CHAPTER 1

Background material

1.1. Notation and conventions

Throughout this survey, (X, d) and (Z, d) will denote metric spaces. We often abuse notation by dropping explicit reference to the metric. We write $B(z, r)$, resp. $\overline{B}(z, r)$, for the open, resp. closed, ball in Z with center z and radius r. For $B = B(z, r)$ we let $tB = B(z, tr)$; we emphasize that the set tB depends on the choice of center and radius for B (which are not necessarily unique in metric spaces). We write $\operatorname{diam} A$ for the diameter of a set $A \subset Z$, and $\operatorname{dist}(A, B)$ for the distance between sets $A, B \subset Z$. Note that

$$\operatorname{diam} \overline{B}(z, r) \leq 2r \tag{1.1.1}$$

for all metric balls $B(z, r)$ in Z.

Unless otherwise specified, we equip the Euclidean space $\mathbb{R}^N$ of dimension N with the Euclidean (ℓ^2) norm and metric. We will occasionally denote the Euclidean metric in $\mathbb{R}^N$ by d_{eucl}. We denote Lebesgue measure in $\mathbb{R}^N$ by $\mathcal{L}^N$. The cardinality of a finite set S will be written $\#S$.

As is customary, we let $C, c, \ldots$ denote various finite positive constants. These constants may depend on auxiliary data a, b, etc.; we indicate this by writing $C(a, b)$ or $c(a, b)$. We will write C if we wish to emphasize that a certain constant is finite, and we will write c or $1/C$ if we wish to emphasize that the constant is positive.

1.2. Mappings

In this section, we review a number of important classes of mappings between metric spaces. The most important of these, from the perspective of this survey, is the class of *quasisymmetric maps* (Definition 1.2.9), which underlies the notions of conformal gauge and conformal dimension. The study of quasiconformal and quasi-Möbius maps has a long history in geometric function theory: the theory of these classes of maps generalizes classical aspects of complex analysis.

We begin with some simpler classes: Lipschitz and Hölder maps.

1.2.1. Lipschitz and Hölder maps. Lipschitz functions are arguably the most important class of functions in analysis in metric spaces. They play a role there analogous to that played by smooth functions in classical Euclidean analysis.

DEFINITION 1.2.1. Let (Z, d) and (Z', d') be metric spaces and let $L > 0$. A map $f : Z \to Z'$ is said to be *L-Lipschitz* if

$$d'(f(z_1), f(z_2)) \leq Ld(z_1, z_2)$$

for all $z_1, z_2 \in Z$. We say that f is *Lipschitz* if it is L-Lipschitz for some $L < \infty$.

When $Z' = \mathbb{R}$ we say that f is a *Lipschitz function.*

The Lipschitz condition is an absolute metric distortion requirement: distances in the image should not exceed distances in the source by more than a fixed, universal multiplicative constant. The importance of Lipschitz functions in analysis rests largely on results such as the following two theorems.

THEOREM 1.2.2 (Rademacher Differentiation Theorem). *Every Lipschitz function defined on a subset of $\mathbb{R}^N$ is differentiable $\mathcal{L}^N$-almost everywhere.*

THEOREM 1.2.3 (McShane Extension Theorem). *Let Z be any metric space and let $f : A \to \mathbb{R}$ be an L-Lipschitz function defined on a set $A \subset Z$. Then there exists an L-Lipschitz extension of f to Z, i.e., there exists an L-Lipschitz function $g : Z \to \mathbb{R}$ so that $g|_A = f$.*

Note that, as a consequence of Theorem 1.2.3, we may without loss of generality assume that the function in Theorem 1.2.2 is defined in all of $\mathbb{R}^N$. Indeed, suppose that $f : A \to \mathbb{R}$, $A \subset \mathbb{R}^N$, is a Lipschitz function. By McShane's Theorem 1.2.3, there is a Lipschitz extension $g : \mathbb{R}^N \to \mathbb{R}$ of f. By Rademacher's Theorem 1.2.2,

$$Z := \{x \in \mathbb{R}^N : g \text{ is not differentiable at } x\}$$

is a set of measure zero in $\mathbb{R}^N$. Since

$$\{x \in A : f \text{ is not differentiable at } x\}$$

is a subset of $Z \cap A$, it also has measure zero. Thus f is differentiable $\mathcal{L}^N$-a.e. in A, as asserted.

We note the following elementary characterization of Lipschitz maps in terms of metric balls. A generalization of this observation to quasisymmetric maps will be useful later on.

LEMMA 1.2.4. *Let $f : Z \to Z'$. Then f is L-Lipschitz if and only $f(B(z, r)) \subset B(f(z), Lr)$ for all balls $B(z, r)$ in Z.*

We recall some basic examples. The distance function $d : Z \to \mathbb{R}$ to a point $z_0 \in Z$, $d(z) := d(z, z_0)$, is a 1-Lipschitz function; this follows directly from the triangle inequality. More generally, for any closed set $F \subset Z$, the function $d(z) = \operatorname{dist}(z, F)$ is 1-Lipschitz. The class of Lipschitz functions on a metric space Z is closed under addition and the operations of taking maxima and minima. Compositions of Lipschitz maps are Lipschitz.

The projection maps

$$\pi_j : \mathbb{R}^N \to \mathbb{R}, \qquad j = 1, \ldots, N,$$

or

$$\pi_j : \ell^\infty \to \mathbb{R}, \qquad j = 1, 2, \ldots,$$

given by $\pi_j(x) = x_j$ are 1-Lipschitz maps. Thus $\pi_j \circ f$ is Lipschitz whenever $f : Z \to \mathbb{R}^N$ or $f : Z \to \ell^\infty$ is Lipschitz. Conversely, if $\pi_j \circ f$ is L-Lipschitz for all j, then f is $L\sqrt{N}$-Lipschitz (if the target is $\mathbb{R}^N$) or is L-Lipschitz (if the target is ℓ^∞).

DEFINITION 1.2.5. A homeomorphism $f : (Z, d) \to (Z', d')$ is *L-bi-Lipschitz* if both f and f^{-1} are L-Lipschitz. We say that f is *bi-Lipschitz* if it is L-bi-Lipschitz for some finite L.

In many contexts, bi-Lipschitz equivalent spaces are indistinguishable. For instance, all of the most well-studied notions of metric dimension (see section 1.4) are invariant under bi-Lipschitz maps. The problem of recognizing when two metric spaces are bi-Lipschitz equivalent is in general quite challenging. For instance, there is still no simple list of intrinsic metric properties which characterize bi-Lipschitz images of $\mathbb{R}^N$, except in the simple case $N = 1$.

Products of bi-Lipschitz maps are always bi-Lipschitz. Given two metric spaces (Z_1, d_1) and (Z_2, d_2) we consider the product space $Z_3 = Z_1 \times Z_2$. We may equip Z_3 with the ℓ^2 product metric

$$d_3((z_1, z_2), (z_1', z_2')) = \sqrt{d_1(z_1, z_1')^2 + d_2(z_2, z_2')^2}.$$

Alternatively, we may use the ℓ^p product metric

$$((z_1, z_2), (z_1', z_2')) \mapsto (d_1(z_1, z_1')^p + d_2(z_2, z_2')^p)^{1/p}$$

for any $1 \leq p < \infty$ or the corresponding ℓ^∞ product metric

$$((z_1, z_2), (z_1', z_2')) \mapsto \max\{d_1(z_1, z_1'), d_2(z_2, z_2')\}.$$

All of these metrics are bi-Lipschitz equivalent.

PROPOSITION 1.2.6. *Let $f_1 : (Z_1, d_1) \to (Z_1', d_1')$ and $f_2 : (Z_2, d_2) \to (Z_2', d_2')$ be bi-Lipschitz maps. Then $(f_1, f_2) : (Z_3, d_3) \to (Z_3', d_3')$ is bi-Lipschitz, where $Z_3 = Z_1 \times Z_2$ and $Z_3' = Z_1' \times Z_2'$.*

The concept of a Lipschitz map admits the following natural generalization.

DEFINITION 1.2.7. Let Z and Z' be as in Definition 1.2.1 and let $L, \alpha > 0$. A map $f : Z \to Z'$ is said to be *(L, α)-Hölder* if

$$d'(f(z_1), f(z_2)) \leq L d(z_1, z_2)^\alpha$$

for all $z_1, z_2 \in Z$. We say that f is *α-Hölder* if it is (L, α)-Hölder for some $L < \infty$, and that f is *Hölder* if it is α-Hölder for some $\alpha > 0$.

