Harmonic maps and associated energy functionals

Ognjen Tošić
Keble College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2024

Acknowledgements

Firstly, I would like to thank my advisor Vladimir Marković for his advice, insight and encouragement. None of the work in this thesis would be possible without his help.

I am also grateful to Nathaniel Sagman for many interesting conversations on harmonic maps and gauge theory, both in Oxford and Luxembourg, and at several conferences.

Finally, I would like to thank all my fellow DPhil students from the topology group, for all the reading groups, junior seminars and subsequent lighthearted discussions at the Rickety Press, that have made it very easy to stay enthusiastic about mathematics over the years.

Abstract

In this thesis, we study the existence and applications of harmonic maps in negative curvature settings. We cover three separate, but related strands of research.

Firstly, we study the existence of harmonic maps between complete simply connected negatively curved manifold. There have been many results in this vein [49, 48, 43, 3, 5, 6], culminating in the result of Benoist–Hulin [5] asserting the existence of a harmonic map at a finite distance from an arbitrary quasi-isometry between pinched Hadamard manifolds (i.e. complete simply connected Riemannian manifolds with sectional curvatures bounded between two negative constants). Nearest-point projections to convex sets are natural objects of study that fall outside the scope of this work. Given a pinched Hadamard manifold M and a convex subset $K \subseteq M$, under certain conditions on K, we show that there exists a harmonic map $h\colon M \to M$ that is at a bounded distance from the nearest-point retraction $r\colon M \to K$. In particular, when M is n-dimensional hyperbolic space, we show this existence when K is the convex hull of (1) a quasicircle in the sphere at infinity of M, or (2) an open subset of the sphere at infinity with Lipschitz boundary.

Secondly, we apply harmonic maps to study the Putman–Wieland conjecture. This is an algebraic conjecture stating that, given a finite cover of closed surfaces $p : \tilde{\Sigma} \to \Sigma$ and a cohomology class $\chi \in H^1(\tilde{\Sigma}, \mathbb{Z}) \setminus \{0\}$, the orbit of χ under the group of mapping classes on Σ that lift via p to $\tilde{\Sigma}$ is infinite. Given such a p and χ, we define the energy functional $\mathrm{E} : \mathcal{T}(\Sigma) \to \mathbb{R}$ defined on the Teichmüller space $\mathcal{T}(\Sigma)$ of the Riemann surface structures on Σ as follows: given an $X \in \mathcal{T}(\Sigma)$, we can lift X to a Riemann surface Y marked by $\tilde{\Sigma}$ such that p is homotopic to a holomorphic immersion $Y \to X$. We then set $\mathrm{E}(X)$ to be the energy of the harmonic form on Y that corresponds to χ. Our main result is that χ is a counterexample to the Putman–Wieland conjecture if and only if E is

constant on $\mathcal{T}(\Sigma)$. As an application, we show that the Putman–Wieland conjecture holds for covers satisfying a suitable expansion property. This part of the thesis was obtained in joint work with my advisor Vladimir Marković [50].

Finally, this leads us naturally to study energy functionals on $\mathcal{T}(\Sigma)$ associated to representations into higher-rank Lie groups. Specifically, given a closed surface Σ and a completely reducible representation $\rho : \pi_1(\Sigma) \to \mathrm{GL}(n, \mathbb{C})$, by the non-abelian Hodge theorem, for any $X \in \mathcal{T}(\Sigma)$, there exists a harmonic ρ-equivariant map $f : \tilde{X} \to \mathrm{GL}(n, \mathbb{C})/\mathrm{U}(n)$, where $\tilde{X}$ is the universal cover of X. The Dirichlet energy of f over a fundamental domain in $\tilde{X}$ for the action of $\pi_1(\Sigma)$ then defines a function $\mathrm{E}_\rho : \mathcal{T}(\Sigma) \to \mathbb{R}$. It is a classical result of Toledo [66] that E_ρ is plurisubharmonic. For $n = 1$, this fact plays a crucial role in our study of the Putman–Wieland conjecture. More generally, we study the directions in $\mathcal{T}(\Sigma)$ along which this plurisubharmonicity of E_ρ is not strict. We give a classification of such directions in terms of the Higgs bundle obtained from ρ by the non-abelian Hodge correspondence. In particular, we show that for a generic representation ρ and a Riemann surface $X \in \mathcal{T}(\Sigma)$, the map E_ρ is strictly plurisubharmonic at X. We are also able to construct, for any $X \in \mathcal{T}(\Sigma)$, a representation ρ such that E_ρ is not strictly plurisubharmonic at X. Finally, we relate the points and directions where E_ρ is not strictly plurisubharmonic to the Hitchin fibration and the $\mathbb{C}^*$-action on the moduli space of Higgs bundles, two basic objects in non-abelian Hodge theory.

Contents

Chapter 1

Overview of results

In this thesis, we study harmonic maps in a non-positive curvature context. Our results can naturally be divided into three classes: existence results in negative curvature, applications of Teichmüller theory to mapping class groups, and the study of energy functionals on Teichmüller space associated to symmetric spaces. We discuss each of them in turn.

A Existence of harmonic maps

We first introduce our existence results. A fundamental conjecture in the theory of harmonic maps is the Schoen–Li–Wang conjecture [59, 44].

Conjecture 1.1 (Schoen–Li–Wang conjecture)**.** Let X be a non-compact symmetric space of rank one. Then given any quasi-isometry $f : X \to X$, there exists a unique harmonic map $h : X \to X$ that is at a bounded distance from f, i.e. $\sup_X \operatorname{dist}(f, h) < \infty$.

The uniqueness part of this conjecture was settled by Li–Wang [44]. Various special cases of Conjecture 1.1 were shown by Marković [48] (for 3-dimensional hyperbolic space $\mathbb{H}^3$), Marković [49] (for the hyperbolic plane $\mathbb{H}^2$), Lemm–Marković [43] (for higher dimensional hyperbolic spaces $\mathbb{H}^n, n \geq 3$). Conjecture 1.1 in full was shown by Benoist–Hulin [3]. They then generalized their proof to all pinched Hadamard manifolds (complete simply connected Riemannian manifolds with all sectional curvatures bounded between two negative constants), showing Theorem 1.2 below [5].

Theorem 1.2. *Let M, N be pinched Hadamard manifolds, and let $f : M \to N$ be a quasi-isometry. Then there exists a harmonic map $h : M \to N$ such that* $\operatorname{dist}(f, h)$ *is bounded.*

Our results have the aim of generalizing Theorem 1.2 to maps that are not necessarily quasi-isometries. The initial maps we consider are nearest-point projections to convex sets. Specifically we are able to show the existence of the harmonic map in two different regimes: when the convex set is "sufficiently small" (Chapter 3) and when it is "sufficiently large" (Chapter 4). The most general statements for pinched Hadamard manifolds are somewhat involved, so we defer them to introductory sections of Chapters 3 and 4. Below we state our results for hyperbolic spaces.

Theorem A.1. *Let $K \subseteq \mathbb{H}^n$ be the convex hull of either*

(1) a quasicircle in the boundary at infinity $\mathbb{S}^{n-1}$ of $\mathbb{H}^n$, or

(2) an open set with Lipschitz boundary in $\mathbb{S}^{n-1}$.

Then there exists a harmonic map $h : \mathbb{H}^n \to \mathbb{H}^n$ that is at a finite distance from the nearest-point projection $r : \mathbb{H}^n \to K$.

We will show Theorem A.1(1) in Chapter 3, and Theorem A.1(2) in Chapter 4.

Both parts of Theorem A.1 were shown by the author in [67, 69]. In this thesis, we present these results in a unified way, as well as a slight strengthening of the results of [67].

B Study of mapping class groups using Teichmüller theory

In Chapter 5, we use harmonic maps to study algebraic questions about mapping class groups. Specifically, we look at the Putman–Wieland conjecture on higher Prym representations. We remind the reader that the mapping class group $\mathrm{Mod}_{g,n}$ is the group of connected components (called mapping classes) of the orientation-preserving diffeomorphism group of $\Sigma_{g,n}$ (the surface without boundary of genus g with n punctures), fixing the punctures pointwise.

Definition 1.3. Let $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ be a covering map. Let Γ_π be the finite index subgroup of $\mathrm{Mod}_{g,n+1}$, consisting of mapping classes that lift to diffeomorphisms of $\Sigma_{h,m}$. Then the higher Prym representation is the induced action of Γ_π on the first cohomology $H^1(\Sigma_h, \mathbb{Z})$, where Σ_h is obtained from $\Sigma_{h,m}$ by filling in the punctures.

