}^{d}$ for $\tau=\frac{\beta}{d^{2}}$ from Definition 4.4.4. By Proposition 4.4.6, $\|C(x)\|_{X} \geq\|x\|_{X}-\frac{\beta}{d}>\alpha \cdot 6 \log _{\beta} d$. Therefore, we may assume $x \in \mathbb{R}^{d}$ has $\|x\|_{X}>\alpha \cdot 6 \log _{\beta} d$, and that all non-zero coordinates have absolute values greater than $\frac{\beta}{d^{2}}$. Equivalently, $b_{k}=0$ for all $k \geq$ $2 \log _{\beta} d$. If for some $i \in[d],\left|x_{i}\right| \geq 1$, then the second term in the RHS of (4.2) is non-zero, and $G\left(\left|x_{i}\right|\right) \geq \alpha \cdot 2 \log _{\beta} d$. So we may further assume all coordinates of $x$ lie
in levels $k=0, \ldots, 2 \log _{\beta} d-1$. Note that

$$
\sum_{i=1}^{d} G\left(\left|x_{i}\right|\right)=\sum_{k=0}^{2 \log _{\beta} d-1} \sum_{i=1}^{d} G\left(\left|V_{k}(x)_{i}\right|\right),
$$

and for each $0 \leq k<2 \log _{\beta} d, \sum_{i=1}^{d} G\left(\left|V_{k}(x)_{i}\right|\right)=\frac{b_{k}}{L_{k}}$.
We partition the levels into two groups,

$$
A=\left\{k \left\lvert\, \frac{b_{k}}{L_{k}}<1\right.\right\} \quad \text { and } \quad B=\left\{k \left\lvert\, \frac{b_{k}}{L_{k}} \geq 1\right.\right\} .
$$

For all $k \in B$,

$$
\left\|V_{k}(x)\right\|_{X} \leq\left\lceil\frac{b_{k}}{L_{k}}\right\rceil \leq \frac{2 b_{k}}{L_{k}}
$$

since by the triangle inequality, we can break $V_{k}(x)$ into at most $\left\lceil\frac{b_{k}}{L_{k}}\right\rceil$ vectors with $L_{k}$ coordinates at level $k$ each having norm at least 1 .

Suppose for the sake of contradiction that $\sum_{k \in B} \frac{b_{k}}{L_{k}} \leq \alpha \cdot 2 \log _{\beta} d$. Then

$$
\alpha \cdot 4 \log _{\beta} d \geq \sum_{k \in B} \frac{2 b_{k}}{L_{k}} \geq \sum_{k \in B}\left\|V_{k}(x)\right\|_{X} .
$$

Additionally, since $\|x\|_{X}>\alpha \cdot 6 \log _{\beta} d$, and

$$
\alpha \cdot 6 \log _{\beta} d<\|x\|_{X} \leq \sum_{k \in A}\left\|V_{k}(x)\right\|_{X}+\sum_{k \in B}\left\|V_{k}(x)\right\|_{X},
$$

it follows that

$$
\sum_{k \in A}\left\|V_{k}(x)\right\|_{X}>\alpha \cdot 2 \log _{\beta} d .
$$

However, this is a contradiction for since $|A| \leq 2 \log _{\beta} d$ and $\left\|V_{k}(x)\right\|_{X} \leq 1$.

Lemma 4.8.3. For any $\varepsilon \in(0,1 / 2)$, there exists a data structure for $A N N$ over any symmetric norm $\|\cdot\|_{X}$ which achieves approximation $O\left(\frac{\log d \log \log d}{\varepsilon^{2}}\right)$ using space $O\left(d n^{1+\varepsilon}\right)$ and query time $O\left(d n^{\varepsilon}\right)$.

Proof. We fix $\beta=\frac{3}{2}$. The proof of this lemma follows in the same way as the proof
of Theorem 4.3.2. The one difference is that we rescale the $\ell_{\infty}$ norm by $\frac{1}{2 \log _{\beta} d}$ after applying the embedding.

### 4.8.1 The $\log ^{\Omega(1)} d$-approximation is necessary

Let us remark that we cannot push the technique much further. Namely, any $G(\cdot)$ (even non-convex) requires approximation $\Omega(\sqrt{\log d})$ for the following norm. Define the norm of a vector to be

$$
\|x\|=\max _{1 \leq k \leq d}\left(\frac{x_{1}^{*}+x_{2}^{*}+\ldots x_{k}^{*}}{\sqrt{k}}\right) .
$$

This is the minimal norm for $a_{k}=\sqrt{k}$ (see Section 4.1.2 for the definition). It is not hard to check that an approximation with any $G(\cdot)$ ends up having a distortion $\Omega(\sqrt{\log d})$.

The idea is to consider the following vectors: for every $1 \leq k \leq d$, we consider a vector

$$
(\underbrace{1,1, \ldots, 1}_{k}, 0,0, \ldots, 0)
$$

and besides, we consider a vector

$$
(1, \sqrt{2}-1, \sqrt{3}-\sqrt{2}, \ldots, \sqrt{d}-\sqrt{d-1})
$$

The remaining calculation is a simple exercise.

### 4.9 Appendix: Lower bound for arbitrary metrics via expander graphs

We give an example of a metric that is hard for current approaches to ANN search. The lower bound is based on the notion of robust expansion, which implies all known lower bounds for ANN [146, 24]. In what follows, we will refer to $d=\log N$ as the dimension of a finite metric space of size $N$.

Our example of a hard metric will be the shortest path metric on any spectral expander graph. We note that a similar theorem to the one below is also known for a finite subset of the high-dimensional Earth-Mover Distance [98].

Fix $M$ to be the metric induced by the shortest path distance on a 3-regular expander $G$ on $N$ nodes. In particular, assume that $1-\lambda(G)>c$, where $c$ is an absolute constant, and $\lambda(G) \in(0,1)$ is the second-largest eigenvalue of the normalized adjacency matrix of $G$. Let $d$ be the dimension $d=\log N$.

Theorem 4.9.1. For any approximation $\alpha>1$, and data set size $n \geq 1$ with $d^{\Omega(1)} \leq n \leq N^{O(1)}$, any $\alpha-A N N$ data structure on $n$ points which makes $t$ cell probes (with cells of size at most $\left.w \leq(d \log n)^{O(1)}\right)$, and has success probability at least $\gamma>n^{-1+o(1)}$, must use space $m=\gamma^{\Omega(1 / t)} N^{\Omega(1 /(\alpha t))}=\gamma^{\Omega(1 / t)} 2^{\Omega(d /(\alpha t))}$.

We proceed by introducing a few definitions from [146], and then prove lower bounds on the robust expansion.

Definition 4.9.2 ([146]). In the Graphical Neighbor Search problem (GNS), we are given a bipartite graph $H=(U, V, E)$ where the dataset comes from $U$ and the queries come from $V$. The dataset consists of pairs $P=\left\{\left(p_{i}, x_{i}\right) \mid p_{i} \in U, x_{i} \in\{0,1\}, i \in[n]\right\}$. On query $q \in V$, if there exists a unique $p_{i}$ with $\left(p_{i}, q\right) \in E$, then we want to return $x_{i}$.

One can use the GNS problem to prove lower bounds on $c$-ANN as follows: build a GNS graph $H$ by taking $U=V=[N]$, and connecting two points $u \in U, v \in V$ iff they are at a distance at most $r$ (see details in [146]). We will also need to make sure that in our instances $q$ is not closer than $c r$ to other points except the near neighbor.

We now introduce the notion of robust expansion, used in [146] to prove lower bounds.

Definition 4.9 .3 (Robust Expansion [146]). For a GNS graph $H=(U, V, E)$, fix a distribution $e$ on $E \subset U \times V$, and let $\mu$ be the marginal on $U$ and $\eta$ be the marginal on $V$. For $\delta, \gamma \in(0,1]$, the robust expansion $\Phi_{r}(\delta, \gamma)$ is defined as follows:

$$
\Phi_{r}(\delta, \gamma)=\min _{A \subset V: \eta(A) \leq \delta} \min _{B \subset U: \frac{e(A X B)}{e(A \times V)} \geq \gamma} \frac{\mu(B)}{\eta(A)}
$$

We now prove a lower bound on the robust expansion $\Phi_{r}(\delta, \gamma)$ for a GNS graph arising from the shortest path metric on the expander graph $G$. Fix $r=d / \alpha$. The hard distribution $e$ is defined as follows: pick $p$ at random from $M$ and obtain $q$ by running a random walk of length $r$ starting at $p$. Note that for $n<N^{1 / 4}$ and sufficiently high constant $\alpha$, the distribution satisfies the weak-independence condition required for applying the results in [146].

Fix any sets $A, B \subset M$, where $a=|A| / N$ and $b=|B| / N$. By the expander mixing lemma applied to $G^{r}$, we obtain that:

$$
\left|E_{G^{r}}(A, B)-\frac{|A| \cdot|B|}{3^{r} N}\right| \leq \lambda^{3} 3^{r} \sqrt{|A| \cdot|B|} .
$$

Considering that $\operatorname{Pr}[q \in B \mid p \in A]=\frac{E_{G^{r}(A, B)}^{a N \cdot 3^{r}}}{}$, we have that:

$$
\operatorname{Pr}[q \in B \mid p \in A] \leq b+\lambda^{r} \sqrt{b / a}
$$

Restricting to sets $A, B$ such that $\operatorname{Pr}[q \in B \mid p \in A] \geq \gamma$, for which we must have that $\Phi_{r}=\Phi_{r}(a, \gamma) \geq b / a$ (by definition), we conclude:

$$
\gamma \leq \Phi_{r} \cdot a+\lambda^{r} \sqrt{\Phi_{r}}
$$

Hence, either $\Phi_{r}=\Omega(\gamma / a)$ or $\Phi_{r}=\Omega\left(\gamma^{2} / \lambda^{2 r}\right)$.
Proof of Theorem 4.9.1. Applying Theorem 1.5 from [146], we have that, for $t \geq 1$ cell probes, either:

- $m^{t} w / n \geq \Omega\left(\gamma \cdot m^{t}\right)$, an impossibility;
- or $m^{t} w / n \geq \Omega\left(\gamma^{2} / \lambda^{2 r}\right)$, or $m^{t}=\Omega\left(\frac{n}{w} \gamma^{2} / \lambda^{2 r}\right)$, implying $m=\gamma^{2 / t} N^{\Omega(1 /(\alpha t))}$.

To show how bad the situation for expander metrics is, we state a lower bound on for $\alpha$-ANN on the expander metric described above in the list-of-points model, which captures the hashing-based algorithms of [24] and in the decision tree model
of [81]. The proofs follow from a simple derivation using the robust expansion lower bounds in Section 7 of [24] and a reduction of decision trees to $O(\log m)$-cell-probe data structures similar to Appendix A in [11].

Theorem 4.9.4. Any list-of-points data structure for ( $c, r$ )-ANN for random instances of $n$ points in the expander metric of dimension $d$ (described above) with query time $t$ and space $m$ has either $t=\Omega(n)$, or $m=\exp (\Omega(d))$.

Theorem 4.9.5. Let $d=\Omega\left(\log ^{1+\varepsilon} n\right)$ for some $\varepsilon>0$. Any decision tree of size $m$ and depth $t$ and word size $w$ succeeding with probability $\gamma$ satisfies:

$$
\frac{m^{O(\log m)} t w}{n} \geq \Phi_{r}\left(\frac{1}{m^{O(\log m)}}, \frac{\gamma}{O(\log m)}\right)
$$

In particular, for any $\rho>0$, if $w \leq n^{\rho}$, either

$$
t \geq \widetilde{\Omega}\left(n^{1-\rho}\right)
$$

or

$$
m=\exp \left(\Omega\left(d^{\varepsilon /(1+\varepsilon)}\right)\right) \operatorname{poly}(n)
$$

## Chapter 5

## Sketching and embedding are equivalent for norms

### 5.1 Introduction

One of the most exciting notions in the modern algorithm design is that of sketching, where an input is summarized into a small data structure. Perhaps the most prominent use of sketching is to estimate distances between points, one of the workhorses of similarity search. For example, some early uses of sketches have been designed for detecting duplicates and estimating resemblance between documents [47, 48, 55]. Another example is Nearest Neighbor Search, where many algorithms rely heavily on sketches, under the labels of dimension reduction (like the Johnson-Lindenstrauss Lemma [93]) or Locality-Sensitive Hashing (see e.g. [85, 108, 14]). Sketches see widespread use in streaming algorithms, for instance when the input implicitly defines a high-dimensional vector (via say frequencies of items in the stream), and a sketch is used to estimate the vector's $\ell_{p}$ norm. The situation is similar in compressive sensing, where acquisition of a signal can be viewed as sketching. Sketching - especially of distances such as $\ell_{p}$ norms-was even used to achieve improvements for classical computational tasks: see e.g. recent progress on numerical linear algebra algorithms [178], or dynamic graph algorithms [5, 99]. Since sketching is a crucial primitive that can lead to many algorithmic advances, it is important to understand its power and
limitations.
A primary use of sketches is for distance estimation between points in a metric space $\left(X, d_{X}\right)$, such as the Hamming space. The basic setup here asks to design a sketching function $\mathbf{s k}: X \rightarrow\{0,1\}^{s}$, so that the distance $d_{X}(x, y)$ can be estimated given only the sketches $\mathbf{s k}(x), \mathbf{s k}(y)$. In the decision version of this problem, the goal is to determine whether the inputs $x$ and $y$ are "close" or "far", as formalized by the Distance Threshold Estimation Problem [157], denoted $\operatorname{DTEP}_{r}(X, D)$, where, for a threshold $r>0$ and approximation $D \geq 1$ given as parameters in advance, the goal is to decide whether $d_{X}(x, y) \leq r$ or $d_{X}(x, y)>D r$. Throughout, it will be convenient to omit $r$ from the subscript. ${ }^{1}$ Efficient sketches sk almost always need to be randomized, and hence we allow randomization, requiring (say) $90 \%$ success probability.

The diversity of applications gives rise to a variety of natural and important metrics $M$ for which we want to solve DTEP: Hamming space, Euclidean space, other $\ell_{p}$ norms, the Earth Mover's Distance, edit distance, and so forth. Sketches for Hamming and Euclidean distances are now classic and well-understood [85, 108]. In particular, both are "efficiently sketchable": one can achieve approximation $D=O(1)$ using sketch size $s=O(1)$ (most importantly, independent of the dimension of $X$ ). Indyk [84] extended these results to efficient sketches for every $\ell_{p}$ norm for $p \in(0,2]$. In contrast, for $\ell_{p}$-spaces with $p>2$, efficient sketching (constant $D$ and $s$ ) was proved impossible using information-theoretic arguments [157, 34]. Extensive subsequent work investigated sketching of other important metric spaces, ${ }^{2}$ or refined bounds (like a trade-off between $D$ and $s$ ) for "known" spaces. ${ }^{3}$

These efforts provided beautiful results and techniques for many specific settings. Seeking a broader perspective, a foundational question has emerged [124, Question \#5]:

Question 5.1.1. Characterize metric spaces which admit efficient sketching.

[^27]To focus the question, efficient sketching will mean constant $D$ and $s$ for us. Since its formulation circa 2006, progress on this question has been limited. The only known characterization is by [76] for distances that are decomposable by coordinates, i.e., $d_{X}(x, y)=\sum_{i=1}^{n} \varphi\left(x_{i}, y_{i}\right)$ for some $\varphi$. In particular, they show a number of general conditions on $\varphi$ which imply an $\Omega(n)$ sketching complexity for $d_{X}$.

### 5.1.1 The embedding approach

To address DTEP in various metric spaces more systematically, researchers have undertaken the approach of metric embeddings. A metric embedding of $X$ is a map $f: X \rightarrow Y$ into another metric space $\left(Y, d_{Y}\right)$. The distortion of $f$ is the smallest $D^{\prime} \geq 1$ for which there exists a scaling factor $t>0$ such that

$$
\forall x, y \in X, \quad d_{Y}(f(x), f(y)) \leq t \cdot d_{X}(x, y) \leq D^{\prime} \cdot d_{Y}(f(x), f(y))
$$

If the target metric $Y$ admits sketching with parameters $D$ and $s$, then $X$ admits sketching with parameters $D \cdot D^{\prime}$ and $s$, by the simple composition $\mathbf{s k}^{\prime}: x \mapsto \mathbf{s k}(f(x))$. This approach of "reducing" sketching to embedding has been very successful, including for variants of the Earth Mover's Distance [55, 86, 117, 134, 15], and for variants of edit distance $[37,141,56,16,63,131,62,61]$. The approach is obviously most useful when $Y$ itself is efficiently sketchable, which holds for all $Y=\ell_{p}, p \in(0,2]$ [84] (we note that $\ell_{p}$ for $0<p<1$ is not a metric space, but rather a quasimetric space; the above definitions of embedding and distortion make sense even when $Y$ is a quasimetric, and we will use this extended definition liberally). In fact, the embeddings mentioned above are all into $\ell_{1}$, except for [16] which employs a more complicated target space. We remark that in many cases the distortion $D^{\prime}$ achieved in the current literature is not constant and depends on the "dimension" of $X$.

Extensive research on embeddability into $\ell_{1}$ has resulted in several important distortion lower bounds. Some address the aforementioned metrics [102, 134, 106, 21], while others deal with metric spaces arising in rather different contexts such as Functional Analysis [149, 57, 58], or Approximation Algorithms [114, 30, 104, 103].

Nevertheless, obtaining (optimal) distortion bounds for $\ell_{1}$-embeddability of several metric spaces of interest, are still well-known open questions [122].

Yet sketching is a more general notion, and one may hope to achieve better approximation by bypassing embeddings into $\ell_{1}$. As mentioned above, some limited success in bypassing an $\ell_{1}$-embedding has been obtained for a variant of edit distance [16], albeit with a sketch size depending mildly on the dimension of $X$. Our results disparage these hopes, at least for the case of normed spaces.

### 5.1.2 Our results

Our main contribution is to show that efficient sketchability of norms is equivalent to embeddability into $\ell_{1-\varepsilon}$ with constant distortion. Below we only assert the "sketching $\Longrightarrow$ embedding" direction, as the reverse direction follows from [84], as discussed above.

Theorem 5.1.2. Let $X$ be a finite-dimensional normed space, and suppose that $0<$ $\varepsilon<1 / 3$. If $X$ admits a sketching algorithm for $\operatorname{DTEP}(X, D)$ for approximation $D>1$ with sketch size $s$, then $X$ linearly embeds into $\ell_{1-\varepsilon}$ with distortion $D^{\prime}=O(s D / \varepsilon)$.

One can ask whether it is possible to improve Theorem 5.1 .2 by showing that $X$, in fact, embeds into $\ell_{1}$. Since many non-embeddability theorems are proved for $\ell_{1}$, such a statement would "upgrade" such results to lower bounds for sketches. Indeed, we show results in this direction too. First of all, the above theorem also yields the following statement.

Theorem 5.1.3. Under the conditions of Theorem 5.1.2, $X$ linearly embeds into $\ell_{1}$ with distortion $O(s D \cdot \log (\operatorname{dim} X))$.

Ideally, we would like an even stronger statement: efficient sketchability for norms is equivalent to embeddability into $\ell_{1}$ with constant distortion (i.e., independent of the dimension of $X$ as above). Such a stronger statement in fact requires the resolution of an open problem posed by Kwapien in 1969 (see [95, 41]). To be precise, Kwapien asks whether every finite-dimensional normed space $X$ that embeds into $\ell_{1-\varepsilon}$ for $0<\varepsilon<1$
with distortion $D_{0} \geq 1$ must also embed into $\ell_{1}$ with distortion $D_{1}$ that depends only on $D_{0}$ and $\varepsilon$ but not on the dimension of $X$ (this is a reformulation of the finitedimensional version of the original Kwapien's question). In fact, by Theorem 5.1.2, the "efficient sketching $\Longrightarrow$ embedding into $\ell_{1}$ with constant distortion" statement is equivalent to a positive resolution of the Kwapien's problem. Indeed, for the other direction, consider a potential counter-example to the Kwapien's problem, i.e., a normed space $X$ that embeds into $\ell_{1-\varepsilon}$ with a constant distortion $D_{0} \geq 1$, but every embedding of $X$ into $\ell_{1}$ incurs a distortion $D_{1}=\omega(1)$, where the asymptotics is with the dimension of $X$ (it is really a sequence of normed spaces). Hence, $X$ admits an efficient sketch obtained by combining the embedding into $\ell_{1-\varepsilon}$ with the sketch of [84], but does not embed into $\ell_{1}$ with constant distortion. Thus, if the answer to Kwapien's question is negative, then our desired stronger statement is false.

To bypass the resolution of the Kwapien's problem, we prove the following variant of the theorem using a result of Kalton [95]: efficient sketchability is equivalent to $\ell_{1}$-embeddability with constant distortion for norms that are "closed" under sumproducts. A sum-product of two normed spaces $X$ and $Y$, denoted $X \oplus_{\ell_{1}} Y$, is the normed space $X \times Y$ endowed with $\|(x, y)\|=\|x\|+\|y\|$. It is easy to verify that $\ell_{1}$, the Earth Mover's Distance, and the trace norm are all closed under taking sumproducts (potentially with an increase in the dimension). Again, we only need to show the "sketching $\Longrightarrow$ embedding" direction, as the reverse direction follows from the arguments above - if a normed space $X$ embed into $\ell_{1}$ with constant distortion, we can combine it with the $\ell_{1}$ sketch of [84] and obtain an efficient sketch for $X$. We discuss the application of this theorem to the Earth Mover's Distance in Section 5.1.3.

Theorem 5.1.4. Let $\left(X_{n}\right)_{n=1}^{\infty}$ be a sequence of finite-dimensional normed spaces. Suppose that for every $i_{1}, i_{2} \geq 1$ there exists $m=m\left(i_{1}, i_{2}\right) \geq 1$ such that $X_{i_{1}} \oplus \ell_{1} X_{i_{2}}$ embeds isometrically into $X_{m}$. Assume that every $X_{n}$ admits a sketching algorithm for $\operatorname{DTEP}\left(X_{n}, D\right)$ for fixed approximation $D>1$ with fixed sketch size s (both independent of $n$ ). Then, every $X_{n}$ linearly embeds into $\ell_{1}$ with bounded distortion (independent of $n)$.

Overall, we almost completely characterize the norms that are efficiently sketchable, thereby making a significant progress on Question 5.1.1. In particular, our results suggest that the embedding approach (embed into $\ell_{p}$ for some $p \in(0,2]$, and use the sketch from [84]) is essentially unavoidable for norms. It is interesting to note that for general metrics (not norms) the implication "efficient sketching $\Longrightarrow$ embedding into $\ell_{1}$ with constant distortion" is false: for example the Heisenberg group embeds into $\ell_{2}$-squared (with bounded distortion) and hence is efficiently sketchable, but it is not embeddable into $\ell_{1}[110,57,58]$ (another example of this sort is provided by Khot and Vishnoi [104]). At the same time, we are not aware of any counter-example to the generalization of Theorem 5.1.2 to general metrics.

### 5.1.3 Applications

To demonstrate the applicability of our results to concrete questions of interest, we consider two well-known families of normed spaces, for which we obtain the first non-trivial lower bounds on the sketching complexity.

Trace norm. Let $\mathcal{T}_{n}$ be the vector space $\mathbb{R}^{n \times n}$ (all real square $n \times n$ matrices) equipped with the trace norm (also known as the nuclear norm and Schatten 1-norm), which is defined to be the sum of singular values. It is well-known that $\mathcal{T}_{n}$ embeds into $\ell_{2}$ (and thus also into $\ell_{1}$ ) with distortion $\sqrt{n}$ (observe that the trace norm is within $\sqrt{n}$ from the Frobenius norm, which embeds isometrically into $\ell_{2}$ ). Pisier [149] proved a matching lower bound of $\Omega(\sqrt{n})$ for the distortion of any embedding of $\mathcal{T}_{n}$ into $\ell_{1}$.

This non-embeddability result, combined with our Theorem 5.1.3, implies a sketching lower bound for the trace norm. Before, only lower bounds for specific types of sketches (linear and bilinear) were known [111].

Corollary 5.1.5. For any sketching algorithm for $\operatorname{DTEP}\left(\mathcal{T}_{n}, D\right)$ with sketch size $s$ the following bound must hold:

$$
s D=\Omega\left(\frac{\sqrt{n}}{\log n}\right) .
$$

Earth Mover's Distance. The (planar) Earth Mover's Distance (also known as the transportation distance, Wasserstein-1 distance, and Monge-Kantorovich distance) is the vector space $\mathrm{EMD}_{n}=\left\{p \in \mathbb{R}^{[n]^{2}}: \sum_{i} p_{i}=0\right\}$ endowed with the norm $\|p\|_{\text {EMD }}$ defined as the minimum cost needed to transport the "positive part" of $p$ to the "negative part" of $p$, where the transportation cost per unit between two points in the grid $[n]^{2}$ is their $\ell_{1}$-distance (for a formal definition see [134]). It is known that this norm embeds into $\ell_{1}$ with distortion $O(\log n)$ [86, 55, 134], and that any $\ell_{1}$-embedding requires distortion $\Omega(\sqrt{\log n})$ [134].

We obtain the first sketching lower bound for $\mathrm{EMD}_{n}$, which in particular addresses a well-known open question [124, Question \#7]. Its proof is a direct application of Theorem 5.1.4 (which we can apply, since $\mathrm{EMD}_{n}$ is obviously closed under taking sum-products), to essentially "upgrade" the known non-embeddability into $\ell_{1}[134]$ to non-sketchability.

Corollary 5.1.6. No sketching algorithm for $\operatorname{DTEP}\left(\operatorname{EMD}_{n}, D\right)$ can achieve approximation $D$ and sketch size $s$ that are constant (independent of $n$ ).

The reason we can not apply Theorem 5.1.3 and get a clean quantitative lower bound for sketches of $\mathrm{EMD}_{n}$ is the factor $\log (\operatorname{dim} X)$ in Theorem 5.1.3. Indeed, the lower bound on the distortion of an embedding of $\mathrm{EMD}_{n}$ into $\ell_{1}$ proved in [134] is $\Omega(\sqrt{\log n})$, which is smaller than $\log (\operatorname{dim} X)=\Theta(\log n)$.

We note that $\mathrm{EMD}_{n}$ is a (slight) generalization of the EMD metric version commonly used in computer science applications. In the latter, given two weighted sets $A, B \subset[n]^{2}$ of the same total weight, one has to solve, using only their sketches $\mathbf{s k}(A), \mathbf{s k}(B)$, the DTEP (EMD,$D)$ problem where the EMD distance is the min-cost matching between $A$ and $B$. Observe that the weights used in the sets $A, B \subset[n]^{2}$ are all positive. The slight difference is that in $\operatorname{DTEP}\left(\mathrm{EMD}_{n}, D\right)$, which asks analogously to estimate $\|p-q\|_{\text {EMD }}$, each of $p, q \in \mathbb{R}^{[n]^{2}}$ has both "positive" and "negative" parts. Nevertheless, we show in Section 5.7 that efficient sketching of EMD on weighted sets implies efficient sketching of the $\mathrm{EMD}_{n}$ norm. Hence, the non-sketchability of $\mathrm{EMD}_{n}$ norm applies to EMD on weighted sets as well.

### 5.1.4 Other related work

Another direction for "characterizing tractable metrics" is in the context of streaming algorithms, where the input is an implicit vector $x \in \mathbb{R}^{n}$ given in the form of updates $(i, \delta)$, with the semantics that coordinate $i$ has to be increased by $\delta \in \mathbb{R}$.

There are two known results in this vein. First, [45] characterized the streaming complexity of computing the sum $\sum_{i} \varphi\left(x_{i}\right)$, for some fixed $\varphi$ (e.g., $\varphi(x)=x^{2}$ for $\ell_{2}$ norm), when the updates are positive. They gave a precise property of $\varphi$ that determines whether the complexity of the problem is small. Second, [112] showed that, in certain settings, streaming algorithms may as well be linear, i.e., maintain a sketch $f(x)=A x$ for a matrix $A$, and the size of the sketch is increased by a factor logarithmic in the dimension of $x$.

Furthermore, after the appearance of the conference version of the current article, there has been another characterization result that significantly generalizes and extends [45]. Specifically, for every symmetric norm $\|\cdot\|_{X}$, it is proved in [42] that the sketching (and streaming) complexity of computing $\|x\|_{X}$ is characterized by the norm's (maximum) modulus of concentration, up to polylogarithmic factors in the dimension of $X$. Finally, we mention a related work [25], where an efficient data structure for the Approximate Nearest Neighbor search (ANN) is constructed for every symmetric norm. It is known [85, 108] that efficient sketches imply good data structures for ANN, however, the result of [25] shows that having efficient ANN data structure is a way more general property of an underlying norm.

### 5.1.5 Proof overview

Following common practice, we think of sketching as a communication protocol. In fact, our results hold for protocols with an arbitrary number of rounds (and access to public randomness).

Our proof of Theorem 5.1.2 can be divided into two parts: information-theoretic and analytic. First, we use information-theoretic tools to convert an efficient protocol for $\operatorname{DTEP}(X, D)$ into a so-called threshold map from $X$ to a Hilbert space. Our notion
of a threshold map can be viewed as a very weak definition of embeddability (see Definition 5.4.5 for details). Second, we use techniques from nonlinear functional analysis to convert a threshold map to a linear map into $\ell_{1-\varepsilon}$.

Information-theoretic part. To get a good threshold map from a protocol for $\operatorname{DTEP}(X, D)$, we proceed in three steps. First, using the fact that $X$ is a normed space, we are able to give a good protocol for $\operatorname{DTEP}\left(\ell_{\infty}^{k}(X), D k\right)$ (Lemma 5.4.3). The space $\ell_{\infty}^{k}(X)$ is a product of $k$ copies of $X$ equipped with the norm $\left\|\left(x_{1}, \ldots, x_{k}\right)\right\|=\max _{i}\left\|x_{i}\right\|$. Then, invoking the main result from [20], we conclude non-existence of certain Poincarétype inequalities for $X$ (Theorem 5.4.4, in the contrapositive).

Finally, we use convex duality together with a compactness argument to conclude the existence of a desired threshold map from $X$ to a Hilbert space (Lemma 5.4.6, again in the contrapositive).

Analytic part. We proceed from a threshold map by upgrading it to a uniform embedding (see Definition 5.2.1) of $X$ into a Hilbert space (Theorem 5.4.12). For this we adapt arguments from [94, 151]. We use two tools from nonlinear functional analysis: an extension theorem for $1 / 2$-Hölder maps from a (general) metric space to a Hilbert space [128] (Theorem 5.4.16), and a symmetrization lemma for maps from metric abelian groups to Hilbert spaces [3] (Lemma 5.4.14).