Again, when $Z' = \mathbb{R}$ we say that f is a *Hölder function.*

The obvious analog of Lemma 1.2.4 holds for Hölder functions. We leave the precise formulation to the reader.

Lipschitz and Hölder maps are closely related. To state this more precisely we recall the notion of a snowflake metric space.

DEFINITION 1.2.8. Let (Z, d) be any metric space, let $0 < \epsilon < 1$, and define d^ϵ on Z by

$$d^\epsilon(z_1, z_2) := d(z_1, z_2)^\epsilon.$$

Then d^ϵ is a metric. The metric space (Z, d^ϵ) is called the *snowflake* of (Z, d) *with parameter* ϵ.

The primary example (and the reason for the terminology) is the *von Koch snowflake curve* $K \subset \mathbb{R}^2$, see Figure 1.1. This self-similar fractal set is homeomorphic to the unit interval $[0, 1]$. It has metric dimension $\frac{\log 4}{\log 3} > 1$. (Here we refer to any of the notions of dimension introduced in section 1.4—all of them agree for this set.) Let $\alpha = \frac{\log 3}{\log 4}$. Then (K, d_{eucl}) is bi-Lipschitz equivalent to the snowflaked space $([0, 1], d^\alpha_{eucl})$.

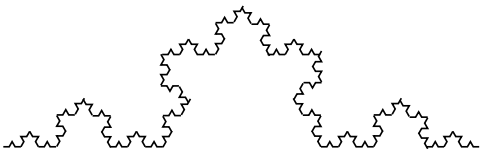

FIGURE 1.1. The von Koch snowflake curve K

It follows immediately from the definitions that a function $f : (Z, d) \to (Z', d')$ is α-Hölder (for some $\alpha \leq 1$) if and only if $f : (Z, d^\alpha) \to (Z', d')$ is Lipschitz. Coupling this with the McShane Extension Theorem, we observe that every real-valued Hölder function defined on a subset of a metric space Z extends to a Hölder function on all of Z, with the same Hölder exponent and coefficient.

Hölder functions with exponent larger than one are of little interest in the classical setting. In fact, if f is a $(1 + \epsilon)$-Hölder function defined on an open subset Ω of $\mathbb{R}^N$, then f is constant. For the proof, we note that f is Lipschitz on compact subdomains of Ω, hence differentiable a.e. in Ω. It is easy to check that the derivative must vanish wherever it is defined. Since f is also absolutely continuous along every line segment in Ω and Ω is connected, we conclude that f is constant. However, by the remarks of the previous paragraph, $(1 + \epsilon)$-Hölder functions defined on more general subsets of $\mathbb{R}^N$ need not be constant.

Note that the product of the snowflakes of (Z_1, d_1) and (Z_2, d_2) with the same parameter ϵ is bi-Lipschitz equivalent to the snowflake of $(Z_1 \times Z_2, d_3)$ with parameter ϵ. Products of snowflakes with different parameters are never bi-Lipschitz equivalent with each other, and are rarely even quasisymmetrically equivalent (for the definition of quasisymmetric maps, see the following section). See the notes to section 4 for more details.

1.2.2. Quasisymmetric, quasiconformal and quasi-Möbius maps. The bi-Lipschitz condition imposes a restriction on absolute metric distortion: distances in the image are uniformly comparable to the corresponding distances in the source. The notion of quasisymmetry, which we now introduce, is more subtle: it imposes a restriction on *relative* metric distortion. Quasisymmetric maps are at the core of this survey. In the Euclidean setting (and in many other contexts), quasisymmetry turns out to be equivalent to the more classical notion of quasiconformality, which arose first in the context of complex analysis. We review the basic facts.

A *distortion function* is a homeomorphism of $[0,\infty)$ onto itself.

DEFINITION 1.2.9. Let η be a distortion function. Let $f : (Z,d) \to (Z',d')$ be a homeomorphism of metric spaces. We say that f is *η-quasisymmetric* if

$$\frac{d'(f(z_1),f(z_2))}{d'(f(z_1),f(z_3))} \leq \eta\left(\frac{d(z_1,z_2)}{d(z_1,z_3)}\right) \tag{1.2.1}$$

for all $z_1,z_2,z_3 \in Z$, $z_1 \neq z_3$. We say that f is *quasisymmetric* if it is η-quasisymmetric for some η, and that (Z,d) and (Z',d') are *quasisymmetrically equivalent* if there exists a quasisymmetry from Z to Z'.

Thus f is quasisymmetric if *ratios* of distances of pairs of points are distorted in a controlled fashion. This three-point condition is extremely useful. We will see numerous applications in later chapters.

EXAMPLES 1.2.10. (1) Every L-bi-Lipschitz map $f : Z \to Z'$ is η-quasisymmetric with distortion function $\eta(t) = L^2 t$.

(2) The identity map from (Z,d) to its snowflake (Z,d^ϵ) with parameter $\epsilon < 1$ is η-quasisymmetric with distortion function $\eta(t) = t^\epsilon$.

(3) If $f_1 : Z \to Z'$ is an η_1-quasisymmetric homeomorphism and $f_2 : Z' \to Z''$ is an η_2-quasisymmetric homeomorphism, then $f_2 \circ f_1 : Z \to Z''$ is $\eta_2 \circ \eta_1$-quasisymmetric. The inverse of an η-quasisymmetric homeomorphism is $\tilde{\eta}$-quasisymmetric with distortion function $\tilde{\eta}(t) = 1/\eta^{-1}(1/t)$. Hence the collection of all quasisymmetric homeomorphisms of Z onto itself forms a group.

Quasisymmetry is a *global* condition: it imposes a uniform requirement on the relative metric distortion of any triple of points. In the Euclidean case, this condition turns out to be equivalent to an infinitesimal version.

DEFINITION 1.2.11. Let $F : \Omega \to \Omega'$ be a homeomorphism between domains in $\mathbb{R}^N$. The *linear dilatation* of F is the function $H_F : \Omega \to \mathbb{R}$ given by

$$H_F(x) := \limsup_{r\to 0} \frac{\sup\{|F(x)-F(y)| : |x-y| = r\}}{\inf\{|F(x)-F(z)| : |x-z| = r\}}. \tag{1.2.2}$$

We say that F is *H-quasiconformal*, $H \geq 1$, if $H_F(x) \leq H$ for all $x \in \Omega$, and that F is *quasiconformal* if it is H-quasiconformal for some $H < \infty$.

It is easy to see from the definitions that if F is η-quasisymmetric, then it is H-quasiconformal with $H = \eta(1)$. In fact, the stronger condition

$$\sup_{r>0} \frac{\sup\{|F(x) - F(y)| : |x - y| = r\}}{\inf\{|F(x) - F(z)| : |x - z| = r\}} \leq \eta(1)$$

holds true. It is a deep fact, central to modern geometric function theory, that the converse also holds true, at least for globally defined mappings. The same result holds true for homeomorphisms between sufficiently nice domains.

THEOREM 1.2.12. *Let $F : \mathbb{R}^N \to \mathbb{R}^N$ be a H-quasiconformal homeomorphism, $N \geq 2$. Then F is η-quasisymmetric for some distortion function η which depends only on H and N.*

REMARK 1.2.13. There are many characterizations of quasiconformal maps. Here is an analytic one. Assume that F is a homeomorphism in the local Sobolev class $W^{1,N}_{loc}(\Omega : \mathbb{R}^N)$ (the latter condition means that the first-order partial derivatives of the components of F exist in the distributional sense as elements of $L^N_{loc}(\Omega)$), and that there exists $K < \infty$ so that the almost everywhere defined differential matrix DF satisfies the following *reverse Hadamard inequality*

$$||DF(x)||^N \leq K \det DF(x) \tag{1.2.3}$$

for a.e. $x \in \Omega$. Geometrically, inequality (1.2.3) asserts that the volume of the image of an infinitesimal ball centered at x is comparable with the volume of the smallest circumscribed ball. By elementary geometry, this condition implies and is implied by the corresponding condition that infinitesimal balls are mapped onto infinitesimal ellipsoids with uniformly bounded eccentricity. Indeed, F is quasiconformal (in the sense of Definition 1.2.11) if and only if it satisfies this condition. The constants H and K depend only on each other and on the dimension N.

REMARK 1.2.14. In the setting of (1.2.3), the regularity assumption $F \in W^{1,N}_{loc}$ is natural, since it is equivalent to the local integrability of the Jacobian $\det DF$ (which is relevant for change of variables arguments). However, it is a deep fact in Euclidean geometric function theory that weaker regularity assumptions are sufficient, and that stronger regularity assumptions are in fact valid in the presence of a distortion condition such as (1.2.3). A fundamental tool is Gehring's higher integrability lemma, which in its simplest form asserts that a $W^{1,N}_{loc}$ mapping F satisfying (1.2.3) for some $K < \infty$ in fact belongs to $W^{1,p}_{loc}$ for all $p < p(N, K)$ for some specific $p(N, K) > N$. See the notes to this section for further remarks.