More precise version of this definition can be found in §2.2.1.1.

Conjecture 1.4 (Putman–Wieland conjecture)**.** In the setting of Definition 1.3, when $g \geq 3$, any non-zero orbit of Γ_π in $H^1(\Sigma_h, \mathbb{Z})$ is infinite.

Conjecture 1.4 was originally studied by Putman–Wieland [57] in relation to a famous question of Ivanov [37]: does the mapping class group have a finite index subgroup that admits a non-zero homomorphism to $\mathbb{Z}$? They show in [57] that a negative answer to the question of Ivanov is essentially equivalent to Conjecture 1.4.

The main result of Chapter 5 is the geometric reformulation of Conjecture 1.4. We state the result below as Theorem B.1, but we first introduce some notation.

Definition 1.5. Given a covering map $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$, any Riemann surface structure on $\Sigma_{g,n}$ can be lifted to $\Sigma_{h,m}$ in a unique way that makes π holomorphic. By filling in the punctures, this provides a map $\sigma_\pi : \mathcal{T}_{g,n} \to \mathcal{T}_h$ that is a holomorphic embedding.

In Definition 1.5 above, $\mathcal{T}_{g,n}$ is the Teichmüller space of marked Riemann surfaces of genus g with n punctures, and $\mathcal{T}_h$ is the Teichmüller space of marked Riemann surfaces of genus h.

Definition 1.6. Given a cohomology class $\chi \in H^1(\Sigma_g, \mathbb{R})$, and a Riemann surface X, there exists a unique harmonic 1-form ω in the de Rham class χ. The energy of this harmonic 1-form is called the Hodge norm of χ, and defines a map

$$\mathrm{E}_\chi : \mathcal{T}_g \to \mathbb{R}.$$

An elaboration on what precisely we mean by energy can be found in §2.2.1.2. We are now ready to state the main result of Chapter 5.

Theorem B.1. *Let $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ be a covering map. Then the class $\chi \in H^1(\Sigma_h, \mathbb{Z})$ has a finite Γ_π orbit if and only if E_χ is constant on the image of σ_π.*

As an application of Theorem B.1, we show the following result.

Definition 1.7. A covering map $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ is called geometrically ε-uniform if there exists a Riemann surface $X \in \mathcal{T}_{g,n}$ such that $\lambda_1(\sigma_\pi(X)) \geq \varepsilon$, where λ_1 denotes the bottom of the spectrum of the Laplacian relative to the hyperbolic metric.

Theorem B.2. *Let $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ be a geometrically $\frac{1}{2(g-1)}$-uniform covering map. Then all non-zero orbits of Γ_π in $H^1(\Sigma_h, \mathbb{Z})$ are infinite.*

All of the results in this section were obtained by the author in joint work with Vladimir Marković, and were previously written up in [50]. They are all shown in Chapter 5 with slightly simplified proofs, available due to subsequent advances in the algebra of the Putman–Wieland conjecture [8].

C Energy functionals assoicated to maps into symmetric spaces

The proof of Theorems B.1 and B.2 depends essentially on the fact that the Hodge norm is plurisubharmonic (meaning subharmonic when restricted to an arbitrary holomorphic disk in Teichmüller space). Note that any harmonic 1-form ω on a marked Riemann surface $X \in \mathcal{T}_g$ in the de Rham class of some $\chi \in H^1(\Sigma_g, \mathbb{R})$ induces a χ-equivariant harmonic map

$$f : \tilde{X} \to \mathbb{R},$$

where $\tilde{X}$ is the universal cover of X and we identify χ with its induced homomorphism $\pi_1(\Sigma_g) \to \mathbb{R}$. Our aim in Chapter 6 is to replace $\mathbb{R}$ with some more general symmetric space of non-positive curvature, and study the associated energy functional.

We first introduce the setting we will work in and define the energy. We suppose that $X = \mathbb{H}^2/\Gamma$ is a closed Riemann surface, where Γ is some Fuchsian group. Given a Riemannian manifold M with isometry group $\mathrm{Isom}(M)$ and a homomorphism $\rho : \Gamma \to \mathrm{Isom}(M)$, we define the energy of a ρ-equivariant C^1 map $f : \mathbb{H}^2 \to M$ as

$$\mathrm{E}(f) = \int_\Phi \|df\|^2 \, d\mathrm{Vol},$$

where Φ is some fundamental domain for the action of Γ on $\mathbb{H}^2$.

Definition 1.8. Given a non-positively curved manifold M with isometry group $\mathrm{Isom}(M)$, and a representation $\rho : \pi_1(\Sigma_g) \to \mathrm{Isom}(M)$, we define the energy functional on Teichmüller space $\mathcal{T}_g$,

$$\mathrm{E}_\rho : \mathcal{T}_g \to \mathbb{R},$$

by associating to a marked Riemann surface $f : \Sigma_g \to X$ the energy of a harmonic map (if it exists) $\tilde{X} \to M$ that is $\rho \circ (f_*)^{-1}$ equivariant.

We give two examples when such harmonic maps are guaranteed to exist,

1. when the image of ρ acts freely, properly discontinuously and cocompactly on M, and

2. when M is a symmetric space of non-compact type associated to the complex Lie group $\mathrm{GL}(n, \mathbb{C})$, and ρ is irreducible.

Existence in (1) follows from the classical results of Eells–Sampson [19], and in (2) from the non-abelian Hodge correspondence [14, 16]. Note that in the general setting of Definition 1.8, the harmonic map may not be unique, but its energy is always well-defined (e.g. [31, Corollary, pp. 675]). This energy functional was originally studied when M is the hyperbolic plane by Tromba [70], and more recently when M is a symmetric space in connection with higher Teichmüller theory [40, 41, 51, 58].

It was shown by Toledo [66] that when M has non-positive Hermitian sectional curvature, the energy E_ρ is plurisubharmonic. However, this plurisubharmonicity is in general not strict (one example is provided by Theorem B.1 and the counterexample to the Putman–Wieland conjecture in genus 2 due to Marković [52]). We study this failure of strict plurisubharmonicity under the stronger condition that the curvature of M is very strongly seminegative in the sense of Siu [62].

We remind the reader that for a marked Riemann surface $X \in \mathcal{T}_g$, the tangent space $T_X\mathcal{T}_g$ is parameterized by sections of $\bar{K}_X \otimes K_X^{-1}$, where K_X is the cotangent (canonical) bundle of X, i.e. expressions of the form $\mu\frac{d\bar{z}}{dz}$ where z is a local holomorphic coordinate. We state our main result on Riemannian manifolds below. Precise description of the setting in which this result applies will be given in §6.1.3.

Theorem C.1. *Suppose that M is a Riemannian manifold with very strongly seminegative curvature, and let $\rho : \pi_1(\Sigma_g) \to \mathrm{Isom}(M)$ be a representation. Let X be a point in $\mathcal{T}_g$ with the associated ρ-equivariant harmonic map $f : \tilde{X} \to M$. Assuming the energy is well-defined and smooth in a neighbourhood of $X \in \mathcal{T}_g$, the Laplacian of E_ρ vanishes at X in the complex direction defined by $\mu \in T_X\mathcal{T}_g$ if and only if*

$$\mu\partial f = \bar{\partial}\xi \text{ and } R^M(\xi, \partial f) = 0,$$

for some section ξ of $\mathbb{C} \otimes f^*TM$, where R^M is the complexified Riemann curvature tensor of M.

We now turn our attention to the case $M = \mathrm{GL}(n,\mathbb{C})/\mathrm{U}(n)$. For a closed Riemann surface X, the non-abelian Hodge correspondence provides bijections

$$\left\{\begin{matrix}\text{representations } \rho \\ \pi_1(X) \to \mathrm{GL}(n,\mathbb{C})\end{matrix}\right\} \to \left\{\begin{matrix}\text{equivariant harmonic maps} \\ f : \tilde{X} \to \mathrm{GL}(n,\mathbb{C})/\mathrm{U}(n)\end{matrix}\right\} \to \left\{\begin{matrix}\text{Higgs bundles} \\ (E,\phi) \text{ over } X\end{matrix}\right\}, \tag{†}$$

A Higgs bundle is a pair (E, ϕ) where E is a holomorphic vector bundle over X, and ϕ is a holomorphic 1-form taking values in $\mathrm{End}(E)$ (for more precise definitions and statements, the reader can consult §2.3). Using Theorem C.1, we relate the directions in which E_ρ is not strictly plurisubharmonic to two classical objects in non-abelian Hodge theory: the $\mathbb{C}^*$-action and the Hitchin fibration. We will introduce both of them briefly here. For more details the reader should consult §2.3 and in particular §2.3.3.