Then we convert a uniform embedding of $X$ into a Hilbert space to a linear embedding into $\ell_{1-\varepsilon}$ by applying the result of Aharoni, Maurey and Mityagin [3] together with the result of Nikishin [138]. A similar argument has been used in [134].

To prove a quantitative version of this step, we examine the proofs from [3] and [138], and obtain explicit bounds on the distortion of the resulting map. We accomplish this in Section 5.5.

Embeddings into $\ell_{1}$. To prove Theorem 5.1.3 (which has dependence on the dimension of $X$ ), we note that it is a simple corollary of Theorem 5.1.2 and a result of Zvavitch [181], which gives a dimension reduction procedure for subspaces of $\ell_{1-\varepsilon}$.

Norms closed under sum-product. Finally, we prove Theorem 5.1.4 - embeddability into $\ell_{1}$ for norms closed under sum-product - by proving and using a finitary
version of the theorem of Kalton [95] (Lemma 5.6.1), instead of invoking Nikishin's theorem as above. We prove the finitary version by reducing it to the original statement of Kalton's theorem via a compactness argument.

Let us point out that Naor and Schechtman [134] showed how to use (the original) Kalton's theorem to upgrade a uniform embedding of $\mathrm{EMD}_{n}$ into a Hilbert space to a linear embedding into $\ell_{1}$ (they used this reduction to exclude uniform embeddability of $\mathrm{EMD}_{n}$ ). Their proof used certain specifics of EMD. In contrast, to get Theorem 5.1.4 for general norms, we seem to need a finitary version of Kalton's theorem.

We also note that in Theorems 5.1.2, 5.1.3 and 5.1.4, we can conclude embeddability into $\ell_{1-\varepsilon}^{d}$ and $\ell_{1}^{d}$ respectively, where $d$ is near-linear in the dimension of the original space. This conclusion uses the known dimension reduction theorems for subspaces from [167, 181].

### 5.2 Preliminaries on functional analysis

We remind a few definitions and standard facts from functional analysis that will be useful for our proofs. A central notion in our proofs is the notion of uniform embeddings, which is a weaker version of embeddability.

Definition 5.2.1. For two metric spaces $X$ and $Y$, we say that a map $f: X \rightarrow Y$ is a uniform embedding, if there exist two non-decreasing functions $L, U: \mathbb{R}_{+} \rightarrow \mathbb{R}_{+}$such that for every $x_{1}, x_{2} \in X$ one has $L\left(d_{X}\left(x_{1}, x_{2}\right)\right) \leq d_{Y}\left(f\left(x_{1}\right), f\left(x_{2}\right)\right) \leq U\left(d_{X}\left(x_{1}, x_{2}\right)\right)$, $U(t) \rightarrow 0$ as $t \rightarrow 0$ and $L(t)>0$ for every $t>0$. The functions $L(\cdot)$ and $U(\cdot)$ are called moduli of the embedding.

Definition 5.2.2. An inner-product space is a real vector space $X$ together with an inner product $\langle\cdot, \cdot\rangle: X \times X \rightarrow \mathbb{R}$, which is a symmetric positive-definite bilinear form. A Hilbert space is an inner-product space $X$ that is complete as a metric space.

Every inner-product space is a normed space: we can set $\|x\|=\sqrt{\langle x, x\rangle}$. For a normed space $X$ we denote by $B_{X}$ its closed unit ball. The main example of a Hilbert space is $\ell_{2}$, the space of all real sequences $\left\{x_{n}\right\}$ with $\sum_{i} x_{i}^{2}<\infty$, where the inner
product is defined as

$$
\langle x, y\rangle=\sum_{i} x_{i} y_{i} .
$$

Definition 5.2.3. For a set $S$, a function $K: S \times S \rightarrow \mathbb{R}$ is called a kernel if $K\left(s_{1}, s_{2}\right)=K\left(s_{2}, s_{1}\right)$ for every $s_{1}, s_{2} \in S$. We say that the kernel $K$ is positive-definite if for every $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{n} \in \mathbb{R}$ and $s_{1}, s_{2}, \ldots, s_{n} \in S$, one has

$$
\sum_{i, j=1}^{n} \alpha_{i} \alpha_{j} K\left(s_{i}, s_{j}\right) \geq 0
$$

We say that $K$ is negative-definite if for every $\alpha_{1}, \ldots, \alpha_{n} \in \mathbb{R}$ with $\alpha_{1}+\alpha_{2}+\ldots+\alpha_{n}=0$ and $s_{1}, s_{2}, \ldots, s_{n} \in S$, one has

$$
\sum_{i, j=1}^{n} \alpha_{i} \alpha_{j} K\left(s_{i}, s_{j}\right) \leq 0 .
$$

The following are standard facts about positive- and negative-definite kernels.
Fact 5.2.4 ([160]). For a kernel $K: S \times S \rightarrow \mathbb{R}$, there exists an embedding $f: S \rightarrow H$, where $H$ is a Hilbert space, such that $K\left(s_{1}, s_{2}\right)=\left\langle f\left(s_{1}\right), f\left(s_{2}\right)\right\rangle_{H}$ for every $s_{1}, s_{2} \in S$, iff $K$ is positive-definite.

Fact 5.2.5 ([161]). For a kernel $K: S \times S \rightarrow \mathbb{R}$, there exists an embedding $f: S \rightarrow H$, where $H$ is a Hilbert space, such that $K\left(s_{1}, s_{2}\right)=\left\|f\left(s_{1}\right)-f\left(s_{2}\right)\right\|_{H}^{2}$ for every $s_{1}, s_{2} \in S$, iff $K(s, s)=0$ for every $s \in S$ and $K$ is negative-definite.

Definition 5.2.6. For an abelian group $G$, we say that a function $f: G \rightarrow \mathbb{R}$ is positive-definite if a kernel $K\left(g_{1}, g_{2}\right)=f\left(g_{1}-g_{2}\right)$ is positive-definite. Similarly, $f$ is said to be negative-definite if $K\left(g_{1}, g_{2}\right)=f\left(g_{1}-g_{2}\right)$ is negative-definite.

The following lemma essentially says that an embedding of an abelian group $G$ into a Hilbert space can be made translation-invariant.

Lemma 5.2.7 (see the proof of Lemma 3.5 in [3]). Suppose that $f$ is a map from an abelian group $G$ to a Hilbert space such that for every $g \in G$ we have

$$
\sup _{g_{1}-g_{2}=g}\left\langle f\left(g_{1}\right), f\left(g_{2}\right)\right\rangle<+\infty .
$$

Then, there exists a map $f^{\prime}$ from $G$ to a Hilbert space such that $\left\langle f^{\prime}\left(g_{1}\right), f^{\prime}\left(g_{2}\right)\right\rangle$ depends only on $g_{1}-g_{2}$ and for every $g_{1}, g_{2} \in G$ we have

$$
\inf _{g_{1}^{\prime}-g_{2}^{\prime}=g_{1}-g_{2}}\left\langle f\left(g_{1}^{\prime}\right), f\left(g_{2}^{\prime}\right)\right\rangle \leq\left\langle f^{\prime}\left(g_{1}\right), f^{\prime}\left(g_{2}\right)\right\rangle \leq \sup _{g_{1}^{\prime}-g_{2}^{\prime}=g_{1}-g_{2}}\left\langle f\left(g_{1}^{\prime}\right), f\left(g_{2}^{\prime}\right)\right\rangle .
$$

Finally, let $\operatorname{dim} X$ denote the dimension of a finite-dimensional vector space $X$.

### 5.3 Preliminaries on communication complexity

Let $X$ be a metric space, on which we would like to solve $\operatorname{DTEP}_{r}(X, D)$ defined as follows for some $r>0$ and $D \geq 1$. Alice has a point $x \in X$, Bob has a point $y \in X$, and they would like to decide between the two cases: $d_{X}(x, y) \leq r$ and $d_{X}(x, y)>D r$. To accomplish this goal, Alice and Bob exchange at most $s$ bits of communication.

There are several types of communication protocols that we consider, depending on the randomness used, which we present below in the order of their power. Our main result applies to the most powerful type. We will later show some connections between the protocols of different types.

- Deterministic protocols. This is a simple two-way communication protocol with no randomness. First, Alice sends a bit to Bob that depends only on $x$. Then, Bob sends a bit to Alice that depends on Alice's first communication bit and on $y$. Then, Alice sends a bit to Bob that depends on $x$ and the two previous communication bits, etc. Finally, whoever sends the $s$-th bit must decide the answer to the DTEP problem. We define a transcript $\Pi_{x, y}$ to be the sequence of $s$ bits sent by the two parties for a given pair of inputs $x$ and $y$.
- Private-coin protocols with bounded number of coins. This is a randomized version of the previous definition. Alice and Bob each have access to an independent random string, denoted $a \in\{0,1\}^{R}$ and $b \in\{0,1\}^{R}$, respectively. Communication bits sent by Alice are allowed to depend on $a$, and those sent by Bob may depend on $b$. We require that for every pair of inputs, the probability
(over the random coins $a, b$ ) of the answer being correct is at least, say, 2/3. Whenever we allow randomness, the transcript $\Pi_{x, y}$ becomes a random variable (depending on $x$ and $y$ ). For fixed $a, b \in\{0,1\}^{R}$, we denote by $\Pi_{x, y}(a, b)$ the (deterministic) transcript for inputs $x$ and $y$ when the random strings are set to be $a$ and $b$, respectively.
- Public-coin protocols with finitely many coins. This is a variant of the previous definition, where Alice and Bob have access to a common random string sampled uniformly from $\{0,1\}^{R}$, and the bits sent by both Alice and Bob can depend on this random string. Again, we require the probability (over the public coins) of the answer being correct to be at least, say, $2 / 3$ for every pair of inputs. Also, we denote by $\Pi_{x, y}(u)$ the deterministic transcript for fixed inputs and public coins $u$.

Clearly, a public-coin protocol with $2 R$ public coins can emulate a private-coin protocol with $R$ random bits for each of Alice and Bob.

- Public-coin protocols with countably many coins. The protocols defined above are standard in the communication complexity literature. However, we need a definition that is stronger: we allow countably many public coins. The reason to consider the stronger notion is that the known protocols for DTEP based on [84] fall into this category. Since we allow infinitely many coins, we need to be careful when defining a class of allowed protocols. A sequence of coin tosses $u$ can be identified with a point in the Cantor space $\Omega=\{0,1\}^{\omega}$ equipped with the standard Lebesgue measure. We require that for every pair of inputs $x, y \in X$, the function $u \mapsto \Pi_{x, y}(u)$ is measurable. This restriction allows us to consider probabilities of the form $\operatorname{Pr}\left[\Pi_{x, y} \in A\right]$, where $A \subseteq\{0,1\}^{s}$ is an arbitrary set of possible transcripts. In particular, the probability of success is well-defined, and we require it, as before, to be at least $2 / 3$.

The results in this paper apply to the most general protocols: public-coin protocols with countably many coins.

We now show some connections between these notions. A crucial tool in our result is a theorem of [20], which is itself based on the tools from [34]. The latter shows a lower bound for private-coin protocols with finitely many coins. We show next how the lower bounds from $[20,34]$ extend to the most general type, public-coin protocols with countably many coins.

### 5.3.1 Information complexity: private-coins vs. public-coins

In general, a lower bound for private-coins protocols does not imply a lower bound for public-coins protocols (without a loss in the parameters). However, such an implication does hold for the particular lower bound technique that we are employing. In particular, we use and exploit the notion of information complexity from [34], defined as follows. Let $(x, y, \lambda)$ be distributed according to a distribution $\mathcal{D}$ over $X \times X \times \Lambda$, where $\Lambda$ is an auxiliary set. We will assume that the support of $\mathcal{D}$ is finite. Then, we can define the information complexity with respect to $\mathcal{D}$, denoted $\operatorname{IC}_{\mathcal{D}}\left(\operatorname{DTEP}_{r}(X, D)\right.$ ), to be the infimum of $I\left(x, y: \Pi_{x, y} \mid \lambda\right)$ over all private-coin protocols for $\operatorname{DTEP}_{r}(X, D)$, which succeed on every valid input with probability at least $2 / 3$, where $I(\cdot: \cdot \mid \cdot)$ is the (conditional) mutual information.

It is a standard fact that $\mathrm{IC}_{\mathcal{D}}\left(\operatorname{DTEP}_{r}(X, D)\right)$ is a lower bound on the communication complexity of $\mathrm{DTEP}_{r}(X, D)$ with private-coin protocols since

$$
I\left(x, y: \Pi_{x, y} \mid \lambda\right) \leq \sup _{x, y, a, b}\left|\Pi_{x, y}(a, b)\right| .
$$

However, we are interested in using the information complexity (as defined above) to lower bound the communication complexity of $\operatorname{DTEP}_{r}(X, D)$ for public-coin protocols with finite number of coins. It turns out that $\mathrm{IC}_{\mathcal{D}}\left(\mathrm{DTEP}_{r}(X, D)\right)$ is a valid lower bound for this case as well, as argued in the claim below.

Lemma 5.3.1. The communication complexity of $\operatorname{DTEP}_{r}(X, D)$ for public-coin protocols with finite number of coins is at least $\mathrm{IC}_{\mathcal{D}}\left(\operatorname{DTEP}_{r}(X, D)\right)$.

Proof. Consider any protocol with public randomness, denoted $\Pi_{x, y}(u)$, where $x, y$ are
the two inputs and $u$ is the public random string. Then

$$
\left.\sup _{x, y, u} \mid \Pi_{x, y}(u)\right) \mid \geq H\left(\Pi_{x, y}(u) \mid \lambda, u\right) \geq I\left(x, y: \Pi_{x, y}(u) \mid \lambda, u\right) .
$$

Now consider the following private-coins protocol $\Pi_{x, y}^{\prime}(a, b)$, where $a, b$ are the two private random strings of Alice and Bob, respectively. In the first round, Alice sends $a$ to Bob to be used as public randomness $u=a$. Then they run $\Pi_{x, y}(u)$. In other words, the transcript of $\Pi_{x, y}^{\prime}(a, b)$ is $\left\langle a, \Pi_{x, y}(a)\right\rangle$. We claim that

$$
I\left(x, y: \Pi_{x, y}^{\prime}(a, b) \mid \lambda\right)=I\left(x, y: \Pi_{x, y}(u) \mid \lambda, u\right)
$$

Indeed, by definition of $\Pi^{\prime}$,

$$
I\left(x, y: \Pi_{x, y}^{\prime}(a, b) \mid \lambda\right)=I\left(x, y: a, \Pi_{x, y}(a) \mid \lambda\right)
$$

and using the chain rule for mutual information,

$$
I\left(x, y: a, \Pi_{x, y}(a) \mid \lambda\right)=I(x, y: a \mid \lambda)+I\left(x, y: \Pi_{x, y}(a) \mid \lambda, a\right)
$$

The first term is exactly zero since $x, y$ and $a$ are independent (conditioned on $\lambda$ ). The remaining term gives the equality we are looking for, and proves the lemma. In particular, we see that the length of a public-coin protocol is at least the information complexity $I\left(x, y: \Pi_{x, y}^{\prime}(a, b) \mid \lambda\right)$ of any private-coin protocol $\Pi^{\prime}$.

### 5.3.2 From countable to finite number of coins

We now observe that if we focus only on a finite number of possible inputs to our DTEP problem, then the existence of a protocol with countably-many coins implies the existence of a protocol with bounded number of coins. This claim will be sufficient to generalize our theorem to the most general type of protocols-public-coin protocols with countably many coins: see the remark after Theorem 5.4.4.

Claim 5.3.2. Fix a public-coin protocol with countably many coins and $s$ bits of
communication. Let $\left(x^{(1)}, y^{(1)}\right),\left(x^{(2)}, y^{(2)}\right), \ldots,\left(x^{(N)}, y^{(N)}\right)$ be $N$ fixed pairs of inputs for the DTEP problem, and let $\varepsilon>0$ be a positive parameter. Then there exists a public-coin protocol with $R=R(s, N, \varepsilon)<\infty$ coins and s bits of communication, such that for every $1 \leq i \leq N$, the success probabilities of the original and the new protocols on $\left(x^{(i)}, y^{(i)}\right)$ differ by at most $\varepsilon$.

Proof. Since we care about correctness of the protocol only on the inputs $\left(x^{(i)}, y^{(i)}\right)$, we can think of the protocol as a distribution over a bounded (as a function of $s$ and $N$ ) number of deterministic protocols (there is only a finite number of distinct protocol transcripts). Then, we can approximate this distribution within a statistical distance $\varepsilon$ using a bounded number of public coins.

### 5.4 From sketches to uniform embeddings

Our main technical result shows that, for a finite-dimensional normed space $X$, good sketches for $\operatorname{DTEP}(X, D)$ imply a good uniform embedding of $X$ into a Hilbert space (Definition 5.2.1). Below is the formal statement.

Theorem 5.4.1. Suppose a finite-dimensional normed space $X$ admits a public-coin randomized communication protocol for $\operatorname{DTEP}(X, D)$ of size s for approximation $D>1$. Then, there exists a map $f: X \rightarrow H$ to a Hilbert space such that for all $x_{1}, x_{2} \in X$,

$$
\min \left\{1, \frac{\left\|x_{1}-x_{2}\right\|_{X}}{s \cdot D}\right\} \leq\left\|f\left(x_{1}\right)-f\left(x_{2}\right)\right\|_{H} \leq K \cdot\left\|x_{1}-x_{2}\right\|_{X}^{1 / 2}
$$

where $K>1$ is an absolute constant.
Theorem 5.4.1 implies a qualitative version of Theorem 5.1.2 using the results of Aharoni, Maurey, and Mityagin [3] and Nikishin [138] (see Theorem 5.4.2).

Theorem 5.4.2 ([3, 138]). For every fixed $0<\varepsilon<1$, any finite-dimensional normed space $X$ that is uniformly embeddable into a Hilbert space is linearly embeddable into $\ell_{1-\varepsilon}$ with a distortion that depends only on $\varepsilon$ and the moduli of the assumed uniform embedding.

To prove the full (quantitative) versions of Theorems 5.1.2 and 5.1.3, we adapt the proofs from [3] and [138] in Section 5.5 to get an explicit bound on the distortion.

In the rest of this section, we prove Theorem 5.4.1 according to the outline in Section 5.1.5, putting the pieces together in Section 5.4.4.

### 5.4.1 Sketching implies the absence of Poincaré inequalities

Sketching is often viewed from the perspective of a two-party communication complexity. Alice receives input $x$, Bob receives $y$, and they need to communicate to solve the DTEP problem. In particular, a sketch of size $s$ implies a communication protocol that transmits $s$ bits: Alice just sends her sketch $\mathbf{s k}(x)$ to Bob, who computes the output of DTEP (based on that message and his sketch $\mathbf{s k}(y)$ ). We assume here a public-coins model, i.e., Alice and Bob have access to a common (public) random string that determines the sketch function sk.

To characterize sketching protocols, we build on results of Andoni, Jayram and Pǎtraşcu [20, Sections 3 and 4]. This works in two steps: first, we show that a protocol for $\operatorname{DTEP}(X, D)$ implies a sketching algorithm for $\operatorname{DTEP}\left(\ell_{\infty}^{k}(X), k D\right)$, with a loss of factor $k$ in approximation (Lemma 5.4.3, see the proof in the end of the section). As usual, $\ell_{\infty}^{k}(X)$ is a normed space derived from $X$ by taking the vector space $X^{k}$ and letting the norm of a vector $\left(x_{1}, \ldots x_{k}\right) \in X^{k}$ be the maximum of the norms of its $k$ components. The second step is to apply a result from [20] (Theorem 5.4.4), which asserts that sketching for $\ell_{\infty}^{k}(X)$ precludes certain Poincaré inequalities for the space $X$.

Lemma 5.4.3. Let $X$ be a finite-dimensional normed space that for some $D \geq 1$ admits a communication protocol for $\operatorname{DTEP}(X, D)$ of size $s$. Then for every integer $k$, the space $\ell_{\infty}^{k}(X)$ admits sketching with approximation $k D$ and sketch size $s^{\prime}=O(s)$.

Proof. Fix a threshold $t>0$, and recall that we defined the success probability of sketching to be 0.9 . By our assumption, there is a sketching function sk for $X$ that achieves approximation $D$ and sketch size $s$ for threshold $k t$. Now define a "sketching" function sk' for $\ell_{\infty}^{k}(X)$ by choosing random signs $\varepsilon_{1}, \ldots, \varepsilon_{k} \in\{ \pm 1\}$, letting
$\mathbf{s k}^{\prime}: x \mapsto \mathbf{s k}\left(\sum_{i=1}^{k} \varepsilon_{i} x_{i}\right)$, and using the same decision procedure used by $\mathbf{s k}$ (for $X$ ).
Now to examine the performance of $\mathbf{s k}^{\prime}$, consider $x, y \in \ell_{\infty}^{k}(X)$. If their distance is at most $t$, then we always have that $\left\|\sum_{i=1}^{k} \varepsilon_{i} x_{i}-\sum_{i=1}^{k} \varepsilon_{i} y_{i}\right\| \leq \sum_{i=1}^{k}\left\|x_{i}-y_{i}\right\| \leq k t$ (i.e., for every realization of the random signs). Thus with probability at least 0.9 the sketch will declare that $x, y$ are "close".

If the distance between $x, y$ is greater than $k D \cdot t$, then for some coordinate, say $i=1$, we have $\left\|x_{1}-y_{1}\right\|>k D \cdot t$. Letting $z=\sum_{i \geq 2} \varepsilon_{i}\left(x_{i}-y_{i}\right)$, we can write $\left\|\sum_{i=1}^{k} \varepsilon_{i} x_{i}-\sum_{i=1}^{k} \varepsilon_{i} y_{i}\right\|=\left\|\varepsilon_{1}\left(x_{1}-y_{1}\right)+z\right\|=\left\|\left(x_{1}-y_{1}\right)+\varepsilon_{1} z\right\|$. The last term must be at least $\left\|x_{1}-y_{1}\right\|$ under at least one of the two possible realizations of $\varepsilon_{1}$, because by the triangle inequality $2\left\|x_{1}-y_{1}\right\| \leq\left\|\left(x_{1}-y_{1}\right)+z\right\|+\left\|\left(x_{1}-y_{1}\right)-z\right\|$. We see that with probability $1 / 2$ we have $\left\|\sum_{i=1}^{k} \varepsilon_{i} x_{i}-\sum_{i=1}^{k} \varepsilon_{i} y_{i}\right\| \geq\left\|x_{1}-y_{1}\right\|>D \cdot k t$, and thus with probability at least $1 / 2 \cdot 0.9=0.45$ the sketch will declare that $x, y$ are "far". This last guarantee is not sufficient for $\mathbf{s k}^{\prime}$ to be called a sketch, but it can easily be amplified.

The final sketch $\mathbf{s} \mathbf{k}^{\prime \prime}$ for $\ell_{\infty}^{k}(X)$ is obtained by $O(1)$ independent repetitions of $\mathbf{s k}^{\prime}$, and returning "far" if at least 0.3 -fraction of the repetitions come up with this decision. These repetitions amplify the success probability to 0.9 , while increasing the sketch size to $O(s)$.

We now state a slight modification of the theorem from [20]. We will actually use its contrapositive, to conclude the absence of Poincaré inequalities.

Theorem 5.4.4 (modification of [20]). Let $X$ be a metric space, and fix $r>0, D \geq 1$. Suppose there are $\alpha>0, \beta \geq 0$, and two symmetric probability measures $\mu_{1}, \mu_{2}$ on $X \times X$ such that

- The support of $\mu_{1}$ is finite and is only on pairs with distance at most r;
- The support of $\mu_{2}$ is finite and is only on pairs with distance greater than Dr; and
- For every $f: X \rightarrow B_{\ell_{2}}$ (where $B_{\ell_{2}}$ is the unit ball of $\ell_{2}$ ),

$$
\underset{(x, y) \sim \mu_{1}}{\mathbb{E}}\|f(x)-f(y)\|^{2} \geq \alpha \cdot \underset{(x, y) \sim \mu_{2}}{\mathbb{E}}\|f(x)-f(y)\|^{2}-\beta .
$$

Then for every integer $k$, the communication complexity of $\operatorname{DTEP}\left(\ell_{\infty}^{k}(X), D\right)$ for protocols with countably many public coins (see Section 5.3 for precise definitions) and with probability of error $\delta_{0}>0$ is at least $\Omega(k) \cdot\left(\alpha\left(1-2 \sqrt{\delta_{0}}\right)-\beta\right)$.

In [20], almost the same theorem is proved with only one difference: the protocols for $\operatorname{DTEP}\left(\ell_{\infty}^{k}(X), D\right)$ are only allowed to use finitely many private coins. Here we use Claims 5.3.1 and 5.3.2 to generalize their theorem to Theorem 5.4.4.

Indeed, because the "hard distributions" $\mu_{1}$ and $\mu_{2}$ are finitely-supported, an inspection of the proofs from [20] shows that there is a finite set of inputs $\mathcal{I}$ such that any private-coin protocol for $\operatorname{DTEP}\left(\ell_{\infty}^{k}(X), D\right)$ that is correct on $\mathcal{I}$ with probability at least $1-\delta_{0}$ must have information complexity at least $\Omega(k) \cdot\left(\alpha\left(1-2 \sqrt{\delta_{0}}\right)-\beta\right)$. But by Claim 5.3.1, we get that any protocol with bounded number of public coins correct on $\mathcal{I}$ must have communication complexity at least $\Omega(k) \cdot\left(\alpha\left(1-2 \sqrt{\delta_{0}}\right)-\beta\right)$. Finally, Claim 5.3.2 implies that the same is true for protocols with countably many public coins that are correct on all valid inputs with probability at least $1-\delta_{0}$.

### 5.4.2 The absence of Poincaré inequalities implies threshold maps

We proceed to prove that non-existence of Poincaré inequalities implies the existence a "threshold map", as formalized in Lemma 5.4.6 below. The proof is similar to duality arguments that one often encounters in embedding theory: for instance, see Proposition 15.5.2 in [121]. First we define the notion of threshold maps.

Definition 5.4.5. A map $f: X \rightarrow Y$ between metric spaces $\left(X, d_{X}\right)$ and $\left(Y, d_{Y}\right)$ is said to be $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2}, \tau_{3}\right)$-threshold for $0<s_{1}<s_{2}, 0<\tau_{1}<\tau_{2}<\tau_{3}$, if for all $x_{1}, x_{2} \in X:$

- if $d_{X}\left(x_{1}, x_{2}\right) \leq s_{1}$, then $d_{Y}\left(f\left(x_{1}\right), f\left(x_{2}\right)\right) \leq \tau_{1}$;
- if $d_{X}\left(x_{1}, x_{2}\right) \geq s_{2}$, then $d_{Y}\left(f\left(x_{1}\right), f\left(x_{2}\right)\right) \geq \tau_{2}$; and
- $d_{Y}\left(f\left(x_{1}\right), f\left(x_{2}\right)\right) \leq \tau_{3}$.

We now provide the main lemma of this section, stated in the contrapositive: the non-existence of threshold maps implies a Poincaré inequality.

Lemma 5.4.6. Suppose $X$ is a metric space that has no $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2},+\infty\right)$-threshold map to a Hilbert space. Then, for every $\delta>0$ there exist two symmetric probability measures $\mu_{1}, \mu_{2}$ on $X \times X$ such that

- The support of $\mu_{1}$ is finite and is only on pairs with distance at most $s_{1}$;
- The support of $\mu_{2}$ is finite and is only on pairs with distance at least $s_{2}$; and
- For every $f: X \rightarrow B_{\ell_{2}}$,

$$
\begin{equation*}
\underset{(x, y) \sim \mu_{1}}{\mathbb{E}}\|f(x)-f(y)\|^{2} \geq\left(\frac{\tau_{1}}{\tau_{2}}\right)^{2} \cdot \underset{(x, y) \sim \mu_{2}}{\mathbb{E}}\|f(x)-f(y)\|^{2}-\delta \tag{5.1}
\end{equation*}
$$

We prove Lemma 5.4.6 via the following three claims. The first one uses standard arguments about embeddability of finite subsets (see, e.g., Proposition 8.12 in [41], or Lemma 1.1 from [31]). We note that this claim requires a finite value for $\tau_{3}$, as opposed to $\tau_{3}=+\infty$, which is the only reason the definition of a threshold embedding (Definition 5.4.5) needs the parameter $\tau_{3}$. In the following claims, we denote by $\binom{X}{2}$ the set of all unordered pairs $\{x, y\}$ with $x, y \in X, x \neq y$.

Claim 5.4.7. For every metric space $X$ and every $0<s_{1}<s_{2}, 0<\tau_{1}<\tau_{2}<\tau_{3}$ there exists an $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2}, \tau_{3}\right)$-threshold map of $X$ to a Hilbert space iff the same is true for every finite subset of $X$.

The proof of Claim 5.4.7 uses standard definitions and facts from general topology: product topology, Tychonoff's theorem, as well as convergence and accumulation points along nets. These definitions can be found in a general topology textbook (see, e.g., [130]).

Proof. The "only if" direction is obvious, so let us turn to the "if" part. Consider the topological space

$$
U=\prod_{\{x, y\} \in\binom{X}{2}}\left[-\tau_{3}^{2}, \tau_{3}^{2}\right] .
$$

By Tychonoff's theorem $U$ is compact. For every finite $X^{\prime} \subset X$ there exists an $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2}, \tau_{3}\right)$-threshold map $f_{X^{\prime}}$ from $X^{\prime}$ to a Hilbert space. It gives rise to a point $u_{X^{\prime}} \in U$ whose coordinates are given by

$$
\left(u_{X^{\prime}}\right)_{x, y}= \begin{cases}\left\|f_{X^{\prime}}(x)-f_{X^{\prime}}(y)\right\|^{2}, & \text { if } x, y \in X^{\prime} \\ 0, & \text { otherwise }\end{cases}
$$

Since $U$ is compact, $u_{X^{\prime}}$ has an accumulation point $u^{*} \in U$ along the net of finite subsets of $X$. Let us reformulate what it means.

Claim 5.4.8. For every $\left\{x_{1}, y_{1}\right\},\left\{x_{2}, y_{2}\right\}, \ldots,\left\{x_{k}, y_{k}\right\} \in\binom{X}{2}$ and every $\varepsilon>0$, there exists a finite set $A \subset X$ such that for all $1 \leq i \leq k$, both $x_{i}, y_{i} \in A$ and $\mid\left(u^{*}\right)_{x_{i}, y_{i}}-$ $\left|\left|f_{A}\left(x_{i}\right)-f_{A}\left(y_{i}\right) \|^{2}\right|<\varepsilon\right.$.

