Introduction

This thesis deals with geometric problems in functional analysis. It is quite hard to distinguish between geometric and non-geometric problems in analysis. Typically, an analyst deals with questions that have to do with functions, with particular emphasis on the internal structure of the mappings involved. Geometric analysis tackles similar problems differently. It turns out that it is sometimes preferable to think of a given function as a member (or a "point") of some larger ambient geometric space, which allows one to draw on geometric methods and intuition to answer questions which had no a priori geometric flavor. Of the many examples of the power of the geometric point of view we point out the applications of Bair's category theorem to the construction of exotic functions and the weak topology methods in the theory of partial differential equations. As it often happens in mathematics, although the original motivation for the development of geometric techniques were their usefulness in "real" analytic problems, the internal richness and beauty of the geometry involved shifted the attention of many mathematicians to the study of these structures in their own right. This is the approach of the present work. We will deal with geometric questions that have to do with metric spaces and Banach spaces since we find the subject appealing and beautiful in itself. In spite of this, our geometric work does occasionally have applications in other fields, such as probability, combinatorics and approximation theory. In this sense our work is a direct continuation of the intensive development that geometric analysis has undergone in the 20'th century: from Banach's seminal work in the beginning of the century, through the development of the structure theory of Banach spaces in the 60's and 70's, the deep developments in the local theory of normed spaces in the 70's and 80's and the more recent work on finite metric spaces and the Lipschitz and uniform classification of Banach spaces.

This work is comprised of two parts: linear problems that have to do with with finite dimensional Banach spaces and convex bodies, and non-linear problems involving Lipschitz embeddings and extensions of Lipschitz/Hölder functions. Before passing to a detailed description of the results obtained in each of the above subjects we will briefly explain the common points between the techniques that were involved in the research of both areas. There are two themes which appear throughout this thesis. The first is that this work has to do with problems of a *finite* nature. Of course, the linear problems we studied deal with finite dimensional normed spaces. Moreover, when studying non-linear questions we concentrated solely on problems that have to do with finite metric spaces. Even seemingly infinite theorems such as extension results for Hölder functions are actually based on results concerning a finite number of points. The second theme of this work is the use of probabilistic techniques. There is nothing novel about the deep connection between probability and geometry. The probabilistic method is a powerful theory that has developed greatly since the 60's. Our work relies heavily on probabilistic intuition and techniques. In fact, some of our results can be viewed as probabilistic theorems with a geometric flavor.

The first part of this work deals with non-linear problems that have to do with Lipschitz embeddings of finite metric spaces in normed spaces and extensions of Lischitz/Hölder maps. Recall that if (X, d) and (Y, d) are metric spaces, then a function $f : X \to Y$ is called Lipschitz with constant K if for every $x, y \in X$, $d(f(x), f(y)) \le Kd(x, y)$. The least such constant K is denoted by $Lip(f)$. Lipschitz functions between metric spaces have been intensively investigated in the 20'th century. The study of Lipschitz functions can be traced back to many fundamental works in the first half of the 20'th century: Banach's fixed point theorem, the Mazur-Ulam theorem on onto isometries between Banach spaces, Kirzbraun's work on extensions of Lipschitz functions between Hilbert spaces and many more. Recently, more sophisticated results were obtained, such as differentiability theorems for Lipschitz functions between Banach spaces and the Lipschitz classification of normed spaces. We can safely say that the theory of Lipschitz functions is now a deep and rich subject. In spite of this, there are numerous fundamental questions concerning Lipschitz functions, and we strongly believe that this topic is still in its (late stages of) infancy. We refer to the book [BL] for a detailed account of the state of the art in this subject (up to approximately the year 2000).

The first chapter of this work deals with Lipschitz embeddings of graphs in normed spaces, particularly in Hilbert space. This subject has been intensively investigated in the past two decades. Apart from the intrinsic geometric interest in these questions, this research has been fueled by its possible applications to theoretical computer science. When one is given a complicated data in the form of a finite metric space (most often this data is actually a graph equipped with the natural graph metric), it can be very helpful to represent this data as a subset of a well understood and familiar space (Hilbert space and more generally ℓ_p spaces are prime examples, but embeddings into other simple spaces such as trees have also been investigated and applied to concrete computing problems). Unfortunately, such embeddings seldom preserve distances. One is therefore forced to compromise by looking at embeddings which distort the geometry as little as possible. How can we quantify this distortion of the geometry of a metric space? This is where the Lipschitz condition comes in. The Lipschitz constant of a function measures to what extent the function expands distances. On the same token, the Lipschitz constant of the inverse of a function measures how much it shrinks distances. It is therefore natural to define the distortion of a one to one mapping $f : X \to Y$ as:

$$\operatorname{dist}(f) = Lip(f) \cdot Lip(f^{-1}) = \sup_{x,y \in X} \frac{d(f(x), f(y))}{d(x, y)} \cdot \sup_{x,y \in X} \frac{d(x, y)}{d(f(x), f(y))}.$$