We first introduce the $\mathbb{C}^*$-action. Note that there is an obvious $\mathbb{C}^*$-action on the space of Higgs bundles, given by $\lambda \cdot (E, \phi) = (E, \lambda\phi)$. Using the non-abelian Hodge correspondence (†), this action can be transported to the space of representations $\mathrm{Rep}(\pi_1(X), \mathrm{GL}(n, \mathbb{C}))$, where this action depends on the underlying Riemann surface X. The action of $i \in \mathbb{C}^*$ then defines a map

$$\mathcal{R} : \mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) \times \mathcal{T}_g \to \mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})),$$

such that $\mathcal{R}(-, X)$ is an order 4 automorphism of $\mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$. Our main result on $\mathcal{R}$ is the following theorem.

Theorem C.2. *Let $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ be an irreducible representation. Then the space of directions in $T\mathcal{T}_g$ in which E_ρ is not strictly plurisubharmonic is exactly the space of directions annihilated by the derivative of $\mathcal{R}_\rho$.*

We now introduce the Hitchin fibration. Denote the moduli space of Higgs bundles by $\mathcal{M}_{\mathrm{Higgs}}(X)$. Then we get a map, called the Hitchin integrable system

$$H : \mathcal{M}_{\mathrm{Higgs}}(X) \to \bigoplus_{i=1}^{n} H^0(X, K_X^{\otimes i}),$$

by mapping (E, ϕ) to the coefficients of the characteristic polynomial of ϕ. Using Theorem C.1 and the non-abelian Hodge correspondence, we show the following.

Theorem C.3. *Suppose that $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ is an irreducible representation, and let X be a Riemann surface. Then if the Laplacian of E_ρ vanishes in the direction defined by some non-zero $\mu \in T_X\mathcal{T}_g$, then the point that corresponds to ρ under the non-abelian Hodge correspondence is a critical point for the Hitchin fibration H.*

Finally, we show that Theorems C.2 and C.3 are non-vacuous, by showing that there exist pairs (X, ρ) where E_ρ is strictly plurisubharmonic at X, and pairs (X, ρ) where it is not.

Theorem C.4. *Let X be a marked Riemann surface of genus $g \geq 3$, and let $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ be a generically chosen representation. Then the energy E_ρ is strictly plurisubharmonic at X. Conversely, if in addition $g \geq 4$, for any $n \geq 2$ there exists a representation $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ such that E_ρ is not strictly plurisubharmonic at X.*

All of the results in this section are shown in Chapter 6, and were previously written up by the author in the paper [68].

Chapter 2

Preliminaries

2.1 Harmonic maps

In this section, we give some background on pinched Hadamard manifolds, harmonic maps, and a shared outline of the main results of Chapter 3 and Chapter 4.

For a smooth map $h : M \to N$ between Riemannian manifolds, we denote by Dh its derivative, and by D^2h its Hessian. In particular Dh is a h^*TN-valued 1-form on M, and that D^2h is a h^*TN-valued symmetric bilinear form on M.

Definition 2.1. For a smooth map $h : X \to Y$ between Riemannian manifolds, we define its tension field to be $\tau(h) = \mathrm{tr}\ D^2h$. The function h is harmonic if $\tau(h) = 0$. When $N = \mathbb{R}$, we denote $\Delta h = \tau(h)$.

2.1.1 Pinched Hadamard manifolds

The results we prove on the existence of harmonic maps are all in the setting of pinched Hadamard manifolds.

Definition 2.2. A Hadamard manifold is a complete simply connected Riemannian manifold with non-positive sectional curvature. A pinched Hadamard manifold is a Hadamard manifold with all sectional curvatures bounded between two negative constants.

We will mainly use the pinching assumption through the comparison theorem below (proofs can be found in [53, Theorem 2.2] and [12, Theorem II.1A.6, Theorem II.1A.7, Proposition II.1.7]).

Theorem 2.3. *Let M be a pinched Hadamard manifold with sectional curvatures bounded between $-b^2$ and $-a^2$. Let $x, y, z \in M$ be distinct points, and let $\overline{xyz}, \underline{xyz}$ be* *triangles in the hyperbolic plane with*

$$\mathrm{dist}(\overline{x},\overline{y}) = a\mathrm{dist}(x,y), \qquad \mathrm{dist}(\underline{x},\underline{y}) = b\mathrm{dist}(x,y), \tag{2.1.1}$$
$$\mathrm{dist}(\overline{y},\overline{z}) = a\mathrm{dist}(y,z), \qquad \mathrm{dist}(\underline{y},\underline{z}) = b\mathrm{dist}(y,z), \tag{2.1.2}$$
$$\mathrm{dist}(\overline{z},\overline{x}) = a\mathrm{dist}(z,x), \qquad \mathrm{dist}(\underline{z},\underline{x}) = b\mathrm{dist}(z,x). \tag{2.1.3}$$

Then

$$\measuredangle_{\underline{x}}(\underline{y},\underline{z}) \leq \measuredangle_x(y,z) \leq \measuredangle_{\overline{x}}(\overline{y},\overline{z}).$$

We will often use without mention the following simple corollary of Theorem 2.3.

Corollary 2.4. *Let M be a pinched Hadamard manifold with sectional curvatures bounded between $-b^2$ and $-a^2$. Let $x,y,z \in M$ be distinct points, and let $\overline{xyz}, \underline{xyz}$ be triangles in the hyperbolic plane such that (2.1.1) and (2.1.3) hold, and such that*

$$\measuredangle_{\overline{x}}(\overline{y},\overline{z}) = \measuredangle_{\underline{x}}(\underline{y},\underline{z}) = \measuredangle_x(y,z).$$

Then

$$a^{-1}\mathrm{dist}(\overline{y},\overline{z}) \leq \mathrm{dist}(y,z) \leq b^{-1}\mathrm{dist}(\underline{y},\underline{z}).$$

2.1.1.1 Visual metrics on the boundary at infinity

Denote by $\partial_\infty M$ the boundary at infinity of M, that is the set of geodesic rays in M up to the equivalence relation of having finite Hausdorff distance (for a more detailed account of the theory of boundaries of negatively curved spaces, the reader may wish to consult [38]).

We equip $\partial_\infty M$ with the family of visual path metrics $\mathrm{dist}_x^{\mathrm{vis}}(\cdot,\cdot)$ indexed by $x \in M$, such that there exists a constant $C > 0$ with

$$C^{-1}e^{-a\mathrm{dist}(x,[y,z])} \leq \mathrm{dist}_x^{\mathrm{vis}}(y,z) \leq Ce^{-a\mathrm{dist}(x,[y,z])}.$$

Remark 2.5. Note that for general Gromov hyperbolic metric spaces, the visual metrics can only be defined such that $\mathrm{dist}(y,z)e^{\kappa\mathrm{dist}(x,[y,z])}$ is bounded, for some $\kappa > 0$ small enough. However since M is a $\mathrm{CAT}(-a^2)$ space, such a metric exists whenever $0 < \kappa \leq a$, [9, §2.4].

2.1.2 Cheng's lemma and nonlinear Schauder estimates

Here we collect some estimates on harmonic maps between pinched Hadamard manifolds. Our first result is due to Cheng [13, equation (2.9)] (a simplified version is stated in [5, Lemma 3.4]). Denote by $B_R(x)$ the metric ball of radius R centered at x belonging to some metric space.

Lemma 2.6 (Cheng's lemma)**.** *Let M, N be Hadamard manifolds with sectional curvatures between $-b^2$ and 0. Then for any $R > \varepsilon > 0$, there exists a constant C that depends only on $R, \varepsilon, b, \dim M, \dim N$, such that for any harmonic map $h : B_R(x) \to N$ with $x \in M$, we have*

$$\|Dh\|_{L^\infty(B_{R-\varepsilon}(x))} \leq C \operatorname{diam}\Big(h\left(B_R(x)\right)\Big).$$

Our second result follows from Schauder elliptic estimates [56, Theorem 70, pp. 303] for linear elliptic operators of second order. We want to apply these results to harmonic maps, that are solutions to a second order semilinear elliptic equation, so a slight modification is required. This modification is well-known, but we include a brief proof for completeness.