Now we define a kernel $K: X \times X \rightarrow \mathbb{R}$, given by (recall Definition 5.2.3):

$$
K(x, y)= \begin{cases}0, & \text { if } x=y \\ \left(u^{*}\right)_{x, y}, & \text { otherwise }\end{cases}
$$

Claim 5.4.9. The kernel $K(\cdot, \cdot)$ is negative-definite.

Proof. Suppose that $K$ is not negative-definite. It means that there exist real numbers $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{n} \in \mathbb{R}$ with $\sum_{i} \alpha_{i}=0$, and $t_{1}, t_{2}, \ldots, t_{n} \in X$ such that

$$
\sum_{i, j=1}^{n} \alpha_{i} \alpha_{j} K\left(t_{i}, t_{j}\right)=\gamma>0
$$

There exists $\varepsilon>0$ such that for every $\left(a_{i j}\right)_{i, j=1}^{n}$ with $\left|a_{i j}-K\left(t_{i}, t_{j}\right)\right|<\varepsilon$ one has

$$
\begin{equation*}
\sum_{i, j=1}^{n} \alpha_{i} \alpha_{j} a_{i j} \geq \gamma / 2>0 \tag{5.2}
\end{equation*}
$$

Now apply Claim 5.4 .8 to get a finite set $A \subset X$ that contains all $s_{i}$ 's such that
$\left\|f_{A}\left(t_{i}\right)-f_{A}\left(t_{j}\right)\right\|^{2}$ is within $\varepsilon$ from $K\left(t_{i}, t_{j}\right)$ for every $i, j$. But by (5.2), it means that

$$
\sum_{i, j=1}^{n} \alpha_{i} \alpha_{j}\left\|f_{A}\left(t_{i}\right)-f_{A}\left(t_{j}\right)\right\|^{2} \geq \gamma / 2>0
$$

which contradicts Fact 5.2.5. This proves Claim 5.4.9.

Thus, by Fact 5.2 .5 , there exists a map $f: X \rightarrow H$ to a Hilbert space $H$ such that for every $x, y \in X$ one has $\|f(x)-f(y)\|^{2}=K(x, y)$. The final step is to verify that $f$ is indeed a required $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2}, \tau_{3}\right)$-map (according to Definition 5.4.5). This can be done exactly the same way as in the proof of Claim 5.4.9. This completes the proof of Claim 5.4.7.

Claim 5.4.10. Suppose that $\left(X, d_{X}\right)$ is a finite metric space and $0<s_{1}<s_{2}$, $0<\tau_{1}<\tau_{2}<\tau_{3}$. Assume that there is no $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2}, \tau_{3}\right)$-threshold map of $X$ to $\ell_{2}$. Then, there exist two symmetric probability measures $\mu_{1}, \mu_{2}$ on $X \times X$ such that

- $\mu_{1}$ is supported only on pairs with distance at most $s_{1}$, while $\mu_{2}$ is supported only on pairs with distance at least $s_{2}$; and
- for every $f: X \rightarrow \ell_{2}$,

$$
\begin{align*}
& \underset{(x, y) \sim \mu_{1}}{\mathbb{E}}\|f(x)-f(y)\|^{2} \\
& \geq\left(\frac{\tau_{1}}{\tau_{2}}\right)^{2} \cdot \underset{(x, y) \sim \mu_{2}}{\mathbb{E}}\|f(x)-f(y)\|^{2}-\left(\frac{2 \tau_{1}}{\tau_{3}}\right)^{2} \cdot \sup _{x \in X}\|f(x)\|^{2} \tag{5.3}
\end{align*}
$$

Proof. Let $\mathcal{L}_{2} \subset \mathbb{R}^{\binom{X}{2}}$ be the cone of squared Euclidean metrics (also known as negative-type distances) on $X$. Let $\mathcal{K} \subset \mathbb{R}^{\binom{X}{2}}$ be the polytope of non-negative functions $l:\binom{X}{2} \rightarrow \mathbb{R}_{+}$such that for every $x, y \in X$ we have

- $l(\{x, y\}) \leq \tau_{3}^{2}$;
- if $d_{X}(x, y) \leq s_{1}$, then $l(\{x, y\}) \leq \tau_{1}^{2}$;
- if $d_{X}(x, y) \geq s_{2}$, then $l(\{x, y\}) \geq \tau_{2}^{2}$.

Notice that $\mathcal{L}_{2} \cap \mathcal{K}=\emptyset$, as otherwise $X$ allows an $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2}, \tau_{3}\right)$-threshold map to $\ell_{2}$. We will need the following claim, which is just a variant of the Hyperplane Separation Theorem.

Claim 5.4.11. There exists $a \in \mathbb{R}^{\binom{X}{2}}$ such that

$$
\begin{array}{ll}
\forall l \in \mathcal{L}_{2}, & \langle a, l\rangle \leq 0 \\
\forall l \in \mathcal{K}, & \langle a, l\rangle>0 \tag{5.5}
\end{array}
$$

Proof. Since both $\mathcal{L}_{2}$ and $\mathcal{K}$ are convex and closed, and, in addition, $\mathcal{K}$ is compact, there exists a separating (affine) hyperplane between $\mathcal{L}_{2}$ and $\mathcal{K}$. Specifically, there is a non-zero $a$ such that for every $l \in \mathcal{L}_{2}$ one has $\langle a, l\rangle \leq \eta$, and for every $l \in \mathcal{K}$ one has $\langle a, l\rangle>\eta$. Since $\mathcal{L}_{2}$ is a cone, one can assume without loss of generality that $\eta=0$. Indeed, the case $\eta<0$ is impossible because $0 \in \mathcal{L}_{2}$, so suppose that $\eta>0$. If for all $l \in \mathcal{L}_{2}$ we have $\langle a, l\rangle \leq 0$, then we are done. Otherwise, take any $l \in \mathcal{L}_{2}$ such that $\langle a, l\rangle>0$, and scale it by sufficiently large $C>0$ to get a point $C l \in \mathcal{L}_{2}$ so that $\langle a, C l\rangle=C\langle a, l\rangle>\eta$, arriving to a contradiction.

We now continue the proof of Claim 5.4.10. We may assume without loss of generality that

$$
\begin{equation*}
\forall\{x, y\} \in\binom{X}{2}, \quad \text { if } d_{X}(x, y)<s_{2} \text { then } a_{\{x, y\}} \leq 0 . \tag{5.6}
\end{equation*}
$$

To see this, let us zero every such $a_{\{x, y\}}>0$, and denote the resulting point $\hat{a}$. Then for every $l \in \mathcal{L}_{2}$ (which clearly has non-negative coordinates), $\langle\hat{a}, l\rangle \leq\langle a, l\rangle \leq 0$. And for every $l \in \mathcal{K}$, let $\hat{l}$ be equal to $l$ except that we zero the same coordinates where we zero $a$ (which in particular satisfy $d_{X}(x, y)<s_{2}$ ); observe that also $\hat{l} \in \mathcal{K}$, and thus $\langle\hat{a}, l\rangle=\langle a, \hat{l}\rangle>0$. We get that $\hat{a}$ separates $\mathcal{K}$ and $\mathcal{L}_{2}$ and also satisfies (5.6).

Now we define non-negative functions $\widetilde{\mu}_{1}, \widetilde{\mu}_{2}, \widetilde{\mu}_{3}:\binom{X}{2} \rightarrow \mathbb{R}_{+}$as follows:

$$
\begin{aligned}
& \widetilde{\mu}_{1}(\{x, y\})=-a(\{x, y\}) \mathbf{1}_{\left\{d_{X}(x, y) \leq s_{1}\right\}} \\
& \widetilde{\mu}_{2}(\{x, y\})=a(\{x, y\}) \mathbf{1}_{\left\{d_{X}(x, y) \geq s_{2} \text { and } a_{x, y} \geq 0\right\}} \\
& \widetilde{\mu}_{3}(\{x, y\})=-a(\{x, y\}) \mathbf{1}_{\left\{d_{X}(x, y)>s_{1} \text { and } a_{x, y}<0\right\}} .
\end{aligned}
$$

By (5.6), these $\widetilde{\mu}_{i}$ "cover" all cases, i.e.,

$$
\forall\{x, y\} \in\binom{X}{2}, \quad a(\{x, y\})=-\widetilde{\mu}_{1}(\{x, y\})+\widetilde{\mu}_{2}(\{x, y\})-\widetilde{\mu}_{3}(\{x, y\}) .
$$

For $i \in\{1,2,3\}$ define $\lambda_{i}=\sum_{\{x, y\}} \widetilde{\mu}_{i}(\{x, y\})$ and $\mu_{i}(\{x, y\})=\widetilde{\mu}_{i}(\{x, y\}) / \lambda_{i}$. We argue that $\mu_{1}$ and $\mu_{2}$ are as required by Claim 5.4.10, and indeed the only non-trivial property to check is the second item. From the condition that $\langle a, l\rangle \leq 0$ for every $l \in \mathcal{L}_{2}$ we get that for every map $f: X \rightarrow \ell_{2}$,

$$
\begin{aligned}
0 \geq \sum_{\{x, y\}} a(\{x, y\}) & \cdot\|f(x)-f(y)\|^{2} \\
& =\sum_{\{x, y\}}\left[-\widetilde{\mu}_{1}(\{x, y\})+\widetilde{\mu}_{2}(\{x, y\})-\widetilde{\mu}_{3}(\{x, y\})\right] \cdot\|f(x)-f(y)\|^{2},
\end{aligned}
$$

which, in turn, implies

$$
\begin{equation*}
\lambda_{1} \cdot \underset{(x, y) \sim \mu_{1}}{\mathbb{E}}\|f(x)-f(y)\|^{2} \geq \lambda_{2} \cdot \underset{(x, y) \sim \mu_{2}}{\mathbb{E}}\|f(x)-f(y)\|^{2}-4 \lambda_{3} \cdot \sup _{x}\|f(x)\|^{2} \tag{5.7}
\end{equation*}
$$

Consider the point $l \in \mathcal{K}$ with value $\tau_{1}^{2}$ on $\operatorname{supp}\left(\mu_{1}\right)$, value $\tau_{2}^{2}$ on $\operatorname{supp}\left(\mu_{2}\right)$, value $\tau_{3}^{2}$ on $\operatorname{supp}\left(\mu_{3}\right)$, and 0 otherwise; the condition $\langle a, l\rangle>0$ gives

$$
-\lambda_{1} \tau_{1}^{2}+\lambda_{2} \tau_{2}^{2}-\lambda_{3} \tau_{3}^{2}>0
$$

which implies $\lambda_{1}<\lambda_{2} \cdot \tau_{2}^{2} / \tau_{1}^{2}$ and $\lambda_{3}<\lambda_{2} \cdot \tau_{2}^{2} / \tau_{3}^{2}$ (in particular, $\lambda_{2}>0$ ). Plugging into (5.7), we get the inequality required for Claim 5.4.10.

We are now ready to prove Lemma 5.4.6.

Proof of Lemma 5.4.6. Consider a metric space $X$, for which there is no threshold map with parameters $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2},+\infty\right)$. We prove that this implies the Poincaré inequality (5.1).

Indeed, $X$ has no $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2}, \tau_{3}\right)$-threshold map to a Hilbert space for any finite value $\tau_{3}$. We set $\tau_{3}>\tau_{2}$ be sufficiently large so that $\left(2 \tau_{1} / \tau_{3}\right)^{2}<\delta$. Then, by Claim 5.4.7 there exists a finite subset $X^{\prime} \subset X$ that has no $\left(s_{1}, s_{2}, \tau_{1}, \tau_{2}, \tau_{3}\right)$-threshold map to a Hilbert space (which without loss of generality can be chosen to be $\ell_{2}$, since $X^{\prime}$ is finite). Now, using Claim 5.4.10, we obtain finitely-supported probability measures $\mu_{1}$ and $\mu_{2}$, which satisfy (5.3). This concludes the proof of Lemma 5.4.6, since its statement only considers $f$ such that the image of $f$ is the unit ball of $\ell_{2}$, and, thus, $\sup _{x \in X}\|f(x)\|^{2} \leq 1$. Note that the measures $\mu_{1}, \mu_{2}$ depend on the value of $\tau_{3}$ (and, as a result, on $\delta$ ).

### 5.4.3 Threshold maps imply uniform embeddings

We now prove that threshold embeddings imply uniform embeddings, formalized as follows.

Theorem 5.4.12. Suppose that $X$ is a finite-dimensional normed space such that there exists a map to a Hilbert space, which is (1, D, $\left.\tau_{1}, \tau_{2},+\infty\right)$-threshold for some $D>1$ and for some $0<\tau_{1}<\tau_{2}$ with $\tau_{2}>8 \tau_{1}$. Then there exists a map $h$ of $X$ into a Hilbert space such that for every $x_{1}, x_{2} \in X$,

$$
\begin{equation*}
\left(\tau_{2}^{1 / 2}-\left(8 \tau_{1}\right)^{1 / 2}\right) \cdot \min \left\{1, \frac{\left\|x_{1}-x_{2}\right\|}{2 D+4}\right\} \leq\left\|h\left(x_{1}\right)-h\left(x_{2}\right)\right\| \leq\left(2 \tau_{1}\left\|x_{1}-x_{2}\right\|\right)^{1 / 2} \tag{5.8}
\end{equation*}
$$

In particular, $h$ is a uniform embedding of $X$ into a Hilbert space with moduli that depend only on $\tau_{1}, \tau_{2}$ and $D$.

Let us point out that in [94, 151], Johnson and Randrianarivony prove that for a Banach space coarse embeddability into a Hilbert space is equivalent to uniform embeddability. Our definition of a threshold map is weaker than that of a coarse embedding (for the latter see [94] say), but we show that we can adapt the proof of
$[94,151]$ to our setting as well (at least whenever the gap between $\tau_{1}$ and $\tau_{2}$ is large enough). Since we only need one direction of the equivalence, we present a part of the argument from [94] with one (seemingly new) addition: Claim 5.4.19. The resulting proof is arguably simpler than the combination of [94] and [151], and yields a clean quantitative bound (5.8).

Intuition. Let us provide some very high-level intuition of the proof of Theorem 5.4.12. We start with a threshold map $f$ from $X$ to a Hilbert space. First, we show that $f$ is Lipschitz on pairs of points that are sufficiently far. In particular, $f$, restricted on a sufficiently crude net $N$ of $X$, is Lipschitz. This allows us to use a certain extension theorem to extend the restriction of $f$ on $N$ to a Lipschitz function on the whole $X$, while preserving the property that $f$ does not contract too much distances that are sufficiently large. Then, we get a required uniform embedding by performing a certain symmetrization step.

The actual proof is different in a number of details; in particular, instead of being Lipschitz the actual property we will be trying to preserve is different.

Useful facts. To prove Theorem 5.4.12, we need the following three results.
Lemma 5.4.13 ([159]). For a set $S$ and a map $f$ from $S$ to a Hilbert space, there exists a map $g$ from $S$ to a Hilbert space such that $\left\|g\left(x_{1}\right)-g\left(x_{2}\right)\right\|=\left\|f\left(x_{1}\right)-f\left(x_{2}\right)\right\|^{1 / 2}$ for every $x_{1}, x_{2} \in S$.

Lemma 5.4.14 (essentially Lemma 3.5 from [3]). Suppose that $f$ is a map from an abelian group $G$ to a Hilbert space such that for every $g \in G$ we have

$$
\sup _{g_{1}-g_{2}=g}\left\|f\left(g_{1}\right)-f\left(g_{2}\right)\right\|<+\infty .
$$

Then, there exists a map $f^{\prime}$ from $G$ to a Hilbert space such that $\left\|f^{\prime}\left(g_{1}\right)-f^{\prime}\left(g_{2}\right)\right\|$ depends only on $g_{1}-g_{2}$ and for every $g_{1}, g_{2} \in G$ we have

$$
\begin{equation*}
\inf _{g_{1}^{\prime}-g_{2}^{\prime}=g_{1}-g_{2}}\left\|f\left(g_{1}^{\prime}\right)-f\left(g_{2}^{\prime}\right)\right\| \leq\left\|f^{\prime}\left(g_{1}\right)-f^{\prime}\left(g_{2}\right)\right\| \leq \sup _{g_{1}^{\prime}-g_{2}^{\prime}=g_{1}-g_{2}}\left\|f\left(g_{1}^{\prime}\right)-f\left(g_{2}^{\prime}\right)\right\| \tag{5.9}
\end{equation*}
$$

Proof. This lemma is similar to Lemma 5.2 .7 with one twist: in the statement, we
now have distances instead of dot products. The proof of Lemma 5.2.7 relies on the characterization from Fact 5.2.4. If instead we use Fact 5.2.5, we can reuse the proof of Lemma 3.5 from [3] verbatim to prove the present lemma.

Let us sketch here the symmetrization procedure. Let $B(G)$ be the vector space of bounded functions $h: G \rightarrow \mathbb{R}$. Then, one can show that there exists a finitely additive invariant mean $M: B(G) \rightarrow \mathbb{R}$ : a linear functional such that

- for every $h \in B(G)$ such that $h \geq 0$ one has $M h \geq 0$;
- for every $h \in B(G)$ and $g \in G$ one has $M h=M(x \mapsto h(x+g))$;
- $M(x \mapsto 1)=1$.

The existence of such $M$ is non-trivial and requires the axiom of choice (see, e.g., Theorem 17.5 from [78]).

Let us now consider a map $f$ from the statement of the lemma and consider the kernel $K\left(g_{1}, g_{2}\right)=\left\|f\left(g_{1}\right)-f\left(g_{2}\right)\right\|^{2}$. Let us define a new function $K^{\prime}\left(g_{1}, g_{2}\right)$ as follows:

$$
K^{\prime}\left(g_{1}, g_{2}\right)=M\left(x \mapsto K\left(x+g_{1}-g_{2}, x\right)\right) .
$$

Now we need to check that:

- $K^{\prime}$ is a kernel (that is, it is non-negative and symmetric) and $K^{\prime}(g, g)=0$ for every $g \in G$;
- $K^{\prime}$ is negative-definite (see Definition 5.2.3), assuming that $K$ is negative-definite (which is true by Fact 5.2.5);
- for every $g_{1}, g_{2} \in G$ one has

$$
\inf _{g_{1}^{\prime}-g_{2}^{\prime}=g_{1}-g_{2}}\left\|f\left(g_{1}^{\prime}\right)-f\left(g_{2}^{\prime}\right)\right\|^{2} \leq K^{\prime}\left(g_{1}, g_{2}\right) \leq \sup _{g_{1}^{\prime}-g_{2}^{\prime}=g_{1}-g_{2}}\left\|f\left(g_{1}^{\prime}\right)-f\left(g_{2}^{\prime}\right)\right\|^{2}
$$

assuming (5.9).

This can be done exactly the same way as in the proof of Lemma 3.5 from [3]. Finally, we observe that $K^{\prime}\left(g_{1}, g_{2}\right)$ depends only on $g_{1}-g_{2}$ and via Fact 5.2 .5 gives a map $f^{\prime}$ from $G$ to a Hilbert space with the required properties.

Definition 5.4.15. We say that a map $f: X \rightarrow Y$ between metric spaces is $1 / 2$-Hölder with constant $C$, if for every $x_{1}, x_{2} \in X$ one has $d_{Y}\left(f\left(x_{1}\right), f\left(x_{2}\right)\right) \leq C \cdot d_{X}\left(x_{1}, x_{2}\right)^{1 / 2}$.

Theorem 5.4.16 ([128], see also Theorem 19.1 in [174]).). Let ( $X, d_{X}$ ) be a metric space and let $H$ be a Hilbert space. Suppose that $f: S \rightarrow H$, where $S \subset X$, is a $1 / 2$-Hölder map with a constant $C>0$. Then there exists a map $g: X \rightarrow H$ that coincides with $f$ on $S$ and is 1/2-Hölder with the constant $C$.

We are now ready to prove Theorem 5.4.12.

Proof of Theorem 5.4.12. We prove the theorem via the following sequence of claims. Suppose that $X$ is a finite-dimensional normed space. Let $f$ be a ( $\left.1, D, \tau_{1}, \tau_{2},+\infty\right)$ threshold map to a Hilbert space.

The first claim is well-known and is a variant of Proposition 1.11 from [41].
Claim 5.4.17. For every $x_{1}, x_{2} \in X$ we have

$$
\left\|f\left(x_{1}\right)-f\left(x_{2}\right)\right\| \leq \max \left\{1,2 \cdot\left\|x_{1}-x_{2}\right\|\right\} \cdot \tau_{1} .
$$

Proof. If $\left\|x_{1}-x_{2}\right\| \leq 1$, then $\left\|f\left(x_{1}\right)-f\left(x_{2}\right)\right\| \leq \tau_{1}$, and we are done. Otherwise, let us take $y_{0}, y_{1}, \ldots, y_{l} \in X$ such that $y_{0}=x_{1}, y_{l}=x_{2},\left\|y_{i}-y_{i+1}\right\| \leq 1$ for every $i$, and $l=\left\lceil\left\|x_{1}-x_{2}\right\|\right\rceil$. In particular, we can take $y_{i}=x_{1}+i \cdot \frac{x_{1}-x_{2}}{\left\|x_{1}-x_{2}\right\|}$ for $i=0,1, \ldots l-1$, and $y_{l}=x_{2}$. We have

$$
\left\|f\left(x_{1}\right)-f\left(x_{2}\right)\right\| \leq \sum_{i=0}^{l-1}\left\|f\left(y_{i}\right)-f\left(y_{i+1}\right)\right\| \leq l \tau_{1}=\left\lceil\left\|x_{1}-x_{2}\right\|\right\rceil \cdot \tau_{1} \leq 2\left\|x_{1}-x_{2}\right\| \cdot \tau_{1}
$$

where the first step is by the triangle inequality, the second step follows from $\| y_{i}-$ $y_{i+1} \| \leq 1$, and the last step follows from $\left\|x_{1}-x_{2}\right\| \geq 1$.

The proof of the next claim essentially appears in [94].

Claim 5.4.18. There exists a map $g$ from $X$ to a Hilbert space such that for every $x_{1}, x_{2} \in X$,

- $\left\|g\left(x_{1}\right)-g\left(x_{2}\right)\right\| \leq\left(2 \tau_{1} \cdot\left\|x_{1}-x_{2}\right\|\right)^{1 / 2} ;$
- if $\left\|x_{1}-x_{2}\right\| \geq D+2$, then $\left\|g\left(x_{1}\right)-g\left(x_{2}\right)\right\| \geq \tau_{2}^{1 / 2}-\left(8 \tau_{1}\right)^{1 / 2}$;

Proof. From Claim 5.4.17 and Lemma 5.4.13 we can get a map $g^{\prime}$ from $X$ to a Hilbert space such that for every $x_{1}, x_{2} \in X$

- $\left\|g^{\prime}\left(x_{1}\right)-g^{\prime}\left(x_{2}\right)\right\| \leq \max \left\{1,\left(2\left\|x_{1}-x_{2}\right\|\right)^{1 / 2}\right\} \cdot \tau_{1}^{1 / 2} ;$
- if $\left\|x_{1}-x_{2}\right\| \geq D$, then $\left\|g^{\prime}\left(x_{1}\right)-g^{\prime}\left(x_{2}\right)\right\| \geq \tau_{2}^{1 / 2}$.

Let $N \subset X$ be a 1-net of $X$ such that all the pairwise distances between points in $N$ are more than 1 . The map $g^{\prime}$ is $1 / 2$-Hölder on $N$ with a constant $\left(2 \tau_{1}\right)^{1 / 2}$, so we can apply Theorem 5.4.16 and get a map $g$ that coincides with $g^{\prime}$ on $N$ and is $1 / 2$-Hölder on the whole $X$ with a constant $\left(2 \tau_{1}\right)^{1 / 2}$. That is, for every $x_{1}, x_{2} \in X$ we have

- $\left\|g\left(x_{1}\right)-g\left(x_{2}\right)\right\| \leq\left(2 \tau_{1} \cdot\left\|x_{1}-x_{2}\right\|\right)^{1 / 2} ;$
- if $x_{1} \in N, x_{2} \in N$ and $\left\|x_{1}-x_{2}\right\| \geq D$, then $\left\|g\left(x_{1}\right)-g\left(x_{2}\right)\right\| \geq \tau_{2}^{1 / 2}$.

To conclude that $g$ is as required, let us lower bound $\left\|g\left(x_{1}\right)-g\left(x_{2}\right)\right\|$ whenever $\left\|x_{1}-x_{2}\right\| \geq D+2$. Suppose that $x_{1}, x_{2} \in X$ are such that $\left\|x_{1}-x_{2}\right\| \geq D+2$. Let $u_{1} \in N$ be the closest net point to $x_{1}$ and, similarly, let $u_{2} \in N$ be the closest net point to $x_{2}$. Observe that

$$
\left\|u_{1}-u_{2}\right\| \geq\left\|x_{1}-x_{2}\right\|-\left\|x_{1}-u_{1}\right\|-\left\|x_{2}-u_{2}\right\| \geq(D+2)-1-1=D
$$

We have
$\left\|g\left(x_{1}\right)-g\left(x_{2}\right)\right\| \geq\left\|g\left(u_{1}\right)-g\left(u_{2}\right)\right\|-\left\|g\left(u_{1}\right)-g\left(x_{1}\right)\right\|-\left\|g\left(u_{2}\right)-g\left(x_{2}\right)\right\| \geq \tau_{2}^{1 / 2}-2\left(2 \tau_{1}\right)^{1 / 2}$,
as required, where the second step follows from the inequality $\left\|g\left(u_{1}\right)-g\left(u_{2}\right)\right\| \geq \tau_{2}^{1 / 2}$, which is true, since $u_{1}, u_{2} \in N$, and that $g$ is $1 / 2$-Hölder with a constant $\left(2 \tau_{1}\right)^{1 / 2}$.

The following claim completes the proof of Theorem 5.4.12.
Claim 5.4.19. There exists a map $h$ from $X$ to a Hilbert space such that for every $x_{1}, x_{2} \in X:$

- $\left\|h\left(x_{1}\right)-h\left(x_{2}\right)\right\| \leq\left(2 \tau_{1} \cdot\left\|x_{1}-x_{2}\right\|\right)^{1 / 2} ;$
- $\left\|h\left(x_{1}\right)-h\left(x_{2}\right)\right\| \geq\left(\tau_{2}^{1 / 2}-\left(8 \tau_{1}\right)^{1 / 2}\right) \cdot \min \left\{1,\left\|x_{1}-x_{2}\right\| /(2 D+4)\right\}$.

Proof. We take the map $g$ from Claim 5.4.18 and apply Lemma 5.4.14 to it. Let us call the resulting map $h$. The first desired condition for $h$ follows from a similar condition for $g$ and Lemma 5.4.14. Let us prove the second one.

If $x_{1}=x_{2}$, then there is nothing to prove. If $\left\|x_{1}-x_{2}\right\| \geq D+2$, then by Claim 5.4.18 and Lemma 5.4.14, $\left\|h\left(x_{1}\right)-h\left(x_{2}\right)\right\| \geq \tau_{2}^{1 / 2}-\left(8 \tau_{1}\right)^{1 / 2}$, and we are done. Otherwise, let us consider points $y_{0}, y_{1}, \ldots, y_{l} \in X$ such that $y_{0}=0, y_{i}-y_{i-1}=x_{1}-x_{2}$ for every $i$, and $l=\left\lceil\frac{D+2}{\left\|x_{1}-x_{2}\right\|}\right\rceil$. Since $\left\|y_{l}-y_{0}\right\|=\left\|l\left(x_{1}-x_{2}\right)\right\|=l\left\|x_{1}-x_{2}\right\| \geq D+2$, we have

$$
\begin{aligned}
\tau_{2}^{1 / 2}-\left(8 \tau_{1}\right)^{1 / 2} \leq\left\|h\left(y_{l}\right)-h\left(y_{0}\right)\right\| & \leq \sum_{i=1}^{l}\left\|h\left(y_{i}\right)-h\left(y_{i-1}\right)\right\| \\
& =l \cdot\left\|h\left(x_{1}\right)-h\left(x_{2}\right)\right\| \leq \frac{2 D+4}{\left\|x_{1}-x_{2}\right\|} \cdot\left\|h\left(x_{1}\right)-h\left(x_{2}\right)\right\|,
\end{aligned}
$$

where the equality follows from the conclusion of Lemma 5.4.14.

Finally, observe that Theorem 5.4.12 is merely a reformulation of Claim 5.4.19.

### 5.4.4 Putting it all together

We now show that Theorem 5.4.1 follows by applying Lemma 5.4.3, Theorem 5.4.4, Lemma 5.4.6, and Theorem 5.4.12, in this order, with an appropriate choice of parameters.

Proof of Theorem 5.4.1. Suppose $\operatorname{DTEP}(X, D)$ admits a protocol of size $s$. By setting $k=C s$ in Lemma 5.4.3 ( $C$ is a large absolute constant, to be chosen later), we conclude that $\operatorname{DTEP}\left(\ell_{\infty}^{C s}(X), C s D\right)$ admits a protocol of size $s^{\prime}=O(s)$.

Now choosing $C$ large enough and applying Theorem 5.4.4 (in contrapositive), we conclude that $X$ has no Poincaré inequalities for distance scales 1 and $C s D$, with $\alpha=0.01$ and $\beta=0.001$.

Applying Lemma 5.4.6 (in contrapositive), we conclude that $X$ allows a threshold map to a Hilbert space with parameters $(1, C s D, 1,10,+\infty)$.

Using Theorem 5.4.12, it follows that there is a map $h$ from $X$ to a Hilbert space, such that for all $x_{1}, x_{2} \in X$,

$$
\min \left\{1, \frac{\left\|x_{1}-x_{2}\right\|}{s \cdot D}\right\} \leq\left\|h\left(x_{1}\right)-h\left(x_{2}\right)\right\| \leq K \cdot\left\|x_{1}-x_{2}\right\|^{1 / 2}
$$

where $K>1$ is an absolute constant, and this proves the theorem.

Remark: Instead of applying Lemma 5.4.3 and Theorem 5.4.4, we could have attempted to apply the reduction from [21] to get a threshold map from $X$ to a Hilbert space directly. That approach is much simpler technically, but has two fatal drawbacks. First, we end up with a threshold map with a gap between $\tau_{1}$ and $\tau_{2}$ being arbitrarily close to 1 , and thus, we are unable to invoke Theorem 5.4.12, which requires the gap to be more than 8 . Second, the parameters of the resulting threshold map are exponential in the number of bits in the communication protocol, which is bad for the quantitative bounds from Section 5.5.