We denote by $c_Y(X)$ the least distortion with which X can be embedded into Y. For $p \ge 1$ we write for simplicity: $c_p(X) = c_{\ell_p}(X)$. The parameter $c_2(X)$ is of particular importance, and is called the Euclidean distortion of X. It turns out that there is a formula for $c_2(X)$ when X is a finite metric space. Let $\mathcal{B}_n$ be the set of all positive semi-definite $n \times n$ matrices Q such that $Q\vec{1} = 0$. Then, thinking of X as a metric defined on the set $\{1, \ldots, n\}$, it was proved in [LLR] that:

$$c_2(X) = \inf_{Q \in \mathcal{B}_n} \sqrt{\frac{\sum_{i,j:Q_{ij}>0} d(i,j)^2 Q_{ij}}{\sum_{ij:Q_{ij}<0} d(i,j)^2 |Q_{ij}|}}.$$

Unfortunately, this formula is hard to work with, and guessing the optimal $Q \in \mathcal{B}_n$ requires a lot of effort. It is therefore not surprising that almost all the known calculations and estimates for $c_2(\cdot)$ were based on other geometric ideas, and not on the above formula.

In [Bou1], Bourgain proved that if X is an n-point metric space then $c_2(X) = O(\log n)$. The proof is based on an elegant averaging technique which can be implemented as a simple embedding algorithm (see [LLR]). As we shall see below, Bourgain's upper bound was eventually shown to be tight. Bourgain himself proved that there exist n-point metric spaces with Euclidean distortion $\Omega\left(\frac{\log n}{\log\log n}\right)$. Is this the end of the story? Of course not. In practical cases more restrictions on the metric space are known, and it would be very interesting to link the value of $c_2(X)$ with other geometric and combinatorial properties of X. At this point we depart from the embedding problem for general metric spaces and restrict our attention to the important family of graph metrics. Every connected graph $G = (V, E)$, induces a natural metric on its vertex set V, which is denoted by d_G. For any two vertices $u, v \in V$, $d_G(u, v)$ is defined as the length of the shortest path joining u and v (for obvious reasons, d_G is often called the geodesic metric). There is a trivial bound for the Euclidean distortion of every graph metric: $c_2(G, d_G) \leq \text{diam}(G)$ (here $\text{diam}(G)$ is the diameter of G). Indeed, this bound is achieved by embedding G as simplex in Hilbert space. In [LLR] it was shown that there exist arbitrarily large graphs G with $c_2(G) = \Omega(\log |G|)$, i.e. Bourgain's bound is tight. What is perhaps more striking is that for these graphs we also have that $\text{diam}(G) = O(\log |G|)$. This means that the trivial bound described above is tight. In other words, there exist complicated graphs such that the best way to represent them as subsets of Hilbert space without distorting the distances too much is to ignore their structure altogether and embed them as if they were a clique! One of the consequences of our work described in the first chapter is a new proof of this intriguing phenomenon.

What are these pathological graphs for which the trivial bound is tight? They are the so called expanders. A graph $G = (V, E)$ is called an expander with conductance Φ if for every $A \subset V$ such that $|A| \leq |V|/2$, $|\{[u, v] \in E; u \in A, v \notin A\}| \geq \Phi|A|$. In other words, a graph is an expander if for every not too large subset of vertices, its edge boundary is large. On a first glance one might doubt that expanders exist (to be more precise, it isn't clear whether there are arbitrarily large graphs with constant degree and constant conductance. Cliques are obviously great expanders, but their degree is huge). It is a classical application of the probabilistic method that constant degree expanders do exist (in fact, in an appropriate probabilistic model most random graphs are expanders). Now that we know that for a constant degree expander G, $c_2(G)$ is of order $\log |G|$, what can be said about $c_p(G)$? Of course, $C_\infty(G) = 1$, but it turns out that $c_p(G)$ decays rather slowly as $p \to \infty$. In [Mat] Matoušek proved that $c_p(G) \sim 1 + \frac{1}{p} \log |G|$ for every $p \geq 1$.

Having recognized that there are "exotic" graphs for which the trivial bound is tight, we will briefly describe the Euclidean distortion of more familiar graphs: discrete cubes, trees and planar graphs. Let $D_n = \{-1, 1\}^n$ be the discrete cube, equipped with the Hamming metric (i.e. the ℓ_1 metric). Enflo proved in [En1] that for $1 \leq p \leq 2$, $c_p(D_n) \sim n^{1-\frac{1}{p}}$. For $2 < p < \infty$ the situation is more complicated: Bourgain, Milman and Wolfson proved in [BMW] that for every $\epsilon >$, $c_p(D_n) \geq c(\epsilon) n^{\frac{1}{2}-\epsilon}$. In [Pi], Pisier later improved this estimate to $c_p(D_n) = \Omega\left(\sqrt{\frac{n}{\log n}}\right)$. In the third chapter of this thesis we will show that in fact, $c_p(D_n) \sim \sqrt{n}$ for $p \geq 2$. If we denote by T_n the full binary tree of depth n, then Bourgain proved in [Bou2] that for $1 < p < \infty$, $c_p(T_n) \sim (\log n)^{\min\{1/2, 1/p\}}$. Finally, it was proved by Rao [Rao] that if G is a planar graph then $c_2(G) = O(\sqrt{\log |G|})$. In chapter 1 we deal with the Euclidean distortion of regular graphs with large girth. Recall that the girth of a graph $G = (V, E)$ is the length of the shortest closed path in G. If G is a k-regular graph, $k \geq 3$ with girth g, then G obviously contains trees of depth $\frac{g}{2} - 1$ (in fact, every ball of radius $\frac{g}{2} - 1$ in G is isometric to such a tree). Bourgain's lower bound for the Lipschitz distortion of trees yields that $c_2(G) = \Omega(\sqrt{\log g})$. The main result of chapter 1 is to improve this lower bound exponentially, by showing that $c_2(G) = \Omega(\sqrt{g})$. We can actually refine this estimate by taking into account the spectral gap of G. If we denote by A the adjacency matrix of G then since G is k-regular, its largest eigenvalue is k. The difference between k and the second largest eigenvalue of A is called the spectral gap of G. In chapter 1 we show that if G is a k-regular graph, $k \geq 3$, with girth g and spectral gap ϵ then $c_2(G) = \Omega(g/\sqrt{\min\{g, k/\epsilon\}})$. Since there are arbitrarily large graphs with a constant spectral gap and which are k-regular, $k \geq 3$, with girth g and diameter $O(g)$, this gives a new proof of the fact that the trivial upper bound for the Euclidean distortion of graphs is tight.