REMARK 1.2.15. Theorem 1.2.12 admits the following local version. Let us say that $F : \Omega \to \Omega'$ is *locally quasisymmetric* if it is quasisymmetric on every compact set $S \subset \Omega$, with distortion function η depending only on S (and N). Then every quasiconformal mapping $F : \Omega \to \Omega'$ between arbitrary domains $\Omega, \Omega' \subset \mathbb{R}^N$, $N \geq 2$, is locally quasisymmetric.

REMARK 1.2.16. The restriction to dimensions $N \geq 2$ is essential. Every C^1 homeomorphism $F : \mathbb{R} \to \mathbb{R}$ is quasiconformal, but such a homeomorphism need not be quasisymmetric.

EXAMPLE 1.2.17. Every 1-quasiconformal map defined between domains in $\mathbb{R}^2$ is a conformal map (in the sense of complex analysis). Every 1-quasiconformal map defined between domains in $\mathbb{R}^N$, $N \geq 3$, is the restriction of a globally defined Möbius transformation of $\mathbb{R}^N$: this is a consequence of the *Gehring–Liouville theorem*.

We remark on an interesting corollary of Theorem 1.2.12: *inverses of quasiconformal maps of* $\mathbb{R}^N$ *(*$N \geq 2$*) are quasiconformal*. This follows immediately from Theorem 1.2.12 and Example 1.2.10(3). It is not at all easy to see this fact directly from the definition of quasiconformality.

The following lemma recasts quasisymmetry in terms of the distortion of metric annuli. It is the analog of Lemma 1.2.4 from the Lipschitz setting.

LEMMA 1.2.18. *A homeomorphism* $f : Z \to Z'$ *is* η*-quasisymmetric if and only if whenever* $B_1 = B(z, r_1) \subset B_2 = B(z, r_2)$ *are concentric balls in* Z*, then there exist concentric balls* $B_1' = B'(f(z), r_1') \subset B_2' = B'(f(z), r_2')$ *so that* $B_1' \subset f(B_1) \subset f(B_2) \subset B_2'$ *and* $r_2'/r_1' \leq \eta(r_2/r_1)$.

We note however an important distinction between quasiconformal maps and quasisymmetric maps. Quasiconformal maps need not preserve boundedness. (Indeed, even conformal maps of the plane can change bounded domains into unbounded domains.) By way of contrast, quasisymmetries always preserve boundedness. More generally, we have the following well-known result on the distortion of relative diameter by quasisymmetric maps.

LEMMA 1.2.19 (Tukia–Väisälä). *Let* $f : Z \to Z'$ *be an* η*-quasisymmetric homeomorphism and assume that* $A \subset B \subset Z$ *are such that* $0 < \operatorname{diam} A \leq \operatorname{diam} B < \infty$. *Then* $\operatorname{diam} f(B)$ *is finite and*

$$\frac{1}{2\eta(\frac{\operatorname{diam} B}{\operatorname{diam} A})} \leq \frac{\operatorname{diam} f(A)}{\operatorname{diam} f(B)} \leq \eta\left(\frac{2\operatorname{diam} A}{\operatorname{diam} B}\right).$$

It is for this reason sometimes advantageous to work with a different class of mappings: the so-called *quasi-Möbius maps*. The quasi-Möbius condition is a four-point condition. One can recover quasisymmetric maps on locally compact metric spaces from quasi-Möbius maps on the one-point (Alexandrov) compactifications by specializing one of the points to the point at infinity.

DEFINITION 1.2.20. The *(real) cross ratio* of four points z_1, z_2, z_3, z_4 in a metric space (Z, d) is

$$[z_1, z_2, z_3, z_4] := \frac{d(z_1, z_3)d(z_2, z_4)}{d(z_1, z_4)d(z_2, z_3)}.$$

A map $F : (Z, d) \to (Z', d')$ is called *θ-quasi-Möbius*, where θ is a distortion function, if

$$[f(z_1), f(z_2), f(z_3), f(z_4)] \leq \theta([z_1, z_2, z_3, z_4])$$

for all quadruples $z_1, z_2, z_3, z_4 \in Z$, $z_1 \neq z_4$, $z_2 \neq z_3$.

Quasisymmetric maps are always quasi-Möbius, while quasi-Möbius maps are quasisymmetric provided the source and target space are simultaneously bounded or simultaneously unbounded.

1.3. Quasisymmetric invariants

In order to distinguish spaces up to some notion of equivalence, one requires invariants. By Example 1.2.10(3), quasisymmetry is an equivalence relation on the class of metric spaces. In this section, we indicate a few important quasisymmetric invariants: the doubling property, uniform perfectness, and uniform disconnectedness.

1.3.1. Doubling spaces.

DEFINITION 1.3.1. A metric space (Z, d) is *doubling* if there exists a constant $N < \infty$ so that any ball $\overline{B}(z, r)$ in Z can be covered by at most N balls of radius $r/2$. We say that (Z, d) is N-doubling.

The doubling condition is a finite-dimensionality hypothesis which controls the growth of the cardinalities of separated subsets of balls at all scales and locations. (This remark is made precise in Definition 1.4.14.) Examples of doubling spaces include Euclidean space of any dimension and arbitrary subsets thereof.

1.3.2. Uniformly perfect spaces.

DEFINITION 1.3.2. A metric space (Z, d) is *uniformly perfect* if there exists a constant $c > 0$ so that $\overline{B}(z, r) \setminus B(z, cr) \neq \emptyset$ for all $z \in Z$ and $0 < r < \operatorname{diam} Z$. We say that (Z, d) is c-uniformly perfect.

Uniform perfectness is a quantitative, scale-invariant version of topological perfectness: it asserts the nonexistence of separating annuli of large modulus. Every connected metric space is uniformly perfect. The Cantor set is uniformly perfect.

1.3.3. Uniformly disconnected spaces.

DEFINITION 1.3.3. Let (Z, d) be a metric space. A finite sequence of points $z_0, z_1, \ldots, z_N$ in X is called a *δ-chain* if $d(z_{i-1}, z_i) \leq \delta d(z_0, z_N)$ for all $i = 1, \ldots, N$. We do not require that the points be distinct. Such a chain is called *trivial* if all of its points coincide. We say that (Z, d) is *uniformly disconnected* if there exists $\delta > 0$ so that there are no nontrivial δ-chains in Z. In this case we also say that (Z, d) is *δ-uniformly disconnected*.

Chains should be viewed as finite analogs of curves. Indeed, in any path connected space, every pair of points can be joined by δ-chains with arbitrarily small δ. Uniform disconnectedness is a quantitative, scale-invariant version of total disconnectedness: it asserts the existence of separating annuli of a definite modulus at all scales and locations. The Cantor set is uniformly disconnected for some specific positive value of δ.

THEOREM 1.3.4. *Let (Z_1, d_1) have property P, where P is one of the following conditions:*

- *the doubling property,*
- *uniform perfectness,*
- *uniform disconnectedness.*

Let $f : Z_1 \to Z_2$ be a quasisymmetric homeomorphism of Z_1 onto another metric space (Z_2, d_2). Then Z_2 has property P. The statement is quantitative.

1.4. Dimensions and measures

The subject of this survey is another quasisymmetric invariant: conformal dimension. Versions of conformal dimension are defined (see section 2.2) by infimizing certain metrically defined notions of dimension over all quasisymmetric images. In this section, we review several notions of dimension for metric spaces. We begin with the most well-known and useful of these notions, the Hausdorff dimension.

1.4.1. Hausdorff measure and dimension. The Hausdorff measures on a metric space arise as instances of a general construction due to Carathéodory. As we will make use of this construction in several different ways, we begin with a general review.

Carathéodory's construction. Let (Z, d) be a metric space. The data associated to Carathéodory's construction are a family $\mathcal{F}$ of subsets of Z and a function $\phi : \mathcal{F} \to [0, +\infty]$ (called a *gauge function*).

Typical examples of families $\mathcal{F}$ which we will use are the full power set of Z, the family of all metric balls in Z, or (in the case $Z = \mathbb{R}^N$) the family of dyadic cubes. We usually consider gauge functions ϕ which depend only on the diameter of the argument, although this is not essential and there are situations where more general gauges are relevant. See, for example, the proof of Lemma 6.2.8.