### 5.5 Quantitative bounds

In this section we prove the quantitative version of our results, namely Theorem 5.1.2 and Theorem 5.1.3, for which we will reuse Theorem 5.4.1. In particular, we prove the following theorem.

Theorem 5.5.1. For a finite-dimensional normed space $X$ and $\Delta>1$, assume we have a map $f: X \rightarrow H$ to a Hilbert space $H$, such that, for an absolute constant $K>0$ and for every $x_{1}, x_{2} \in X$ :

- $\left\|f\left(x_{1}\right)-f\left(x_{2}\right)\right\|_{H} \leq K \cdot\left\|x_{1}-x_{2}\right\|_{X}^{1 / 2} ;$ and
- if $\left\|x_{1}-x_{2}\right\|_{X} \geq \Delta$, then $\left\|f\left(x_{1}\right)-f\left(x_{2}\right)\right\|_{H} \geq 1$.

Then, for any $\varepsilon \in(0,1 / 3)$, the space $X$ linearly embeds into $\ell_{1-\varepsilon}$ with distortion $O(\Delta / \epsilon)$.

Note that Theorem 5.1.2 now follows from applying Theorem 5.4.1 together with Theorem 5.5.1 for $\Delta=s D$. We can further prove Theorem 5.1.3 by using the following result of Zvavitch from [181].

Lemma 5.5.2 ([181]). Every d-dimensional subspace of $L_{1-\varepsilon}$ embeds linearly into $\ell_{1-\varepsilon}^{d \cdot \text { poly }(\log d)}$ with distortion $O(1)$.

Indeed, applying Lemma 5.5.2 together with Theorem 5.5.1, we get that for every $0<\varepsilon<1 / 3$ the space $X$ linearly embeds into $\ell_{1-\varepsilon}^{\operatorname{poly}(\operatorname{dim} X)}$ with distortion $O(\Delta / \varepsilon)$. Thus, $X$ is embeddable into $\ell_{1}$ with distortion

$$
O\left(\Delta \cdot(\operatorname{dim} X)^{O(\varepsilon)} / \varepsilon\right) .
$$

Setting $\varepsilon=\Theta(1 / \log (\operatorname{dim} X))$, we obtain Theorem 5.1.3.
It remains to prove Theorem 5.5.1. Its proof proceeds by adjusting the arguments from [3] and [138].

Proof of Theorem 5.5.1. Fix $X, \Delta>0$, and the corresponding map $f: X \rightarrow H$. We first prove the following lemma.

Lemma 5.5.3. There exists a probability measure $\mu$ on $\mathbb{R}^{\operatorname{dim} X}$ symmetric around the origin such that its (real-valued) characteristic function $\varphi: X \rightarrow \mathbb{R}$ has the following properties for every $x \in X$ :

- $\varphi(x) \geq e^{-\widetilde{K} \cdot\|x\|_{X}} ;$ and
- if $\|x\|_{X} \geq \Delta$, then $\varphi(x) \leq 1 / e$.

Here $\widetilde{K}>0$ is an absolute constant.
Proof. It is known from [161] that for a Hilbert space $H$ the function $g: h \mapsto e^{-\|h\|_{H}^{2}}$ is positive-definite. Thus, there exists a function $\widetilde{g}: H \rightarrow \widetilde{H}$ to a Hilbert space $\widetilde{H}$ such
that for every $h_{1}, h_{2} \in H$ one has $\left\langle\widetilde{g}\left(h_{1}\right), \widetilde{g}\left(h_{2}\right)\right\rangle_{\tilde{H}}=e^{-\left\|h_{1}-h_{2}\right\|_{H}^{2}}$. Setting $\widetilde{f}=\widetilde{g} \circ f$, we get a function $\widetilde{f}: X \rightarrow \widetilde{H}$ to a Hilbert space such that for an absolute constant $\widetilde{K}>0$ for every $x_{1}, x_{2} \in X$, we have:

- $\left\|\widetilde{f}\left(x_{1}\right)\right\|_{\widetilde{H}}=1$;
- $\left\langle\tilde{f}\left(x_{1}\right), \widetilde{f}\left(x_{2}\right)\right\rangle_{\widetilde{H}} \geq e^{-\widetilde{K} \cdot\left\|x_{1}-x_{2}\right\|_{x}}$; and
- if $\left\|x_{1}-x_{2}\right\|_{X} \geq \Delta$, then $\left\langle\widetilde{f}\left(x_{1}\right), \widetilde{f}\left(x_{2}\right)\right\rangle_{\widetilde{H}} \leq 1 / e$.

Applying Lemma 5.2.7 and Lemma 5.2.4, we obtain a positive-definite function $\varphi: X \rightarrow \mathbb{R}$ such that:

- $\varphi(0)=1$;
- for every $x \in X$ one has $\varphi(x) \geq e^{-\widetilde{K} \cdot\|x\|_{X}}$; and
- if $\|x\|_{X} \geq \Delta$, then $\varphi(x) \leq 1 / e$.

We can now use Bochner's theorem, which is the following characterization of continuous positive-definite functions, via the Fourier transform.

Theorem 5.5.4 (Bochner's theorem, see, e.g., [74]). If a map $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is positivedefinite, continuous in zero, and $f(0)=1$, then there exists a probability measure $\mu$ on $\mathbb{R}^{d}$ such that $f$ is the $\mu$ 's characteristic function. That is, for every $x \in \mathbb{R}^{d}$,

$$
f(x)=\int_{\mathbb{R}^{d}} e^{i\langle x, v\rangle} \mu(d v)
$$

In particular, note that we have that $\varphi(0)=1, \varphi$ is positive-definite and is continuous at zero. Hence, by Bochner's theorem, we get a probability measure $\mu$ over $\mathbb{R}^{\operatorname{dim} X}$ whose characteristic function equals to $\varphi$. That is, for every $x \in X$ we get

$$
\varphi(x)=\int_{\mathbb{R}^{\operatorname{dim} X}} e^{i\langle x, v\rangle} \mu(d v)
$$

where $\langle\cdot, \cdot\rangle$ is the standard dot product in $\mathbb{R}^{\operatorname{dim} X}$. Clearly, $\mu$ is symmetric around the origin, since $\varphi$ is real-valued. This completes the proof of Lemma 5.5.3.

Our next goal is to show that $\mu$ gives rise to a one-measurement linear sketch for $X$ with approximation $O(\Delta)$ and a certain additional property that will be useful to us. The following lemma contains two standard facts about one-dimensional characteristic functions (see, e.g., [143]). We include the proof for completeness.

Lemma 5.5.5. Let $\nu$ be a symmetric probability measure over the real line, and let

$$
\psi(t)=\int_{\mathbb{R}} e^{i v t} \nu(d v)
$$

be its characteristic function (which is real-valued due to the symmetry of $\nu$ ). Then,

- if for some $R>0$ and $0<\varepsilon<1$ we have $|\psi(R)| \leq 1-\varepsilon$, then

$$
\begin{equation*}
\nu\left(\left\{v \in \mathbb{R}:|v| \geq \Omega_{\varepsilon}(1 / R)\right\}\right) \geq \Omega_{\varepsilon}(1) \tag{5.10}
\end{equation*}
$$

- for every $\delta>0$ one has

$$
\begin{equation*}
\nu(\{v \in \mathbb{R}:|v| \geq 1 / \delta\}) \leq O(1 / \delta) \cdot \int_{-\delta}^{\delta}(1-\psi(t)) d t \tag{5.11}
\end{equation*}
$$

Proof. Let us start with proving the first claim. We have for every $\alpha>0$

$$
\begin{aligned}
& 1-\varepsilon \geq|\psi(R)| \geq \int_{\mathbb{R}} \cos (v R) \nu(d v) \\
& \qquad \begin{aligned}
& \geq \cos \alpha \cdot \nu(\{v \in \mathbb{R}:|v R| \leq \alpha\})-\nu(\{v \in \mathbb{R}:|v R|>\alpha\}) \\
&=(1+\cos \alpha) \cdot \nu(\{v \in \mathbb{R}:|v R| \leq \alpha\})-1,
\end{aligned}
\end{aligned}
$$

where the second step uses the fact that $\psi$ is real-valued. Thus, we have

$$
\nu(\{v \in \mathbb{R}:|v R| \leq \alpha\}) \leq \frac{2-\varepsilon}{1+\cos \alpha}
$$

Setting $\alpha=\Theta(\sqrt{\varepsilon})$, we get the desired bound.

Now let us prove the second claim. We have, for every $\delta>0$,

$$
\begin{aligned}
& \int_{-\delta}^{\delta}(1-\psi(t)) d t=\int_{-\delta}^{\delta} \int_{\mathbb{R}}\left(1-e^{i v t}\right) \nu(d v) d t=2 \delta \cdot \int_{\mathbb{R}}\left(1-\frac{\sin (\delta v)}{\delta v}\right) \nu(d v) \\
& \geq 2(1-\sin 1) \cdot \delta \cdot \nu(\{v \in \mathbb{R}:|\delta v| \geq 1\}),
\end{aligned}
$$

where we use that $(1-\sin y / y)>(1-\sin 1)$ for every $y$ such that $|y|>1$.

Now we will show that the probability measure $\mu$ from Lemma 5.5.3 gives a good linear sketch for $X$. To see this we use the conditioning on the characteristic function of $\mu$, namely, we exploit them using the above Lemma 5.5.5. In order to do this, we look at the one-dimensional projections of $\mu$ as follows. Let $x \in X$ be a fixed vector. For a measurable subset $A \subseteq \mathbb{R}$ we define

$$
\nu(A)=\mu\left(\left\{v \in \mathbb{R}^{\operatorname{dim} X}:\langle x, v\rangle \in A\right\}\right)
$$

It is immediate to check that the characteristic function $\psi$ of $\nu$ is as follows: $\psi(t)=$ $\varphi(t \cdot x)$ (recall that $\varphi$ is the characteristic function of $\mu$ ). Next we apply Lemma 5.5.5 to $\psi$ and use the properties of $\varphi$ from the conclusion of Lemma 5.5.3. Namely, we get for every $x \in X$ :

$$
\begin{equation*}
\mu\left(\left\{v \in \mathbb{R}^{\operatorname{dim} X}:|\langle x, v\rangle| \geq \Omega\left(\|x\|_{X} / \Delta\right)\right\}\right)=\Omega(1) \tag{5.12}
\end{equation*}
$$

and for every $t>0$,

$$
\begin{equation*}
\mu\left(\left\{v \in \mathbb{R}^{\operatorname{dim} X}:|\langle x, v\rangle| \geq t \cdot\|x\|_{X}\right\}\right) \leq O(1 / t) \tag{5.13}
\end{equation*}
$$

Indeed, (5.12) follows from the bound $\varphi(x) \leq 1 / e$ whenever $\|x\|_{X} \geq \Delta$ and (5.10). The inequality (5.13) follows from the estimate $\varphi(x) \geq e^{-\widetilde{K} \cdot\|x\|_{X}}$ and (5.11) (for
$\left.1 / \delta=t\|x\|_{X}\right)$ that together give

$$
\begin{aligned}
\mu\left(\left\{v \in \mathbb{R}^{\operatorname{dim} X}:|\langle x, v\rangle| \geq t \cdot\|x\|_{X}\right\}\right) & =\nu\left(\left\{r \in \mathbb{R}:|r| \geq t\|x\|_{X}\right\}\right) \\
& \leq O\left(t\|x\|_{X}\right) \int_{-1 / t\|x\|_{X}}^{1 / t\|x\|_{X}}\left(1-e^{-\tilde{K} s\|x\|_{X}}\right) d s
\end{aligned}
$$

as well as from the inequality

$$
t \int_{-1 / t}^{1 / t}\left(1-e^{-C s}\right) d s \leq t \int_{-1 / t}^{1 / t}(1-(1-C s)) d s=C / t
$$

Hence, $\mu$ gives rise to a one-measurement linear sketch of $X$ with approximation $O(\Delta)$, whose "upper tail" is not too heavy.

Finally, we are ready to describe a desired linear embedding of $X$ into $L_{1-\varepsilon}$; we map $X$ into $L_{1-\varepsilon}(\mu)$ as follows: $x \mapsto(v \mapsto\langle x, v\rangle)$. The following Lemma states that the distortion of this embedding is $O(\Delta / \varepsilon)$, as required.

Lemma 5.5.6. For $0<\varepsilon<1 / 3$ and every $x \in X$,

$$
\Omega\left(\|x\|_{X} / \Delta\right) \leq\|v \mapsto\langle x, v\rangle\|_{L_{1-\varepsilon}(\mu)} \leq O\left(\|x\|_{X} / \varepsilon\right) .
$$

Proof. The lower bound is straightforward:

$$
\|v \mapsto\langle x, v\rangle\|_{L_{1-\varepsilon}(\mu)}^{1-\varepsilon}=\int_{\mathbb{R}^{\operatorname{dim} X}}|\langle x, v\rangle|^{1-\varepsilon} \mu(d v) \geq \Omega(1) \cdot \Omega\left(\|x\|_{X} / \Delta\right)^{1-\varepsilon},
$$

where the last step follows from (5.12).
For the upper bound, we have for every $\alpha>0$,

$$
\begin{aligned}
\|v \mapsto\langle x, v\rangle\|_{L_{1-\varepsilon}(\mu)}^{1-\varepsilon} & =\int_{\mathbb{R}^{\operatorname{dim} X}}|\langle x, v\rangle|^{1-\varepsilon} \mu(d v) \\
& =\int_{0}^{\infty} \mu\left(\left\{v \in \mathbb{R}^{\operatorname{dim} X}:|\langle x, v\rangle|^{1-\varepsilon} \geq s\right\}\right) d s \\
& \leq \alpha+O(1) \cdot\|x\|_{X} \cdot \int_{\alpha}^{\infty} s^{-\frac{1}{1-\varepsilon}} d s \leq \alpha+O(1) \cdot \frac{\|x\|_{X}}{\varepsilon} \cdot \alpha^{-\frac{\varepsilon}{1-\varepsilon}},
\end{aligned}
$$

where the third step follows from (5.13). Choosing $\alpha=\|x\|_{X}^{1-\varepsilon} / \varepsilon^{1-\varepsilon}$, we get

$$
\|v \mapsto\langle x, v\rangle\|_{L_{1-\varepsilon}(\mu)} \leq O(1) \cdot 2^{\frac{1}{1-\varepsilon}} \cdot \frac{\|x\|_{X}}{\varepsilon}
$$

This concludes the proof of Theorem 5.5.1.

### 5.6 Embedding into $\ell_{1}$ via sum-products

Finally, we prove Theorem 5.1.4: good sketches for norms closed under the sum-product imply embeddings into $\ell_{1}$ with constant distortion. First we invoke Theorem 5.4.1 and get a sequence of good uniform embeddings into a Hilbert space, whose moduli depend only on the sketch size and the approximation. Then, we use the main result of this section: Lemma 5.6.1. Before stating the lemma, let us remind a few notions. For a metric space $X$, recall that the metric space $\ell_{1}^{k}(X)=\bigoplus_{\ell_{1}}^{k} X_{n}$ is the direct sum of $k$ copies of $X$, with the associated distance defined as a sum-product ( $\ell_{1}$-product) over the $k$ copies. We define $\ell_{1}(X)$ similarly. We also denote $X \oplus_{\ell_{1}} Y$ the sum-product of $X$ and $Y$.

Lemma 5.6.1. Let $\left(X_{n}\right)_{n=1}^{\infty}$ be a sequence of finite-dimensional normed spaces. Suppose that for every $i_{1}, i_{2} \geq 1$ there exists $m=m\left(i_{1}, i_{2}\right) \geq 1$ such that $X_{i_{1}} \oplus_{\ell_{1}} X_{i_{2}}$ is isometrically embeddable into $X_{m}$. If every $X_{n}$ admits a uniform embedding into a Hilbert space with moduli independent of $n$, then every $X_{n}$ is linearly embeddable into $\ell_{1}$ with distortion independent of $n$.

Note that Theorem 5.1.4 just follows from combining Lemma 5.6.1 with Theorem 5.4.1.

Before proving Lemma 5.6.1, we state the following two useful theorems. The first one (Theorem 5.6.2) follows from the fact that uniform embeddability into a Hilbert space is determined by embeddability of finite subsets [41]. The second one (Theorem 5.6.3) follows by composing results of Aharoni, Maurey, and Mityagin [3] and Kalton [95].

Theorem 5.6.2 (Proposition 8.12 from [41]). Let $A_{1} \subset A_{2} \subset \ldots$ be metric spaces and let $A=\bigcup_{i} A_{i}$. If every $A_{n}$ is uniformly embeddable into a Hilbert space with moduli independent of $n$, then the whole $A$ is uniformly embeddable into a Hilbert space.

Theorem 5.6.3 ([3, 95]). A Banach space $X$ is linearly embeddable into $L_{1}$ iff $\ell_{1}(X)$ is uniformly embeddable into a Hilbert space.

We are now ready to proceed with the proof of Lemma 5.6.1.

Proof of Lemma 5.6.1. Let $X=X_{1} \oplus_{\ell_{1}} X_{2} \oplus_{\ell_{1}} \ldots$. More formally,

$$
X=\left\{\left(x_{1}, x_{2}, \ldots\right): x_{i} \in X_{i}, \sum_{i}\left\|x_{i}\right\|<\infty\right\}
$$

where the norm is defined as follows:

$$
\left\|\left(x_{1}, x_{2}, \ldots\right)\right\|=\sum_{i}\left\|x_{i}\right\|
$$

We claim that the space $\ell_{1}(X)$ embeds uniformly into a Hilbert space. To see this, consider $U_{p}=\ell_{1}^{p}\left(X_{1} \oplus_{\ell_{1}} X_{2} \oplus_{\ell_{1}} \ldots \oplus_{\ell_{1}} X_{p}\right)$, which can be naturally seen as a subspace of $\ell_{1}(X)$. Then, $U_{1} \subset U_{2} \subset \ldots \subset U_{p} \subset \ldots \subset \ell_{1}(X)$ and $\bigcup_{p} U_{p}$ is dense in $\ell_{1}(X)$. By the assumption of the lemma, $U_{p}$ is isometrically embeddable into $X_{m}$ for some $m$, thus, $U_{p}$ is uniformly embeddable into a Hilbert space with moduli independent of $p$. Now, by Theorem 5.6.2, $\bigcup_{p} U_{p}$ is uniformly embeddable into a Hilbert space. Since $\bigcup_{p} U_{p}$ is dense in $\ell_{1}(X)$, the same holds also for the whole $\ell_{1}(X)$, as claimed.

Finally, since $\ell_{1}(X)$ embeds uniformly into a Hilbert space, we can apply Theorem 5.6.3 and conclude that $X$ is linearly embeddable into $L_{1}$. The lemma follows since $X$ contains every $X_{i}$ as a subspace.

### 5.7 Appendix: EMD reduction

Recall that $\mathrm{EMD}_{n}$ is a normed space on all signed measures on $[n]^{2}$ (that sum up to zero). We also take the view that a weighted set in $[n]^{2}$ is in fact a measure on $[n]^{2}$.

Lemma 5.7.1. Suppose the EMD metric between non-negative measures (of the same total measure) admits a sketching algorithm sk with approximation $D>1$ and sketch size $s$. Then the normed space $\mathrm{EMD}_{n}$ admits a sketching algorithm $\mathbf{s k}^{\prime}$ with approximation $D$ and sketch size $O(s)$.

Proof. The main idea is that if $x, y$ are signed measures and we add a sufficiently large term $M>0$ to all of their coordinates, then the resulting vectors $x^{\prime}=x+M \cdot \overrightarrow{1}$ and $y^{\prime}=y+M \cdot \overrightarrow{1}$ are measures (all their coordinates are non-negative) of the same total mass, and $\|x-y\|_{\text {EMD }}$ is equal to the EMD distance between measures $x^{\prime}, y^{\prime}$. The trouble is in identifying a large enough $M$. We use the values of $x$ and $y$ themselves to agree on $M$. Details follow.

Without loss of generality we can fix the DTEP threshold to be $r=1$.
We design the sketch $\mathbf{s k}^{\prime}$ as follows. First choose a hash function $h: \mathbb{N} \rightarrow\{0,1\}^{9}$ (using public randomness). Fix an input $x \in \mathbb{R}^{n^{2}}$ of total measure zero, i.e., $\sum_{i} x_{i}=0$. Let $m(x)=\min _{i} x_{i}$, and let $b(x)$ be the largest multiple of 2 that is smaller than $m(x)$. Since $x$ has total measure zero, $b(x)<m(x) \leq 0$. Now let $b^{(1)}(x)=b(x)$ and $b^{(2)}(x)=b(x)-2$, and then $x^{(q)}=x-b^{(q)}(x) \cdot \overrightarrow{1}$ for $q=1,2$. Notice that in both cases $x^{(q)}>x \geq 0$ (component-wise). Now let the sketch $\mathbf{s k}^{\prime}(x)$ be the concatenation of $\mathbf{s k}\left(x^{(q)}\right), h\left(b^{(q)}(x)\right)$ for $q=1,2$.

The distinguisher works as follows, given two sketches

$$
\mathbf{s k}^{\prime}(x)=\left(\mathbf{s k}\left(x^{(q)}\right), h\left(b^{(q)}(x)\right)\right)_{q=1,2}
$$

and

$$
\mathbf{s k}^{\prime}(y)=\left(\mathbf{s k}\left(y^{(q)}\right), h\left(b^{(q)}(x)\right)\right)_{q=1,2} .
$$

If there are $q_{x}, q_{y} \in\{1,2\}$ whose hashes agree $h\left(b^{(q)}(x)\right)=h\left(b^{(q)}(y)\right)$ (breaking ties arbitrarily if there are multiple possible agreements), then output whatever the EMD metric distinguisher would output on $\operatorname{sk}\left(x^{\left(q_{x}\right)}\right), \mathbf{s k}\left(y^{\left(q_{y}\right)}\right)$. Otherwise output "far" (i.e., that $\left.\|x-y\|_{\text {EMD }}>D\right)$.

To analyze correctness, consider the case when $\|x-y\|_{\text {EMD }} \leq 1$. Without loss of generality, suppose $m(x) \geq m(y)$. Then $m(x)-m(y) \leq 1$ (otherwise $x, y$ are further
away in EMD norm than 1). Hence either $b(x)=b(y)$ or $b(x)=b(y)+2$. Then there exists a corresponding $q \in\{1,2\}$ for which the hashes agree $h\left(b^{(1)}(x)\right)=h\left(b^{(q)}(y)\right)$. By properties of the hash function, with sufficiently large constant probability the hashes match only when the $b$ 's match, in which case the values $q_{x}, q_{y}$ used by the distinguisher satisfy $b^{\left(q_{x}\right)}(x)=b^{\left(q_{y}\right)}(y)$. In this case, $\|x-y\|_{\mathrm{EMD}}=d_{\mathrm{EMD}}\left(x^{\left(q_{x}\right)}, y^{\left(q_{y}\right)}\right)$, and the correctness now depends on sk, and the distinguisher for the EMD metric.

Otherwise, if $\|x-y\|_{\text {EMD }}>D$, either the $b$-values coincide for some $q_{x}, q_{y}$ and then the above argument applies again, or with sufficiently large constant probability the hashes will not agree and the distinguisher outputs (correctly) "far".

There is a small loss in success probability due to use of the hash function, but that can be amplified back by independent repetitions.

Notice that the above lemma assumes a sketching algorithm for the EMD metric between any non-negative measures of the same total measure, and not only in the case where the total measure is 1 . The proof can be easily modified so that any non-negative measure being used always has a fixed total measure (say 1, by simply scaling the inputs), which translates to scaling the threshold $r$ of the DTEP problem. This is acceptable because, under standard definitions, a metric space is called sketchable if it admits a sketching scheme for every threshold $r>0$.

## Chapter 6

## Hardness results for the ANN problem

### 6.1 Introduction

Lower bounds for NNS and ANN have received considerable attention. Such lower bounds are ideally obtained in the cell-probe model [127, 126], where one measures the number of memory cells the query algorithm accesses. Despite a number of success stories, high cell-probe lower bounds are notoriously hard to prove. In fact, there are few techniques for proving high cell-probe lower bounds, for any (static) data structure problem. For ANN in particular, we have no viable techniques to prove $\omega(\log n)$ query time lower bounds. Due to this state of affairs, one may rely on restricted models of computation, which nevertheless capture existing algorithmic approaches.

Early lower bounds for NNS were obtained for data structures in exact or deterministic settings [44,52, 35, 115, 88, 53, 150, 180]. [53, 116] obtained an almost tight cell-probe lower bound for the randomized Approximate Nearest Neighbor Search under the $\ell_{1}$ distance. In that problem, there is no distance threshold $r$, and instead the goal is to find a data point that is not much further than the closest data point. This twist is the main source of hardness, so the result is not applicable to the ANN problem as introduced above.

There are few results that show lower bounds for randomized data structures for

ANN. The first such result [19] shows that any data structure that solves ( $1+\varepsilon, r$ )-ANN for $\ell_{1}$ or $\ell_{2}$ using $t$ cell probes requires space $n^{\Omega\left(1 / t \varepsilon^{2}\right)} .{ }^{1}$ This result shows that the algorithms of $[85,108]$ are tight up to constants in the exponent for $t=O(1)$.

In [146] (following up on [145]), the authors introduce a general framework for proving lower bounds for ANN under any metric. They show that lower bounds for ANN are implied by the robust expansion of the underlying metric space. Using this framework, [146] show that ( $c, r)$-ANN using $t$ cell probes requires space $n^{1+\Omega(1 / t c)}$ for the Manhattan distance and $n^{1+\Omega\left(1 / t c^{2}\right)}$ for the Euclidean distance (for every $c>1$ ).

Lower bounds have also been obtained for other metrics. For the $\ell_{\infty}$ distance, [11] show a lower bound for deterministic ANN data structures. This lower bound was later generalized to randomized data structures [146, 98]. A recent result [2] adapts the framework of [146] to Bregman divergences.

To prove higher lower bounds, researchers resorted to lower bounds for restricted models. These examples include: decision trees [11] (the corresponding upper bound from [81] is in the same model), LSH [129, 140, 17] and data-dependent LSH [27].

### 6.1.1 Our results

We show new cell-probe and restricted lower bounds for $(c, r)$-ANN matching the upper bounds from Chapter 2. All our lower bounds rely on a certain canonical hard distribution for the Hamming space (defined later in Section 6.2.1). Via a standard reduction [114], we obtain similar hardness results for $\ell_{p}$ with $1<p \leq 2$ (with $c$ being replaced by $\left.c^{p}\right)$.

## One cell probe

First, we show a tight lower bound on the space needed to solve ANN for a random instance, for query algorithms that use a single cell probe. More formally, we prove the following theorem:

Theorem 6.1.1 (see Section 6.3.2). Any data structure that:

[^28]- solves ( $c, r$ )-ANN for the random instance defined in Section 6.2.1 with probability at least $2 / 3$,
- operates on memory cells of size $n^{o(1)}$,
- for each query, looks up a single cell,
must use at least $n^{\left(\frac{c}{c-1}\right)^{2}-o(1)}$ words of memory.
The space lower bound matches:
- The upper bound from Chapter 2 for random instances that can be made single-probe;
- The same upper bound for worst-case instances with query time $n^{o(1)}$.

The previous best lower bound from [146] for a single probe are weaker by a polynomial factor.

We prove Theorem 6.1.1 by computing tight bounds on the robust expansion of a hypercube $\{-1,1\}^{d}$ as defined in [146]. Then, we invoke a result from [146], which yields the desired cell probe lower bound. We obtain estimates on the robust expansion via a combination of the hypercontractivity inequality and Hölder's inequality [139]. Equivalently, one could obtain the same bounds by an application of the Generalized Small-Set Expansion Theorem for $\{-1,1\}^{d}$ from [139].

## Two cell probes

To state our results for two cell probes, we first define the decision version of ANN (first introduced in [146]). Suppose that with every data point $p \in P$ we associate a bit $x_{p} \in\{0,1\}$. A new goal is: given a query $q \in\{-1,1\}^{d}$ which is within distance $r$ from a data point $p \in P$, if $P \backslash\{p\}$ is at distance at least $c r$ from $q$, return $x_{p}$ with probability at least $2 / 3$. It is easy to see that any algorithm for $(c, r)$-ANN would solve this decision version.

We prove the following lower bound for data structures making only two cell probes per query.

Theorem 6.1.2 (see Section 6.5). Any data structure that:

- solves the decision ANN for the random instance (Section 6.2.1) with probability $2 / 3$,
- operates on memory cells of size $o(\log n)$,
- accesses at most two cells for each query,
must use at least $n^{\left(\frac{c}{c-1}\right)^{2}-o(1)}$ words of memory.

Informally speaking, Theorem 6.1.2 shows that the second cell probe cannot improve the space bound by more than a subpolynomial factor. To the best of our knowledge, this is the first lower bound on the space of any static data structure problem without a polynomial gap between $t=1$ and $t \geq 2$ cell-probes. Previously, the highest ANN lower bound for two queries was weaker by a polynomial factor [146]. This remains the case even if we plug the tight bound on the robust expansion of a hypercube into the framework of [146]. Thus, in order to obtain a higher lower bound for $t=2$, we must depart from the framework of [146].

Our proof establishes a connection between two-query data structures (for the decision version of ANN), and two-query locally-decodable codes (LDC) [179]. A possibility of such a connection was suggested in [145]. In particular, we show that any data structure violating the lower bound from Theorem 6.1.2 implies a too-good-to-be-true two-query LDC, which contradicts known LDC lower bounds from [101, 39].

The first lower bound for unrestricted two-query LDCs was proved in [101] via a quantum argument. Later, the argument was simplified and made classical in [39]. It turns out that, for our lower bound, we need to resort to the original quantum argument of [101] since it has a better dependence on the noise rate a code is able to tolerate. During the course of our proof, we do not obtain a full-fledged LDC, but rather an object which can be called an $L D C$ on average. For this reason, we are unable to use [101] as a black box but rather adjust their proof to the average case.

Finally, we point out an important difference with Theorem 6.1.1: in Theorem 6.1.2 we allow words to be merely of size $o(\log n)$ (as opposed to $n^{o(1)}$ ). Nevertheless, for
the decision version of ANN for random instances our upper bounds hold even for such "tiny" words. In fact, our techniques do not allow us to handle words of size $\Omega(\log n)$ due to the weakness of known lower bounds for two-query LDC for large alphabets. In particular, our argument can not be pushed beyond word size $2^{\widetilde{\Theta}(\sqrt{\log n})}$ in principle, since this would contradict known constructions of two-query LDCs over large alphabets [71]!