The methods with which we obtained the results of chapter 1 are of particular interest. The proof of the lower bound $c_2(G) = \Omega(\sqrt{g})$, is based on the notion of Markov type (this notion will resurface in chapters 2 and 3 as well). Let X be a metric space. A stochastic process $\{Z_k\}_{k=0}^{\infty}$ is called a symmetric Markov chain on X if there are $x_1, \ldots, x_m \in X$ such that each Z_k takes values in $\{x_1, \ldots, x_m\}$, there is a symmetric stochastic $m \times m$ matrix A such that for each k, $P(Z_{k+1} = x_j | Z_k = x_i) = a_{ij} = a_{ji}$, and Z_0 is uniformly distributed (i.e. $P(Z_0 = x_i) = \frac{1}{m}$ for each i). X is said to have Markov type $p > 0$ if there is a constant $K > 0$ such that for every symmetric Markov chain on X, $\{Z_k\}_{k=0}^{\infty}$, and for every integer n, $\mathbb{E}d(Z_n, Z_0)^p \leq K^p n \mathbb{E}d(Z_1, Z_0)^p$. Roughly speaking, X has Markov type p if every symmetric Markov chain in X, at time n, is not expected to wander further than $Kn^{1/p}$ times its average step. The notion of Markov type was introduced by K. Ball, in connection with extension problems for Lipschitz maps (we will return to Ball's results later in this introduction). Ball showed that Hilbert space has type 2. In chapter 1 we show that graphs with large girth behave badly in terms of their Markov type, and use this observation to prove the aforementioned lower bound for the Euclidean distortion.

To better understand the proof of the estimates in terms of the spectral gap, we recall the basic approach to the proofs of all the existing lower bounds for the Euclidean distortion of graphs. The simplest geometric way to distinguish between two metric spaces is by the "strength" of the triangle inequality. Usually one shows that a metric space is far from Euclidean since it contains "too many" triples for which the triangle inequality holds as a (near) equality. One natural way to measure this phenomenon is via Poincaré type inequalities. By a Poincaré inequality for a function f from the vertexes of a graph $G = (V, E)$ into some metric space (X, d) we mean a bound for the average size of $\{d(f(u), f(v))\}_{u,v \in V}$ in terms of f's average "gradient" $\{d(f(u), f(v))\}_{[u,v] \in E}$. In proving the lower bound for the distortion required to embed an n-point, k-regular expander $G = (V, E)$ with conductance Φ in ℓ_p, $p \geq 1$, Matoušek used the following Poincaré inequality for functions $f : V \to \ell_p$:

$$\sum_{u,v \in V} \|f(u) - f(v)\|_p \leq k(n-1) \left(\frac{4pk}{\Phi}\right)^p \sum_{[u,v] \in E} \|f(u) - f(v)\|_p^p.$$

For the embeddability of the Hamming cube D_n in a metric space (X, d), the following Poincaré inequality for functions $f : D_n \to X$ is relevant:

$$\sum_{d(u,v)=n} d(f(u), f(v))^2 \leq K^2 n^{\frac{2}{p}-1} \sum_{[u,v] \in E(D_n)} d(f(u), f(v))^2.$$

When such an inequality holds, for some $p > 0$, for every n with K independent of n and f, we say that (X, d) has Non-Linear type p. This notion was introduced by Bourgain, Milman and Wolfson in [BMW], and we will return to it when describing the results of chapter 3. We refer to chapter 1 for the Poincaré inequality used for obtaining the lower bound of the Euclidean distortion of trees. Our main result in the first chapter is based on the following new Poincaré inequality: if G is a k regular graph, $k \geq 3$, with girth g and spectral gap ϵ, then for every $1 < s < g/2$ and $f : V(G) \to \ell_2$:

$$\sum_{d_G(u,v)=s} \|f(u) - f(v)\|^2 \leq C(k-1)^s \cdot \frac{1 - e^{-C\epsilon s/k}}{\epsilon} \cdot \sum_{[u,v]\in E(G)} \|f(u) - f(v)\|^2,$$

where C is an absolute numerical constant. The proof of this inequality is based on the analysis of an important sequence of polynomials associated with the graph G: the so-called Geronimus polynomials.