Theorem 2.7 (Nonlinear Schauder elliptic estimates)**.** *Let M, N be pinched Hadamard manifolds, and let $\Omega_0 \subset \Omega \subset M$ be open sets with compact closures, such that $\bar{\Omega}_0 \subset \Omega_1$. Suppose $h : \Omega \to N$ is a harmonic map with bounded image. Then for any $\alpha \in (0, 1)$, we have*

$$\|h\|_{C^{2,\alpha}(\Omega_0)} \leq C = C\Big(\Omega, \Omega_0, N, \operatorname{diam}\big(h\left(\Omega\right)\big)\Big)$$

Proof. Let B be a closed ball containing $h(\Omega)$ of radius comparable to $\operatorname{diam}\big(h(\Omega)\big)$. Let $\Psi : \operatorname{int}(B) \to \mathbb{R}^{\dim N}$ be an embedding with the properties

$$\left\|D\Psi^{\pm 1}\right\|_\infty, \left\|D^2\Psi^{\pm 1}\right\|_\infty < c_0.$$

Such coordinates exist by [5, Lemma 5.2], and here c_0 depends only on curvature bounds and dimension of N, and $\operatorname{diam}\big(h(\Omega)\big)$. We write the harmonic map equation in the coordinates given by Ψ. The Riemannian metric only depends on the first derivative of Ψ^{-1}, and the Christoffel symbols only on the first two derivatives of Ψ^{-1}, so in particular we obtain a pointwise bound on both.

Pick arbitrary local coordinates for Ω. We denote by $\mu = 1, 2, ..., \dim N$ indices that refer to coordinates on N, and by $i = 1, 2, ..., \dim M$ indices that refer to coordinates on M. We also set h_i^μ to be the derivative in the i-th direction of the μ-component of h, and by $\left(h_{ij}^\mu\right)_{i,j=1,2,...,\dim M}$ the second derivative of the μ-component.

The harmonic map equation is

$$\Delta\left(h^{\mu}\right)+g^{ij}h_i^{\nu}h_j^{\eta}\Gamma_{\nu\eta}^{\mu}=0,$$

where g_{ij} is the Riemannian metric on M, g^{ij} is its inverse, and $\Gamma_{\nu\eta}^{\mu}$ are Christoffel symbols on N. Note that by Lemma 2.6, we have a bound on the derivative of h. Since $\Gamma_{\nu\eta}^{\mu}$ is bounded, and since the Laplacian is elliptic, by the standard Schauder estimates [56, Theorem 70, pp. 303] we get a bound on the $C^{2,\alpha}$-norm of h. □

2.1.3 General outline for showing existence of harmonic maps

Here we present the rough outline of proofs of existence of harmonic maps in Chapters 3 and 4. This same outline could apply to [5].

In [5], Chapter 3 and Chapter 4, we start with a Lipschitz map $f: M \to N$ between pinched Hadamard manifolds, with the goal of constructing a harmonic map $h: M \to N$ such that $\sup_M \operatorname{dist}(h, f) < \infty$.

We first exhaust M by closed subsets with smooth boundary $\Omega_1 \subset \Omega_2 \subset ...$ with $\bigcup_{n=1}^{\infty} \Omega_n = M$. We then use the classical result of Hamilton [29].

Theorem 2.8. *Let M be a compact Riemannian manifold with boundary, and let N be a non-positively curved Riemannian manifold. Then for any map $f: \partial M \to N$, if there exists a continuous extension of f to M, there exists a harmonic extension, i.e. a harmonic map $h: M \to N$ such that $h|_{\partial M} = f$.*

Remark 2.9. In [29], Theorem 2.8 is only stated for N compact with convex boundary. The version stated here follows from applying [29, Theorem, pp. 6] to a large ball containing the image of f.

Using Theorem 2.8, we construct harmonic maps $h_n: \Omega_n \to N$ such that $h_n|_{\partial\Omega_n} = f|_{\partial\Omega_n}$. By the arguments of Benoist–Hulin [5], if $\sup_{\Omega_n} \operatorname{dist}(h_n, f)$ is bounded in n, we can finish the proof by using Lemma 2.6, Theorem 2.7 and the Arzela–Ascoli theorem. For completeness, we explain the argument in the proposition below.

Proposition 2.10. *Let M, N be pinched Hadamard manifolds, and let $\Omega_1 \subset \Omega_2 \subset ...$ be an exhaustion of M by relatively compact open sets. Let $h_n: \Omega_n \to N$ be harmonic maps such that*

$$\sup_n \sup_{\Omega_n} \operatorname{dist}(f, h_n) < \infty.$$

for some continuous map $f: M \to N$. Then there exists a harmonic map $h: M \to N$ such that $\sup_M \operatorname{dist}(h, f) < \infty$.

Proof. For any fixed compact set $K \subset M$, we have

$$\operatorname{diam}(h_n(K)) \leq 2 \sup_{\Omega_n} \operatorname{dist}(h_n, f) + \operatorname{diam}(f(K)),$$

for n large enough such that $K \subset \Omega_n$. Hence $\operatorname{diam}(h_n(K))$ is bounded, and hence by Cheng's lemma ([13] or Lemma 2.6)

$$\sup_n \|Dh_n\|_{L^\infty(K)} < \infty.$$

By the Arzela–Ascoli theorem, we may pass to a subsequence and extract a limit $h_n \to h$, that is uniform on compact subsets of M.

From the fact that $\sup_n \|Dh_n\|_{L^\infty(K)} < \infty$ for any compact set $K \subset M$, we see that for $\alpha \in (0,1)$, we have

$$\sup_n \|h_n\|_{C^\alpha(K)} < \infty.$$

From Theorem 2.7 and the fact that h_n is harmonic, we see that $\sup_n \|h_n\|_{C^{2,\alpha}(K)} < \infty$ for any compact $K \subset M$. Applying Arzela–Ascoli again, we may extract a further subsequence such that $D^2 h_n \to H$. It is easy to see that $H = D^2 h$, so in particular h is harmonic.

Finally, we have

$$\sup_M \operatorname{dist}(h, f) \leq \sup_n \sup_{\Omega_n} \operatorname{dist}(h_n, f) < \infty,$$

which concludes the proof of Proposition 2.10. □

Therefore, in both Chapters 3 and 4, we show existence results by showing the bound $\sup_n \sup_{\Omega_n} \operatorname{dist}(h_n, f) < \infty$, and appealing to Proposition 2.10.

2.1.4 Notation

When A, B are quantities of geometric interest in our proofs, we write $A \lesssim B$ when there exists an absolute constant $C > 0$ such that $A \leq CB$. We similarly write $A \gtrsim B$ when $B \lesssim A$, and $A \approx B$ when $A \lesssim B \lesssim A$. We often allow C to depend on the setting of the theorem we are trying to prove (e.g. pinching constants and dimension of the ambient pinched Hadamard manifold, or properties of the convex set we started with). Write $A \lesssim_{x_1,x_2,\ldots,x_n} B$ for $A \leq CB$ when C is allowed to depend on quantities x_i, and define $\gtrsim_{x_1,x_2,\ldots,x_n}$ and $\approx_{x_1,x_2,\ldots,x_n}$ analogously.

We collect below some pieces of notation that appear throughout Chapters 3 and 4 for the reader's convenience,

- given a Riemannian manifold M, the distance function $\text{dist} : M \times M \to \mathbb{R}_+ = \{x \in \mathbb{R} : x \geq 0\}$ always refers to the path metric induced by the Riemannian metric on M,

- we denote by $B_R(x)$ the ball of radius R centered at x, under the metric given by dist,

- we denote by $\sigma_{x,R}$ the harmonic measure on the sphere $\partial B_R(x)$, as seen from x, i.e. the measure defined by the equality

$$h(x) = \int_{\partial B_R(x)} h(y) d\sigma_{x,R}(y)$$

  for all bounded harmonic functions $h : B_R(x) \to \mathbb{R}$,

- when M is a pinched Hadamard manifold, we denote by $\partial_\infty M$ the visual boundary at infinity of M (as in §2.1.1.1),

- for $x, y \in M \cup \partial_\infty M$, we denote by $[x, y]$ the geodesic segment joining x and y,

- for $a \in M, b, c \in M \cup \partial_\infty M \setminus \{a\}$, we denote by $\measuredangle_a(b, c)$ the angle at a between the geodesics $[a, b]$ and $[a, c]$,

- for $x \in M, \xi \in M \cup \partial_\infty M \setminus \{x\}$ and $\theta > 0$, we denote by $\text{Cone}(x\xi, \theta)$ the set of points $y \in M \cup \partial_\infty M$ such that $\measuredangle_x(\xi, y) < \theta$,

- we denote by $\mathbb{H}^n$ the n-dimensional hyperbolic space, and by $\partial_\infty \mathbb{H}^n = \mathbb{S}^{n-1}$ the $(n-1)$-dimensional sphere at infinity,

- we denote by $\|f\|_\infty$ the supremum of some function f (if f is a section of some vector bundle equipped with a natural metric, we still denote by $\|f\|_\infty$ the supremum of the norm of f),

- for a set $S \subseteq M \cup \partial_\infty M$, we denote by $\text{CH}(S)$ its convex hull, that is the intersection of all convex sets containing S.