## The general time-space trade-off

Finally, we prove conditional lower bound on the entire time-space trade-off from Chapter 2 matching our upper bounds that up to $n^{o(1)}$ factors. Note that-since we show polynomial query time lower bounds - proving similar lower bounds unconditionally is far beyond the current reach of techniques. Any such statement would constitute a major breakthrough in cell probe lower bounds.

Our lower bounds are proved in the following model, which can be loosely thought of comprising all hashing-based frameworks we are aware of:

Definition 6.1.3. A list-of-points data structure for the ANN problem is defined as follows:

- We fix (possibly random) sets $A_{i} \subseteq\{-1,1\}^{d}$, for $1 \leq i \leq m$; also, with each possible query point $q \in\{-1,1\}^{d}$, we associate a (random) set of indices $I(q) \subseteq[m] ;$
- For a given dataset $P$, we maintain $m$ lists of points $L_{1}, L_{2}, \ldots, L_{m}$, where $L_{i}=P \cap A_{i} ;$
- On query $q$, we scan through each list $L_{i}$ for $i \in I(q)$ and check whether there exists some $p \in L_{i}$ with $\|p-q\|_{1} \leq c r$. If it exists, return $p$.

The total space is defined as $s=m+\sum_{i=1}^{m}\left|L_{i}\right|$ and the query time is $t=|I(q)|+$ $\sum_{i \in I(q)}\left|L_{i}\right|$.

For this model, we prove the following theorem.

Theorem 6.1.4 (see Section 6.4). Consider any list-of-points data structure for the $(c, r)-A N N$ problem for random instances of $n$ points in the d-dimensional Hamming space with $d=\omega(\log n)$, which achieves a total space of $n^{1+\rho_{u}}$, and has query time $n^{\rho_{q}-o(1)}$, for $2 / 3$ success probability. Then it must hold that:

$$
\begin{equation*}
c \sqrt{\rho_{q}}+(c-1) \sqrt{\rho_{u}} \geq \sqrt{2 c-1} \tag{6.1}
\end{equation*}
$$

### 6.2 Preliminaries

We introduce a few techniques and concepts to be used primarily for our lower bounds. We start by defining the approximate nearest neighbor search problem.

Definition 6.2.1. The goal of the $(c, r)$-approximate nearest neighbor problem with failure probability $\delta$ is to construct a data structure over a set of points $P \subset\{-1,1\}^{d}$ supporting the following query: given any point $q$ such that there exists some $p \in P$ with $\|q-p\|_{1} \leq r$, report some $p^{\prime} \in P$ where $\left\|q-p^{\prime}\right\|_{1} \leq c r$ with probability at least $1-\delta$.

### 6.2.1 Random Hamming instances

Here we define a random instance that will be the hard distribution for all the further lower bounds. It corresponds to the $(c, r)$-ANN problem over the Hamming space $\{0,1\}^{d}$ and is similar to Euclidean random instances from Section 2.6 that were crucial for the algorithms from Chapter 2.

- A dataset $P \subset\{0,1\}^{d}$ is given by $n$ points, where each vector is drawn independently and uniformly at random from $\{0,1\}^{d}$. We assume that $d=\omega(\log n)$.
- A query $q \in\{0,1\}^{d}$ is drawn by first choosing a dataset point $p \in P$ uniformly at random, and then choosing $q$ by flipping each bit of $p$ independently with probability $\frac{1}{2 c}$.
- The goal of the data structure is to preprocess $P$ in order to recover the data point $p$ from the query point $q$ quickly.

Any $\left(c-o(1),(1+o(1)) \cdot \frac{d}{2 c}\right)$-ANN data structure for the unit sphere must be able to find $p$ given $q$ with high probability. Indeed, the distance between $p$ and $q$ is at most $(1+o(1)) \cdot \frac{d}{2 c}$ with high probability, while all the distances from $q$ to other data points are at least $(1-o(1)) \cdot \frac{d}{2}$ with high probability, where in both cases we use that $d=\omega(\log n)$.

### 6.2.2 Graphical neighbor search and robust expansion

We introduce a few definitions from [146] to setup the lower bounds for the ANN.
Definition 6.2.2 ([146]). In the Graphical Neighbor Search problem (GNS), we are given a bipartite graph $G=(U, V, E)$ where the dataset comes from $U$ and the queries come from $V$. The dataset consists of pairs $P=\left\{\left(p_{i}, x_{i}\right) \mid p_{i} \in U, x_{i} \in\{0,1\}, i \in[n]\right\}$. On query $q \in V$, if there exists a unique $p_{i}$ with $\left(p_{i}, q\right) \in E$, then we want to return $x_{i}$.

We will sometimes use the GNS problem to prove lower bounds on $(c, r)$-ANN as follows: we build a GNS graph $G$ by taking $U=V=\{-1,1\}^{d}$, and connecting two points $u \in U, v \in V$ iff their Hamming distance most $r$ (see details in [146]). We will also ensure $q$ is not closer than $c r$ to other points apart from the near neighbor.
[146] showed lower bounds for ANN are intimately tied to the following quantity of a metric space.

Definition 6.2.3 (Robust Expansion [146]). Consider a GNS graph $G=(U, V, E)$, and fix a distribution $e$ on $E \subset U \times V$, where $\mu$ is the marginal distribution on $U$ and $\eta$ is the marginal distribution on $V$. For $\delta, \gamma \in(0,1]$, the robust expansion $\Phi_{r}(\delta, \gamma)$ is:

$$
\Phi_{r}(\delta, \gamma)=\min _{A \subset V: \eta(A) \leq \delta} \min _{B \subset U: \frac{e(A X B)}{e(A \times V)} \geq \gamma} \frac{\mu(B)}{\eta(A)}
$$

### 6.2.3 Locally-decodable codes (LDC)

Our 2-probe lower bounds uses results on Locally-Decodable Codes (LDCs). We present the standard definitions and results on LDCs below, although in Section 6.5, we will use a weaker definition of LDCs for our 2-query lower bound.

Definition 6.2.4. A $(t, \delta, \varepsilon)$ locally-decodable code (LDC) encodes $n$-bit strings $x \in\{0,1\}^{n}$ into $m$-bit codewords $C(x) \in\{0,1\}^{m}$ such that, for each $i \in[n]$, the bit $x_{i}$ can be recovered with probability $\frac{1}{2}+\varepsilon$ while making only $t$ queries into $C(x)$, even if the codeword is arbitrarily modified (corrupted) in $\delta m$ bits.

We will use the following lower bound on the size of the LDCs.

Theorem 6.2.5 (Theorem 4 from [101]). If $C:\{0,1\}^{n} \rightarrow\{0,1\}^{m}$ is $a(2, \delta, \varepsilon)-L D C$, then

$$
m \geq 2^{\Omega\left(\delta \varepsilon^{2} n\right)}
$$

### 6.3 One-probe data structures

### 6.3.1 Robust expansion of the Hamming space

The goal of this section is to compute tight bounds for the robust expansion $\Phi_{r}(\delta, \gamma)$ in the Hamming space of dimension $d$, as defined in the preliminaries. We use these bounds for all of our lower bounds in the subsequent sections.

We use the following model for generating dataset points and queries corresponding to the random instance of Section 6.2.1.

Definition 6.3.1. For any $x \in\{-1,1\}^{n}, N_{\sigma}(x)$ is a probability distribution over $\{-1,1\}^{n}$ representing the neighborhood of $x$. We sample $y \sim N_{\sigma}(x)$ by choosing $y_{i} \in\{-1,1\}$ for each coordinate $i \in[d]$ independently; with probability $\sigma$, we set $y_{i}=x_{i}$, and with probability $1-\sigma, y_{i}$ is set uniformly at random.

Given any Boolean function $f:\{-1,1\}^{n} \rightarrow \mathbb{R}$, the function $T_{\sigma} f:\{-1,1\}^{n} \rightarrow \mathbb{R}$ is

$$
\begin{equation*}
T_{\sigma} f(x)=\underset{y \sim N_{\sigma}(x)}{\mathbb{E}}[f(y)] \tag{6.2}
\end{equation*}
$$

In the remainder of this section, will work solely on the Hamming space $V=$ $\{-1,1\}^{d}$. We let

$$
\sigma=1-\frac{1}{c} \quad d=\omega(\log n)
$$

and $\mu$ will refer to the uniform distribution over $V$.
A query is generated as follows: we sample a dataset point $x$ uniformly at random and then generate the query $y$ by sampling $y \sim N_{\sigma}(x)$. From the choice of $\sigma$ and $d, d(x, y) \leq \frac{d}{2 c}(1+o(1))$ with high probability. In addition, for every other point in the dataset $x^{\prime} \neq x$, the pair $\left(x^{\prime}, y\right)$ is distributed as two uniformly random points (even though $y \sim N_{\sigma}(x)$, because $x$ is randomly distributed). Therefore, by taking a union-bound over all dataset points, we can conclude that with high probability, $d\left(x^{\prime}, y\right) \geq \frac{d}{2}(1-o(1))$ for each $x^{\prime} \neq x$.

Given a query $y$ generated as described above, we know there exists a dataset point $x$ whose distance to the query is $d(x, y) \leq \frac{d}{2 c}(1+o(1))$. Every other dataset point lies at a distance $d\left(x^{\prime}, y\right) \geq \frac{d}{2}(1-o(1))$. Therefore, the two distances are a factor of $c-o(1)$ away.

The following lemma is the main result of this section, and we will reference this lemma in subsequent sections.

Lemma 6.3.2 (Robust expansion). Consider the Hamming space equipped with the Hamming norm. For any $p, q \in[1, \infty)$ where $(q-1)(p-1)=\sigma^{2}$, any $\gamma \in[0,1]$, and $m \geq 1$,

$$
\Phi_{r}\left(\frac{1}{m}, \gamma\right) \geq \gamma^{q} m^{1+\frac{q}{p}-q} .
$$

The robust expansion comes from a straight forward application from small-set expansion. In fact, one can easily prove tight bounds on robust expansion via the following lemma:

Theorem 6.3.3 (Generalized Small-Set Expansion Theorem, [139]). Let $0 \leq \sigma \leq 1$. Let $A, B \subset\{-1,1\}^{n}$ have volumes $\exp \left(-\frac{a^{2}}{2}\right)$ and $\exp \left(-\frac{b^{2}}{2}\right)$ and assume $0 \leq \sigma a \leq b \leq$ a. Then

$$
\operatorname{Pr}_{\substack{(x, y) \\ \sigma-\text { correlated }}}[x \in A, y \in B] \leq \exp \left(-\frac{1}{2} \frac{a^{2}-2 \sigma a b+b^{2}}{1-\sigma^{2}}\right)
$$

We compute the robust expansion via an application of the Bonami-Beckner Inequality and Hölder's inequality. This computation gives us more flexibility with
respect to parameters which will become useful in subsequent sections. We now recall the necessary tools.

Theorem 6.3.4 (Bonami-Beckner Inequality [139]). Fix $1 \leq p \leq q$ and $0 \leq \sigma \leq$ $\sqrt{(p-1) /(q-1)}$. Any Boolean function $f:\{-1,1\}^{n} \rightarrow \mathbb{R}$ satisfies

$$
\left\|T_{\sigma} f\right\|_{q} \leq\|f\|_{p}
$$

Theorem 6.3.5 (Hölder's Inequality). Let $f:\{-1,1\}^{n} \rightarrow \mathbb{R}$ and $g:\{-1,1\}^{n} \rightarrow \mathbb{R}$ be arbitrary Boolean functions. Fix $s, t \in[1, \infty)$ where $\frac{1}{s}+\frac{1}{t}=1$. Then

$$
\langle f, g\rangle \leq\|f\|_{s}\|g\|_{t}
$$

We will let $f$ and $g$ be indicator functions for two sets $A$ and $B$, and use a combination of the Bonami-Beckner Inequality and Hölder's Inequality to lower bound the robust expansion. The operator $T_{\sigma}$ applied to $f$ will measure the neighborhood of set $A$. We compute an upper bound on the correlation of the neighborhood of $A$ and $B$ (referred to as $\gamma$ ) with respect to the volumes of $A$ and $B$, and the expression will give a lower bound on robust expansion.

We also need the following lemma.

Lemma 6.3.6. Let $p, q \in[1, \infty)$, where $(p-1)(q-1)=\sigma^{2}$ and $f, g:\{-1,1\}^{d} \rightarrow \mathbb{R}$ be two Boolean functions. Then

$$
\left\langle T_{\sigma} f, g\right\rangle \leq\|f\|_{p}\|g\|_{q} .
$$

Proof. We first apply Hölder's Inequality to split the inner-product into two parts, apply the Bonami-Beckner Inequality to each part.

$$
\left\langle T_{\sigma} f, f\right\rangle=\left\langle T_{\sqrt{\sigma}} f, T_{\sqrt{\sigma}} g\right\rangle \leq\left\|T_{\sqrt{\sigma}} f\right\|_{s}\left\|T_{\sqrt{\sigma}} g\right\|_{t} .
$$

We pick the parameters $s=\frac{p-1}{\sigma}+1$ and $t=\frac{s}{s-1}$, so $\frac{1}{s}+\frac{1}{t}=1$. Note that $p \leq s$
because $\sigma<1$ and $p \geq 1$ because $(p-1)(q-1)=\sigma^{2} \leq \sigma$. We have

$$
q \leq \frac{\sigma}{p-1}+1=t
$$

In addition,

$$
\sqrt{\frac{p-1}{s-1}}=\sqrt{\sigma} \quad \sqrt{\frac{q-1}{t-1}}=\sqrt{(q-1)(s-1)}=\sqrt{\frac{(q-1)(p-1)}{\sigma}}=\sqrt{\sigma} .
$$

We finally apply the Bonami-Beckner Inequality to both norms to obtain

$$
\left\|T_{\sqrt{\sigma}} f\right\|_{s}\left\|T_{\sqrt{\sigma}} g\right\|_{t} \leq\|f\|_{p}\|g\|_{q}
$$

We are now ready to prove Lemma 6.3.2.

Proof of Lemma 6.3.2. We use Lemma 6.3.6 and the definition of robust expansion. For any two sets $A, B \subset V$, let $a=\frac{1}{2^{d}}|A|$ and $b=\frac{1}{2^{d}}|B|$ be the measure of set $A$ and $B$ with respect to the uniform distribution. We refer to $\chi_{A}:\{-1,1\}^{d} \rightarrow\{0,1\}$ and $\chi_{B}:\{-1,1\}^{d} \rightarrow\{0,1\}$ as the indicator functions for $A$ and $B$. Then,

$$
\begin{equation*}
\gamma=\operatorname{Pr}_{x \sim \mu, y \sim N_{\sigma}(x)}[x \in B \mid y \in A]=\frac{1}{a}\left\langle T_{\sigma} \chi_{A}, \chi_{B}\right\rangle \leq a^{\frac{1}{p}-1} b^{\frac{1}{q}} . \tag{6.3}
\end{equation*}
$$

Therefore, $\gamma^{q} a^{q-\frac{q}{p}} \leq b$. Let $A$ and $B$ be the minimizers of $\frac{b}{a}$ satisfying (6.3) and $a \leq \frac{1}{m}$.

$$
\Phi_{r}\left(\frac{1}{m}, \gamma\right)=\frac{b}{a} \geq \gamma^{q} a^{q-\frac{q}{p}-1} \geq \gamma^{q} m^{1+\frac{q}{p}-q} .
$$

### 6.3.2 One-probe data structures

In this section, we prove Theorem 6.1.1. Our proof relies on the main result of [146] for the GNS problem:

Theorem 6.3.7 (Theorem 1.5 [146]). There exists an absolute constant $\gamma$ such that the following holds. Any randomized cell-probe data structure making $t$ probes and using $m$ cells of $w$ bits for a weakly independent instance of GNS which is correct with probability greater than $\frac{1}{2}$ must satisfy

$$
\frac{m^{t} w}{n} \geq \Phi_{r}\left(\frac{1}{m^{t}}, \frac{\gamma}{t}\right)
$$

Proof of Theorem 6.1.1. The lower bound follows from applying Lemma 6.3.2 directly to Theorem 6.3.7. Setting $t=1$ in Theorem 6.3.7, we obtain

$$
m w \geq n \cdot \Phi_{r}\left(\frac{1}{m}, \gamma\right) \geq n \gamma^{q} m^{1+\frac{q}{p}-q}
$$

for some $p, q \in[1, \infty)$ and $(p-1)(q-1)=\sigma^{2}$. Rearranging the inequality and letting $p=1+\frac{\log \log n}{\log n}$, and $q=1+\sigma^{2} \frac{\log n}{\log \log n}$, we obtain

$$
m \geq \frac{\gamma^{\frac{p}{p-1}} n^{\frac{p}{p q-q}}}{w^{\frac{p}{p q-q}}} \geq n^{\frac{1}{\sigma^{2}}-o(1)} .
$$

Since $\sigma=1-\frac{1}{c}$ and $w=n^{o(1)}$, we obtain the desired result.

Corollary 6.3.8. Any 1 cell probe data structure with cell size $n^{o(1)}$ for $c$-approximate nearest neighbors on the sphere in $\ell_{2}$ needs $n^{1+\frac{2 c^{2}-1}{\left(c^{2}-1\right)^{2}}-o(1)}$ many cells.

Proof. Each point in the Hamming space $\{-1,1\}^{d}$ (after scaling by $\frac{1}{\sqrt{d}}$ ) can be thought of as lying on the unit sphere. If two points are a distance $r$ apart in the Hamming space, then they are $2 \sqrt{r}$ apart on the sphere with $\ell_{2}$ norm. Therefore a data structure for a $c^{2}$-approximation on the sphere gives a data structure for a $c$-approximation in the Hamming space.

### 6.4 List-of-points data structures

In this section, we prove Theorem 6.1.4, i.e., a tight lower bound against data structure that fall inside the "list-of-points" model, as defined in Def. 6.1.3.

Theorem 6.4.1 (Restatement of Theorem 6.1.4). Let $D$ be a list-of-points data structure which solves $(c, r)-A N N$ for $n$ points in the d-dimensional Hamming space with $d=\omega(\log n)$. Suppose $D$ is specified by a sequence of $m$ sets $\left\{A_{i}\right\}_{i=1}^{m}$ and a procedure for outputting a subset $I(q) \subset[m]$ using expected space $s=n^{1+\rho_{u}}$, and expected query time $n^{\rho_{q}-o(1)}$ with success probability $\frac{2}{3}$. Then

$$
c \sqrt{\rho_{q}}+(c-1) \sqrt{\rho_{u}} \geq \sqrt{2 c-1}
$$

We will prove the lower bound by giving a lower bound on list-of-points data structures which solve the random instance for the Hamming space defined in Section 6.2.1. The dataset consists of $n$ points $\left\{u_{i}\right\}_{i=1}^{n}$ where each $u_{i} \sim V$ drawn uniformly at random, and a query $v$ is drawn from the neighborhood of a random dataset point. Thus, we may assume $D$ is a deterministic data structure.

Fix a data structure $D$, where $A_{i} \subset V$ specifies which dataset points are placed in $L_{i}$. Additionally, we may define $B_{i} \subset V$ which specifies which query points scan $L_{i}$, i.e., $B_{i}=\{v \in V \mid i \in I(v)\}$. Suppose we sample a random dataset point $u \sim V$ and then a random query point $v$ from the neighborhood of $u$. Let

$$
\gamma_{i}=\operatorname{Pr}\left[v \in B_{i} \mid u \in A_{i}\right]
$$

represent the probability that query $v$ scans the list $L_{i}$, conditioned on $u$ being in $L_{i}$. Additionally, we write $s_{i}=\mu\left(A_{i}\right)$ as the normalized size of $A$. The query time for $D$ is given by the following expression:

$$
\begin{align*}
T & =\sum_{i=1}^{m} \chi_{B_{i}}(v)\left(1+\sum_{j=1}^{n} \chi_{A_{i}}\left(u_{j}\right)\right) \\
\mathbb{E}[T] & =\sum_{i=1}^{m} \mu\left(B_{i}\right)+\sum_{i=1}^{m} \gamma_{i} \mu\left(A_{i}\right)+(n-1) \sum_{i=1}^{m} \mu\left(B_{i}\right) \mu\left(A_{i}\right) \\
& \geq \sum_{i=1}^{m} \Phi_{r}\left(s_{i}, \gamma_{i}\right) s_{i}+\sum_{i=1}^{m} s_{i} \gamma_{i}+(n-1) \sum_{i=1}^{m} \Phi_{r}\left(s_{i}, \gamma_{i}\right) s_{i}^{2} \tag{6.4}
\end{align*}
$$

Since the data structure succeeds with probability $\gamma$,

$$
\begin{equation*}
\sum_{i=1}^{m} s_{i} \gamma_{i} \geq \gamma=\operatorname{Pr}_{j \sim[n], v \sim N\left(u_{j}\right)}\left[\exists i \in[m]: v \in B_{i}, u_{j} \in A_{i}\right] \tag{6.5}
\end{equation*}
$$

Additionally, since $D$ uses at most $s$ space,

$$
\begin{equation*}
n \sum_{i=1}^{m} s_{i} \leq O(s) \tag{6.6}
\end{equation*}
$$

Using the two constraints in (6.5) and (6.6), we will use the estimates of robust expansion in order to find a lower bound for (6.4). From Lemma 6.3.2, for any $p, q \in[1, \infty)$ where $(p-1)(q-1)=\sigma^{2}$ where $\sigma=1-\frac{1}{c}$,

$$
\begin{aligned}
\mathbb{E}[T] & \geq \sum_{i=1}^{m} s_{i}^{q-\frac{q}{p}} \gamma_{i}^{q}+(n-1) \sum_{i=1}^{m} s_{i}^{q-\frac{q}{p}+1} \gamma_{i}^{q}+\gamma \\
\gamma & \leq \sum_{i=1}^{m} s_{i} \gamma_{i} \\
O\left(\frac{s}{n}\right) & \geq \sum_{i=1}^{m} s_{i}
\end{aligned}
$$

We set $S=\left\{i \in[m]: s_{i} \neq 0\right\}$ and for $i \in S$, we write $v_{i}=s_{i} \gamma_{i}$. Then

$$
\begin{equation*}
\mathbb{E}[T] \geq \sum_{i \in S} v_{i}^{q}\left(s_{i}^{-\frac{q}{p}}+(n-1) s_{i}^{-\frac{q}{p}+1}\right) \geq \sum_{i \in S}\left(\frac{\gamma}{|S|}\right)^{q}\left(s_{i}^{-\frac{q}{p}}+(n-1) s_{i}^{-\frac{q}{p}+1}\right) \tag{6.7}
\end{equation*}
$$

where we used the fact $q \geq 1$. Consider

$$
\begin{equation*}
F=\sum_{i \in S}\left(s_{i}^{-\frac{q}{p}}+(n-1) s_{i}^{-\frac{q}{p}+1}\right) \tag{6.8}
\end{equation*}
$$

We analyze three cases separately:

- $0<\rho_{u} \leq \frac{1}{2 c-1}$
- $\frac{1}{2 c-1}<\rho_{u} \leq \frac{2 c-1}{(c-1)^{2}}$
- $\rho_{u}=0$.

For the first two cases, we let

$$
\begin{equation*}
q=1-\sigma^{2}+\sigma \beta \quad p=\frac{\beta}{\beta-\sigma} \quad \beta=\sqrt{\frac{1-\sigma^{2}}{\rho_{u}}} \tag{6.9}
\end{equation*}
$$

Since $0<\rho_{u} \leq \frac{2 c-1}{(c-1)^{2}}$, one can verify $\beta>\sigma$ and both $p$ and $q$ are at least 1 .

Lemma 6.4.2. When $\rho_{u} \leq \frac{1}{2 c-1}$, and $s=n^{1+\rho_{u}}$,

$$
\mathbb{E}[T] \geq \Omega\left(n^{\rho_{q}}\right)
$$

where $\rho_{q}$ and $\rho_{u}$ satisfy Equation 6.1.

Proof. In this setting, $p$ and $q$ are constants, and $q \geq p$. Therefore, $\frac{q}{p} \geq 1 . F$ can be viewed as consisting of the contributions of each $s_{i}$ 's in Equation 6.8, constrained by (6.6). One can easily verify that $F$ minimized when $s_{i}=O\left(\frac{s}{n|S|}\right)$, so substituting in (6.7),

$$
\mathbb{E}[T] \geq \Omega\left(\frac{\gamma^{q} S^{-q / p+1} n^{q / p}}{|S|^{q-q / p}}\right) \geq \Omega\left(\gamma^{q} S^{1-q} n^{q / p}\right)
$$

since $q-q / p>0$ and $|S| \leq s$. In addition, $p, q$ and $\gamma$ are constants, and note the fact $s=n^{1+\rho_{u}}$, and (6.9), we let $n^{\rho_{q}}$ be the best query time we can achieve. Combining these facts, along with the lower bound for $\rho_{q}$ in (6.4), we obtain the following relationship between $\rho_{q}$ and $\rho_{u}$ :

$$
\begin{aligned}
\rho_{q} & =\left(1+\rho_{u}\right)(1-q)+\frac{q}{p} \\
& =\left(1+\rho_{u}\right)\left(\sigma^{2}-\sigma \beta\right)+\frac{\left(1-\sigma^{2}+\sigma \beta\right)(\beta-\sigma)}{\beta} \\
& =\left(\sqrt{1-\sigma^{2}}-\sqrt{\rho_{u}} \sigma\right)^{2} \\
& =\left(\frac{\sqrt{2 c-1}}{c}-\sqrt{\rho_{u}} \cdot \frac{(c-1)}{c}\right)^{2} .
\end{aligned}
$$

Lemma 6.4.3. When $\rho_{u}>\frac{1}{2 c-1}$,

$$
\mathbb{E}[T] \geq \Omega\left(n^{\rho_{q}}\right)
$$

where $\rho_{q}$ and $\rho_{u}$ satisfy Equation 6.1.
Proof. We follow a similar pattern to Lemma 6.4.2.

$$
\frac{\partial F}{\partial s_{i}}=\left(-\frac{q}{p}\right) s_{i}^{-\frac{q}{p}-1}+\left(-\frac{q}{p}+1\right)(n-1) s_{i}^{-\frac{q}{p}}
$$

Consider the case when each $\frac{\partial F}{\partial s_{i}}\left(s_{i}\right)=0$, by setting $s_{i}=\frac{q}{(p-q)(n-1)}$. Since $q<p$, this value is positive and $\sum_{i \in S} s_{i} \leq O\left(\frac{m}{n}\right)$ for large enough $n$. Thus, $F$ is minimized at this point, and $\mathbb{E}[T] \geq\left(\frac{\gamma}{|S|}\right)^{q}|S|\left(\frac{q}{(p-q)(n-1)}\right)^{-\frac{q}{p}}$. Since $q \geq 1$ and $|S| \leq s$,

$$
\mathbb{E}[T] \geq\left(\frac{\gamma}{s}\right)^{q} s\left(\frac{q}{(p-q)(n-1)}\right)^{-\frac{q}{p}}
$$

Since $p, q$ and $\gamma$ are constants, $\mathbb{E}[T] \geq \Omega\left(n^{\rho_{q}}\right)$,

$$
\rho_{q}=\left(1+\rho_{u}\right)(1-q)+\frac{q}{p}
$$

which is the same expression for $\rho_{q}$ as in Lemma 6.4.2.
Lemma 6.4.4. When $\rho_{u}=0$ (so $s=O(n)$ ),

$$
\mathbb{E}[T] \geq n^{\rho_{q}-o(1)}
$$

where $\rho_{q}=\frac{2 c-1}{c^{2}}=1-\sigma^{2}$.
Proof. In this case, we let

$$
q=1+\sigma^{2} \cdot \frac{\log n}{\log \log n} \quad p=1+\frac{\log \log n}{\log n}
$$

Since $q>p$, we have

$$
\mathbb{E}[T]=\Omega\left(\gamma^{q} s^{1-q} n^{\frac{q}{p}}\right)=n^{1-\sigma^{2}-o(1)}
$$

which is the desired expression.

### 6.5 Two-probe data structures

In this section, we prove the cell probe lower bound for $t=2$ cell probes stated in Theorem 6.1.2.

We follow the framework in [146] and prove lower bounds for GNS when $U=V$ with measure $\mu$ (see Def. 6.2.2). We assume there is an underlying graph $G$ with vertex set $V$. For any point $p \in V$, we write $p$ 's neighborhood, $N(p)$, as the set of points with an edge incident on $p$ in $G$.

In the 2-probe GNS problem, we are given a dataset $P=\left\{p_{i}\right\}_{i=1}^{n} \subset V$ of $n$ points as well as a bit-string $x \in\{0,1\}^{n}$. The goal is to build a data structure supporting the following types of queries: given a point $q \in V$, if there exists a unique neighbor $p_{i} \in N(q) \cap P$, return $x_{i}$ with probability at least $\frac{2}{3}$ after making two cell-probes.

We let $D$ denote a data structure with $m$ cells of $w$ bits each. $D$ will depend on the dataset $P$ as well as the bit-string $x$. We will prove the following theorem.

Theorem 6.5.1. There exists a constant $\gamma>0$ such that any non-adaptive $G N S$ data structure holding a dataset of $n \geq 1$ points which succeeds with probability $\frac{2}{3}$ using two cell probes and $m$ cells of $w$ bits satisfies

$$
\frac{m \log m \cdot 2^{O(w)}}{n} \geq \Omega\left(\Phi_{r}\left(\frac{1}{m}, \gamma\right)\right) .
$$

Theorem 6.1.2 will follow from Theorem 6.5.1 together with the robust expansion bound from Lemma 6.3.2 for the special case of non-adaptive probes. We will later show how to reduce adaptive algorithms losing a sub-polynomial factor in the space for $w=o(\log n)$ in Section 6.5.6. We now proceed to proving Theorem 6.5.1.

At a high-level, we show that a "too-good-to-be-true", 2-probe data structure implies a weaker notion of 2-query locally-decodable code (LDC) with small noise rate using the same amount of space ${ }^{2}$. Even though our notion of LDC is weaker

[^29]than Def. 6.2.4, we adapt the tools for showing 2-query LDC lower bounds from [101]. These arguments, using quantum information theory, are very robust and work well with the weaker 2-query LDC we construct.

We note that [145] was the first to suggest the connection between ANN and LDCs. This work represents the first concrete connection which gives rise to better lower bounds.

Proof structure. The proof of Theorem 6.5.1 proceeds in six steps.