The second chapter of this work deals with extension problems for Hölder functions. Recall that if X and Y are metric spaces, then a function $f : X \to Y$ is called α Hölder with constant C if for every $x, y \in X$, $d(f(x), f(y)) \leq Cd(x, y)^\alpha$. We denote By $\mathcal{A}(X, Y)$ the set of all $\alpha > 0$ such that for all $D \subset X$ and for all α Hölder $f : D \to Y$ there is an $\bar{f} : X \to Y$ which is α Hölder with the same constant as that of f and the restriction of $\bar{f}$ to D is f. Such an $\bar{f}$ is called an isometric extension of f. Analogously, we denote by $\mathcal{B}(X, Y)$ the set of all $\alpha > 0$ such that there is a constant $K > 0$ such that for all $D \subset X$ and for any α Hölder function $g : D \to Y$ with constant C there is an α Hölder function with constant KC, $\bar{g} : X \to Y$ which extends g. Such a $\bar{g}$ is called an isomorphic extension of g. Note that there is an inherent difference between the construction of isometric and isomorphic extensions. Since the Hölder constant in the isometric extension problem remains fixed, it suffices to extend a given function one point at a time, in which case a simple Zorn lemma argument shows that a maximal extension must be defined on all of X. This observation goes back to Kirszbraun's solution of the isometric extension problem for Hilbert spaces [K]. The isomorphic extension problem poses different challenges: there is no choice but to extend a function to an arbitrarily large number of extra points, and then use some sort of compactness argument.

The first part of chapter 2 deals with the monotonicity of $\mathcal{B}(X, Y)$, i.e. we ask whether it is true that $\alpha \in \mathcal{B}(X, Y)$ and $0 < \beta \leq \alpha$ implies that $\beta \in \mathcal{B}(X, Y)$. For $\alpha > 0$ we denote by $K_\alpha(X, Y)$ the infimum of all $K > 0$ such that for all $D \subset X$ and every $f : D \to Y$ which is α Hölder with constant C, there is an extension of f to X which is α Hölder with constant KC. We prove that if we assume that Y is complete, $\alpha \in \mathcal{B}(X, Y)$ and $0 < \beta \leq \alpha$ then for every $t > 0$, $K > K_\alpha(X, Y)$, $D \subset X$ and $f : D \to Y$ which is β Hölder with constant C there is a function $\bar{f} : X \to Y$ which extends f and satisfies :

$$d(x, y) \leq t \Longrightarrow \rho(\bar{f}(x), \bar{f}(y)) \leq \frac{21K^2\alpha}{\beta} Cd(x, y)^\beta.$$

The main point here is that we got estimates that are independent of t. The proof of this result is based on an iterative approximation procedure, in which we restrict f to appropriate discrete sets, extend this restriction to an α Hölder function, and take special care to glue the functions obtained in such a way that they will converge to the desired extension. Since the estimate we got does not depend on t, a simple weak* compactness argument gives that

$\mathcal{B}(X,Y)$ is monotone whenever Y is a dual Banach spaces. We posed the question whether this result is true for Banach spaces Y without the extra duality assumption. This problem was subsequently solved by Brudnyi and Shvartsman [BS] who applied deep results in interpolation theory (namely the K-divisibility theorem) to prove that $\mathcal{B}(X,Y)$ is monotone for arbitrary Banach spaces Y (they proved in fact a more general result on the stability of the Lipschitz extension problem under metric transforms, see [BS] for more details. However, their results rely heavily on the linear structure of Y, and do not seem to yield a result for complete metric spaces as above).

The second part of chapter 2 is devoted to the proof of the following result: if $1<p,q<2$ then:
1) $\mathcal{A}(L_p,L_q)=\left(0,\frac{p}{q^*}\right]$ (here $q^*=q/(q-1)$).
2) $\mathcal{B}(L_p,L_q)=\left(0,\frac{p}{2}\right]$.
3) For any $\frac{p}{2}<\alpha\le 1$ there is an α Hölder function from a subset of the unit ball of L_p to L_q which cannot be extended to an α Hölder function defined on all of the unit ball of L_p.

Since when $1<q<2$, $q^*>2$, we get a clear phase transition between the isometric and isomorphic extension problems: $K_\alpha(L_p,L_q)$, as a function of α, remains constantly 1 for $\alpha\le p/q^*$, for $p/q^*<\alpha\le p/2$ it is strictly larger than 1 but finite, and for every $\alpha>p/2$, $K_\alpha(L_p,L_q)=\infty$. Part 1) was proved by Wells and Williams in [W-W]. In chapter 2 we give a simplified proof of this result. Part 2) is proved using a fundamental result of K. Ball [B]: if $1<q<2$ and X is a metric space with Markov type 2 then $1\in\mathcal{B}(X,L_q)$. This theorem was Ball's original motivation for the definition of the notion of Markov type. In the proof of part 3) we construct an unextendable function. The proof of its unextendability uses an averaging argument. As a byproduct of our construction we answer negatively a question posed by K. Ball in [B] by showing that it is not true that all Lipschitz maps from subsets of Hilbert space into normed spaces extend to the whole of Hilbert space. Finally, we show that the above phase transition does not always occur: if $1<p\le 2\le q<\infty$ then $\mathcal{A}(L_p,L_q)=\mathcal{B}(L_p,L_q)=\left(0,\frac{p}{q}\right]$.