Associated to (Z, d), $\mathcal{F}$ and ϕ, and for each $0 < \delta \leq \infty$, we define the following set function:

$$\Phi_\delta(A) := \inf \left\{ \sum_{i=1}^{\infty} \phi(A_i) : A \subset \bigcup_{i=1}^{\infty} A_i, \operatorname{diam} A_i < \delta, A_i \in \mathcal{F} \right\}, \qquad A \subset Z.$$

Some comments are in order. We have imposed no further assumptions on $\mathcal{F}$, thus, it may happen that the infimum in question is taken over the empty

set for certain choices of δ and A. In this case, we assign $\Phi_\delta(A) = +\infty$. In case $A = \emptyset$ we allow the trivial (empty) covering (one way to do this is to include the empty set as an element of $\mathcal{F}$ and assign $\phi(\emptyset) = 0$), thus $\Phi_\delta(\emptyset) = 0$ for all δ. With these remarks in mind, the value of $\Phi_\delta(A)$ is well-defined for all δ and A.

We call $(\Phi_\delta)_{\delta>0}$ the family of *Carathéodory premeasures* associated to $\mathcal{F}$ and ϕ. We also call Φ_∞ the *Carathéodory content*.

For each $0 < \delta \leq \infty$, the function Φ_δ is an outer measure on Z. Indeed, it is clearly monotone and countably subadditive, and by convention it vanishes on the empty set. It is not necessarily a Borel measure, however. The following definition overcomes this shortcoming.

DEFINITION 1.4.1. The *Carathéodory measure* associated to the data $\mathcal{F}$ and ϕ is

$$\Phi(A) := \lim_{\delta\to 0} \Phi_\delta(A) = \sup_{\delta>0} \Phi_\delta(A),$$

where $A \subset Z$.

The existence of the limit and the equality of these two formulas is an easy consequence of the monotonicity of the function $\delta \mapsto \Phi_\delta(A)$.

We use the following criterion, also due to Carathéodory, to verify that Φ is a Borel measure.

THEOREM 1.4.2. *Let μ be an outer measure on a metric space (Z, d). Then μ is a Borel measure if and only if $\mu(A \cup B) = \mu(A) + \mu(B)$ for every pair of sets $A, B \subset Z$ with $\operatorname{dist}(A, B) > 0$.*

COROLLARY 1.4.3. *Carathéodory measures are Borel.*

PROOF. It suffices to verify the condition in Theorem 1.4.2. For fixed $A, B \subset Z$ with $\operatorname{dist}(A, B) > 0$, we have $\Phi_\delta(A\cup B) = \Phi_\delta(A)+\Phi_\delta(B)$ whenever $\delta \leq \frac{1}{2}\operatorname{dist}(A, B)$. Taking the limit as $\delta \to 0$ completes the proof. □

Hausdorff measure and content. We now specialize to the most important example. Let $\mathcal{F}$ be the power set of Z, let $s \geq 0$, and let $h_s : \mathcal{F} \to [0, \infty]$ be given by

$$h_s(A) := (\operatorname{diam} A)^s.$$

DEFINITION 1.4.4. The *Hausdorff premeasures*, resp. *Hausdorff content*, resp. *Hausdorff measure*, in dimension s are the functions $(\mathcal{H}^s_\delta)_{\delta>0}$, $\mathcal{H}^s_\infty$, $\mathcal{H}^s$, given by Carathéodory's construction for the gauge function $\phi = h_s$ on $\mathcal{F}$.

By the preceding remarks, $\mathcal{H}^s$ is a Borel measure on Z for any $s \geq 0$.

EXAMPLES 1.4.5. On any metric space, the Hausdorff measure in dimension zero coincides with the *counting measure*

$$\mathcal{H}^0(A) = \#A.$$

When $Z = \mathbb{R}^N$ and $s = m \le N$ is an integer, $\mathcal{H}^s$ coincides with the usual surface area measure for m-dimensional smooth submanifolds in $\mathbb{R}^N$, up to a positive multiple. In particular, $\mathcal{H}^N$ coincides with the Lebesgue measure $\mathcal{L}^N$ up to a multiple. For intervals $I \subset \mathbb{R}$, $\mathcal{L}^1(I) = \mathcal{H}^1(I) = \mathcal{H}^1_\infty(I) = \operatorname{diam} I$.

REMARK 1.4.6. Let $\mathcal{F}$ be the family of all metric balls in (Z, d), let $\phi_s(A) = (\operatorname{diam} A)^s$ for $A \in \mathcal{F}$, and let Φ^s_δ be the corresponding family of Carathéodory premeasures. Then

$$\mathcal{H}^s_\delta(A) \le \Phi^s_\delta(A) \le 2^s \mathcal{H}^s_\delta(A) \tag{1.4.1}$$

for all $A \subset Z$ and all $\delta \ge 0$. The left hand inequality is obvious; for the right hand inequality we note that any nonempty subset A of Z is contained inside a ball of radius $\operatorname{diam} A$ and use (1.1.1). These measures Φ^s are sometimes known as the *spherical Hausdorff measures* on Z.

REMARK 1.4.7. Let $Z = \mathbb{R}^N$ and let $\mathcal{F}$ be the family of dyadic cubes, i.e., closed cubes with sides parallel to the coordinate axes and vertices at consecutive dyadic rational coordinates. More precisely, $A \in \mathcal{F}$ if and only if

$$A = [k_1 2^{-m}, (k_1 + 1)2^{-m}] \times \cdots \times [k_N 2^{-m}, (k_N + 1)2^{-m}]$$

for integers m and $k_1, k_2, \ldots, k_N$. Let ϕ_s be as in Remark 1.4.6 and let Φ^s_δ be the corresponding family of Carathéodory premeasures. Then

$$\mathcal{H}^s_\delta(A) \le \Phi^s_\delta(A) \le (4\sqrt{N})^s \mathcal{H}^s_\delta(A) \tag{1.4.2}$$

for all $A \subset Z$ and all $\delta \ge 0$. The argument is similar to the previous one. These measures Φ^s are sometimes known as the *dyadic Hausdorff measures* on $\mathbb{R}^N$.

LEMMA 1.4.8. *Let $0 \le s < t < \infty$ and $A \subset Z$. If $\mathcal{H}^s(A) < \infty$, then $\mathcal{H}^t(A) = 0$.*

PROOF. For each $0 < \delta < \infty$ it follows from the definitions that $\mathcal{H}^t_\delta(A) \le \delta^{t-s} \mathcal{H}^s_\delta(A) \le \delta^{t-s} \mathcal{H}^s(A)$. Let $\delta \to 0$. □

DEFINITION 1.4.9. The *Hausdorff dimension* of a set $A \subset Z$ is

$$\dim A = \inf\{s \ge 0 : \mathcal{H}^s(A) = 0\} = \sup\{s \ge 0 : \mathcal{H}^s(A) = +\infty\}.$$

Note that the equality of these two formulas follows immediately upon considering the contrapositive in Lemma 1.4.8. Throughout this survey, we will denote by dim the Hausdorff dimension. On occasion, if we wish to distinguish it from other notions of dimension, we may write $\dim_H$ for the Hausdorff dimension.

Hausdorff dimension is monotonic:

$$\dim A \le \dim B \qquad \text{if } A \subset B$$

and countably stable:

$$\dim \left(\bigcup_{i=1}^{\infty} A_i \right) = \sup_i \dim A_i \qquad \text{for } A_1, A_2, \ldots \text{ in } Z.$$

We will see many examples of dimensions of specific sets and metric spaces throughout this text. For now we merely observe that the Hausdorff dimension of $\mathbb{R}^N$ (or any open set in $\mathbb{R}^N$) is equal to N. Similarly, the Hausdorff dimension of any smooth m-dimensional submanifold of $\mathbb{R}^N$ is equal to m.

REMARK 1.4.10. By (1.4.1) and (1.4.2), the value of $\dim A$ is unchanged if we replace the Hausdorff measures in its definition with either the spherical Hausdorff measures or the dyadic Hausdorff measures (in $\mathbb{R}^N$).

Estimates for the Hausdorff dimension of a set A in Z arise primarily via measure-theoretic arguments. It suffices to find Borel measures supported on A with suitable volume growth properties. Two primary results along these lines are the *Mass Distribution Principle* and its converse, *Frostman's lemma*.