2.2 Putman–Wieland conjecture and classical Teichmüller theory

In this section, we state some background results on mapping class groups and those aspects of Teichmüller theory and complex geometry that will be used in Chapter 5.

2.2.1 Definitions

We first give definitions of terms briefly introduced in Chapter 1. We denote by $\Sigma_{g,n}$ the closed surface of genus g with n (labelled) points removed.

Definition 2.11. The mapping class group of $\Sigma_{g,n}$, denoted $\mathrm{Mod}_{g,n}$, is the group of orientation-preserving diffeomorphisms $\Sigma_{g,n} \to \Sigma_{g,n}$ that fix the set of punctures pointwise, up to homotopy through diffeomorphisms fixing the set of punctures pointwise.

2.2.1.1 Higher Prym representations and homomorphisms between mapping class groups

We now elaborate on Definition 1.3. Recall the Birman exact sequence

$$1 \to \pi_1(\Sigma_{g,n}) \xrightarrow{\iota_{g,n}} \mathrm{Mod}_{g,n+1} \xrightarrow{\mathcal{F}_{g,n}} \mathrm{Mod}_{g,n} \to 1. \tag{2.2.1}$$

The second non-trivial map in this sequence is denoted $\mathcal{F}_{g,n} : \mathrm{Mod}_{g,n+1} \to \mathrm{Mod}_{g,n}$ and called the forgetful homomorphism. It is obtained by filling in one of the punctures. We denote the resulting point on $\Sigma_{g,n}$ by x_0. The first map realizes $\pi_1(\Sigma_{g,n}, x_0)$ as the point-pushing subgroup of $\mathrm{Mod}_{g,n+1}$. This is a normal subgroup, and the conjugation action of $\mathrm{Mod}_{g,n+1}$ on $\pi_1(\Sigma_{g,n}, x_0)$ (induced by the embedding in (2.2.1)) is precisely the induced action of mapping classes on the fundamental group $\pi_1(\Sigma_{g,n}, x_0)$. In symbols, for any orientation-preserving diffeomorphism $\psi : \Sigma_{g,n+1} \to \Sigma_{g,n+1}$ fixing the punctures, by filling in x_0, we get a map

$$\psi_* : \pi_1(\Sigma_{g,n}, x_0) \to \pi_1(\Sigma_{g,n}, x_0).$$

We then have $\iota_{g,n}(\psi_*(\gamma)) = \psi\iota_{g,n}(\gamma)\psi^{-1}$. In most of the thesis, we will identify $\pi_1(\Sigma_{g,n})$ with the point-pushing subgroup of $\mathrm{Mod}_{g,n+1}$ without mention. The reader can consult [22, §4.2] for more details.

Given a finite cover $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$, we let K be the image of π_* in $\pi_1(\Sigma_{g,n}, x_0)$ (after fixing the basepoint $y_0 \in \pi^{-1}(x_0)$ for $\Sigma_{h,m}$ arbitrarily). Then $[\pi_1(\Sigma_{g,n}, x_0) : K] < \infty$, and hence there is some finite index subgroup $\Gamma_\pi \leq \mathrm{Mod}_{g,n+1}$ of mapping classes that preserve K under the conjugacy action on $\pi_1(\Sigma_{g,n}, x_0)$ induced by (2.2.1). This is exactly the set of mapping classes that lift via π using x_0 as a basepoint. We denote this lifting map by

$$\Lambda_\pi : \Gamma_\pi \to \mathrm{Mod}_h.$$

Definition 2.12. The mapping class group Mod_h acts symplectically on $H^1(\Sigma_h, \mathbb{Z})$. The induced action by Γ_π via the homomorphism Λ_π is called a higher Prym representation.

We will usually use Λ_π without mention to act by Γ_π on $H^1(\Sigma_h, \mathbb{R})$.

2.2.1.2 The Hodge norm

We now elaborate on Definition 1.6. Given a Riemann surface S, we recall that, by Hodge theory, any de Rham class $\chi \in H^1(S, \mathbb{R})$ contains a unique harmonic 1-form ω. For completeness, we define below what we mean in Definition 1.6 by the "energy" of a 1-form.

Definition 2.13. The energy of a (continuous) 1-form ω is defined as

$$\int_S \omega \wedge \star\omega,$$

where $\star$ is the Hodge star.

Remark 2.14. In particular $\mathrm{E}_\chi(S)$ is the least possible energy of a smooth 1-form in the class χ.

Finally we express the Hodge norm of χ explicitly in terms of holomorphic forms representing χ.

Claim 2.15. The Hodge norm of $\chi \in H^1(\Sigma_g, \mathbb{R})$ is given by

$$\mathrm{E}_\chi(S) = \frac{i}{2} \int_S \phi \wedge \bar{\phi},$$

for $S \in \mathcal{T}_g$, where ϕ is the unique holomorphic 1-form on S such that the de Rham class of $\mathrm{Re}(\phi)$ is χ.

Proof. It is well known that $\phi = \omega + i\star\omega$, where ω is the harmonic 1-form in the class χ. This follows easily from the conformal invariance of the Hodge star on 1-forms and a direct computation. Thus

$$\frac{i}{2} \int_S (\omega + i \star \omega) \wedge (\omega - i \star \omega) = \int_S \omega \wedge \star\omega = E_\chi(S).$$

□

2.2.2 Teichmüller and moduli spaces

In this section we recall the definitions of Teichmüller space and the moduli space, and explain those aspects of the local deformation theory of Riemann surfaces that will be used in Chapter 5. All of the material in this section is very standard and can be found in many books, e.g. [23, 36] for Teichmüller spaces and [30] for moduli spaces.

Definition 2.16. The Teichmüller space $\mathcal{T}_{g,n}$ of the surface $\Sigma_{g,n}$ is the set of marked Riemann surfaces, i.e. homeomorphisms $f : \Sigma_{g,n} \to X$, where X is a closed Riemann surface with n points removed, up to the following equivalence relation. The marked Riemann surfaces $f : \Sigma_{g,n} \to X$ and $g : \Sigma_{g,n} \to Y$ are considered equivalent if $g \circ f^{-1}$ is isotopic to a biholomorphism $X \to Y$.

There is a natural action of the mapping class group $\mathrm{Mod}_{g,n}$ on $\mathcal{T}_{g,n}$, given by

$$\begin{aligned} \mathrm{Mod}_{g,n} \times \mathcal{T}_{g,n} &\longrightarrow \mathcal{T}_{g,n} \\ A \cdot f &\longrightarrow f \circ A^{-1}. \end{aligned}$$

Definition 2.17. The moduli space of $\Sigma_{g,n}$ is the quotient $\mathcal{M}_{g,n} = \mathcal{T}_{g,n}/\mathrm{Mod}_{g,n}$.

The moduli space $\mathcal{M}_{g,n}$ is a complex orbifold and a quasiprojective variety (meaning a Zariski open subset of a projective variety). More generally, for a subgroup $\Gamma \leq \mathrm{Mod}_{g,n}$, we denote by $\mathcal{M}^{\Gamma}_{g,n} = \mathcal{T}_{g,n}/\Gamma$ the (étale) covering space of $\mathcal{M}_{g,n}$ that corresponds to the subgroup Γ of its fundamental group.

Local theory

Here we describe without proof the tangent and cotangent spaces to Teichmüller space. Fix a Riemann surface $X \in \mathcal{T}_g$, and denote by K_X its cotangent bundle. This is a holomorphic bundle of complex dimension 1.

A tangent vector at X in Teichmüller space is a *Beltrami differential*, i.e. a section μ of the (smooth, but non-holomorphic) complex line bundle $\bar{K}_X \otimes K_X^{-1}$. In local coordinate z on X, this section takes the form

$$\mu(z)\frac{d\bar{z}}{dz}.$$

Given any such μ, with $\sup_X|\mu| < 1$, there exists a Riemann surface X^μ with a map $f : X \to X^\mu$, with the property that $f_{\bar{z}} = \mu f_z$ in the local coordinate z. Then for any bounded μ, the path $t \to X^{t\mu}$ is tangent to a vector $T_X\mathcal{T}_g$ that we denote $[\mu]$.