In [B], K. Ball posed the question whether for $2<p<\infty$, L_p has Markov type 2. It follows from the contents of chapters 1 and 2 that a positive answer to this problem would have several interesting applications. The motivation for chapter 3 grew out of an attempt to calculate the Markov type of L_p, $p>2$. In order to better understand the results we briefly overview the notions of type that are used in geometric analysis. The notion of (Rademacher) type is of fundamental importance in modern Banach space theory. A Banach space X is said to have type $p\ge 1$ if there is a constant $T>0$ such that for every n and for every $x_1,\dots,x_n\in X$,

$$\left(\mathbb{E}\left\|\sum_{i=1}^n \epsilon_i x_i\right\|^2\right)^{1/2}\le T\left(\sum_{i=1}^n \|x_i\|^p\right)^{1/p}, \tag{1}$$

where $\epsilon_1,\dots\epsilon_n$ are i.i.d. random variables which take the values ±1 with probability 1/2. Since a subspace of a space with type p also has type p, this notion is particularly useful in proving that a Banach space does not linearly embed in another space. When passing to the Lipschitz category, one would like to define an analogous notion of type which is purely non-linear (and of course, one would like to define this notion in such a way that a substantial part of the linear theory of type would have appropriate analogues in the non-linear theory). Such definitions were given by several authors: Enflo, Bourgain, Milman and Wolfson, and as described above,

K. Ball. Let X be a metric space. A subset $C = \{x_\epsilon\}_{\epsilon \in \{1,-1\}^n} \subset X$ is called an n-dimensional cube in X. An unordered pair $\{x_\epsilon, x_{\epsilon'}\}$ is called an edge of C if ϵ and ϵ' differ in exactly one coordinate. The set of all edges of C is denoted by $edge(C)$. Similarly, the set of diagonals of C is defined by $diag(C) = \{\{x_\epsilon, x_{-\epsilon}\}\}_{\epsilon \in \{1,-1\}^n}$. Now, the first step in defining non-linear type is to notice that (1) can be interpreted as an inequality on "linear" cubes in X. Indeed, if we define for $\epsilon \in \{1,-1\}^n$, $x_\epsilon = \sum_{i=1}^n \epsilon_i x_i$ and $C = \{x_\epsilon\}_{\epsilon \in \{1,-1\}^n}$ then it is easy to deduce from (1) that:

$$\left(\sum_{\{a,b\} \in diag(C)} \|a-b\|^2 \right)^{1/2} \leq K n^{\frac{1}{p}-\frac{1}{2}} \left(\sum_{\{u,v\} \in edge(C)} \|u-v\|^2 \right)^{1/2}. \tag{2}$$

Similarly, (1) and Kahane's inequality give:

$$\sum_{\{a,b\} \in diag(C)} \|a-b\|^p \leq K \sum_{\{u,v\} \in edge(C)} \|u-v\|^p. \tag{3}$$

Now comes the bold new step. Enflo [En2] defined what is called today Enflo-type as follows: a metric space has Enflo-type $p > 0$ if (3) holds for an arbitrary cube $C \subset X$. Similarly, as we have mentioned before, Bourgain, Milman and Wolfson [BMW] defined a metric space X as having Non-Linear type $p > 0$ if (2) holds for an arbitrary cube $C \subset X$. On the face of it, when X is a Banach space, these notions seem substantially stronger than the notion of type. Bourgain, Milman and Wolfson proved that if X has type $p > 0$ then X has Non-Linear type q for every $0 < q < p$. Similarly, Pisier [Pi] proved that if X has type $p > 0$ then X has Enflo-type q for every $0 < q < p$. The starting point of chapter 3 is the fact that for any metric space, Markov type p implies Enflo type p. Having realized that a positive solution to the Markov type 2 problem would imply that L_p, $2 < p < \infty$ has Enflo type 2, it was somewhat discouraging that the only space that was known to have Enflo type 2 was Hilbert space. The purpose of the third chapter was to remedy this. We prove that the class of spaces with Enflo type 2 includes UMD spaces with linear type 2 (A Banach space X is called UMD if martingale differences are unconditional in X. We refer to chapter 3 for a precise definition, but point out that the class of UMD spaces contains all the classical reflexive spaces, particularly L_p, $2 < p < \infty$). Unfortunately, the repertoire of spaces known to have Markov type 2 is presently restricted to Hilbert spaces. The proof of the above result is based on showing that for UMD spaces the $\log n$ term in an inequality due to Pisier [Pi] can be replaced by a constant independent of n. We refer to chapter 3 for a precise description of this and related inequalities. We just point out that chapter 3 draws on techniques from Harmonic analysis. Finally, we apply the calculation of the Non-Linear type of L_p, $p > 2$, to close the gap that was mentioned in the beginning of this introduction concerning the distortion required to embed the discrete cube in this space.

The second part of this work, chapters 4-7, deals with applications of the cone measure to convex geometry. The study of finite dimensional convex bodies originated from several fundamental works in the end of the 19'th century and the beginning of the 20'th century. We point out the ground-breaking contributions of Brunn, Minkowski, Balschke, Alexandrov and others. The motivation for this research was twofold. On the one hand, convex bodies and their geometry naturally attract geometers due to their beauty and their apparent (and misleading) simplicity. On the other hand, mathematicians such as Minkowski discovered deep connections between convex geometry and number theory. In the 1970's and 1980's, the intense development of the local theory of Banach spaces emphasized the importance of convex geometry- after all there is a natural correspondence between norms on $\mathbb{R}^n$ and centrally symmetric convex bodies in $\mathbb{R}^n$.