THEOREM 1.4.11. *(1) (Mass Distribution Principle) If there exists a positive measure μ on A satisfying $\mu(B(z,r)) \le r^s$ for all $B(z,r) \subset Z$, then $\mathcal{H}^s(A) > 0$.*

(2) (Frostman's lemma) If A is compact and $\mathcal{H}^s(A) > 0$, then there exists a positive Borel measure μ on A satisfying $\mu(B(z,r)) \le r^s$ for all $B(z,r) \subset Z$.

The proof of (1) is trivial: $\mathcal{H}^s(A) \ge \mathcal{H}^s_\infty(A) \ge \mu(A) > 0$. The proof of Frostman's lemma is significantly more involved; see the notes to this section for specific references.

Every uniformly perfect metric space has positive Hausdorff dimension. Indeed, by iterating the uniform perfectness condition one can construct a subset which is bi-Lipschitz equivalent to a Cantor set of positive dimension.

1.4.2. Topological dimension. We will denote by $\dim_T Z$ the *topological dimension* (or *small inductive dimension*) of Z. This is defined as follows:

$$\dim_T \emptyset = -1$$

and $\dim_T Z \le n$ if Z admits a neighborhood basis of sets $\{A_i\}$ whose boundaries satisfy $\dim_T \partial A_i \le n-1$. It is clear that $\mathbb{R}^N$ has topological dimension N. Any compact totally disconnected metric space has topological dimension zero. Note that the topological dimension is always an integer, unlike Hausdorff dimension which may assume any nonnegative real value.

The main fact which we want to highlight is the following.

THEOREM 1.4.12. *For any metric space Z, the inequality*

$$\dim_T Z \le \dim Z \tag{1.4.3}$$

holds.

There is a beautiful "Fubini-type" proof for Theorem 1.4.12 which reduces it to the following implication: *if $\mathcal{H}^{n+1}(Z) = 0$, then $\mathcal{H}^n(\partial B(z,r)) = 0$*

for each $z \in Z$ and for $\mathcal{L}^1$-a.e. $r > 0$. Since the metric balls are a neighborhood basis in Z, Theorem 1.4.12 follows by induction on the topological dimension.

1.4.3. Ahlfors regularity. Many natural spaces which arise in geometric function theory, dynamics and rigidity exhibit strong self-similarity properties. This self-similarity implies that the space looks similar at all locations and scales. Ahlfors regularity provides one quantitative measurement of such self-similarity, from the perspective of measure theory. A metric space is Ahlfors regular, roughly speaking, if the measures of balls of a fixed radius are all comparable, and scale like an appropriate power of the radius. Here is the precise definition.

DEFINITION 1.4.13. A metric space (Z, d) is *Ahlfors s-regular* for some $s > 0$ if there exists a Borel measure μ on Z and a constant $C < \infty$ so that

$$\frac{1}{C} r^s \leq \mu(\overline{B}(z, r)) \leq C r^s \tag{1.4.4}$$

for all balls $\overline{B}(z, r)$ in Z with $r < \operatorname{diam} Z$. We say that (Z, d) is *Ahlfors regular* if it is Ahlfors s-regular for some $s > 0$. We also say that μ is an *Ahlfors regular measure* on (Z, d).

If (Z, d) is Ahlfors s-regular with respect to some Borel measure μ, then it is Ahlfors s-regular with respect to the Hausdorff measure $\mathcal{H}^s$ (possibly with a different constant C) and μ and $\mathcal{H}^s$ are comparable. Indeed, for any covering of a ball $B(z, r)$ with sets A_i, we compute

$$\frac{1}{C} r^s \leq \mu(B(z, r)) \leq \sum_i \mu(A_i) \leq \sum_i (\operatorname{diam} A_i)^s$$

(where we included each set A_i in a ball of radius equal to its diameter), from which it follows that

$$\mathcal{H}^s(B(z, r)) \geq \mathcal{H}^s_\infty(B(z, r)) \geq \frac{1}{C} r^s.$$

The remaining direction is similar, but uses also a suitable covering theorem (see Lemma 4.1.1).

For each $s > 0$ there exist Ahlfors s-regular metric spaces. Examples can be selected from among the self-similar Cantor subsets of Euclidean spaces. In fact, all self-similar subsets of Euclidean space generated by iterated function systems satisfying the so-called open set condition are Ahlfors regular. See section 8.3 for more details.

Every Ahlfors regular space is uniformly perfect. Indeed, suppose that (Z, d) satisfies (1.4.4) for some $s > 0$, some C and some measure μ. If $\overline{B}(z, r) \setminus B(z, cr)$ is empty for some positive r, then $r^s/C \leq \mu(\overline{B}(z, r)) = \mu(\overline{B}(z, cr)) \leq C(cr)^s$ and $c \geq C^{-2/s}$. Thus Z is $\frac{1}{2}C^{-2/s}$-uniformly perfect.

Similarly, every Ahlfors s-regular space with $s < 1$ is uniformly disconnected. For a more general statement, see section 1.4.4.

1.4.4. Assouad dimension. In a metric space (Z, d), call a subset S *ϵ-separated* if $d(z_1, z_2) \geq \epsilon$ for all distinct $z_1, z_2 \in S$. Let $N(S, \epsilon)$ be the maximal cardinality of an ϵ-separated subset of S. Then Z is doubling (as in Definition 1.3.1) if and only if there exists $C < \infty$ so that

$$N(B(z, r), r/2) \leq C \tag{1.4.5}$$

for all balls $B(z, r)$ in Z with $r < \operatorname{diam} Z$. An easy inductive argument shows that if (1.4.5) holds, then there exists C' and s depending only on C so that

$$N(B(z, r), \epsilon r) \leq C' \epsilon^{-s} \tag{1.4.6}$$

for all balls $B(z, r)$ in Z with $r < \operatorname{diam} Z$ and all $0 < \epsilon < 1$. For instance, we may take $s = \log_2 C$.

Definition 1.4.14. The *Assouad dimension*, $\dim_A Z$, of a doubling metric space (Z, d) is the infimum of those $s > 0$ for which (1.4.6) holds for some $C' < \infty$, for all balls $B(z, r)$ and $0 < \epsilon \leq 2$.

Thus (Z, d) is doubling precisely when it has finite Assouad dimension. The inequality

$$\dim Z \leq \dim_A Z \tag{1.4.7}$$

between Hausdorff and Assouad dimension is valid for all spaces (Z, d); this is an easy exercise. Every nonempty open set in $\mathbb{R}^N$ has Assouad dimension N. Assouad dimension is monotonic: $\dim_A S \leq \dim_A Z$ whenever $S \subset Z$.

Let (Z, d) be an Ahlfors s-regular space. Then (1.4.6) holds for some C' depending only on the constant C from (1.4.4) and on s. To see this, note that if $z_1, \ldots, z_N$ are points in a ball $B(z, r)$ which are all at distance ϵr from each other ($\epsilon \leq 2$), then the balls $B(z_i, \frac{1}{2}\epsilon r)$ are all disjoint and contained in $B(z, 2r)$. The estimate $N \leq C^2 4^s \epsilon^{-s}$ follows from volume considerations. From this and the previous remarks we conclude that

$$Z \text{ Ahlfors } s\text{-regular} \quad \Longrightarrow \quad \dim Z = \dim_A Z = s. \tag{1.4.8}$$

In Chapter 5 we will prove that every metric space with Assouad dimension strictly less than one is uniformly disconnected. It follows that Ahlfors s-regular spaces with $s < 1$ are uniformly disconnected.

Several warnings are in order. Assouad dimension is not countably stable (Example 1.4.15). Also, the Assouad dimension of a metric space Z agrees with the Assouad dimension of its completion, thus the rationals $\mathbb{Q}$ have Assouad dimension one. These properties illustrate that Assouad dimension behaves quite differently from its Hausdorff counterpart.

Example 1.4.15. For each $p > 0$, the set $\{1, 2^{-p}, 3^{-p}, 4^{-p}, \ldots\}$ has Assouad dimension equal to one.

In Chapter 6 we will give an alternate derivation of Example 1.4.15 by showing that Assouad dimension does not increase upon passing to weak

tangents. It is not hard to verify that $\{1, 2^{-p}, 3^{-p}, 4^{-p}, \ldots\}$ has the set $[0, \infty)$ as a tangent space at the origin. Then

$$1 = \dim_A [0, \infty) \leq \dim_A \{1, 2^{-p}, 3^{-p}, 4^{-p}, \ldots\} \leq \dim_A \mathbb{R} = 1$$

and equality holds throughout.

A fundamental property of doubling metric spaces is Assouad's Embedding Theorem.