All tangent vectors to Teichmüller space arise in this way, but there exist μ for which the curve $t \to X^{t\mu}$ has vanishing derivative at $t = 0$. To describe which μ have this property, we need the cotangent space of $\mathcal{T}_g$.

A cotangent vector at X is a holomorphic section of $K_X^{\otimes 2}$, i.e. a holomorphic quadratic differential. In the local coordinate z, this is an object of the form

$$\Phi(z)dz^2,$$

where Φ is a holomorphic function. The pairing between Φ and μ is

$$\langle \Phi, \mu \rangle := \int_X \Phi\mu.$$

Moreover we have $[\mu] = [\nu]$ for μ, ν Beltrami forms on X, if and only if $\langle \Phi, \mu \rangle = \langle \Phi, \nu \rangle$ for all holomorphic quadratic differentials Φ.

2.2.3 Complex geometry

The central Theorem 5.1 of §5.1 is that the Hodge norm E_χ is plurisubharmonic.

Definition 2.18. Let M be a complex manifold and $f : M \to \mathbb{R}$ be a C^2 function. For $x \in M$, we say that f is plurisubharmonic, resp. strictly plurisubharmonic at x, if the matrix

$$L = \left(\frac{\partial^2 f}{\partial z_i \partial \bar{z}_j}(x) \right)_{1 \leq i,j \leq n}$$

is positive semidefinite, resp. positive definite, for local holomorphic coordinates $(z_1, ..., z_n)$ near x.

To derive Theorem B.1 from the plurisubharmonicity of the Hodge norm, we will use the following standard fact from complex geometry.

Lemma 2.19. *Let M be a compact connected complex manifold, and let $N \subseteq M$ be a submanifold of (complex) codimension* 1*. Then any bounded plurisubharmonic function on $M \setminus N$ is constant.*

Proof. Let $f : M \setminus N \to \mathbb{R}$ be bounded and plurisubharmonic. By a standard result in complex analysis, that can be found e.g. in the book of Demailly [15, Theorem (5.24)], the map f has a unique plurisubharmonic extension to $\bar{f} : M \to \mathbb{R}$. Since M is compact, we pick a point $x \in M$ at which $\bar{f}$ achieves its maximum. Let $(z_1, z_2, ..., z_n)$ be local holomorphic coordinates near x, that identify x with $(0, 0, ..., 0)$.

Then $f_i = \bar{f}(\underbrace{0, ..., 0}_{i-1}, z, 0, ..., 0)$ is a subharmonic function on a small disk around $0 \in \mathbb{C}$, that achieves its maximum at 0. Hence by the strong maximum principle, f_i is constant for all $0 \leq i < n$. It follows that $\bar{f}$ is constant in a neighborhood of x. By connectedness of M, the function $\bar{f}$ is constant everywhere. □

2.3 Non-abelian Hodge theorem

Non-abelian Hodge theorem defines a correspondence between the holomorphic data of Higgs bundles and the set of harmonic maps from a Riemann surface into a symmetric space. In this thesis, we focus on harmonic maps into the space $\mathrm{GL}(n, \mathbb{C})/\mathrm{U}(n)$, so while Higgs bundles can be defined for more general Lie groups, we will specialize to the case $\mathrm{GL}(n, \mathbb{C})$.

Definition 2.20. A rank n Higgs bundle on a Riemann surface S is a pair (E, ϕ), where E is a holomorphic rank n vector bundle on S and ϕ is a holomorphic 1-form that takes values in $\mathrm{End}(E)$.

The non-abelian Hodge correspondence provides a bijection between the set of completely reducible (also known as semisimple) representations $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ and the set of polystable Higgs bundles, where irreducible representations correspond to stable Higgs bundles. We first define the notions of stability and semistability.

Recall that the degree of a holomorphic rank n vector bundle E, denoted $\deg(E)$, is the number of self-intersections of the zero section in $\bigwedge^n E$. Define the slope of a holomorphic vector bundle E to be $\mu(E) = \frac{\deg(E)}{\mathrm{rank}(E)}$.

Definition 2.21. A Higgs bundle (E, ϕ) is stable if for any proper ϕ-invariant holomorphic subbundle $F \leq E$, we have $\mu(F) < \mu(E)$. It is polystable if it is a direct sum of stable bundles of equal slope.

The non-abelian Hodge correspondence goes through equivariant harmonic maps into the symmetric space $\mathrm{GL}(n, \mathbb{C})/\mathrm{U}(n)$, and the Higgs field ϕ is the derivative of the appropriate harmonic map.

We first state some preliminary results on symmetric spaces in §2.3.1. We will then describe without proof how a harmonic map yields a Higgs bundle in §2.3.2. Finally, we introduce the $\mathbb{C}^*$-action and the Hitchin integrable system appearing in Theorems C.2 and C.3 in §2.3.3.

2.3.1 Symmetric spaces and the Maurer–Cartan form

In this subsection, we let G be a semisimple Lie group, and K be its maximal compact subgroup. Then $X = G/K$ is a symmetric space. We will describe the tangent bundle to X in terms of its Lie algebra, that we denote $\mathfrak{g}$. Most proofs in this subsection will be omitted, for more background and proofs we refer the reader to the books [39, 32].

We denote the Lie subalgebra of $\mathfrak{g}$ corresponding to K by $\mathfrak{k}$, and the orthogonal complement to $\mathfrak{k}$ under the Killing form by $\mathfrak{p}$. Note that we have

$$\begin{aligned}[\mathfrak{k},\mathfrak{k}] + [\mathfrak{p},\mathfrak{p}] &\leq \mathfrak{k},\\ [\mathfrak{k},\mathfrak{p}] &\leq \mathfrak{p}.\end{aligned} \tag{2.3.1}$$

We let $X = G/K$ be equipped with some left-invariant metric, such that $\mathfrak{p} = \mathfrak{k}^\perp$ with respect to this metric. We will identify TX with a $G \times_K \mathfrak{p}$, and describe the Levi–Civita connection in terms of this identification. Assume that X has non-positive curvature.

Let $\omega \in \Omega^1(G, \mathfrak{g})$ be the Maurer–Cartan form of G, that is the unique 1-form with the following properties

1. $L_g^*\omega = \omega$, and
2. $R_g^*\omega = \mathrm{Ad}_{g^{-1}}\omega$.

In particular, ω identifies all tangent spaces of G to $\mathfrak{g}$, in a left-invariant manner. Denote by $\omega^\mathfrak{k}, \omega^\mathfrak{p}$ the results of composing ω with the projections $\mathfrak{g} \to \mathfrak{k}, \mathfrak{g} \to \mathfrak{p}$, respectively, given by the orthogonal splitting $\mathfrak{g} = \mathfrak{k} + \mathfrak{p}$.

The form $\omega^\mathfrak{p}$ descends to a form on $X = G/K$, taking values in the bundle $G \times_K \mathfrak{p}$ over X, where K acts on $\mathfrak{p}$ by the restriction of the adjoint action of G. It defines an isomorphism $TX \to G \times_K \mathfrak{p}$. Similarly, the form $\omega^\mathfrak{k}$ defines a connection on the principal K-bundle $G \to X$. There is a connection ∇ on $G \times_K \mathfrak{p}$ induced from $\omega^\mathfrak{k}$. The following result is well-known in the theory of symmetric spaces. Since an explicit reference is unknown to the author, we include a proof for completeness. Note that in a slightly less general setting, the same result was shown by Slegers [63, Lemma 2.2].

Proposition 2.22. *If V is a vector field on X, then* $\omega^\mathfrak{p}(\nabla V) = \nabla\left(\omega^\mathfrak{p}(V)\right)$.

Proof. Note that the first ∇ refers to the Levi–Civita connection on X, and that the second ∇ refers to the $\omega^\mathfrak{k}$-induced connection on $G \times_K \mathfrak{p}$. We equip G with a left-invariant metric such that $G \to G/K$ is a Riemannian submersion. Denote by ad_ξ^* the adjoint to the map $\mathrm{ad}_\xi = [\xi, -] : \mathfrak{g} \to \mathfrak{g}$ with respect to this metric.

Claim 2.23. Let V be a vector field on G such that $\omega(V)$ takes values in $\mathfrak{p}$. Then for a left-invariant vector field ξ corresponding to an element of $\mathfrak{p}$, the function $\omega(\nabla_\xi V) - d_\xi(\omega(V))$ takes values in $\mathfrak{k}$.