Let $K \subset \mathbb{R}^n$ be a centrally symmetric convex body. The (normalized) cone measure on ∂K is defined as follows: for every $A \subset \partial K$,

$$\mu_K(A) = \frac{\mathrm{v}ol(\{tx; x \in A \quad \text{and} \quad 0 \le t \le 1\})}{\mathrm{v}ol(K)}.$$

In other words, $\mu_K(A)$ is the normalized volume of the cone with base A and vertex at 0. In the particular case $K = B_p^n = \left\{x \in \mathbb{R}^n; \|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}\right\}$ we write $\mu_{B_p^n} = \mu_p^n$. We also denote by σ_p^n the normalized surface area measure on $\partial B_p^n = S_p^n$. Note that $\sigma_p^n = \mu_p^n$ if and only if $p \in \{1, 2, \infty\}$. It follows from a combination of results from chapters 4 and 5, that when the dimension is large, μ_p^n and σ_p^n are close in the following sense: for every $p \ge 1$ and every measurable $A \subset S_p^n$,

$$|\mu_p^n(A) - \sigma_p^n(A)| \le C\left(1 - \frac{1}{p}\right)\left|\frac{1}{p} - \frac{1}{2}\right|\frac{\sqrt{np}}{n+p} = O\left(\frac{1}{\sqrt{n}}\right), \tag{4}$$

where C is a universal constant. In fact, in chapter 5 we prove the following stronger estimate (except possibly for the dependence of the constants on p):

$$\left|\frac{\sigma_p^n(A)}{\mu_p^n(A)} - 1\right| \le \frac{C_p}{\sqrt{n}}\left[\log\left(\frac{C_p}{\mu_p^n(A)}\right)\right]^{1-\min\left\{\frac{1}{p},\frac{1}{2}\right\}}. \tag{5}$$

In chapters 4 and 5 several applications for these estimates were obtained, which we will now describe.

We begin with the results of chapter 4. In 1987, Diaconis and Freedman [DF] proved that as long as $k = o(n)$, the total variation distance between the distribution of the first k coordinates of a vector chosen randomly with respect to the surface measure on the n dimensional Euclidean sphere, and the k dimensional Gaussian distribution, tends to zero as n tends to infinity. Since for $1 < p < \infty$, $p \neq 2$, there are two natural probability measures on the sphere of ℓ_p^n, namely the cone measure and the surface measure, in trying to find an ℓ_p version of the Diaconis Freedman theorem one is faced with the problem how to properly choose a random vector on the sphere. For reasons that will become clear later, the measure μ_p^n is much easier to work with. Rachev and Rüschendorf proved in [RaR] that for $k = o(n)$, the distribution of the first k coordinates of a vector chosen with respect to μ_p^n is close in total variation distance to the distribution on $\mathbb{R}^k$ whose density is proportional to $e^{-\|x\|_p^p}$. Using a much more complicated argument, Mogul'skiĭ [Mo] later proved the same statement for random vectors chosen according to σ_p^n. Using the estimate (4) we see that the result for σ_p^n actually follows from the much simpler result for μ_p^n. This is the general theme along which the estimates (4) and (5) are applied: one proves a statement for the cone measure, which is usually rather easy to handle, and then transfers it to a similar statement for the surface measure.

Another result in chapter 4 shows that in contrast to the ℓ_p version of the Diaconis Freedman theorem, if we project σ_p^n onto a *random* direction, we will get with very large probability a measure which is close the the Gaussian measure on $\mathbb{R}$. This is the central limit property of the surface measure on the unit ball of ℓ_p^n. The classical central limit theorem is a cornerstone of modern probability theory, and is one of the deepest *natural* phenomena discovered in mathematics. It states that if $\{X_i\}_{i=1}^{\infty}$ are i.i.d. random variable with mean zero and second moment 1, then $(X_1 + \ldots + X_n)/\sqrt{n}$ converges in distribution to the standard Gaussian distribution. If we denote by P the distribution of the random vector $X = (X_1, \ldots, X_n) \in \mathbb{R}^n$, then another way to say the same thing is that the distribution of the orthogonal projection of P onto the direction of the vector $(1, \ldots, 1)$ tends to the Gaussian. When stating the central limit theorem in this fashion, it becomes unclear why should the particular direction $(1, \ldots, 1)$ be singled out among all other possible directions in $\mathbb{R}^n$. Of course, some directions are trivially ruled out, such as the coordinate direction $(1, 0, \ldots, 0)$, but one might expect that the projection of P onto most directions is close to the Gaussian distribution. This is indeed the case, as proved by Romik in [Ro] (such a result also follows from the classical Lindberg-Feller theorem). As is common in probability theory, the next step is to allow some dependence between the coordinates of a distribution on $\mathbb{R}^n$, such that similar results still hold. The Central Limit Conjecture for convex bodies states that if μ is the normalized volume measure on a centrally symmetric convex body $K \subset \mathbb{R}^n$ in isotropic position, then the orthogonal projection of μ onto most directions is close to the Gaussian distribution (here "most" and "close" should be uniform over all such bodies K). A convex body K is said to be in isotropic position if its covariance matrix $\left(\int_K x_i x_j dx\right)$ is the identity (the isotropicity assumption corresponds to the normalization of the second moment in the scalar case). Chapter 4 presents a positive solution of the Central Limit Conjecture for μ_p^n and σ_p^n (from which the central limit property for volume measure on B_p^n follows formally. The result for the volume measure was also proved in [ABP]).