THEOREM 1.4.16 (Assouad). *If (Z, d) is a doubling metric space and $\epsilon < 1$, then there exists $N < \infty$ so that the snowflaked space (Z, d^ϵ) admits a bi-Lipschitz embedding in $\mathbb{R}^N$.*

The conclusion in Theorem 1.4.16 need not hold if $\epsilon = 1$. Non-abelian Carnot groups, equipped with their standard Carnot-Carathéodory (sub-Riemannian) metric, provide counterexamples. We discuss the sub-Riemannian geometry of Carnot groups in more detail in chapter 3.

1.4.5. Behavior of dimensions under mappings. Lipschitz and Hölder maps behave well with respect to Hausdorff dimensions and Ahlfors regularity, but can behave rather pathologically with respect to Assouad dimension.

THEOREM 1.4.17. *(1) Let $f : Z \to Z'$ be L-Lipschitz. Then $\mathcal{H}^s_{L\delta}(f(A)) \leq L^s \mathcal{H}^s_\delta(A)$ for all $A \subset Z$ and all $0 \leq \delta \leq \infty$.*

(2) Let $f : Z \to Z'$ be (L, α)-Hölder. Then $\mathcal{H}^{s/\alpha}_{L\delta^\alpha}(f(A)) \leq L^{s/\alpha} \mathcal{H}^s_\delta(A)$ for all $A \subset Z$ and all $0 \leq \delta \leq \infty$.

The proof is obvious.

COROLLARY 1.4.18. *(1) If $f : Z \to Z'$ is Lipschitz, then $\dim f(A) \leq \dim A$ for all $A \subset Z$. If $f : Z \to Z'$ is a bi-Lipschitz homeomorphism, then*

$$\dim f(A) = \dim A$$

for all $A \subset Z$.

(2) If $f : Z \to Z'$ is α-Hölder, then $\dim f(A) \leq \alpha^{-1} \dim A$ for all $A \subset Z$. If $f : Z \to Z'$ is an (α, β)-bi-Hölder homeomorphism (i.e., f is α-Hölder and f^{-1} is β-Hölder) for some α, β with $\alpha\beta \leq 1$, then

$$\beta \dim A \leq \dim f(A) \leq \alpha^{-1} \dim A$$

for all $A \subset Z$.

(3) For any $\epsilon < 1$ and any metric space (Z, d), $\dim(Z, d^\epsilon) = \epsilon^{-1} \dim Z$.

It is also easy to see that the image of an Ahlfors s-regular metric space under a bi-Lipschitz homeomorphism is again s-regular. The image of an Ahlfors regular metric space under an (α, β)-bi-Hölder homeomorphism need not be Ahlfors regular, but the snowflake of an s-regular metric space Z with parameter $\epsilon < 1$ is s/ϵ-regular.

Bi-Hölder homeomorphisms can also be ill-behaved with respect to the Assouad dimension.

EXAMPLE 1.4.19. There exists a countable doubling metric space (Z, d) and a Lipschitz homeomorphism $f : Z \to Z'$ onto a non-doubling space (Z', d'). In fact, for each $\beta < 1$ we construct a $(1, \beta)$-bi-Hölder homeomorphism from a space Z with $\dim_A Z = 1$ onto a space Z' with $\dim_A Z' = \infty$.

By Example 1.4.15 the Assouad dimension of $\{n^{-p} : n \in \mathbb{N}\}$ is equal to one, for each $p > 0$. By the finite stability of the Assouad dimension, $\{n^{-p} : n \geq n_0\}$ has Assouad dimension equal to one for any (fixed) $n_0 \in \mathbb{N}$. Let us consider $Z = \{n \in \mathbb{N} : n \geq n_0\}$ with the metric $d_{(p)}$ defined by $d_{(p)}(m, n) = |m^{-p} - n^{-p}|$; then $(Z, d_{(p)})$ is isometric with $\{n^{-p} : n \geq n_0\}$.

Let $\beta < 1$ and define $d' : \mathbb{N} \times \mathbb{N} \to \infty$ by $d'(m, n) = \max\{m^{-1/\beta}, n^{-1/\beta}\}$ if $m \neq n$ and $d'(m, n) = 0$ if $m = n$. We leave it to the reader to check that d' is a metric on $\mathbb{N}$.

If we choose $p > 0$ such that $p < \frac{1}{\beta} - 1$ and choose $n_0 = n_0(p) \in \mathbb{N}$ appropriately, then the identity map from $(Z, d_{(p)})$ to (Z, d') satisfies

$$\frac{1}{C} d_{(p)}(m, n)^{1/\beta} \leq d'(m, n) \leq C d_{(p)}(m, n)$$

for all $m, n \in Z$ and a suitable constant $C = C(p) < \infty$, i.e., the identity map is a $(1, \beta)$-bi-Hölder homeomorphism. In particular, it is Lipschitz.

We show that (Z, d') is not doubling. Again by the finite stability of the Assouad dimension, it suffices to show that $(\mathbb{N}, d')$ is not doubling. Since

$$\lim_{m \to \infty} \operatorname{diam}_{d'}\{m, m+1, m+2, \ldots\} = \lim_{m \to \infty} m^{-1/\beta} = 0,$$

(m) is a Cauchy sequence in $(\mathbb{N}, d')$. There is a unique limit point for this Cauchy sequence in the completion of $(\mathbb{N}, d')$; we denote this point by 0. The metric d' extends to $\{0\} \cup \mathbb{N}$ and $d'(0, m) = m^{-1/\beta}$. We will show that

$$\lim_{r \to 0} N(B(0, r), \tfrac{1}{2} r) = \infty$$

which implies that $(\mathbb{N}, d')$ is not doubling.

Observe that $B(0, r) = \{0\} \cup \{m \in \mathbb{N} : m > r^{-\beta}\}$ and that the points $m, m+1, \ldots, m+p \in B(0, r)$ are all $\frac{1}{2}r$-separated provided

$$\left(\frac{1}{m+p}\right)^{1/\beta} \geq \frac{r}{2},$$

i.e., $m + p \leq (2/r)^{\beta}$. Choose $p = [(2^{\beta} - 1)r^{-\beta}] - 1$ and $m_0 = [r^{-\beta} + 1]$, where $[\cdot]$ denotes the integer part function. Then $m_0, \ldots, m_0 + p$ are all $(\frac{1}{2}r)$-separated and

$$N(B(0, r), \tfrac{1}{2} r) \geq p + 1 = [(2^{\beta} - 1)r^{-\beta}]$$

which tends to infinity as $r \to 0$.

However, Assouad dimension is invariant under bi-Lipschitz maps.

We conclude this chapter with an application of Theorem 1.4.17.

PROPOSITION 1.4.20. *For any connected metric space (Z, d), we have* $\operatorname{diam} Z \leq \mathcal{H}^1_{\infty}(Z)$.

PROOF. Suppose first that Z is unbounded. We claim that $\mathcal{H}^1_\infty(Z) = +\infty$. Fix $z_0 \in Z$ and consider the map $f : Z \to \mathbb{R}$ given by $f(z) = d(z, z_0)$. This map is 1-Lipschitz, so Theorem 1.4.17 implies that

$$\mathcal{H}^1_\infty(Z) \geq \mathcal{H}^1_\infty(f(Z)) = \mathcal{H}^1_\infty([0, +\infty)) = \mathcal{L}^1([0, +\infty)) = \infty.$$

Here we used the fact that $f(Z)$ is an unbounded interval in the nonnegative real axis containing 0, hence is equal to $[0, \infty)$.

Now suppose that Z is bounded. Let $\epsilon > 0$ and fix $z_0 \in Z$ so that

$$\sup_{z \in Z} d(z, z_0) \geq \operatorname{diam} Z - \epsilon.$$

Let f be as above. Again, $f(Z)$ is an interval in the nonnegative real axis containing 0, with

$$\mathcal{L}^1(f(Z)) \geq \sup_{z \in Z} |f(z) - f(z_0)| \geq \operatorname{diam} Z - \epsilon.$$

As before we find

$$\mathcal{H}^1_\infty(Z) \geq \mathcal{H}^1_\infty(f(Z)) = \mathcal{L}^1(f(Z)) \geq \operatorname{diam} Z - \epsilon.$$

Now let $\epsilon \to 0$. □

As a corollary, note that every space of Hausdorff dimension strictly less than one is totally disconnected. (If such a space Z contained a nontrivial connected component, we would conclude that $\mathcal{H}^1(Z) > 0$.)