Proof. Note that by the Koszul formula, if ξ, η, θ are left-invariant vector fields on G, we have

$$\begin{aligned} 2\langle \nabla_\xi \eta, \theta \rangle &= \langle [\xi, \eta], \theta \rangle - \langle [\xi, \theta], \eta \rangle - \langle [\eta, \theta], \xi \rangle \\ &= \langle [\xi, \eta] - \mathrm{ad}^*_\xi \eta - \mathrm{ad}^*_\eta \xi, \theta \rangle \end{aligned}$$

We write $V = \sum_i a_i \xi_i$, where ξ_i are left-invariant vector fields corresponding to a basis of $\mathfrak{p}$, and $a_i : G \to \mathbb{R}$ are functions. Then

$$\begin{aligned} \omega(\nabla_\xi V) &= \omega \left(\sum_i d_\xi(a_i)\xi_i + a_i \nabla_\xi \xi_i \right) \\ &= d_\xi \big(\omega(V)\big) + \sum_i \frac{a_i}{2} \omega \Big([\xi, \xi_i] - \mathrm{ad}^*_\xi \xi_i - \mathrm{ad}^*_{\xi_i} \xi \Big) . \end{aligned}$$

We observe that $\xi, \xi_i \in \mathfrak{p}$, and hence by (2.3.1), we have $[\xi, \xi_i] \in \mathfrak{k}$ and since $\mathfrak{k}, \mathfrak{p}$ are orthogonal, we have

$$\mathrm{ad}^*_\xi \xi_i, \mathrm{ad}^*_{\xi_i} \xi \in \mathfrak{k}.$$

□

Let W be a vector field on X such that $\omega^{\mathfrak{p}}(W) = V$. We let $\bar{W}$ be the vector field on G such that $\omega(\bar{W}) = V$. Then $\bar{W}$ projects to W under the natural quotient map $G \to X$. Since $\omega(\bar{W}(g)) \in \mathfrak{p}$, we see that $\bar{W}(g) \in (L_g)_\star \mathfrak{p}$. In particular, $\bar{W}$ is perpendicular to the fibres of $G \to X$.

Given any section V' of $G \times_K \mathfrak{p}$, introduce vector fields $W', \bar{W}'$ analogously. It follows by standard results on Riemannian submersions [56, Proposition 13 (4), pp. 94], that $\nabla_{W'} W$ is the projection of $\nabla_{\bar{W}'} \bar{W}$ to X. Thus, using (2.3.1),

$$\omega^{\mathfrak{p}}(\nabla_{W'} W) = \omega^{\mathfrak{p}}(\nabla_{\bar{W}'} \bar{W}) = d_{\bar{W}'} V.$$

The result now follows since $\omega^{\mathfrak{k}}(\bar{W}') = 0$. □

2.3.2 Harmonic maps and the Hitchin equation

We are now in a position to describe the correspondence (†) for a marked Riemann surface $S \in \mathcal{T}_g$.

$$\left\{\begin{matrix}\text{representations } \rho \\ \pi_1(S) \to \mathrm{GL}(n,\mathbb{C})\end{matrix}\right\} \overset{(1)}{\to} \left\{\begin{matrix}\text{equivariant harmonic maps} \\ f : \tilde{S} \to \mathrm{GL}(n,\mathbb{C})/\mathrm{U}(n)\end{matrix}\right\} \overset{(2)}{\to} \left\{\begin{matrix}\text{Higgs bundles} \\ (E,\phi) \text{ over } S\end{matrix}\right\}. \tag{2.3.2}$$

In the non-abelian Hodge correspondence, we consider only completely reducible (also known as semisimple) representations ρ up to conjugacy, i.e. the first set in the equation above is the character variety

$$\mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n,\mathbb{C})) = \mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n,\mathbb{C})) /\!\!/ \mathrm{GL}(n,\mathbb{C}).$$

Here "$/\!\!/$" refers to the GIT (geometric invariant theory) quotient, and $\mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n,\mathbb{C}))$ is the space of completely reducible homomorphisms $\pi_1(\Sigma_g) \to \mathrm{GL}(n,\mathbb{C})$. The harmonic maps in (2.3.2) are thus only considered up to a global left translate by an element in $\mathrm{GL}(n,\mathbb{C})$. Finally, the Higgs bundles in (2.3.2) are all polystable of degree 0 and rank n, and moreover, irreducible representations ρ correspond to stable Higgs bundles.

Remark 2.24. When ρ is irreducible, both $\mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n,\mathbb{C}))$ and $\mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n,\mathbb{C}))$ are smooth near ρ (by [27, §1.2, Proposition] and [42, Theorem 3], respectively).

We now describe the maps (1), (2) and their inverses in (2.3.2).

The map (1) is provided by the results of Donaldson [16] and Corlette [14], showing that for a conjugacy class of completely reducible representations $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n,\mathbb{C})$, there exists a ρ-equivariant harmonic map $f : \tilde{S} \to \mathrm{GL}(n,\mathbb{C})/\mathrm{U}(n)$.

We now describe the map (2). The pullback by f of the principal $\mathrm{U}(n)$-bundle $\mathrm{GL}(n,\mathbb{C}) \to X_n$ descends to a principal $\mathrm{U}(n)$-bundle P_f over S. We consider the associated bundle $E = P_f \times_{\mathrm{U}(n)} \mathbb{C}^n$ which is a complex vector bundle equipped with a Hermitian metric. The projected Maurer–Cartan form $\omega^{\mathfrak{u}(n)}$ defines a connection d_A on E that preserves the metric. Since f is harmonic, the triple $(E, d_A^{0,1}, \omega^{\mathfrak{p}}(\partial f))$ forms a Higgs bundle.

We now turn to the inverse maps. Note that any Higgs bundle $(E, d_A^{0,1}, \omega^{\mathfrak{p}}(\partial f))$ obtained as above comes equipped with a Hermitian metric h, with the following property.

Definition 2.25. Given a Higgs bundle (E, ϕ) over S, a Hermitian metric h on E is called harmonic if $\nabla^h + \phi + \phi^{*h}$ is flat, where ∇^h is the Chern connection on E, or equivalently if h is a solution to the Hitchin equation

$$F_{\nabla^h} + [\phi, \phi^{*h}] = 0, \tag{2.3.3}$$

where F_{∇^h} is the curvature of the Chern connection ∇^h.

Conversely, given a rank n degree 0 Higgs bundle (E, ϕ) over S, any Hermitian metric h on E defines a reduction of the structure group of E to $\mathrm{U}(n)$. The developing map $f : \tilde{S} \to \mathrm{GL}(n, \mathbb{C})/\mathrm{U}(n)$ then defines a map into X_n, which is harmonic if and only if the metric is harmonic. The other direction of the non-abelian Hodge theorem, due to Hitchin [33], states that any degree 0 polystable Higgs bundle (E, ϕ) admits a solution to the Hitchin equation (2.3.3), and thus a harmonic map that produces (E, ϕ) under the map (2). Moreover, the representation ρ can be recovered as the holonomy of $\nabla^h + \phi + \phi^{*h}$.

We encapsulate the above discussion as Theorem 2.26 below. Denote by $\mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S)$ the set of (isomorphism classes of) polystable degree 0 Higgs bundles over S.

Theorem 2.26. *Given a Riemann surface S, the following map is a bijection*

$$\begin{aligned} \mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) &\longrightarrow \mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S) \\ \rho &\longrightarrow \mathrm{Higgs}(\rho, S) := (E, d_A^{0,1}, \omega^{\mathfrak{p}}(\partial f)), \end{aligned}$$

where $f : \tilde{S} \to X$ is the a ρ-equivariant harmonic map.

2.3.3 $\mathbb{C}^*$-action and the Hitchin fibration

2.3.3.1 $\mathbb{C}^*$-action

Here we define the $\mathbb{C}^*$-action that appears in Theorem C.2. Let $S \in \mathcal{T}_g$ be a marked Riemann surface.

Definition 2.27. Given a Riemann surface S, the $\mathbb{C}^*$-action on $\mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S)$ is defined by

$$\lambda \cdot (E, \phi) = (E, \lambda\phi).$$

Note in particular that the $\mathbb{C}^*$-action preserves both stability and polystability, so by the non-abelian Hodge correspondence we get a $\mathbb{C}^*$-action on $\mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$ that preserves the open subset consisting of irreducible representations.

2.3.3.2 The Hitchin fibration and spectral curves

Hitchin fibration was introduced by Hitchin [34] to parameterize Higgs bundles for which the resulting representation has the image in the real form $\mathrm{GL}(n, \mathbb{R})$. Given a Riemann surface S, let K_S be its cotangent bundle.