1. First we use Yao's principle to focus on deterministic non-adaptive data structures for GNS with two cell-probes. We provide distributions over $n$-point datasets $P$, as well as bit-strings $x$ and a query $q$, and assume the existence of a deterministic data structure succeeding with probability at least $\frac{2}{3}$.
2. We simplify the deterministic data structure in order to get "low-contention" data structures. These are data structures which do not rely on any single cell too much (similar to Def. 6.1 in [146]).
3. We use ideas from [146] to understand how queries neighboring particular dataset points probe various cells of the data structure. We fix an $n$-point dataset $P$ with a constant fraction of the points satisfying the following condition: many possible queries in the neighborhood of these points probe disjoint pairs of cells.
4. For the fixed dataset $P$, we show that we can recover a constant fraction of bits of $x$ with significant probability even if we corrupt the contents of some cells.
5. We reduce to data structures with 1-bit words in order to apply the LDC arguments from [101].
6. Finally, we design an LDC with weaker guarantees and use the arguments in [101] to prove lower bounds on the space of the weak LDC.
though there is a slight ambiguity with the data structure notion of query, we say "2-query LDCs" in order to be consistent with the LDC literature.

### 6.5.1 Deterministic data structures

Definition 6.5.2. A non-adaptive randomized algorithm $R$ for the GNS problem making two cell-probes is an algorithm specified by the following two components:

1. A procedure which preprocess a dataset $P=\left\{p_{i}\right\}_{i=1}^{n}$ of $n$ points, as well as a bit-string $x \in\{0,1\}^{n}$ in order to output a data structure $D \in\left(\{0,1\}^{w}\right)^{m}$.
2. An algorithm $R$ that given a query $q$, chooses two indices $(i, j) \in[m]^{2}$ and specifies a function $f_{q}:\{0,1\}^{w} \times\{0,1\}^{w} \rightarrow\{0,1\}$.

We require the data structure $D$ and the algorithm $R$ satisfy

$$
\operatorname{Pr}_{R, D}\left[f_{q}\left(D_{j}, D_{k}\right)=x_{i}\right] \geq \frac{2}{3}
$$

whenever $q \in N\left(p_{i}\right)$ and $p_{i}$ is the unique such neighbor.
Note that the procedure which outputs the data structure does not depend on the query $q$, and that the algorithm $R$ does not depend on the dataset $P$ or bit-string $x$.

Definition 6.5.3. We define the following distributions:

- Let $\mathcal{P}$ be the uniform distribution supported on $n$-point datasets from $V$.
- Let $\mathcal{X}$ be the uniform distribution over $\{0,1\}^{n}$.
- Let $\mathcal{Q}(P)$ be the distribution over queries given by first drawing a dataset point $p \in P$ uniformly at random and then drawing $q \in N(p)$ uniformly at random.

Lemma 6.5.4. Assume $R$ is a non-adaptive randomized algorithm for GNS using two cell-probes. Then, there exists a non-adaptive deterministic algorithm A for GNS using two cell-probes succeeding with probability at least $\frac{2}{3}$ when the dataset $P \sim \mathcal{P}$, the bit-string $x \sim \mathcal{X}$, and $q \sim \mathcal{Q}(P)$.

Proof. We apply Yao's principle to the success probability of the algorithm. By assumption, there exists a distribution over algorithms which can achieve probability of success at least $\frac{2}{3}$ for any single query. Therefore, for the fixed distributions $\mathcal{P}, \mathcal{X}$,
and $\mathcal{Q}$, there exists a deterministic algorithm achieving at least the same success probability.

In order to simplify notation, we let $A^{D}(q)$ denote output of the algorithm $A$. We assume that $A(q)$ outputs a pair of indices $(j, k)$ as well as the function $f_{q}:\{0,1\}^{w} \times$ $\{0,1\}^{w} \rightarrow\{0,1\}$, and thus, we use $A^{D}(q)$ as the output of $f_{q}\left(D_{j}, D_{k}\right)$. For any fixed dataset $P=\left\{p_{i}\right\}_{i=1}^{n}$ and bit-string $x \in\{0,1\}^{n}$,

$$
\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[A^{D}(q)=x_{i}\right]=\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[f_{q}\left(D_{j}, D_{k}\right)=x_{i}\right] .
$$

This notation allows us to succinctly state the probability of correctness when the query is a neighbor of $p_{i}$.

For the remainder of the section, we let $A$ denote a non-adaptive deterministic algorithm succeeding with probability at least $\frac{2}{3}$ using $m$ cells of width $w$. The success probability is taken over the random choice of the dataset $P \sim \mathcal{P}, x \sim \mathcal{X}$ and $q \sim \mathcal{Q}(P)$.

### 6.5.2 Making low-contention data structures

For any $t \in\{1,2\}$ and $j \in[m]$, let $A_{t, j}$ be the set of queries which probe cell $j$ at the $t$-th probe of algorithm $A$. Since $A$ is deterministic, the indices $(i, j) \in[m]^{2}$ which $A$ outputs are completely determined by two collections $\mathcal{A}_{1}=\left\{A_{1, j}\right\}_{j \in[m]}$ and $\mathcal{A}_{2}=\left\{A_{2, j}\right\}_{j \in[m]}$ which independently partition the query space $V$. On query $q$, if $q \in A_{1, i}$ and $q \in A_{2, j}$, algorithm $A$ outputs the indices $(i, j)$.

We now define the notion of low-contention data structures, which requires the data structure to not rely on any one particular cell too much by ensuring no $A_{t, j}$ is too large.

Definition 6.5.5. A deterministic non-adaptive algorithm $A$ using $m$ cells has low contention if every set $\mu\left(A_{t, j}\right) \leq \frac{1}{m}$ for $t \in\{1,2\}$ and $j \in[m]$.

We now use the following lemma to argue that up to a small increase in space, a data structure can be made low-contention.

Lemma 6.5.6. Let $A$ be a deterministic non-adaptive algorithm for GNS making two cell-probes using $m$ cells. There exists a deterministic non-adaptive algorithm $A^{\prime}$ for GNS making two cell-probes using 3 m cells which has low contention and succeeds with the same probability.

Proof. Suppose $\mu\left(A_{t, j}\right) \geq \frac{1}{m}$ for some $j \in[m]$. We partition $A_{t, j}$ into enough sets $\left\{A_{t, k}^{(j)}\right\}_{k}$ of measure $\frac{1}{m}$ and at most one set with measure between 0 and $\frac{1}{m}$. For each of set $A_{t, k}^{(j)}$, we make a new cell $j_{k}$ with the same contents as cell $j$. When a query lies inside $A_{t, k}^{(j)}$ the $t$-th probe is made to the new cell $j_{k}$ instead of cell $j$.

We apply the above transformation on all sets with $\mu\left(A_{t, j}\right) \geq \frac{1}{m}$. In the resulting data structure, in each partition $\mathcal{A}_{1}$ and $\mathcal{A}_{2}$, there can be at most $m$ cells of measure $\frac{1}{m}$ and at most $m$ sets with measure less than $\frac{1}{m}$. Therefore, the transformed data structure has at most $3 m$ cells. Since the contents remain the same, the data structure succeeds with the same probability.

Given Lemma 6.5.6, we assume that $A$ is a deterministic non-adaptive algorithm for GNS with two cell-probes using $m$ cells which has low contention. The extra factor of 3 in the number of cells is absorbed in the asymptotic notation.

### 6.5.3 Datasets which shatter

We fix some $\gamma>0$ to be a sufficiently small constant.

Definition 6.5.7 (Weak-shattering [146]). We say a partition $A_{1}, \ldots, A_{m}$ of $V(K, \gamma)$ weakly shatters a point $p$ if

$$
\sum_{i \in[m]}\left(\mu\left(A_{i} \cap N(p)\right)-\frac{1}{K}\right)^{+} \leq \gamma
$$

where the operator $(\cdot)^{+}: \mathbb{R} \rightarrow \mathbb{R}^{+}$is the identity on positive real numbers and zero otherwise.

Lemma 6.5.8 (Shattering [146]). Let $A_{1}, \ldots, A_{k}$ be a collection of disjoint subsets of
measure at most $\frac{1}{m}$. Then

$$
\underset{p \sim \mu}{\operatorname{Pr}}[p \text { is }(K, \gamma) \text {-weakly shattered }] \geq 1-\gamma
$$

for $K=\Phi_{r}\left(\frac{1}{m}, \frac{\gamma^{2}}{4}\right) \cdot \frac{\gamma^{3}}{16}$.

For the remainder of the section, we let

$$
K=\Phi_{r}\left(\frac{1}{m}, \frac{\gamma^{2}}{4}\right) \cdot \frac{\gamma^{3}}{16} .
$$

We are interested in dataset points which are shattered with respect to the collections $\mathcal{A}_{1}$ and $\mathcal{A}_{2}$. Intuitively, queries which are near-neighbors of these dataset points probe various disjoint cells in the data structure, so their corresponding bit is stored in many cells.

Definition 6.5.9. Let $p \in V$ be a dataset point which is ( $K, \gamma$ )-weakly shattered by $\mathcal{A}_{1}$ and $\mathcal{A}_{2}$. Let $\beta_{1}, \beta_{2} \subset N(p)$ be arbitrary subsets where each $j \in[m]$ satisfies

$$
\mu\left(A_{1, j} \cap N(p) \backslash \beta_{1}\right) \leq \frac{1}{K}
$$

and

$$
\mu\left(A_{2, j} \cap N(p) \backslash \beta_{2}\right) \leq \frac{1}{K}
$$

Since $p$ is $(K, \gamma)$-weakly shattered, we can pick $\beta_{1}$ and $\beta_{2}$ with measure at most $\gamma$ each. We will refer to $\beta(p)=\beta_{1} \cup \beta_{2}$.

For a fixed dataset point $p \in P$, we refer to $\beta(p)$ as the set holding the slack in the shattering of measure at most $2 \gamma$. For a given collection $\mathcal{A}$, let $S(\mathcal{A}, p)$ be the event that the collection $\mathcal{A}(K, \gamma)$-weakly shatters $p$. Note that Lemma 6.5.8 implies that $\operatorname{Pr}_{p \sim \mu}[S(\mathcal{A}, p)] \geq 1-\gamma$.

Lemma 6.5.10. With high probability over the choice of $n$-point dataset, at most $4 \gamma n$ points do not satisfy $S\left(\mathcal{A}_{1}, p\right)$ and $S\left(\mathcal{A}_{2}, p\right)$.

Proof. The expected number of points $p$ which do not satisfy $S\left(\mathcal{A}_{1}, p\right)$ and $S\left(\mathcal{A}_{2}, p\right)$ is at most $2 \gamma n$. Therefore via a Chernoff bound, the probability that more than $4 \gamma n$ points do not satisfy $S\left(\mathcal{A}_{1}, p\right)$ and $S\left(\mathcal{A}_{2}, p\right)$ is at most $\exp \left(-\frac{2 \gamma n}{3}\right)$.

We call a dataset good if there are at most $4 \gamma n$ dataset points which are not ( $K, \gamma$ )-weakly shattered by $\mathcal{A}_{1}$ and $\mathcal{A}_{2}$.

Lemma 6.5.11. There exists a good dataset $P=\left\{p_{i}\right\}_{i=1}^{n}$ where

$$
\operatorname{Pr}_{x \sim \mathcal{X}, q \sim \mathcal{Q}(P)}\left[A^{D}(q)=x_{i}\right] \geq \frac{2}{3}-o(1)
$$

Proof. For any fixed dataset $P=\left\{p_{i}\right\}_{i=1}^{n}$, let

$$
\mathbf{P}=\operatorname{Pr}_{x \sim \mathcal{X}, q \sim Q(p)}\left[A^{D}(q)=x_{i}\right] .
$$

Then,

$$
\begin{aligned}
\frac{2}{3} & \leq \underset{P \sim \mathcal{P}}{\mathbb{E}}[\mathbf{P}] \\
& =(1-o(1)) \cdot \underset{P \sim \mathcal{P}}{\mathbb{E}}[\mathbf{P} \mid P \text { is good }]+o(1) \cdot \underset{P \sim \mathcal{P}}{\mathbb{E}}[\mathbf{P} \mid P \text { is not good }] \\
\frac{2}{3}-o(1) & \leq(1-o(1)) \cdot \underset{P \sim \mathcal{P}}{\mathbb{E}}[\mathbf{P} \mid P \text { is good }] .
\end{aligned}
$$

Therefore, there exists a dataset which is not shattered by at most $4 \gamma n$ and

$$
\operatorname{Pr}_{x \sim \mathcal{X}, q \sim \mathcal{Q}(P)}\left[A^{D}(y)=x_{i}\right] \geq \frac{2}{3}-o(1)
$$

### 6.5.4 Corrupting some cell contents of shattered points

In the rest of the proof, we fix the dataset $P=\left\{p_{i}\right\}_{i=1}^{n}$ satisfying the conditions of Lemma 6.5.11, i.e., $P$ satisfies

$$
\begin{equation*}
\operatorname{Pr}_{x \sim \mathcal{X}, q \sim \mathcal{Q}(P)}\left[A^{D}(q)=x_{i}\right] \geq \frac{2}{3}-o(1) \tag{6.10}
\end{equation*}
$$

We now introduce the notion of corruption of the data structure cells $D$, which parallels the notion of noise in locally-decodable codes. Remember that, after fixing some bit-string $x$, the algorithm $A$ produces some data structure $D \in\left(\{0,1\}^{w}\right)^{m}$.

Definition 6.5.12. We call $D^{\prime} \in\left(\{0,1\}^{w}\right)^{m}$ a corrupted version of $D$ at $k$ cells if $D$ and $D^{\prime}$ differ on at most $k$ cells, i.e., if $\left|\left\{i \in[m]: D_{i} \neq D_{i}^{\prime}\right\}\right| \leq k$.

Definition 6.5.13. For a fixed $x \in\{0,1\}^{n}$, let

$$
\begin{equation*}
c_{x}(i)=\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[A^{D}(q)=x_{i}\right] \tag{6.11}
\end{equation*}
$$

denote the recovery probability of bit $i$. Note that from the definitions of $\mathcal{Q}(P)$, $\mathbb{E}\left[c_{x}(i)\right] \geq \frac{2}{3}-o(1)$, where the expectation is taken over $x \sim \mathcal{X}$ and $i \in[n]$.

In this section, we show there exist a subset $S \subset[n]$ of size $\Omega(n)$ where each $i \in S$ has constant recovery probability averaged over $x \sim \mathcal{X}$, even if the algorithm probes a corrupted version of data structure. We let $\varepsilon>0$ be a sufficiently small constant.

Lemma 6.5.14. Fix a vector $x \in\{0,1\}^{n}$, and let $D \in\left(\{0,1\}^{w}\right)^{m}$ be the data structure that algorithm $A$ produces on dataset $P$ and bit-string $x$. Let $D^{\prime}$ be a corruption of $D$ at $\varepsilon K$ cells. For every $i \in[n]$ where events $S\left(\mathcal{A}_{1}, p_{i}\right)$ and $S\left(\mathcal{A}_{2}, p_{i}\right)$ occur,

$$
\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[A^{D^{\prime}}(q)=x_{i}\right] \geq c_{x}(i)-2 \gamma-2 \varepsilon
$$

Proof. Note that $c_{x}(i)$ represents the probability that algorithm $A$ run on a uniformly chosen query from the neighborhood of $p_{i}$ returns the correct answer, i.e. $x_{i}$. We denote the subset $C_{1} \subset N(p)$ of queries that when run on $A$ return $x_{i}$; so, $\mu\left(C_{1}\right)=c_{x}(i)$ by definition.

By assumption, $p_{i}$ is $(K, \gamma)$-weakly shattered by $\mathcal{A}_{1}$ and $\mathcal{A}_{2}$, so by Def. 6.5.9, we specify some $\beta(p) \subset N(p)$ where $\mu\left(C_{1} \cap \beta(p)\right) \leq \mu(\beta(p)) \leq 2 \gamma$. Let $C_{2}=C_{1} \backslash \beta(p)$, where $\mu\left(C_{2}\right) \geq c_{i}(x)-2 \gamma$. Again, by assumption that $p_{i}$ is $(K, \gamma)$-weakly shattered, each $j \in[m]$ and $t \in\{1,2\}$ satisfy $\mu\left(C_{2} \cap A_{t, j}\right) \leq \frac{1}{K}$. Let $\Delta \subset[m]$ be the set of $\varepsilon K$ cells where $D$ and $D^{\prime}$ differ, and let $C_{3} \subset C_{2}$ be given by

$$
C_{3}=C_{2} \backslash\left(\bigcup_{j \in \Delta}\left(A_{1, j} \cup A_{2, j}\right)\right) .
$$

Thus,

$$
\mu\left(C_{3}\right) \geq \mu\left(C_{2}\right)-\sum_{j \in \Delta}\left(\mu\left(C_{2} \cap A_{1, j}\right)+\mu\left(C_{2} \cap A_{2, j}\right)\right) \geq c_{i}(x)-2 \gamma-2 \varepsilon
$$

If $q \in C_{3}$, then on query $q$, algorithm $A$ probes cells outside of $\Delta$, so $A^{D^{\prime}}(q)=A^{D}(q)=$ $x_{i}$.

Lemma 6.5.15. There exists a set $S \subset[n]$ of size $\Omega(n)$ with the following property. If $i \in S$, then events $S\left(\mathcal{A}_{1}, p_{i}\right)$ and $S\left(\mathcal{A}_{2}, p_{i}\right)$ occur, and

$$
\underset{x \sim \mathcal{X}}{\mathbb{E}}\left[c_{x}(i)\right] \geq \frac{1}{2}+\nu
$$

where $\nu$ is a constant. ${ }^{3}$

Proof. For $i \in[n]$, let $E_{i}$ be the event that $S\left(\mathcal{A}_{1}, p_{i}\right)$ and $S\left(\mathcal{A}_{2}, p_{i}\right)$ occur and $\mathbb{E}_{x \sim \mathcal{X}}\left[c_{x}(i)\right] \geq \frac{1}{2}+\nu$. Additionally, let

$$
\mathbf{P}=\operatorname{Pr}_{i \in[n]}\left[E_{i}\right] .
$$

[^30]We set $S=\left\{i \in[n] \mid E_{i}\right\}$, so it remains to show that $\mathbf{P}=\Omega(1)$. To this end,

$$
\begin{array}{rlr}
\frac{2}{3}-o(1) & \leq \underset{x \sim \mathcal{X}, i \in[n]}{\mathbb{E}}\left[c_{x}(i)\right] & \text { (by Equations } 6.10 \text { and 6.11) } \\
& \leq 4 \gamma+\mathbf{P}+\left(\frac{1}{2}+\nu\right) \cdot(1-\mathbf{P}) & \text { (since } P \text { is good) } \\
\frac{1}{6}-o(1)-4 \gamma-\nu & \leq \mathbf{P} \cdot\left(\frac{1}{2}-\nu\right)
\end{array}
$$

Fix the set $S \subset[n]$ satisfying the conditions of Lemma 6.5.15. We combine Lemma 6.5.14 and Lemma 6.5.15 to obtain the following condition on the dataset.

Lemma 6.5.16. Whenever $i \in S$,

$$
\underset{x \sim \mathcal{X}}{\mathbb{E}}\left[\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[A^{D^{\prime}}(q)=x_{i}\right]\right] \geq \frac{1}{2}+\eta
$$

where $\eta=\nu-2 \gamma-2 \varepsilon$ and $D^{\prime}$ differs from $D$ in $\varepsilon K$ cells.
Proof. Whenever $i \in S, p_{i}$ is $(K, \gamma)$-weakly shattered. By Lemma 6.5.15, $A$ outputs $x_{i}$ with probability $\frac{1}{2}+\nu$ on average when probing the data structure $D$ on input $q \sim N\left(p_{i}\right)$, i.e

$$
\underset{x \sim \mathcal{X}}{\mathbb{E}}\left[\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[A^{D}(q)=x_{i}\right]\right] \geq \frac{1}{2}+\nu
$$

Therefore, from Lemma 6.5 .14 , if $A$ probes $D^{\prime}$ which is a corruption of $D$ in any $\varepsilon K$ cells, $A$ will recover $x_{i}$ with probability at least $\frac{1}{2}+\nu-2 \gamma-2 \varepsilon$ averaged over all $x \sim \mathcal{X}$ where $q \sim N\left(p_{i}\right)$. In other words,

$$
\underset{x \sim \mathcal{X}}{\mathbb{E}}\left[\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[A^{D^{\prime}}(q)=x_{i}\right]\right] \geq \frac{1}{2}+\nu-2 \gamma-2 \varepsilon
$$

Summarizing the results of the section, we conclude with the following theorem.
Theorem 6.5.17. There exists a two-probe algorithm and a subset $S \subseteq[n]$ of size $\Omega(n)$, satisfying the following property. When $i \in S$, we can recover $x_{i}$ with probability
at least $\frac{1}{2}+\eta$ over a random choice of $x \sim \mathcal{X}$, even if we probe a corrupted version of the data structure at $\varepsilon K$ cells.

Proof. We describe how one can recover bit $x_{i}$ from a data structure generated by algorithm $A$. In order to recover $x_{i}$, we generate a random query $q \sim N\left(p_{i}\right)$ and probe the data structure at the cells specified by $A$. From Lemma 6.5.16, there exists a set $S \subset[n]$ of size $\Omega(n)$ for which the described algorithm recovers $x_{i}$ with probability at least $\frac{1}{2}+\eta$, where the probability is taken on average over all possible $x \in\{0,1\}^{n}$.

Since we fixed the dataset $P=\left\{p_{i}\right\}_{i=1}^{n}$ satisfying the conditions of Lemma 6.5.11, we will abuse a bit of notation, and refer to algorithm $A$ as the algorithm which recovers bits of $x$ described in Theorem 6.5.17. We say that $x \in\{0,1\}^{n}$ is an input to algorithm $A$ in order to initialize the data structure with dataset $P=\left\{p_{i}\right\}_{i=1}^{n}$ and $x_{i}$ is the bit associated with $p_{i}$.

### 6.5.5 Decreasing the word size

In order to apply the lower bounds of 2-query locally-decodable codes, we reduce to the case when the word size $w$ is one bit.

Lemma 6.5.18. There exists a deterministic non-adaptive algorithm $A^{\prime}$ which on input $x \in\{0,1\}^{n}$ builds a data structure $D^{\prime}$ using $m \cdot 2^{w}$ cells of 1 bit. For any $i \in S$ as well as any corruption $C$ which differs from $D^{\prime}$ in at most $\varepsilon K$ cells satisfies

$$
\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[A^{\prime C}(q)=x_{i}\right]\right] \geq \frac{1}{2}+\frac{\eta}{2^{2 w}} .
$$

Proof. Given algorithm $A$ which constructs the data structure $D \in\left(\{0,1\}^{w}\right)^{m}$ on input $x \in\{0,1\}^{n}$, construct the following data structure $D^{\prime} \in(\{0,1\})^{m \cdot 2^{w}}$. For each cell $D_{j} \in\{0,1\}^{w}$, make $2^{w}$ cells containing all parities of the $w$ bits in $D_{j}$. This procedure increases the size of the data structure by a factor of $2^{w}$.

Fix $i \in S$ and $q \in N\left(p_{i}\right)$ be a query. If the algorithm $A$ produces a function $f_{q}:\{0,1\}^{w} \times\{0,1\}^{w} \rightarrow\{0,1\}$ which succeeds with probability at least $\frac{1}{2}+\zeta$ over $x \in\{0,1\}^{n}$, then there exists a signed parity on some input bits which equals $f_{q}$ in at
least $\frac{1}{2}+\frac{\zeta}{2^{2 w}}$ inputs $x \in\{0,1\}^{n}$. Let $S_{j}$ be the parity of the bits of cell $j$ and $S_{k}$ be the parity of the bits of cell $k$. Let $f_{q}^{\prime}:\{0,1\} \times\{0,1\} \rightarrow\{0,1\}$ denote the parity or the negation of the parity which equals $f_{q}$ on $\frac{1}{2}+\frac{\zeta}{2^{2 w}}$ possible input strings $x \in\{0,1\}^{n}$.

Algorithm $A^{\prime}$ will evaluate $f_{q}^{\prime}$ at the cell containing the parity of the $S_{j}$ bits in cell $j$ and the parity of $S_{k}$ bits in cell $k$. Let $I_{S_{j}}, I_{S_{k}} \in\left[m \cdot 2^{w}\right]$ be the indices of these cells. If $C^{\prime}$ is a sequence of $m \cdot 2^{w}$ cells which differ in $\varepsilon K$ many cells from $D^{\prime}$, then

$$
\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[f_{q}^{\prime}\left(C_{I_{S_{j}}}, C_{I_{S_{k}}}\right)=x_{i}\right]\right] \geq \frac{1}{2}+\frac{\eta}{2^{2 w}}
$$

whenever $i \in S$.

For the remainder of the section, we will prove a version of Theorem 6.5.1 for algorithms with 1-bit words. Given Lemma 6.5.18, we will modify the space to $m \cdot 2^{w}$ and the probability to $\frac{1}{2}+\frac{\eta}{2^{2 w}}$ to obtain the answer. So for the remainder of the section, assume algorithm $A$ has 1 bit words.

### 6.5.6 Connection to locally-decodable codes

To complete the proof of Theorem 6.5.1, it remains to prove the following lemma.
Lemma 6.5.19. Let $A$ be a non-adaptive deterministic algorithm which makes 2 cell probes to a data structure $D$ of $m$ cells of 1 bit and recover $x_{i}$ with probability $\frac{1}{2}+\eta$ on random input $x \in\{0,1\}^{n}$ even after $\varepsilon K$ cells are corrupted whenever $i \in S$ for some fixed $S$ of size $\Omega(n)$. Then the following must hold:

$$
\frac{m \log m}{n} \geq \Omega\left(\varepsilon K \eta^{2}\right)
$$

The proof of the lemma uses [101] and relies heavily on notions from quantum computing. In particular, quantum information theory applied to LDC lower bounds.

## Crash course in quantum computing

We introduce a few concepts from quantum computing that are necessary in our subsequent arguments. The quantum state of a qubit is described by a unit-length
vector in $\mathbb{C}^{2}$. We write the quantum state as a linear combination of the basis states $\binom{1}{0}=|0\rangle$ and $\binom{0}{1}=|1\rangle$. The quantum state $\alpha=\binom{\alpha_{1}}{\alpha_{2}}$ can be written

$$
|\alpha\rangle=\alpha_{1}|0\rangle+\alpha_{2}|1\rangle
$$

where we refer to $\alpha_{1}$ and $\alpha_{2}$ as amplitudes and $\left|\alpha_{1}\right|^{2}+\left|\alpha_{2}\right|^{2}=1$. The quantum state of an $m$-qubit system is a unit vector in the tensor product $\mathbb{C}^{2} \otimes \cdots \otimes \mathbb{C}^{2}$ of dimension $2^{m}$. The basis states correspond to all $2^{m}$ bit-strings of length $m$. For $j \in\left[2^{m}\right]$, we write $|j\rangle$ as the basis state $\left|j_{1}\right\rangle \otimes\left|j_{2}\right\rangle \otimes \cdots \otimes\left|j_{m}\right\rangle$ where $j=j_{1} j_{2} \ldots j_{m}$ is the binary representation of $j$. We will write the $m$-qubit quantum state $|\phi\rangle$ as unit-vector given by linear combination over all $2^{m}$ basis states. So $|\phi\rangle=\sum_{j \in\left[2^{m}\right]} \phi_{j}|j\rangle$. As a shorthand, $\langle\phi|$ corresponds to the conjugate transpose of a quantum state.

A mixed state $\left\{p_{i},\left|\phi_{i}\right\rangle\right\}$ is a probability distribution over quantum states. In this case, we the quantum system is in state $\left|\phi_{i}\right\rangle$ with probability $p_{i}$. We represent mixed states by a density matrix $\sum p_{i}\left|\phi_{i}\right\rangle\left\langle\phi_{i}\right|$.

A measurement is given by a family of Hermitian positive semi-definite operators which sum to the identity operator. Given a quantum state $|\phi\rangle$ and a measurement corresponding to the family of operators $\left\{M_{i}^{*} M_{i}\right\}_{i}$, the measurement yields outcome $i$ with probability $\| M_{i}|\phi\rangle \|^{2}$ and results in state $\frac{M_{i}|\phi\rangle}{\| M_{i}|\phi\rangle \|}$, where the norm $\|\cdot\|$ is the $\ell_{2}$ norm. We say the measurement makes the observation $M_{i}$.

Finally, a quantum algorithm makes a query to some bit-string $y \in\{0,1\}^{m}$ by starting with the state $|c\rangle|j\rangle$ and returning $(-1)^{c \cdot y_{j}}|c\rangle|j\rangle$. One can think of $c$ as the control qubit taking values 0 or 1 ; if $c=0$, the state remains unchanged by the query, and if $c=1$ the state receives a $(-1)^{y_{j}}$ in its amplitude. The queries may be made in superposition to a state, so the state $\sum_{c \in\{0,1\}, j \in[m]} \alpha_{c j}|c\rangle|j\rangle$ becomes $\sum_{c \in\{0,1\}, j \in[m]}(-1)^{c \cdot y_{j}} \alpha_{c j}|c\rangle|j\rangle$.

## Weak quantum random access codes from GNS algorithms

Definition 6.5.20. $C:\{0,1\}^{n} \rightarrow\{0,1\}^{m}$ is a $(2, \delta, \eta)$-LDC if there exists a randomized decoding algorithm making at most 2 queries to an $m$-bit string $y$ non-adaptively,
and for all $x \in\{0,1\}^{n}, i \in[n]$, and $y \in\{0,1\}^{m}$ where $d(y, C(x)) \leq \delta m$, the algorithm can recover $x_{i}$ from the two queries to $y$ with probability at least $\frac{1}{2}+\eta$.

In their paper, [101] prove the following result about 2-query LDCs.

Theorem 6.5.21 (Theorem 4 in [101]). If $C:\{0,1\}^{n} \rightarrow\{0,1\}^{m}$ is a $(2, \delta, \eta)-L D C$, then $m \geq 2^{\Omega\left(\delta \eta^{2} n\right)}$.

The proof of Theorem 6.5.21 proceeds as follows. They show how to construct a 1-query quantum-LDC from a classical 2 -query LDC. From a 1 -query quantum-LDC, [101] constructs a quantum random access code which encodes $n$-bit strings in $O(\log m)$ qubits. Then they apply a quantum information theory lower bound due to Nayak [135]:

Theorem 6.5.22 (Theorem 2 stated in [101] from Nayak [135]). For any encoding

$$
x \mapsto \rho_{x}
$$

of $n$-bit strings into m-qubit states, such that a quantum algorithm, given query access to $\rho_{x}$, can decode any fixed $x_{i}$ with probability at least $1 / 2+\eta$, it must hold that $m \geq(1-H(1 / 2+\eta)) n$.