We have in fact a higher dimensional version of the above statement. To introduce it, we define the following distance between two probability measures, P and Q on $\mathbb{R}^k$: $T(P,Q) = \sup_H |P(H) - Q(H)|$, where the supremum is taken over all affine half-spaces in $\mathbb{R}^k$. If we denote by $\lambda_{n,k}$ the normalized Haar measure on the Grassmanian manifold $G(n,k)$ of all k dimensional subspaces of $\mathbb{R}^n$, then the result in chapter 4 reads as follows: Let $p \geq 1$ and $\tilde{\sigma_p^n}$ be the measure on $\mathbb{R}^n$ given by $\tilde{\sigma_p^n}(A) = \sigma(A/a_{n,p})$, where $a_{n,p} = \left[\frac{\Gamma(1/p)}{\Gamma(3/p)} \cdot \frac{\Gamma((n+2)/p)}{\Gamma(n/p)}\right]^{1/2}$. Then for every $\epsilon > 0$ and $k \leq c\epsilon^4 n$:

$$\lambda_{n,k}\left[\left\{E \in G(n,k) : T(\mathrm{Proj}_E(\tilde{\sigma_p^n}), \gamma_k) \geq \epsilon\right\}\right] \leq C\exp(-cn\epsilon^4), \tag{6}$$

where $c, C > 0$ are universal constants. Here $\mathrm{Proj}_E(\tilde{\sigma_p^n})$ is the orthogonal projection of the measure $\tilde{\sigma_p^n}$ onto the subspace E, and γ_k is the standard Gaussian measure on $\mathbb{R}^k$. The proof of this result uses measure concentration techniques and the special properties of μ_p^n to prove (6) for the measure $\tilde{\mu_p^n}$ (which is defined analogously to $\tilde{\sigma_p^n}$), and then inequality (4) to transfer the result to $\tilde{\sigma_p^n}$.

In chapter 5, which is rather technical, inequality (5) is used to transfer concentration inequalities for μ_p^n due to Gromov and Milman [GM] and Schechtman and Zinn [SZ2] to similar results for σ_p^n. We refer to chapter 5 for a precise formulation of these results. We would like to point out that among other results, the probabilistic properties of μ_p^n are studied in chapter 5, in particular it is shown to satisfy the so-called "negative correlation property". Additionally, a precise concentration inequality for the ℓ_q^n norm on the ℓ_p^n sphere (with respect to μ_p^n) is obtained, a result which generalizes a theorem of Schechtman and Zinn [SZ2].

The reason the measure μ_p^n is so useful is the following concrete realization of it (due to Schechtman and Zinn, [SZ1]). Let $g_1, \dots, g_n$ be i.i.d random variables with density proportional to $e^{-|t|^p}$. If we denote by G the random vector $(g_1, \dots, g_n) \in \mathbb{R}^n$ then $G/\|G\|_p$ generates the measure μ_p^n. Moreover, $\|G\|_p$ and $G/\|G\|_p$ are independent random variables. This result plays a crucial role in both the proofs of the above statements and the geometric results in chapter 6. In chapter 6 we study the maximal and minimal volume of the projection of B_p^n onto $n-1$ dimensional hyperplanes. If we denote $H_k = \left\{x \in \mathbb{R}^n; \sum_{i=1}^k x_i = 0\right\}$ then the main result of this chapter states that for every hyperplane $H \subset \mathbb{R}^n$, if $p > 2$ then $vol(P_{H_1}(B_p^n)) \leq vol(P_H(B_p^n)) \leq vol(P_{H_n}(B_p^n))$, and if $1 < p < 2$ then $vol(P_H(B_p^n)) \leq vol(P_{H_1}(B_p^n))$, where P_H denotes the orthogonal projection onto H. Moreover, we prove an analogue of the Meyer-Pajor theorem [MP]: for every hyperplane H, $\mathrm{vol}(P_H(B_p^n))/\mathrm{vol}(B_p^{n-1})$ is an increasing function of $p \geq 1$ (Meyer and Pajor proved the same monotonicity statement for volumes of sections of B_p^n). The main point in the proof of these results is to use the Schechtman-Zinn representation of μ_p^n and the calculation of $\frac{d\mu_p^n}{d\sigma_p^n}$ that was done in chapter 4 to prove the following formula for volumes of projections: let $X_1, ..., X_n$ be i.i.d. random variables with density proportional to $|t|^{(2-p)/(p-1)} \exp\left(-|t|^{p/(p-1)}\right)$, $p > 1$. Then for every $a \in S^{n-1}$:

$$\frac{\mathrm{vol}(P_{a^\perp} B_p^n)}{\mathrm{vol}(B_p^{n-1})} = \frac{\mathbb{E}\left|\sum_{i=1}^n a_i X_i\right|}{\mathbb{E}|X_1|}.$$

This formula allows us to apply results from probability theory and Fourier analysis to prove the above statements (namely, we use the Choquet ordering between random variables and the theory of completely monotonic functions).