1.5. Notes

Rademacher's theorem 1.2.2 is a remarkable fact, highly indicative of the geometry of the underlying Euclidean space. It asserts that Lipschitz functions are intrinsically equipped with a powerful infinitesimal rigidity property (local linearity), at least away from a null set. Rademacher's original paper appeared in 1919 [**144**] and has stimulated a wealth of subsequent generalizations, reformulations and applications which we cannot hope to describe in detail here. We mention only the striking extension of Theorem 1.2.2 by Cheeger [**44**] to the category of arbitrary doubling metric measure spaces satisfying a Poincaré inequality. This generalization extends the Rademacher theorem to a variety of new spaces with nonsmooth infinitesimal geometry, such as Carnot groups and other Carnot-Carathéodory manifolds (section 3.3.5) and the boundaries of certain hyperbolic buildings (section 3.5). It is now clear, as a consequence of the results of [**44**] that the validity of a Rademacher-type differentiation theorem has strong implications for the geometry of the underlying space.

McShane's Extension Theorem 1.2.3 has also opened new lines of investigation which have expanded dramatically in recent years. The existence of **a** Lipschitz extension is quite easy to prove. In fact, given any L-Lipschitz function $f : A \to \mathbb{R}$ defined on a nonempty set $A \subset Z$, it is easy to check that the formulas

$$g_+(z) := \inf\{f(a) + Ld(a, z) : a \in A\}$$

and

$$g_-(z) := \sup\{f(a) - Ld(a,z) : a \in A\}$$

define L-Lipschitz extensions of f to all of Z. Furthermore, $g_-(z) \leq g_+(z)$ for all $z \in Z$. The existence of *absolute minimizing Lipschitz extensions* (AMLE's), i.e., extensions which minimize the Lipschitz constant locally, was first shown by Aronsson [**2**] in the Euclidean setting. The study of such extensions involves significant use of tools from nonlinear PDE, especially the theory of the ∞-Laplace equation. The uniqueness of AMLE's was later established by Jensen [**86**]. See [**3**] for a comprehensive overview of the theory of ∞-harmonic functions and absolutely minimizing Lipschitz extensions in the Euclidean setting. Quite recently, Peres, Schramm, Sheffield and Wilson [**143**] provided a game-theoretic reinterpretation of the ∞-Laplace equation in terms of "tug-of-war" and thereby obtain a probabilistic approach to Jensen's uniqueness theorem valid in extremely general settings. This area of research is still developing rapidly.

Lipschitz extension theorems with other target spaces are also of some interest. Lang, Pavlović and Schroeder [**100**], [**99**] consider the case of target spaces with curvature bounds. In a series of papers [**103**], [**104**], Lee and Naor consider the problem of Lipschitz extension where the subset A is fixed and the superset Z is allowed to vary. They obtain striking results on this *absolute extension problem* for a variety of spaces and targets, including doubling metric spaces, graphs and Banach spaces.

Theorem 1.2.12 and Remark 1.2.15 are folklore, implicit in early work of Gehring [**57**], [**58**] and Väisälä [**174**]. The local quasisymmetry of quasiconformal maps is sometimes known as the *egg yolk principle*; Heinonen [**76**, Chapter 11] attributes the term to McMullen. The precise definition of quasisymmetric maps in the metric space context was first given by Tukia and Väisälä [**166**]. However, the notion was implicit in many earlier works. In particular, quasisymmetric maps on the real line (as the boundary maps of quasiconformal maps of the upper half plane) have been well studied since at least the work of Ahlfors and Beurling [**14**]. The canonical reference for quasiconformal maps in $\mathbb{R}^N$ is Väisälä's monograph [**175**]. For a recent treatment, see Iwaniec–Martin [**84**].

Liouville's theorem is a strong rigidity statement about conformal mappings in Euclidean spaces of sufficiently high dimension. It asserts that any conformal mapping between domains in $\mathbb{R}^N$ is the restriction of a Möbius transformation, provided $N \geq 3$. Liouville's original proof required a suitable a priori regularity assumption. The version of Liouville's theorem stated in Example 1.2.17 is due to Gehring [**58**]. The result is false when $N = 2$; every analytic map with nonvanishing derivative is conformal. Liouville's theorem also holds in more general spaces such as Carnot groups.

Gehring's higher integrability lemma (Remark 1.2.14) is one of the most important tools in geometric function theory; it instigated numerous investigations of self-improving integrability properties which continue to the

present day. The sharp integrability exponent $p(N,K)$ in Remark 1.2.14 is not known. A well-known conjecture asserts that

$$p(N,K) = \frac{NK^{1/(N-1)}}{K^{1/(N-1)}-1}; \tag{1.5.1}$$

this value, if correct, would be sharp as indicated by the radial stretch map $F(x) = |x|^{(1/K)-1}x$. In a seminal work, Astala [**5**] established the conjecture (1.5.1) in the case $N = 2$.

THEOREM 1.5.1 (Astala). *Let $F : \Omega \to \Omega'$ be a $W^{1,2}_{loc}$ homeomorphism between domains $\Omega, \Omega' \subset \mathbb{C}$ satisfying the distortion inequality $||DF||^2 \leq K \det DF$ a.e. Then $F \in W^{1,p}_{loc}$ for each $p < \frac{2K}{K-1}$.*

Gehring's higher integrability lemma and Astala's theorem 1.5.1 have direct applications to the study of Hausdorff dimension distortion by quasiconformal maps. See Chapter 8 for further information. A comprehensive account of the theory of elliptic partial differential equations and quasiconformal mappings in the plane can be found in the recent book of Astala, Iwaniec and Martin [**6**].

An excellent reference for the theory of topological dimension is the book of Hurewicz and Wallman [**81**]. Theorem 1.4.12 can be found there; see also [**76**, Theorem 8.14]. For the theory of Hausdorff dimension, we recommend the books of Rogers [**149**] and Falconer [**52**], [**51**]. A vast amount of information can be found in Federer's monumental treatise [**53**]. For a more recent discussion, see Mattila [**116**].

Assouad dimension was considered (under a different name) by Assouad [**4**]. An extensive development of the history and theory of Assouad dimension was provided by Luukkainen [**107**]. Example 1.4.19 is a slight elaboration on an example in the appendix of [**107**].

Assouad's embedding theorem 1.4.16 [**4**] features in a surprisingly wide spectrum of applications. Bonk and Schramm [**26**] use it to construct coarse embeddings of Gromov hyperbolic spaces, Hajłasz [**71**] uses it to build Whitney-style surjections of high regularity with large critical locus, and Semmes [**153**] uses it to establish a bi-Lipschitz embedding theorem for doubling spaces with target equipped with a metric defined via deformation by an A_∞ weight (compare section 7.1).

CHAPTER 2

Conformal gauges and conformal dimension

2.1. Conformal gauges

The problem of classifying metric spaces up to quasisymmetry arises naturally in many contexts. Conformal gauges provide a convenient language for discussing quasisymmetric invariance. We recall the definition.

DEFINITION 2.1.1. The *conformal gauge* $\mathcal{G}$ of a metric space (Z, d) is the collection of all metric spaces (Z', d') which are quasisymmetrically equivalent to (Z, d).

Any quasisymmetrically invariant property of metric spaces can be understood as a property of the associated conformal gauge. For example, by Theorem 1.3.4, we can speak of a doubling conformal gauge, a uniformly perfect conformal gauge, or a uniformly disconnected conformal gauge.

The following theorem characterizes the conformal gauge of the Cantor set. It is a quantitative version of the well-known topological characterization of the Cantor set as the unique compact, perfect, totally disconnected metric space (up to homeomorphism).

THEOREM 2.1.2. *A metric space is quasisymmetrically equivalent to the Cantor set if and only if it is compact, uniformly perfect and uniformly disconnected.*

In other words, the Euclidean conformal gauge on the Cantor set is the unique conformal gauge which is compact, uniformly perfect and uniformly disconnected.

The proof of Theorem 2.1.2 consists of encoding the Cantor set quasisymmetrically via the abstract symbol space $\{0, 1\}^{\mathbb{N}}$ (equipped with a suitable self-similar metric), and then showing that any space satisfying the assumptions of the theorem admits a similar quasisymmetric coding by the same symbol space.

We note the following consequence of Theorem 2.1.2: if a conformal gauge $\mathcal{G}$ is compact, uniformly perfect and uniformly disconnected, then $\mathcal{G}$ is doubling.

Other spaces whose conformal gauges can be explicitly characterized include the real line $\mathbb{R}$ and the circle $\mathbb{S}^1$. In general, however, it is extremely difficult to write down a simple list of properties characterizing a given conformal gauge. For example, characterizing the conformal gauge of the sphere $\mathbb{S}^N$, $N \geq 2$, is one of the most difficult and important problems in

the area. We will have more to say about this problem later in this survey (see section 7.2).