Our proof will follow a pattern similar to the proof of Theorem 6.5.21. We assume the existence of a GNS algorithm $A$ which builds a data structure $D \in\{0,1\}^{m}$.

Our algorithm $A$ from Theorem 6.5.17 does not satisfy the strong properties of an LDC, preventing us from applying 6.5.21 directly. However, it does have some LDC-ish guarantees. In particular, we can recover bits in the presence of $\varepsilon K$ corruptions to $D$. In the LDC language, this means that we can tolerate a noise rate of $\delta=\frac{\varepsilon K}{m}$. Additionally, we cannot necessarily recover every coordinate $x_{i}$, but we can recover $x_{i}$ for $i \in S$, where $|S|=\Omega(n)$. Also, our success probability is $\frac{1}{2}+\eta$ over the random choice of $i \in S$ and the random choice of the bit-string $x \in\{0,1\}^{n}$. Our proof follows by adapting the arguments of [101] to this weaker setting.

Lemma 6.5.23. Let $r=\frac{2}{\delta a^{2}}$ where $\delta=\frac{\varepsilon K}{m}$ and $a \leq 1$ is a constant. Let $D$ be the data structure from above (i.e., satisfying the hypothesis of Lemma 6.5.19). Then there exists a quantum algorithm that, starting from the $r(\log m+1)$-qubit state with $r$ copies of $|U(x)\rangle$, where

$$
|U(x)\rangle=\frac{1}{\sqrt{2 m}} \sum_{c \in\{0,1\}, j \in[m]}(-1)^{c \cdot D_{j}}|c\rangle|j\rangle
$$

can recover $x_{i}$ for any $i \in S$ with probability $\frac{1}{2}+\Omega(\eta)$ (over a random choice of $x$ ).

Assuming Lemma 6.5.23, we can complete the proof of Lemma 6.5.19.

Proof of Lemma 6.5.19. The proof is similar to the proof of Theorem 2 of [101]. Let $\rho_{x}$ represent the $s$-qubit system consisting of the $r$ copies of the state $|U(x)\rangle$, where $s=r(\log m+1) ; \rho_{x}$ is an encoding of $x$. Using Lemma 6.5.23, we can assume we have a quantum algorithm that, given $\rho_{x}$, can recover $x_{i}$ for any $i \in S$ with probability $\alpha=\frac{1}{2}+\Omega(\eta)$ over the random choice of $x \in\{0,1\}^{n}$.

We will let $H(A)$ be the Von Neumann entropy of $A$, and $H(A \mid B)$ be the conditional entropy and $H(A: B)$ the mutual information.

Let $X M$ be the $(n+s)$-qubit system

$$
\frac{1}{2^{n}} \sum_{x \in\{0,1\}^{n}}|x\rangle\langle x| \otimes \rho_{x} .
$$

The system corresponds to the uniform superposition of all $2^{n}$ strings concatenated with their encoding $\rho_{x}$. Let $X$ be the first subsystem corresponding to the first $n$ qubits and $M$ be the second subsystem corresponding to the $s$ qubits. We have

$$
\begin{aligned}
H(X M) & =n+\frac{1}{2^{n}} \sum_{x \in\{0,1\}^{n}} H\left(\rho_{x}\right) \geq n=H(X) \\
H(M) & \leq s
\end{aligned}
$$

since $M$ has $s$ qubits. Therefore, the mutual information $H(X: M)=H(X)+$ $H(M)-H(X M) \leq s$. Note that $H(X \mid M) \leq \sum_{i=1}^{n} H\left(X_{i} \mid M\right)$. By Fano's inequality,
if $i \in S$,

$$
H\left(X_{i} \mid M\right) \leq H(\alpha)
$$

where we are using the fact that Fano's inequality works even if we can recover $x_{i}$ with probability $\alpha$ averaged over all $x$ 's. Additionally, if $i \notin S, H\left(X_{i} \mid M\right) \leq 1$. Therefore,

$$
\begin{aligned}
s \geq H(X: M) & =H(X)-H(X \mid M) \\
& \geq H(X)-\sum_{i=1}^{n} H\left(X_{i} \mid M\right) \\
& \geq n-|S| H(\alpha)-(n-|S|) \\
& =|S|(1-H(\alpha)) .
\end{aligned}
$$

Furthermore, $1-H(\alpha) \geq \Omega\left(\eta^{2}\right)$ since, and $|S|=\Omega(n)$, we have

$$
\begin{aligned}
\frac{2 m}{a^{2} \varepsilon K}(\log m+1) & \geq \Omega\left(n \eta^{2}\right) \\
\frac{m \log m}{n} & \geq \Omega\left(\varepsilon K \eta^{2}\right) .
\end{aligned}
$$

It remains to prove Lemma 6.5.23, which we proceed to do in the rest of the section. We first show that we can simulate our GNS algorithm with a 1-query quantum algorithm.

Lemma 6.5.24. Fix an $x \in\{0,1\}^{n}$ and $i \in[n]$. Let $D \in\{0,1\}^{m}$ be the data structure produced by algorithm $A$ on input $x$. Suppose $\operatorname{Pr}_{q \sim N\left(p_{i}\right)}\left[A^{D}(q)=x_{i}\right]=\frac{1}{2}+b$ for $b>0$. Then there exists a quantum algorithm which makes one quantum query (to $D$ ) and succeeds with probability $\frac{1}{2}+\frac{4 b}{7}$ to output $x_{i}$.

Proof. We use the procedure in Lemma 1 of [101] to determine the output algorithm $A$ on input $x$ at index $i$. The procedure simulates two classical queries with one quantum query.

Without loss of generality, all quantum algorithms which make 1-query to $D$ can
be specified in the following manner: there is a quantum state $\left|Q_{i}\right\rangle$, where

$$
\left|Q_{i}\right\rangle=\sum_{c \in\{0,1\}, j \in[m]} \alpha_{c j}|c\rangle|j\rangle
$$

which queries $D$. After querying $D$, the resulting quantum state is $\left|Q_{i}(x)\right\rangle$, where

$$
\left|Q_{i}(x)\right\rangle=\sum_{c \in\{0,1\}, j \in[m]}(-1)^{c \cdot D_{j}} \alpha_{c j}|c\rangle|j\rangle .
$$

There is also a quantum measurement $\{R, I-R\}$ such that, after the algorithm obtains the state $\left|Q_{i}(x)\right\rangle$, it performs the measurement $\{R, I-R\}$. If the algorithm observes $R$, it outputs 1 and if the algorithm observes $I-R$, it outputs 0 .

From Lemma 6.5.24, we know there exist a state $\left|Q_{i}\right\rangle$ and a measurement $\{R, I-R\}$ where if algorithm $A$ succeeds with probability $\frac{1}{2}+\eta$ on random $x \sim\{0,1\}^{n}$, then the quantum algorithm succeeds with probability $\frac{1}{2}+\frac{4 \eta}{7}$ on random $x \sim\{0,1\}^{n}$.

In order to simplify notation, we write $p(\phi)$ as the probability of making observation $R$ from state $|\phi\rangle$. Since $R$ is a positive semi-definite matrix, $R=M^{*} M$ and so $p(\phi)=\| M|\phi\rangle \|^{2}$.

In exactly the same way as [101], we can remove parts of the quantum state $\left|Q_{i}(x)\right\rangle$ where $\alpha_{c j}>\frac{1}{\sqrt{\delta m}}=\frac{1}{\sqrt{\varepsilon K}}$. If we let $L=\left\{(c, j) \left\lvert\, \alpha_{c j} \leq \frac{1}{\sqrt{\varepsilon K}}\right.\right\}$, after keeping only the amplitudes in $L$, we obtain the quantum state $\frac{1}{a}\left|A_{i}(x)\right\rangle$, where

$$
\left|A_{i}(x)\right\rangle=\sum_{(c, j) \in L}(-1)^{c \cdot D_{j}} \alpha_{c j}|c\rangle|j\rangle, \quad a=\sqrt{\sum_{(c, j) \in L} \alpha_{c j}^{2}}
$$

Lemma 6.5.25. Fix $i \in S$. The quantum state $\left|A_{i}(x)\right\rangle$ satisfies

$$
\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[\left.p\left(\frac{1}{a} A_{i}(x)\right) \right\rvert\, x_{i}=1\right]-\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[\left.p\left(\frac{1}{a} A_{i}(x)\right) \right\rvert\, x_{i}=0\right] \geq \frac{8 \eta}{7 a^{2}} .
$$

Proof. Note that since $\left|Q_{i}(x)\right\rangle$ and $\{R, I-R\}$ simulate $A$ and succeed with probability
at least $\frac{1}{2}+\frac{4 \eta}{7}$ on a random $x \in\{0,1\}^{n}$, we have that

$$
\frac{1}{2} \underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[p\left(Q_{i}(x)\right) \mid x_{i}=1\right]+\frac{1}{2} \underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[1-p\left(Q_{i}(x)\right) \mid x_{i}=0\right] \geq \frac{1}{2}+\frac{4 \eta}{7},
$$

which we can simplify to say

$$
\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[p\left(Q_{i}(x)\right) \mid x_{i}=1\right]+\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[p\left(Q_{i}(x)\right) \mid x_{i}=0\right] \geq \frac{8 \eta}{7}
$$

Since $\left|Q_{i}(x)\right\rangle=\left|A_{i}(x)\right\rangle+\left|B_{i}(x)\right\rangle$ and $\left|B_{i}(x)\right\rangle$ contains at most $\varepsilon K$ parts, if all probes to $D$ in $\left|B_{i}(x)\right\rangle$ had corrupted values, the algorithm should still succeed with the same probability on random inputs $x$. Therefore, the following two inequalities hold:

$$
\begin{align*}
& \underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[p\left(A_{i}(x)+B(x)\right) \mid x_{i}=1\right]+\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[p\left(A_{i}(x)+B(x)\right) \mid x_{i}=0\right] \geq \frac{8 \eta}{7}  \tag{6.12}\\
& \underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[p\left(A_{i}(x)-B(x)\right) \mid x_{i}=1\right]+\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[p\left(A_{i}(x)-B(x)\right) \mid x_{i}=0\right] \geq \frac{8 \eta}{7} \tag{6.13}
\end{align*}
$$

Note that $p(\phi \pm \psi)=p(\phi)+p(\psi) \pm(\langle\phi| R|\psi\rangle+\langle\psi| D|\phi\rangle)$ and $p\left(\frac{1}{c} \phi\right)=\frac{p(\phi)}{c^{2}}$. One can verify by averaging the two inequalities (6.12) and (6.13) that we get the desired expression.

Lemma 6.5.26. Fix $i \in S$. There exists a quantum algorithm that starting from the quantum state $\frac{1}{a}\left|A_{i}(x)\right\rangle$, can recover the value of $x_{i}$ with probability $\frac{1}{2}+\frac{2 \eta}{7 a^{2}}$ over random $x \in\{0,1\}^{n}$.

Proof. The algorithm and argument are almost identical to Theorem 3 in [101], we just check that it works under the weaker assumptions. Let

$$
q_{1}=\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[\left.p\left(\frac{1}{a} A_{i}(x)\right) \right\rvert\, x_{i}=1\right] \quad q_{0}=\underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[\left.p\left(\frac{1}{a} A_{i}(x)\right) \right\rvert\, x_{i}=0\right] .
$$

From Lemma 6.5.25, we know $q_{1}-q_{0} \geq \frac{8 \eta}{7 a^{2}}$. In order to simplify notation, let $b=\frac{4 \eta}{7 a^{2}}$. So we want a quantum algorithm which starting from state $\frac{1}{a}\left|A_{i}(x)\right\rangle$ can recover $x_{i}$ with probability $\frac{1}{2}+\frac{b}{2}$ on random $x \in\{0,1\}^{n}$. Assume $q_{1} \geq \frac{1}{2}+b$, since otherwise
$q_{0} \leq \frac{1}{2}-b$ and the same argument will work for 0 and 1 flipped. Also, assume $q_{1}+q_{0} \geq 1$, since otherwise simply outputting 1 on observation $R$ and 0 on observation $I-R$ will work.

The algorithm works in the following way: it outputs 0 with probability $1-\frac{1}{q_{1}+q_{0}}$ and otherwise makes the measurement $\{R, I-R\}$ on state $\frac{1}{a}\left|A_{i}(x)\right\rangle$. If the observation made is $R$, then the algorithm outputs 1 , otherwise, it outputs 0 . The probability of success over random input $x \in\{0,1\}^{n}$ is

$$
\begin{align*}
\underset{x \in\{0,1\}^{n}}{\mathbb{E}} & {[\operatorname{Pr}[\text { returns correctly }]] } \\
& =\frac{1}{2} \underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[\operatorname{Pr}[\text { returns } 1] \mid x_{i}=1\right]+\frac{1}{2} \underset{x \in\{0,1\}^{n}}{\mathbb{E}}\left[\operatorname{Pr}\left[\text { returns 0] } \mid x_{i}=0\right] .\right. \tag{6.14}
\end{align*}
$$

When $x_{i}=1$, the probability the algorithm returns correctly is $(1-q) p\left(\frac{1}{a} A_{i}(x)\right)$ and when $x_{i}=0$, the probability the algorithm returns correctly is $q+(1-q)\left(1-p\left(\frac{1}{a} A_{i}(x)\right)\right)$. So simplifying (6.14),

$$
\underset{x \in\{0,1\}^{n}}{\mathbb{E}}[\operatorname{Pr}[\text { returns correctly }]]=\frac{1}{2}(1-q) q_{1}+\frac{1}{2}\left(q+(1-q)\left(1-q_{0}\right)\right) \geq \frac{1}{2}+\frac{b}{2} .
$$

Now we can finally complete the proof of Lemma 6.5.23.

Proof of Lemma 6.5.23. Again, the proof is exactly the same as the finishing arguments of Theorem 3 in [101], and we simply check the weaker conditions give the desired outcome. On input $i \in[n]$ and access to $r$ copies of the state $|U(x)\rangle$, the algorithm applies the measurement $\left\{M_{i}^{*} M_{i}, I-M_{i}^{*} M_{i}\right\}$ where

$$
M_{i}=\sqrt{\varepsilon K} \sum_{(c, j) \in L} \alpha_{c j}|c, j\rangle\langle c, j| .
$$

This measurement is designed in order to yield the state $\frac{1}{a}\left|A_{i}(x)\right\rangle$ on $|U(x)\rangle$ if the measurement makes the observation $M_{i}^{*} M_{i}$. The fact that the amplitudes of $\left|A_{i}(x)\right\rangle$ are not too large makes $\left\{M_{i}^{*} M_{i}, I-M_{i}^{*} M_{i}\right\}$ a valid measurement.

The probability of observing $M_{i}^{*} M_{i}$ is $\langle U(x)| M_{i}^{*} M_{i}|U(x)\rangle=\frac{\delta a^{2}}{2}$, where we used that $\delta=\frac{\varepsilon K}{m}$. So the algorithm repeatedly applies the measurement until observing outcome $M_{i}^{*} M_{i}$. If it never makes the observation, the algorithm outputs 0 or 1 uniformly at random. If the algorithm does observe $M_{i}^{*} M_{i}$, it runs the output of the algorithm of Lemma 6.5.26. The following simple calculation (done in [101]) gives the desired probability of success on random input,

$$
\begin{aligned}
& \underset{x \in\{0,1\}^{n}}{\mathbb{E}}[\operatorname{Pr}[\text { returns correctly }]] \\
& \geq\left(1-\left(1-\delta a^{2} / 2\right)^{r}\right)\left(\frac{1}{2}+\frac{2 \eta}{7 a^{2}}\right)+\left(1-\delta a^{2} / 2\right)^{r} \cdot \frac{1}{2} \geq \frac{1}{2}+\frac{\eta}{7 a^{2}}
\end{aligned}
$$

## On adaptivity

We can extend our lower bounds from the non-adaptive to the adaptive setting.

Lemma 6.5.27. If there exists a deterministic data structure which queries two memory cells adaptively and succeeds with probability at least $\frac{1}{2}+\eta$, there exists a deterministic data structure which makes the two queries non-adaptively and succeeds with probability at least $\frac{1}{2}+\frac{\eta}{2^{w}}$.

Proof. The algorithm guesses the outcome of the first cell probe and simulates the adaptive algorithm with the guess. After knowing which two probes to make, we probe the data structure non-adaptively. If the algorithm guessed the contents of the first cell-probe correctly, then we output the value of the non-adaptive algorithm. Otherwise, we output a random value. This algorithm is non-adaptive and succeeds with probability at least $\left(1-\frac{1}{2^{w}}\right) \cdot \frac{1}{2}+\frac{1}{2^{w}}\left(\frac{1}{2}+\eta\right)=\frac{1}{2}+\frac{\eta}{2^{w}}$.

Applying Lemma 6.5.27, from an adaptive algorithm succeeding with probability $\frac{2}{3}$, we obtain a non-adaptive algorithm succeeding with probability $\frac{1}{2}+\Omega\left(2^{-w}\right)$. This value is lower than the intended $\frac{2}{3}$, but we may still reduce to a weak LDC, where we require $\gamma=\Theta\left(2^{-w}\right)$, $\varepsilon=\Theta\left(2^{-w}\right)$, and $|S|=\Omega\left(2^{-w} n\right)$. With these minor changes to
the parameters in Subsections 6.5.1 through 6.5.6, one can easily verify

$$
\frac{m \log m \cdot 2^{\Theta(w)}}{n} \geq \Omega\left(\Phi_{r}\left(\frac{1}{m}, \gamma\right)\right)
$$

This inequality yields tight lower bounds (up to sub-polynomial factors) for the Hamming space when $w=o(\log n)$.

In the case of the Hamming space, we can compute robust expansion in a similar fashion to Theorem 6.1.1. In particular, for any $p, q \in[1, \infty)$ where $(p-1)(q-1)=\sigma^{2}$, we have

$$
\begin{aligned}
\frac{m \log m \cdot 2^{O(w)}}{n} & \geq \Omega\left(\gamma^{q} m^{1+q / p-q}\right) \\
m^{q-q / p+o(1)} & \geq n^{1-o(1)} \gamma^{q} \\
m & \geq n^{\frac{1-o(1)}{q-q / p+o(1)}} \gamma^{\frac{q}{q-q / p+o(1)}}=n^{\frac{p}{p q-q}-o(1)} \gamma^{\frac{p}{p-1}-o(1)} .
\end{aligned}
$$

Let $p=1+\frac{w f(n)}{\log n}$ and $q=1+\sigma^{2} \frac{\log n}{w f(n)}$ where we require that $w f(n)=o(\log n)$ and $f(n) \rightarrow \infty$ as $n \rightarrow \infty$. Then,

$$
m \geq n^{\frac{1}{\sigma^{2}}-o(1)} 2^{\frac{\log n}{w f(n)}} \geq n^{\frac{1}{\sigma^{2}}-o(1)}
$$

## Bibliography

[1] Twitter. Available as https://twitter.com/.
[2] Amirali Abdullah and Suresh Venkatasubramanian. A Directed Isoperimetric Inequality with Application to Bregman Near Neighbor Lower Bounds. In Proceedings of the $4^{7}$ th ACM Symposium on the Theory of Computing (STOC '2015), pages 509-518, 2015.
[3] I. Aharoni, B. Maurey, and B. S. Mityagin. Uniform Embeddings of Metric Spaces and of Banach Spaces into Hilbert Spaces. Israel Journal of Mathematics, 52(3):251-265, 1985.
[4] Thomas D. Ahle, Rasmus Pagh, Ilya Razenshteyn, and Francesco Silvestri. On the Complexity of Inner Product Similarity Join. In Proceedings of the 35th ACM Symposium on Principles of Database Systems (PODS '2016), pages 151-164, 2016. Available as arXiv:1510.02824.
[5] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing Graph Structure via Linear Measurements. In Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA '2012), pages 459-467, 2012.
[6] Nir Ailon and Bernard Chazelle. The Fast Johnson-Lindenstrauss Transform and Approximate Nearest Neighbors. SIAM Journal on Computing, 39(1):302-322, 2009.
[7] Nir Ailon and Holger Rauhut. Fast and RIP-Optimal Transforms. Discrete and Computational Geometry, 52(4):780-798, 2014.
[8] Zeyuan Allen-Zhu, Yang Yuan, and Karthik Sridharan. Exploiting the Structure: Stochastic Gradient Methods Using Raw Clusters. In Proceedings of Advances in Neural Information Processing Systems 29 (NIPS '2016), pages 1642-1650, 2016. Available as arXiv:1602.02151.
[9] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. Journal of Computer and System Sciences, 58(1):137-147, 1999.
[10] Alexandr Andoni. Nearest Neighbor Search: the Old, the New, and the Impossible. PhD thesis, MIT, 2009.
[11] Alexandr Andoni, Dorian Croitoru, and Mihai Patrascu. Hardness of Nearest Neighbor under L-infinity. In Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2008), pages 424-433, 2008.
[12] Alexandr Andoni, Khanh Do Ba, Piotr Indyk, and David P. Woodruff. Efficient Sketches for Earth-Mover Distance, with Applications. In Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2009), pages 324-330, 2009.
[13] Alexandr Andoni and Piotr Indyk. Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions. In Proceedings of the 47 th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2006), pages 459-468, 2006.
[14] Alexandr Andoni and Piotr Indyk. Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions. Communications of the ACM, 51(1):117-122, 2008.
[15] Alexandr Andoni, Piotr Indyk, and Robert Krauthgamer. Earth Mover Distance over High-Dimensional Spaces. In Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms (SODA '2009), pages 343-352, 2009.
[16] Alexandr Andoni, Piotr Indyk, and Robert Krauthgamer. Overcoming the $\ell_{1}$ Non-Embeddability Barrier: Algorithms for Product Metrics. In Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms (SODA '2009), pages 865-874, 2009.
[17] Alexandr Andoni, Piotr Indyk, Thijs Laarhoven, Ilya Razenshteyn, and Ludwig Schmidt. Practical and Optimal LSH for Angular Distance. In Proceedings of Advances in Neural Information Processing Systems 28 (NIPS '2015), pages 1225-1233, 2015. Available as arXiv:1509.02897.
[18] Alexandr Andoni, Piotr Indyk, Huy L. Nguyen, and Ilya Razenshteyn. Beyond Locality-Sensitive Hashing. In Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA '2014), pages 1018-1028, 2014. Available as arXiv:1306.1547.
[19] Alexandr Andoni, Piotr Indyk, and Mihai Patrascu. On the Optimality of the Dimensionality Reduction Method. In Proceedings of the 47 th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2006), pages 449-458, 2006.
[20] Alexandr Andoni, T. S. Jayram, and Mihai Patrascu. Lower Bounds for Edit Distance and Product Metrics via Poincaré-Type Inequalities. In Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA '2010), pages 184-192, 2010.
[21] Alexandr Andoni and Robert Krauthgamer. The Computational Hardness of Estimating Edit Distance. SIAM Journal on Computing, 39(6):2398-2429, 2010.
[22] Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Polylogarithmic approximation for edit distance and the asymmetric query complexity. In Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS '2010), pages 377-386, 2010.
[23] Alexandr Andoni, Robert Krauthgamer, and Ilya Razenshteyn. Sketching and Embedding are Equivalent for Norms. In Proceedings of the 47 th ACM Symposium on the Theory of Computing (STOC '2015), pages 479-488, 2015. Available as arXiv:1411.2577.
[24] Alexandr Andoni, Thijs Laarhoven, Ilya Razenshteyn, and Erik Waingarten. Optimal Hashing-based Time-Space Trade-offs for Approximate Near Neighbors. In Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA '2017), 2017. Available as arXiv:1608.03580.
[25] Alexandr Andoni, Huy L. Nguyen, Aleksandar Nikolov, Ilya Razenshteyn, and Erik Waingarten. Approximate Near Neighbors for General Symmetric Norms. In Proceedings of the 49th ACM Symposium on the Theory of Computing (STOC '2017), 2017. Available as arXiv:1611.06222.
[26] Alexandr Andoni and Ilya Razenshteyn. Optimal Data-Dependent Hashing for Approximate Near Neighbors. In Proceedings of the 47 th ACM Symposium on the Theory of Computing (STOC '2015), pages 793-801, 2015. Available as arXiv:1501.01062.
[27] Alexandr Andoni and Ilya Razenshteyn. Tight Lower Bounds for Data-Dependent Locality-Sensitive Hashing. In Proceedings of the 32nd International Symposium on Computational Geometry (SoCG '2016), pages 9:1-9:11, 2016. Available as arXiv:1507.04299.
[28] Alexandr Andoni, Ilya Razenshteyn, and Negev Shekel Nosatzki. LSH Forest: Practical Algorithms Made Theoretical. In Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA '2017), 2017.
[29] Andreas Argyriou, Rina Foygel, and Nathan Srebro. Sparse Prediction with the $k$-Support Norm. In Proceedings of Advances in Neural Information Processing Systems 25 (NIPS '2012), pages 944-952, 2012.
[30] Yonatan Aumann and Yuval Rabani. An $O(\log k)$ Approximate Min-Cut MaxFlow Theorem and Approximation Algorithm. SIAM Journal on Computing, 27(1):291-301, 1998.
[31] Keith Ball. Markov Chains, Riesz Transforms and Lipschitz Maps. Geometric and Functional Analysis, 2(2):137-172, 1992.
[32] Keith Ball. An Elementary Introduction to Modern Convex Geometry, volume 31 of MSRI Publications. Cambridge University Press, 1997.
[33] Ziv Bar-Yossef, T. S. Jayram, Robert Krauthgamer, and Ravi Kumar. Approximating Edit Distance Efficiently. In Proceedings of the 45 th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2004), pages 550-559, 2004.
[34] Ziv Bar-Yossef, T.S. Jayram, Ravi Kumar, and D. Sivakumar. An Information Statistics Approach to Data Stream and Communication Complexity. Journal of Computer and System Sciences, 68(4):702-732, 2004.
[35] Omer Barkol and Yuval Rabani. Tighter Lower Bounds for Nearest Neighbor Search and Related Problems in the Cell Probe Model. Journal of Computer and System Sciences, 64(4):873-896, 2002.
[36] Yair Bartal and Lee-Ad Gottlieb. Approximate Nearest Neighbor Search for $\ell_{p}$-Spaces $(2<p<\infty)$ via Embeddings. Available as arXiv:1512.01775, 2015.
[37] Tugkan Batu, Funda Ergun, and Cenk Sahinalp. Oblivious String Embeddings and Edit Distance Approximation. In Proceedings of the 1 17th ACM-SIAM Symposium on Discrete Algorithms (SODA '2006), pages 792-801, 2006.
[38] Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New Directions in Nearest Neighbor Searching with Applications to Lattice Sieving. In Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms (SODA '2016), pages 10-24, 2016.
[39] Avraham Ben-Aroya, Oded Regev, and Ronald de Wolf. A Hypercontractive Inequality for Matrix-Valued Functions with Applications to Quantum Computing and LDCs. In Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2008), pages 477-486, 2008.
[40] Colin Bennett and Robert Sharpley. Interpolation of Operators, volume 129 of Pure and Applied Mathematics. Academic Press, 1988.
[41] Yoav Benyamini and Joram Lindenstrauss. Geometric Nonlinear Functional Analysis. Vol. 1, volume 48 of American Mathematical Society Colloquium Publications. American Mathematical Society, Providence, RI, 2000.
[42] Jaroslaw Blasiok, Vladimir Braverman, Stephen R. Chestnut, Robert Krauthgamer, and Lin F. Yang. Streaming Symmetric Norms via Measure Concentration. In Proceedings of the 49th ACM Symposium on the Theory of Computing (STOC '2017), 2017. Available as arXiv:1511.01111.
[43] Mariusz Bojarski, Anna Choromanska, Krzysztof Choromanski, Francois Fagan, Cedric Gouy-Pailler, Anne Morvan, Nourhan Sakr, Tamas Sarlos, and Jamal Atif. Structured Adaptive and Random Spinners for Fast Machine Learning