In the above notation, the independence of $\|G\|_p$ and $G/\|G\|_p$ plays a central role in many results. One might hope that similar random variables exist for norms other than the ℓ_p norms. The results in chapter 7 show that this in not the case, i.e. that under some mild assumptions, the ℓ_p^n norm is the only norm on $\mathbb{R}^n$ which admits a random vector $X \in \mathbb{R}^n$ with independent coordinates such that $X/\|X\|$ is independent of $\|X\|$. To be more precise, let $\|\cdot\|$ be a norm on $\mathbb{R}^n$ which admits such a random vector X. If X is absolutely continuous and its density is of the form $e^{-V(x)}$, where V is locally integrable, then the norm $\|\cdot\|$ is a weighted ℓ_p norm. Additionally, if X has countable support then, unless the norm $\|\cdot\|$ is a weighted ℓ_∞ norm, $\|X\|$ is a constant random variable (i.e. the independence of $X/\|X\|$ and $\|X\|$ is trivial). When $\|\cdot\|$ is a weighted ℓ_∞ norm then a complete characterization of all such X is also given in chapter 7.

In chapter 8 we briefly discuss open problems and directions of future research which arise from the work presented here.

We end this introduction with an explanation concerning the structure of this thesis. Each chapter is a paper written (with the exception of chapters 2 and 5) with other mathematicians. The names of the authors of each chapter are clearly stated in its heading. I would like to thank my co-authors for their kind permission to include our joint results in my thesis.

References

[ABP] Antilla M.; Ball K.; Perissinaki I. The central limit problem for convex bodies. Trans. Amer. Math. Soc., to appear.

[B] Ball K. Markov Chains, Riesz Transforms and Lipschitz Maps. Geometric and Functional Analysis 2, p. 137-172 (1992).

[Bou1] Bourgain J. On Lipschitz embedding of finite metric spaces in Hilbert space. Israel J. Math. 52, p. 46-52, (1985).

[Bou2] Bourgain J. The metrical interpretation of superreflexivity in Banach spaces. Israel J. Math. p. 56 p. 222 - 230 (1986).

[BL] Benyamini, Y; Lindenstrauss, J. Geometric nonlinear functional analysis. Vol. 1. American Mathematical Society Colloquium Publications, 48. American Mathematical Society, Providence, RI, 2000.

[BMW] Bourgain, J.; Milman, V.; Wolfson, H. On type of metric spaces. Trans. Amer. Math. Soc. 294 (1986), no. 1, 295–317.

[BS] Brudnyi Y.; Shvartsman P. Stability of the Lipschitz Extension Property Under Metric Transforms. To appear in Geometric and Functional Analysis.

[DF] Diaconis P.; Freedman D. A dozen de Finetti-style results in search of a theory. Ann. Inst. Henri Poincare **23** (1987) 397-423.

[En1] Enflo P., On the nonexistence of uniform homeomorphisms between L_p-spaces, *Arkiv För Matematik*, 8(1969), 103-105.

[En2] Enflo, P. On infinite-dimensional topological groups. Seminaire sur la Geometrie des Espaces de Banach (1977–1978), Exp. No. 10–11, 11 pp., Ecole Polytech., Palaiseau, 1978.

[GM] Gromov, M. Milman V.D. Generalization of the spherical isoperimetric inequality to uniformly convex Banach spaces. Compositio Math. 62 (1987), no. 3, 263-282.

[K] Kirszbraun M. D., Über die zussammenziehenden und Lipschitzchen Transformationen. Fund. Math.,**22** (1934), 77-108.

[LLR] Linial N.; London E.; Rabinovich Yu. The geometry of graphs and some of its algorithmic applications, Combinatorica, 15 p. 215 - 245 (1995).

[Mat] Matoušek J., Embedding expanders into ℓ_p, Israel J. Math. 102 p. 189-197 (1997).

[MP] Meyer M.; Pajor A. Sections of the unit ball of l_p^n. J. Funct. Anal., 80:109–123, 1988.

[Mo] Mogul'skiĭ A.A. de Finetti-type results for ℓ_p. (Russian). Sibirsk. Mat. Zh. 32 (1991), no. 4, 88-95. Translation in Siberian Math. J. 32 (1992), no. 4, 609-616.

[Pi] Pisier, G. Probabilistic methods in the geometry of Banach spaces. Probability and analysis (Varenna, 1985), 167–241, Lecture Notes in Math., 1206, Springer, Berlin, 1986.

[RaR] Rachev S.T.; Rüschendorf L. Approximate independence of distributions on spheres and their stability properties. Ann. Probab. 19 (1991), no. 3, 1311-1337.

[Rao] Rao S.B. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In 15th Annual ACM Symposium on Computational Geometry, p. 300-306 (1999).

[Ro] Romik D. Measure concentration techniques for randomized central limit theorems. Preprint (1998).

[SZ1] Schechtman G; Zinn J. On the volume of the intersection of two L_p^n balls. Proc. Amer. Math. Soc. **110** (1990) , 1, 217-224.

[SZ2] Schechtman G.; Zinn J. Concentration on the ℓ_p^n ball. Geometric aspects of functional analysis, Lecture Notes in Math., 1745, Springer Verlag, Berlin, 245-256, 2000.

[W-W] Wells J.H.; Williams L.R. Embeddings and Extensions in Analysis. Ergebnisse **84**, Springer-Verlag (1975).