In the language of conformal gauges, Assouad's embedding theorem 1.4.16 states that a conformal gauge $\mathcal{G}$ embeds in a Euclidean space if and only if it is doubling.

2.2. Conformal dimension

We now recall one of the most important quasisymmetry invariants, the focus of this survey. The following concept was first introduced by Pansu [**133**] in connection with the quasi-isometric classification problem for rank one symmetric spaces.

DEFINITION 2.2.1. The *conformal dimension*, $\operatorname{Cdim}\mathcal{G}$, of a conformal gauge $\mathcal{G}$ is the infimum of the Hausdorff dimensions of all metric spaces in $\mathcal{G}$.

We will abuse notation by writing $\operatorname{Cdim} Z$ for the conformal dimension of the conformal gauge associated to a given metric space (Z, d). Occasionally, we will write $\operatorname{Cdim}_H Z$ in place of $\operatorname{Cdim} Z$ to emphasize that Hausdorff dimension is under consideration.

Let (Z, d) be any metric space. Then

$$\dim_T Z = \inf \dim Z' \tag{2.2.1}$$

where the infimum is taken over all metric spaces (Z', d') which are homeomorphic with Z. In other words, there is never any obstruction to lowering the Hausdorff dimension of a metric space by a homeomorphism (apart from the obvious *a priori* restriction (1.4.3)). On the other hand, bi-Lipschitz homeomorphisms preserve Hausdorff dimension. Quasisymmetric homeomorphisms form a class which is intermediate between bi-Lipschitz homeomorphisms and arbitrary homeomorphisms. By (1.4.3) and (2.2.1), the following *a priori* inequalities hold true:

$$\dim_T Z \leq \operatorname{Cdim} Z \leq \dim Z.$$

It follows from Corollary 1.4.18(3) that the supremum of $\dim Z'$ over all spaces (Z', d') in a conformal gauge $\mathcal{G}$ is equal to either 0 or $+\infty$. Conformal dimension, on the other hand, is a subtle and useful invariant which reveals much about the structure of the gauge. The goal of this survey is to document and justify this claim.

As we will see later, it is of critical importance to determine when the Hausdorff dimension of a space cannot be decreased by quasisymmetries, i.e., when it is minimal for conformal dimension.

DEFINITION 2.2.2. A metric space (Z, d) is *minimal for conformal dimension* if

$$\operatorname{Cdim} Z = \dim Z.$$

EXAMPLES 2.2.3. (1) For each $n \geq 1$, Euclidean space $\mathbb{R}^N$ is minimal for conformal dimension. Indeed, $\operatorname{Cdim} \mathbb{R}^N = \dim \mathbb{R}^N = N$. More generally, any nonfractal metric space (i.e., a metric space whose Hausdorff and topological dimensions agree) is minimal for conformal dimension, for purely topological reasons.

(2) The Cantor set K is not minimal for conformal dimension. Indeed, $\dim K = \frac{\log 2}{\log 3}$ but $\operatorname{Cdim} K = 0$. To prove the latter result, consider a one-parameter family of Cantor sets $K_\lambda \subset [0,1]$, $0 < \lambda < \frac{1}{2}$, where K_λ consists of two subsets, one of which contains 0 and the other contains 1, so that each subset is similar to K by scale factor λ. Note that $K = K_{1/3}$. It is easy to verify that K_λ is doubling, uniformly perfect and uniformly disconnected. By Theorem 2.1.2, K_λ is quasisymmetrically equivalent to K for each $\lambda > 0$. (It is also straightforward to verify directly that the canonical monotone increasing homeomorphism $K \to K_\lambda$ is quasisymmetric.) Thus $\operatorname{Cdim} K \leq \frac{\log 2}{\log 1/\lambda}$ for each $\lambda > 0$ and hence $\operatorname{Cdim} K = 0$. Observe that this infimum is not achieved; every quasisymmetric image of K is uniformly perfect and hence has positive Hausdorff dimension.

In fact, there are no metric spaces with conformal dimension strictly between zero and one. More precisely, the range of conformal dimension is equal to $\{0\} \cup [1, \infty]$. The nonexistence of metric spaces with conformal dimension in the interval $(0,1)$ is due to Kovalev and is one of the major recent advances in the subject. We will discuss Kovalev's theorem further in Chapters 6 and 8.

Conformal dimension is obviously monotonic, but (in contrast with Hausdorff dimension) it is not countably stable. In fact, it is not even finitely stable: there exist sets $A, B \subset [0,1]$ of conformal dimension zero such that $A \cup B = [0,1]$! See Example 5.2.1. In section 5.2 we prove that finite unions of disjoint compact sets have conformal dimension no larger than the maximum of the conformal dimensions of the pieces. See Proposition 5.2.3. We observe the following modest corollary: *the conformal dimension of a compact space Z is not increased by adjoining a finite number of points to Z.*

Variant notions are occasionally more useful. Similar invariants can be built from any of the standard metric notions of dimension: Minkowski dimension, packing dimension, etc. One particularly useful variant is conformal Assouad dimension.

DEFINITION 2.2.4. The *conformal Assouad dimension*, $\operatorname{Cdim}_A \mathcal{G}$, of a conformal gauge $\mathcal{G}$ is the infimum of the Assouad dimensions of all metric spaces in $\mathcal{G}$.

Note that $\operatorname{Cdim}_A \mathcal{G}$ is finite if and only if $\mathcal{G}$ is doubling.

DEFINITION 2.2.5. The *Ahlfors regular conformal dimension*, $\operatorname{ARCdim} \mathcal{G}$, of a conformal gauge $\mathcal{G}$ is the infimum of the dimensions of all Ahlfors regular metric spaces in $\mathcal{G}$.

Note that a conformal gauge may contain both Ahlfors regular and non-Ahlfors regular metrics. However, every complete, doubling, uniformly perfect gauge contains a large supply of Ahlfors regular metrics. We will justify the latter assertion in chapter 7.

We record two elementary inequalities which follow from (1.4.7) and (1.4.8):

$$\operatorname{Cdim}_H \mathcal{G} \leq \operatorname{Cdim}_A \mathcal{G} \leq \operatorname{ARCdim} \mathcal{G}$$

for every conformal gauge $\mathcal{G}$. In fact, the second inequality is an equality for a large class of gauges.

PROPOSITION 2.2.6. *Let $\mathcal{G}$ be a complete, uniformly perfect conformal gauge. Then* $\operatorname{ARCdim} \mathcal{G} = \operatorname{Cdim}_A \mathcal{G}$. *Indeed, for any $s > \operatorname{Cdim}_A \mathcal{G}$ there exists a closed Ahlfors s-regular set Y contained in some Euclidean space $\mathbb{R}^N$ so that $Y \in \mathcal{G}$.*

2.3. Notes

Conformal dimension was introduced by Pansu [**133**] in connection with the quasi-isometric classification of rank one symmetric spaces. Bourdon [**27**], [**29**] extended the method to treat lattices in some hyperbolic buildings. We will review the results of Pansu and Bourdon in chapter 3 and discuss ideas from the proofs of those results in chapter 4.

The language of conformal gauges was suggested by Heinonen [**76**, Chapter 15] as a formalism for the study of quasisymmetrically invariant properties. Theorem 2.1.2 is due to David and Semmes [**47**, Proposition 15.11]. The quasisymmetric characterization of $\mathbb{S}^1$ is well-known [**166**, Theorem 4.9]: *A topological circle (Z, d) is quasisymmetrically equivalent to $\mathbb{S}^1$ if and only if it is doubling and has bounded turning.*

DEFINITION 2.3.1. A connected, locally connected metric space (Z, d) is said to have *bounded turning* if there exists $C < \infty$ so that any two points $z_0, z_1 \in Z$ are included in a connected set $A \subset Z$ with $\operatorname{diam} A \leq C d(z_0, z_1)$.

A topological circle Z has bounded turning if and only if the following condition holds for some constant $C < \infty$: *for all $a, b \in Z$,*

$$\min\{\operatorname{diam} \gamma_1, \operatorname{diam} \gamma_2\} \leq C d(a, b)$$

where γ_1 and γ_2 are the two maximal subarcs of $Z \setminus \{a, b\}$. A planar topological circle has bounded turning if and only if it is a *quasicircle*, i.e., it is the image of $\mathbb{S}^1$ under a quasiconformal map $\mathbb{C} \to \mathbb{C}$ [**14**].

Proposition 2.2.6 was proved by Juha Heinonen in 1996 during a graduate topics course at the University of Michigan and first appeared in print in [**76**, Chapter 14]. We will sketch the proof in chapter 7.