Computations. In Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS '2017), 2017.
[44] Allan Borodin, Rafail Ostrovsky, and Yuval Rabani. Lower Bounds for High Dimensional Nearest Neighbor Search and Related Problems. In Proceedings of the 31st ACM Symposium on the Theory of Computing (STOC '1999), pages 312-321, 1999.
[45] Vladimir Braverman and Rafail Ostrovsky. Zero-One Frequency Laws. In Proceedings of the 42rd ACM Symposium on the Theory of Computing (STOC '2010), pages 281-290, 2010.
[46] Bo Brinkman and Moses Charikar. On the Impossibility of Dimension Reduction in $\ell_{1}$. Journal of the ACM, 52(5):766-788, 2005.
[47] Andrei Z. Broder. On the Resemblance and Containment of Documents. In Proceedings of the Compression and Complexity of Sequences (SEQUENCES '1997), pages 21-29, 1997.
[48] Andrei Z. Broder, Steven C. Glassman, Mark S. Manasse, and Geoffrey Zweig. Syntactic Clustering of the Web. Computer Networks and ISDN Systems, 29(8-13):1157-1166, 1997.
[49] Joshua Brody and Amit Chakrabarti. A Multi-Round Communication Lower Bound for Gap Hamming and Some Consequences. In Proceedings of the 24th Annual IEEE Conference on Computational Complexity (CCC '2009), pages 358-368, 2009.
[50] Joshua Brody, Amit Chakrabarti, Oded Regev, Thomas Vidick, and Ronald de Wolf. Better Gap-Hamming Lower Bounds via Better Round Elimination. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX '2010, and 14th International Workshop, RANDOM '2010, pages 476-489, 2010.
[51] J. Lawrence Carter and Mark N. Wegman. Universal Classes of Hash Functions. Journal of Computer and System Sciences, 18(2):143-154, 1979.
[52] Amit Chakrabarti, Bernard Chazelle, Benjamin Gum, and Alexey Lvov. A Lower Bound on the Complexity of Approximate Nearest-Neighbor Searching on the Hamming Cube. In Proceedings of the 31st ACM Symposium on the Theory of Computing (STOC '1999), pages 305-311, 1999.
[53] Amit Chakrabarti and Oded Regev. An Optimal Randomized Cell Probe Lower Bound for Approximate Nearest Neighbor Searching. SIA M Journal on Computing, 39(5):1919-1940, 2010.
[54] Amit Chakrabarti and Oded Regev. An Optimal Lower Bound on the Communication Complexity of Gap-Hamming-Distance. SIAM Journal on Computing, 41(5):1299-1317, 2012.
[55] Moses Charikar. Similarity Estimation Techniques from Rounding Algorithms. In Proceedings of the 34th ACM Symposium on the Theory of Computing (STOC '2002), pages 380-388, 2002.
[56] Moses Charikar and Robert Krauthgamer. Embedding the Ulam Metric into $\ell_{1}$. Theory of Computing, 2(11):207-224, 2006.
[57] Jeff Cheeger and Bruce Kleiner. Differentiating Maps into $L^{1}$, and the Geometry of BV Functions. Annals of Mathematics, 171(2):1347-1385, 2010.
[58] Jeff Cheeger, Bruce Kleiner, and Assaf Naor. Compression Bounds for Lipschitz Maps from the Heisenberg Group to $L_{1}$. Acta Mathematica, 207(2):291-373, 2011.
[59] Kenneth L. Clarkson. A Randomized Algorithm for Closest-Point Queries. SIAM Journal on Computing, 17(4):830-847, 1988.
[60] Kenneth L. Clarkson and David P. Woodruff. Sketching for $M$-Estimators: A Unified Approach to Robust Regression. In Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms (SODA '2015), pages 921-939, 2015.
[61] Graham Cormode and S. Muthukrishnan. The String Edit Distance Matching Problem with Moves. ACM Transactions on Algorithms, 3(1):2:1-2:19, 2007.
[62] Graham Cormode, S. Muthukrishnan, and Cenk Sahinalp. Permutation Editing and Matching via Embeddings. In Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP '2001), pages 481-492, 2001.
[63] Graham Cormode, Mike Paterson, Suleyman Cenk Sahinalp, and Uzi Vishkin. Communication Complexity of Document Exchange. In Proceedings of the 11th ACM-SIAM Symposium on Discrete Algorithms (SODA '2000), pages 197-206, 2000.
[64] Anirban Dasgupta, Ravi Kumar, and Tamás Sarlós. Fast Locality-Sensitive Hashing. In Proceedings of the 17th ACM SIGKDD Conference on Knowledge Discovery and Data (KDD '2011), pages 1073-1081, 2011.
[65] Sanjoy Dasgupta and Anupam Gupta. An Elementary Proof of a Theorem of Johnson and Lindenstrauss. Random Structures and Algorithms, 22(1):60-65, 2003.
[66] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S. Mirrokni. LocalitySensitive Hashing Scheme Based on p-Stable Distributions. In Proceedings of the 20th ACM Symposium on Computational Geometry (SoCG '2004), pages 253-262, 2004.
[67] Inderjit S. Dhillon, Pradeep Ravikumar, and Ambuj Tewari. Nearest Neighbor Based Greedy Coordinate Descent. In Proceedings of Advances in Neural Information Processing Systems 24 (NIPS '2011), pages 2160-2168, 2011.
[68] Persi Diaconis and David Freedman. A Dozen de Finetti-style Results in Search of a Theory. Annales de l'institut Henri Poincaré, 23(S2):397-423, 1987.
[69] S. J. Dilworth and S. J. Montgomery-Smith. The Distribution of Vector-Valued Rademacher Series. Annals of Probability, 21(4):2046-2052, 1993.
[70] Moshe Dubiner. Bucketing Coding and Information Theory for the Statistical High-Dimensional Nearest-Neighbor Problem. IEEE Transactions on Information Theory, 56(8):4166-4179, 2010.
[71] Zeev Dvir and Sivakanth Gopi. 2-Server PIR with Sub-Polynomial Communication. In Proceedings of the 47 th ACM Symposium on the Theory of Computing (STOC '2015), pages 577-584, 2015.
[72] Kave Eshghi and Shyamsundar Rajaram. Locality Sensitive Hash Functions Based on Concomitant Rank Order Statistics. In Proceedings of the 14th ACM SIGKDD Conference on Knowledge Discovery and Data (KDD '2008), pages 221-229, 2008.
[73] Uriel Feige and Gideon Schechtman. On the Optimality of the Random Hyperplane Rounding Technique for MAX CUT. Random Structures and Algorithms, 20(3):403-440, 2002.
[74] William Feller. An Introduction to Probability Theory and its Applications, Vol. 2. Wiley, New York, 1971.
[75] Matteo Frigo and Steven G. Johnson. The Design and Implementation of FFTW3. Proceedings of the IEEE, 93(2):216-231, 2005.
[76] Sudipto Guha, Piotr Indyk, and Andrew McGregor. Sketching Information Divergences. Journal of Machine Learning Research, 72(1-2):5-19, 2008.
[77] Sariel Har-Peled, Piotr Indyk, and Rajeev Motwani. Approximate Nearest Neighbor: Towards Removing the Curse of Dimensionality. Theory of Computing, 8(1):321-350, 2012.
[78] Edwin Hewitt and Kenneth A. Ross. Abstract Harmonic Analysis: Volume I. Grundlehren der mathematischen Wissenschaften. Springer New York, 1994.
[79] Thomas Hofmann, Aurélien Lucchi, Simon Lacoste-Julien, and Brian McWilliams. Variance Reduced Stochastic Gradient Descent with Neighbors. In Proceedings of Advances in Neural Information Processing Systems 28 (NIPS '2015), pages 2305-2313, 2015.
[80] Piotr Indyk. High-Dimensional Computational Geometry. PhD thesis, Stanford University, 2001.
[81] Piotr Indyk. On Approximate Nearest Neighbors under $\ell_{\infty}$ Norm. Journal of Computer and System Sciences, 63(4):627-638, 2001.
[82] Piotr Indyk. Approximate Nearest Neighbor Algorithms for Fréchet Distance via Product Metrics. In Proceedings of the 18th ACM Symposium on Computational Geometry (SoCG '2002), pages 102-106, 2002.
[83] Piotr Indyk. Approximate Nearest Neighbor under Edit Distance via Product Metrics. In Proceedings of the 15 th ACM-SIAM Symposium on Discrete Algorithms (SODA '2004), pages 646-650, 2004.
[84] Piotr Indyk. Stable Distributions, Pseudorandom Generators, Embeddings, and Data Stream Computation. Journal of the ACM, 53(3):307-323, 2006.
[85] Piotr Indyk and Rajeev Motwani. Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality. In Proceedings of the 30th ACM Symposium on the Theory of Computing (STOC '1998), pages 604-613, 1998.
[86] Piotr Indyk and Nitin Thaper. Fast Color Image Retrieval via Embeddings. Workshop on Statistical and Computational Theories of Vision (at ICCV), 2003.
[87] Piotr Indyk and David Woodruff. Optimal Approximations of the Frequency Moments of Data Streams. In Proceedings of the 35th ACM Symposium on the Theory of Computing (STOC '2005), pages 202-208, 2005.
[88] T. S. Jayram, Subhash Khot, Ravi Kumar, and Yuval Rabani. Cell-Probe Lower Bounds for the Partial Match Problem. Journal of Computer and System Sciences, 69(3):435-447, 2004.
[89] T. S. Jayram, Ravi Kumar, and D. Sivakumar. The One-Way Communication Complexity of Hamming Distance. Theory of Computing, 4(1):129-135, 2008.
[90] T. S. Jayram and David P. Woodruff. The Data Stream Space Complexity of Cascaded Norms. In Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2009), pages 765-774, 2009.
[91] Herve Jégou, Matthijs Douze, and Cordelia Schmid. Product Quantization for Nearest Neighbor Search. IEEE Transactions on Pattern Analysis and Machine Intelligence, 33(1):117-128, January 2011.
[92] Fritz John. Extremum Problems with Inequalities as Subsidiary Conditions. In Studies and Essays Presented to R. Courant on his 60th Birthday, January 8, 1948, pages 187-204. Interscience Publishers, Inc., New York, N. Y., 1948.
[93] William Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In Conference in modern analysis and probability (New Haven, Connecticut, 1982), volume 26 of Contemporary Mathematics, pages 189-206. 1984.
[94] William B. Johnson and N. Lovasoa Randrianarivony. $l_{p}(p>2)$ does not Coarsely Embed into a Hilbert Space. Proceedings of the American Mathematical Society, 134(4):1045-1050, 2006.
[95] N. J. Kalton. Banach Spaces Embedding into $L_{0}$. Israel Journal of Mathematics, 52(4):305-319, 1985.
[96] Bala Kalyanasundaram and Georg Schintger. The Probabilistic Communication Complexity of Set Intersection. SIAM Journal on Discrete Mathematics, 5(4):545557, 1992.
[97] Michael Kapralov. Smooth Tradeoffs Between Insert and Query Complexity in Nearest Neighbor Search. In Proceedings of the 34th ACM Symposium on Principles of Database Systems (PODS '2015), pages 329-342, 2015.
[98] Michael Kapralov and Rina Panigrahy. NNS Lower Bounds via Metric Expansion for $\ell_{\infty}$ and EMD. In Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP '2012), pages 545-556, 2012.
[99] Bruce M. Kapron, Valerie King, and Ben Mountjoy. Dynamic Graph Connectivity in Polylogarithmic Worst Case Time. In Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA '2013), pages 1131-1142, 2013.
[100] Christopher Kennedy and Rachel Ward. Fast Cross-Polytope Locality-Sensitive Hashing. In Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS '2017), 2017. Available as arXiv:1602.06922.
[101] Iordanis Kerenidis and Ronald de Wolf. Exponential Lower Bound for 2-Query Locally Decodable Codes via a Quantum Argument. Journal of Computer and System Sciences, 69(3):395-420, 2004. Available as arXiv:quant-ph/0208062.
[102] Subhash Khot and Assaf Naor. Nonembeddability Theorems via Fourier Analysis. Mathematische Annalen, 334(4):821-852, 2006.
[103] Subhash Khot and Rishi Saket. SDP Integrality Gaps with Local $\ell_{1}$ Embeddability. In Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2009), pages 565-574, 2009.
[104] Subhash Khot and Nisheeth Vishnoi. The Unique Games Conjecture, Integrality Gap for Cut Problems and Embeddability of Negative-Type Metrics into $\ell_{1}$. Journal of the $A C M, 62: 1(8): 8: 1-8: 39,2015$.
[105] Felix Krahmer and Rachel Ward. A Unified Framework for Linear Dimensionality Reduction in L1. Results in Mathematics, 70(1):209-231, 2016.
[106] Robert Krauthgamer and Yuval Rabani. Improved Lower Bounds for Embeddings into $L_{1}$. SIAM Journal on Computing, 38(6):2487-2498, 2009.
[107] Ilan Kremer, Noam Nisan, and Dana Ron. On Randomized One-Round Communication Complexity. Computational Complexity, pages 21-49, 1999.
[108] Eyal Kushilevitz, Rafail Ostrovsky, and Yuval Rabani. Efficient Search for Approximate Nearest Neighbor in High Dimensional Spaces. SIAM Journal on Computing, 30(2):457-474, 2000.
[109] Thijs Laarhoven. Search Problems in Cryptography: From Fingerprinting to Lattice Sieving. PhD thesis, Eindhoven University of Technology, 2015.
[110] James R. Lee and Assaf Naor. $L_{p}$ Metrics on the Heisenberg Group and the Goemans-Linial Conjecture. In Proceedings of the 47 th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2006), pages 99-108, 2006.
[111] Yi Li, Huy L. Nguyên, and David P. Woodruff. On Sketching Matrix Norms and the Top Singular Vector. In Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA '2014), pages 1562-1581, 2014.
[112] Yi Li, Huy L Nguyên, and David P. Woodruff. Turnstile streaming algorithms might as well be linear sketches. In Proceedings of the 46 th ACM Symposium on the Theory of Computing (STOC '2014), pages 174-184, 2014.
[113] Moshe Lichman. UCI Machine Learning Repository, 2013.
[114] Nathan Linial, Eran London, and Yuri Rabinovich. The Geometry of Graphs and Some of its Algorithmic Applications. Combinatorica, 15(2):215-245, 1995.
[115] Ding Liu. A Strong Lower Bound for Approximate Nearest Neighbor Searching. Information Processing Letters, 92(1):23-29, 2004.
[116] Mingmou Liu, Xiaoyin Pan, and Yitong Yin. Randomized Approximate Nearest Neighbor Search with Limited Adaptivity. In Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '2016), pages 23-33, 2016.
[117] Qin Lv, Moses Charikar, and Kai Li. Image Similarity Search with Compact Data Structures. In Proceedings of the 13 th ACM International Conference on Information and Knowledge Management (CIKM '2004), pages 208-217, 2004.
[118] Qin Lv, William Josephson, Zhe Wang, Moses Charikar, and Kai Li. Multi-Probe LSH: Efficient Indexing for High-Dimensional Similarity Search. In Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB '2007), pages 950-961, 2007.
[119] Albert W. Marshall, Ingram Olkin, and Barry C. Arnold. Inequalities: Theory of Majorization and its Applications. Academic Press, 2011.
[120] Jiří Matoušek. On the distortion required for embedding finite metric spaces into normed spaces. Israel Journal of Mathematics, 93:333-344, 1996.
[121] Jiří Matoušek. Lectures on Discrete Geometry. Springer, 2002.
[122] Jiří Matoušek and Assaf Naor. Open Problems on Embeddings of Finite Metric Spaces. Available as http://kam.mff.cuni.cz/~matousek/metrop.ps.gz. Last accessed in November, 2014., 2011.
[123] Andrew M. McDonald, Massimiliano Pontil, and Dimitris Stamos. Spectral $k$ Support Norm Regularization. In Proceedings of Advances in Neural Information Processing Systems 27 (NIPS '2014), pages 3644-3652, 2014.
[124] Andrew McGregor. Open problems in data streams and related topics. IITK Workshop on Algorithms For Data Streams, 2006. Available as http: //sublinear.info.
[125] Stefan Meiser. Point Location in Arrangements of Hyperplanes. Information and Computation, 106(2):286-303, 1993.
[126] Peter Bro Miltersen. Cell Probe Complexity - a Survey. In Advances in Data Structures, 1999.
[127] Peter Bro Miltersen, Noam Nisan, Shmuel Safra, and Avi Wigderson. On Data Structures and Asymmetric Communication Complexity. In Proceedings of the 27th ACM Symposium on the Theory of Computing (STOC '1995), pages 103-111, 1995.
[128] George J. Minty. On the Extension of Lipschitz, Lipschitz-Hölder Continuous, and Monotone Functions. Bulletin of the American Mathematical Society, 76(2):334-339, 1970.
[129] Rajeev Motwani, Assaf Naor, and Rina Panigrahy. Lower Bounds on Locality Sensitive Hashing. SIAM Journal on Discrete Mathematics, 21(4):930-935, 2007.
[130] James R. Munkres. Topology. Prentice Hall, Incorporated, 2000.
[131] S. Muthukrishnan and S. C. Sahinalp. Approximate Nearest Neighbors and Sequence Comparisons with Block Operations. In Proceedings of the 32nd ACM Symposium on the Theory of Computing (STOC '2000), pages 416-424, 2000.
[132] Assaf Naor. A Spectral Gap Precludes Low-Dimensional Embeddings. In Proceedings of the 33rd International Symposium on Computational Geometry (SoCG '2017), 2017.
[133] Assaf Naor and Yuval Rabani. On Approximate Nearest Neighbor Search in $\ell_{p}$, $p>2$. Manuscript, 2006.
[134] Assaf Naor and Gideon Schechtman. Planar Earthmover is not in $L_{1}$. SIAM Journal on Computing, 37(3):804-826, 2007.
[135] Ashwin Nayak. Optimal lower bounds for quantum automata and random access codes. In Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS '1999), pages 369-376, 1999.
[136] Jelani Nelson. Sketching and Streaming High-Dimensional Vectors. PhD thesis, Massachusetts Institute of Technology, 2011.
[137] Huy L. Nguyên. Algorithms for High Dimensional Data. PhD thesis, Princeton University, 2014. Available as http://arks.princeton.edu/ark:/88435/ dsp01b8515q61f.
[138] E. M. Nikišin. A Resonance Theorem and Series in Eigenfunctions of the Laplace Operator. Izvestiya Rossiiskoi Akademii Nauk Seriya Matematicheskaya, 36:795-813, 1972.
[139] Ryan O'Donnell. Analysis of Boolean Functions. Cambridge University Press, 2014.
[140] Ryan O'Donnell, Yi Wu, and Yuan Zhou. Optimal Lower Bounds for LocalitySensitive Hashing (Except When $q$ is Tiny). ACM Transactions on Computation Theory, 6(1):5, 2014.
[141] Rafail Ostrovsky and Yuval Rabani. Low Distortion Embedding for Edit Distance. Journal of the ACM, 54(5):23:1-23:16, 2007.
[142] Mark H. Overmars and Jan van Leeuwen. Some Principles for Dynamizing Decomposable Searching Problems. Information Processing Letters, 12(1):49-53, 1981.
[143] Dmitry Panchenko. Lecture notes in probability theory, 2008.
[144] Rina Panigrahy. Entropy Based Nearest Neighbor Search in High Dimensions. In Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA '2006), pages 1186-1195, 2006.
[145] Rina Panigrahy, Kunal Talwar, and Udi Wieder. A Geometric Approach to Lower Bounds for Approximate Near-Neighbor Search and Partial Match. In Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2008), pages 414-423, 2008.
[146] Rina Panigrahy, Kunal Talwar, and Udi Wieder. Lower Bounds on Near Neighbor Search via Metric Expansion. In Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS '2010), pages 805-814, 2010.
[147] Mihai Patrascu. Unifying the Landscape of Cell-Probe Lower Bounds. SIAM Journal on Computing, 40(3):827-847, 2011.
[148] Jeffrey Pennington, Richard Socher, and Christopher Manning. GloVe: Global Vectors for Word Representation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP '2014), pages 1532-1543, 2014.
[149] Gilles Pisier. Some Results on Banach Spaces Without Local Unconditional Structure. Compositio Mathematica, 37(1):3-19, 1978.
[150] Mihai Pǎtraşcu and Mikkel Thorup. Higher Lower Bounds for Near-Neighbor and Further Rich Problems. In Proceedings of the 47 th Annual IEEE Symposium on Foundations of Computer Science (FOCS '2006), pages 646-654, 2006.
[151] N. Lovasoa Randrianarivony. Characterization of Quasi-Banach Spaces which Coarsely Embed into a Hilbert Space. Proceedings of the American Mathematical Society, 134(5):1315-1317, 2006.
[152] Ilya Razenshteyn and Ludwig Schmidt. FALCONN: Similarity Search over High-Dimensional Data. Available as https://falconn-lib.org/, 2015.
[153] Ilya Razenshteyn and Ludwig Schmidt. FFHT: A C99 Implementation of the Fast Hadamard Transform. Available as https://github.com/FALCONN-LIB/FFHT, 2015.
[154] Ilya Razenshteyn, Zhao Song, and David P. Woodruff. Weighted Low Rank Approximations with Provable Guarantees. In Proceedings of the 48th ACM Symposium on the Theory of Computing (STOC '2016), pages 250-263, 2016.
[155] Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. The Earth Mover's Distance as a Metric for Image Retrieval. International Journal of Computer Vision, 40(2):99-121, 2000.
[156] Walter Rudin. Functional Analysis. McGraw-Hill, Inc., 2nd edition, 1991.
[157] Michael Saks and Xiaodong Sun. Space Lower Bounds for Distance Approximation in the Data Stream Model. In Proceedings of the 34th ACM Symposium on the Theory of Computing (STOC '2002), pages 360-369, 2002.
[158] Ludwig Schmidt, Matthew Sharifi, and Ignacio Lopez Moreno. Large-Scale Speaker Identification. In Proceedings of the 2014 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '2014), 2014.
[159] Isaac J. Schoenberg. On Certain Metric Spaces Arising from Euclidean Spaces by a Change of Metric and their Imbedding in Hilbert Space. Annals of Mathematics, 38(4):787-793, 1937.
[160] Issac J. Schoenberg. Remarks to Maurice Fréchet's article "Sur La Définition Axiomatique D'une Classe D'espace Distanciés Vectoriellement Applicable Sur L'espace De Hilbert'". Annals of Mathematics, 36(3):724-732, 1935.
[161] Issac J. Schoenberg. Metric Spaces and Positive Definite Functions. Transactions of the American Mathematical Society, 44(3):522-536, 1938.
[162] Gregory Shakhnarovich, Trevor Darrell, and Piotr Indyk. Nearest-Neighbor Methods in Learning and Vision: Theory and Practice. MIT Press, 2006.
[163] Alexander A. Sherstov. The Communication Complexity of Gap Hamming Distance. Theory of Computing, 8(8):197-208, 2012.
[164] Malcolm Slaney, Yury Lifshits, and Junfeng He. Optimal Parameters for LocalitySensitive Hashing. Proceedings of the IEEE, 100(9):2604-2623, 2012.
[165] Ryan Spring and Anshumali Shrivastava. Scalable and Sustainable Deep Learning via Randomized Hashing. In Proceedings of the 23rd ACM SIGKDD Conference on Knowledge Discovery and Data (KDD '2017), 2017. Available as arXiv:1602.08194.
[166] Narayanan Sundaram, Aizana Turmukhametova, Nadathur Satish, Todd Mostak, Piotr Indyk, Samuel Madden, and Pradeep Dubey. Streaming Similarity Search over One Billion Tweets Using Parallel Locality-Sensitive Hashing. In Proceedings of the 39th International Conference on Very Large Data Bases (VLDB '2013), pages 1930-1941, 2013.
[167] Michel Talagrand. Embedding Subspaces of $L_{1}$ into $l_{1}^{N}$. Proceedings of the American Mathematical Society, 108(2):363-369, 1990.
[168] Kengo Terasawa and Yuzuru Tanaka. Spherical LSH for Approximate Nearest Neighbor Search on Unit Hypersphere. In Proceedings of the 10th International Workshop on Algorithms and Data Structures (WADS '2007), pages 27-38, 2007.
[169] Gregory Valiant. Finding Correlations in Subquadratic Time, with Applications to Learning Parities and the Closest Pair Problem. Journal of the ACM, 62(2):13, 2015.
[170] Thomas Vidick. A Concentration Inequality for the Overlap of a Vector on a Large Set, with Application to the Communication Complexity of the Gap-Hamming-Distance Problem. Chicago Journal of Theoretical Computer Science, 2012(1), 2012.
[171] Jingdong Wang, Heng Tao Shen, Kingkuan Song, and Jianqiu Ji. Hashing for Similarity Search: a Survey. Available as arXiv:1408.2927, 2014.
[172] Jun Wang, Wei Liu, Sanjiv Kumar, and Shih-Fu Chang. Learning to Hash for Indexing Big Data - A Survey. Available as arXiv:1509.05472, 2015.
[173] Kilian Weinberger, Anirban Dasgupta, John Langford, Alex Smola, and Josh Attenberg. Feature Hashing for Large Scale Multitask Learning. In Proceedings of the 26th International Conference on Machine Learning (ICML '2009), pages 1113-1120, 2009.
[174] J. H. Wells and L. R. Williams. Embeddings and Extensions in Analysis. SpringerVerlag, New York-Heidelberg, 1975.
[175] Ryan Williams. A New Algorithm for Optimal 2-Constraint Satisfaction and its Implications. Theoretical Computer Science, 348(2-3):357-365, 2005.
[176] Przemysław Wojtaszczyk. Banach Spaces for Analysts. Cambridge University Press, 1991.
[177] David P. Woodruff. Optimal Space Lower Bounds for all Frequency Moments. In Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA '2004), pages 167-175, 2004.
[178] David P. Woodruff. Sketching as a Tool for Numerical Linear Algebra. Foundations and Trends in Theoretical Computer Science, 10:1-157, 2014.
[179] Sergey Yekhanin. Locally Decodable Codes, volume 6 of Foundations and Trends in Theoretical Computer Science. Now Publishers, 2012.
[180] Yitong Yin. Simple Average-Case Lower Bounds for Approximate Near-Neighbor from Isoperimetric Inequalities. In Proceedings of the 43 rd International Colloquium on Automata, Languages and Programming (ICALP '2016), pages 84:1-84:13, 2016. Available as arXiv:1602.05391.
[181] A. Zvavitch. More on Embedding Subspaces of $L_{p}$ into $l_{p}^{N}, 0<p<1$. In Geometric aspects of functional analysis, volume 1745 of Lecture Notes in Math., pages 269-280. Springer, Berlin, 2000.


[^0]:    ${ }^{1}$ In particular with their passion for algorithms implementations.

[^1]:    ${ }^{1}$ for $\operatorname{IPv} 4$, or $2^{128}$ for $\operatorname{IPv} 6$

[^2]:    ${ }^{2}$ The $\ell_{1}$ distance is defined as $\|x-y\|_{1}=\sum_{i=1}^{d}\left|x_{i}-y_{i}\right|$, while the $\ell_{2}$ distance is $\sqrt{\sum_{i=1}^{d}\left(x_{i}-y_{i}\right)^{2}}$.
    ${ }^{3}$ The Hamming distance between two binary vectors is defined as the number of positions where they differ.

[^3]:    ${ }^{4}$ Technically, to get this hardness result, we need to modify the reduction from [175] slightly. A variant of this modification can be found in [4], but even before it had been a folklore result.

[^4]:    ${ }^{5}$ One could go even further and argue that if there is no such a gap, then the ANN instance is not very interesting to start with, since the nearest neighbors are not distinguishable from the rest of the points.
    ${ }^{6}$ From now on, $\widetilde{O}(f)$ is defined as $O\left(f \log { }^{O(1)} f\right)$.

[^5]:    ${ }^{7}$ There is a one-to-one correspondence between partitions and hash functions $h: \mathbb{R}^{d} \rightarrow \mathbb{Z}$ if we only care about collisions. We will be using both "languages" interchangeably.

[^6]:    ${ }^{8}$ Given a point, we can quickly find a part of $\mathcal{P}$, which it belongs to.
    ${ }^{9}$ Defined as $\|x-y\|_{\infty}=\max _{i=1}^{d}\left|x_{i}-y_{i}\right|$.

[^7]:    ${ }^{10}$ In case of $\ell_{1}$, it is assumed to be a random subset of $\{0,1\}^{d}$, while for $\ell_{2}$, it must be a random subset of a unit sphere $S^{d-1} \subset \mathbb{R}^{d}$, in both cases queries are generated so that there is one data point that is $c$ times closer than all the other ones.

[^8]:    ${ }^{11}$ In all the known constructions of high-quality partitions, the point location time is proportional to space occupied by a partition.

[^9]:    ${ }^{12}$ Though, see [28], where we make a first step in making such a reduction practical for the Hamming space.

[^10]:    ${ }^{13} \mathrm{~A}$ norm is a function $\|\cdot\|: \mathbb{R}^{d} \rightarrow \mathbb{R}_{+}$that satisfies:

    - $\|x\|=0$ iff $x=0$;
    - $\|\alpha x\|=|\alpha|\|x\| ;$
    - $\|x+y\| \leq\|x\|+\|y\|$.

[^11]:    ${ }^{15} \mathrm{~A}$ metric on a set $X$ is a symmetric function $d: X \times X \rightarrow \mathbb{R}_{+}$such that:

[^12]:    ${ }^{16}$ This can be implemented efficiently using classical hash tables.

[^13]:    ${ }^{1}$ As an extreme example, consider the case of the unit sphere and $c r \geq 2$. Then, the $(c, r)$-ANN problem becomes trivial, since any data point is an answer to any valid query.

[^14]:    ${ }^{2}$ See Figure 2-4 for comparison for the case $c=2$.

[^15]:    ${ }^{3}$ This follows from the fact that two random vectors on a high-dimension sphere are nearly orthogonal with high probability. See, e.g., [32, 121].

[^16]:    ${ }^{4}$ Note that $\eta_{s}$ may not be equal to $\eta_{q}$. It is exactly this discrepancy that will govern the time-space trade-off.

[^17]:    ${ }^{5}$ There will be at most $O\left(\frac{\log n}{\tau}\right)=n^{o(1)}$ of $C_{i}$ 's.

[^18]:    ${ }^{1}$ To get an idea why this is the case, see Section 2.9.3 or the proof of Lemma 2.9.2.

[^19]:    ${ }^{2}$ More quantitatively, we need to set $T=2^{\Omega(\sqrt{\log n})}$ in order for the exponents $\rho_{s}, \rho_{q}$ to deviate from their limits by $O\left(\frac{1}{\sqrt{\log n}}\right)$. See Section 2.8.4 and Section 2.8.5 for the details.

[^20]:    ${ }^{3}$ Note that one can apply Lemma 2 from the arXiv version of [173] to claim that-after such a dimension reduction - the distance between any two points remains sufficiently concentrated for the bounds from Theorem 3.3.1 to still hold (with $d$ replaced by $d^{\prime}$ ).

[^21]:    ${ }^{4}$ The situation is qualitatively similar for other values of $r_{1}$.
    ${ }^{5}$ More specifically, for the "partial" version from Section 3.3.1, since $T$ should be constant, while $d$ grows

[^22]:    ${ }^{6}$ In order to simplify notation, we consider a slightly modified version of the cross-polytope LSH that maps both the standard basis vector $+e_{j}$ and its opposite $-e_{j}$ to the same hash value. It is easy to extend the multiprobe scheme defined here to the "full" cross-polytope LSH from Section 3.3.

[^23]:    ${ }^{1}$ The exact dependence, for $\ell_{2}$, is that one can achieve any $\rho_{u}, \rho_{q} \geq 0$ satisfying $c^{2} \sqrt{\rho_{q}}+\left(c^{2}-\right.$ 1) $\sqrt{\rho_{u}}=\sqrt{2 c^{2}-1}$.

[^24]:    ${ }^{2}$ For every $n, m$, one has $a_{n+m} \leq a_{n}+a_{m}$.

[^25]:    ${ }^{3}$ In a recent work, Naor [132] showed that this approach is impossible. He shows that embedding an $N$-point spectral expander with constant distortion into any normed space requires poly $(N)$ dimensions.

[^26]:    ${ }^{4}$ Here we assume the natural condition that $S(n)$ is increasing, which is, otherwise, easy to guarantee.

[^27]:    ${ }^{1}$ When $X$ is a normed space it suffices to consider $r=1$ by simply scaling the inputs $x, y$.
    ${ }^{2}$ Other metric spaces include edit distance $[33,37,141,21]$ and its variants $[63,131,62,61,56,16]$, the Earth Mover's Distance in the plane or in hypercubes [55, 86, 117, 102, 15, 12], cascaded norms of matrices [90], and the trace norm of matrices [111].
    ${ }^{3}$ These refinements include the Gap-Hamming-Distance problem [177, 89, 49, 50, 54, 163, 170], and LSH in $\ell_{1}$ and $\ell_{2}$ spaces [129, 140].

[^28]:    ${ }^{1}$ The correct dependence on $1 / \varepsilon$ requires the stronger Lopsided Set Disjointness lower bound from [147].

[^29]:    ${ }^{2}$ A 2-query LDC corresponds to LDCs which make two probes to their memory contents. Even

[^30]:    ${ }^{3}$ One can think of $\nu$ as around $\frac{1}{10}$.