Definition 2.28. The Hitchin fibration is given by

$$\begin{aligned} H : \mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S) &\longrightarrow \mathcal{B}(S) = \bigoplus_{i=1}^{n} H^0(S, K_S^i) \\ (E, \phi) &\longrightarrow (p_1(\phi), p_2(\phi), ..., p_n(\phi)), \end{aligned}$$

where p_i is the x^{n-i} coefficient in the characteristic polynomial $\chi_\phi(x) = \det(x\mathrm{id}_E - \phi)$.

There is a section $\Gamma : \mathcal{B}(S) \to \mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S)$, called the Hitchin section, whose image is a connected component of the set of Higgs bundles that correspond to representations $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{R})$ [34].

Definition 2.29. When (E, ϕ) is a Higgs bundle and $a = H(E, \phi)$, we define the spectral curve of a to be the subvariety S_a of the total space of the canonical bundle K_S cut out by the equation

$$\chi_a(x) = x^n + x^{n-1}p_1(\phi) + ... + p_n(\phi) = 0,$$

where $\chi_a(x) = \det(x\mathrm{id}_E - \phi)$ is the characteristic polynomial of ϕ (that depends only on a).

Definition 2.30. We say that $a \in \bigoplus_{i=1}^n H^0(S, K_S^i)$ defines a smooth fibre of the Hitchin fibration when S_a is reduced, irreducible, and smooth.

When $g \geq 3, n \geq 2$, for a generic $a \in \bigoplus_{i=1}^n H^0(S, K_S^i)$, it is well-known that the curve S_a is smooth. This follows e.g. from [47, Proposition 2.1] and the fact that $\deg(K_S^n) \geq 4g - 4 > 2g + 1$, and hence K_S^n is very ample.

Chapter 3

Harmonic projections to thin convex sets

3.1 Introduction

In this chapter, we show that given any pinched Hadamard manifold M, and a closed subset $S \subseteq \partial_\infty M$ of the boundary at infinity with low dimension, there exists a harmonic map $h : M \to M$ at a finite distance from the nearest point retraction $r : M \to \mathrm{CH}(S)$. Our main result is Theorem 3.1 below. It relies on a notion of dimension that we call *invariant upper Minkowski dimension*, related to the Minkowski (i.e. box-counting) dimension in $\partial_\infty M$ with respect to visual metrics based at various points in M. We will define this notion precisely in §3.1.1.

Theorem 3.1. *Let M be a pinched Hadamard manifold of dimension n. Let $S \subseteq \partial_\infty M$ be a closed set in the boundary at infinity of M, with the invariant upper Minkowski dimension less than $n-1$. Then there exists a harmonic map $h : M \to M$ at a bounded distance from the nearest-point retraction to the convex hull of S.*

Combined with the well-known fact [26] that quasicircles have Minkowski dimension less than $n-1$, Theorem 3.1 shows Theorem A.1(1). We will elaborate on this argument in §3.5. We first define the invariant upper Minkowski dimension appearing in Theorem 3.1 in §3.1.1. After that we provide further motivation for Theorem 3.1 in §3.1.2, and then outline how the rest of this chapter is structured in §3.1.3.

3.1.1 Invariant upper Minkowski dimension

The appropriate notion of dimension we will use for subsets of $\partial_\infty M$ is defined below. For a subset S of some metric space (X, d), we denote by $N_d(S, \varepsilon)$ the smallest number of ε-balls needed to cover S. Recall from §2.1.1.1 the definition of the visual metrics

$\mathrm{dist}_x^{\mathrm{vis}}$ on the boundary at infinity $\partial_\infty M$ of a pinched Hadamard manifold M, for $x \in M$.

Definition 3.2. If M is a pinched Hadamard manifold, for $S \subseteq \partial_\infty M$, the invariant upper Minkowski dimension of S, denoted $\overline{\dim} S$, is the infimum of all $\beta \geq 0$ with the property that there exists a constant C such that

$$N_{\mathrm{dist}_x^{\mathrm{vis}}}(S, \varepsilon) \leq C\varepsilon^{-\beta}, \tag{3.1.1}$$

for all $x \in M$ and $\varepsilon > 0$.

Remark 3.3. If we fix some arbitrary basepoint $o \in M$, and write $\mathrm{dist}^{\mathrm{vis}} = \mathrm{dist}_o^{\mathrm{vis}}$, there is an alternative characterization of invariant upper Minkowski dimension that refers only to the metric $\mathrm{dist}^{\mathrm{vis}}$, assuming that M has a cobounded isometry group. Recall that a subset $A \subset X$ of a metric space is cobounded if there exists a constant $C > 0$ such that the C-neighbourhood of A covers X.

Lemma 3.4. *Let M be a pinched Hadamard manifold with an isometry group that acts with cobounded orbits. For a set $S \subseteq \partial_\infty M$, the invariant upper Minkowski dimension is the infimum of all $\beta \geq 0$ such that there exists a constant $C > 0$ with the property that*

$$N_{\mathrm{dist}^{\mathrm{vis}}}(\gamma S, \varepsilon) \leq C\varepsilon^{-\beta} \textit{ for all } \gamma \in \mathrm{Isom}(M) \textit{ and } \varepsilon > 0.$$

3.1.2 Connection to the question of Yau

Yau conjectured in [72, Question 38] that any simply connected, complete Kähler manifold with sectional curvature at most -1 admits a non-constant holomorphic map to the disk. A natural analogue of this question for general (not necessarily Kähler) Riemannian manifolds is as follows.

Question 3.5. Does every pinched Hadamard manifold admit a non-constant harmonic map to the hyperbolic plane $\mathbb{H}^2$?

In this chapter, we give a partial affirmative answer to Question 3.5 utilizing harmonic maps that are close to projections of convex sets. Our strategy broadly consists of two steps.

1. We start with some quasi-isometric embedding $\iota : \mathbb{H}^2 \to M$. This defines a quasicircle S in the boundary at infinity of M. A modification of the nearest-point retraction onto the convex hull of S gives a map $r : M \to \mathbb{H}^2$.

2. We deform the map r to a harmonic map.

We show in particular that for hyperbolic spaces, answers to Question 3.5 are abundant.

Theorem 3.6. *Let $\iota : \mathbb{H}^2 \to \mathbb{H}^n$ be a quasi-isometric embedding. Then there exists a harmonic map $h : \mathbb{H}^n \to \mathbb{H}^2$ such that*

$$\sup_{x \in \mathbb{H}^2} \operatorname{dist}(x, h(\iota(x))) < \infty.$$

Note that ι as in Theorem 3.6 defines a quasi-circle S in the boundary at infinity $\partial_\infty M$. The proof of Theorem 3.6 will simultaneously prove Theorem A.1(1).

3.1.3 Outline

Throughout the chapter we fix a pinched Hadamard manifold M, with sectional curvatures between $-b^2$ and $-a^2$, for some fixed $0 < a \leq b$.

We let S be a set in $\partial_\infty M$, and let K be the convex hull of S. We let $r : M \to K$ be the nearest-point retraction.

We derive Theorems 3.6 and 3.1 from the following slightly more general version of Theorem 3.1. Recall from Definition 2.1 that $\tau(f)$ denotes the tension field of a smooth map $f : M \to N$ between Riemannian manifolds.

Theorem 3.7. *Let M, N be pinched Hadamard manifolds such that M has sectional curvatures at most $-a^2$. Let $f : M \to N$ be a smooth map with*

$$\left\|\tau(f)\right\| \leq Ce^{-a\operatorname{dist}(\cdot,K)},$$

for some constant $C > 0$. If $\overline{\dim} S < n - 1$, there exists a harmonic map $h : M \to N$ at a bounded distance from f.

The proof consists of four steps.

1. We construct a smooth map $\tilde{r} : M \to M$ that is at a bounded distance from r, so that its derivative and Hessian have the property $\|D_x\tilde{r}\|, \left\|D_x^2\tilde{r}\right\| \leq Ce^{-\operatorname{dist}(x,K)}$, for some constant C. We use this both in the proof of Theorem 3.7, and to derive Theorem 3.1 from Theorem 3.7.

   This $\tilde{r}$ is the result of a construction of Benoist and Hulin in [5, §2.2]. Their exact statements do not apply here since $\tilde{r}$ is not a quasi-isometry, and moreover we need slightly stronger conclusions than they do. We summarize their construction in §3.2, and explain how it applies to r.

2. Denote by $N_C(K)$ the C-neighbourhood of K. We show in Lemma 3.15 the bound

$$\mathrm{vol}(N_C(K) \cap B_\rho(x)) \leq C' \exp\left(\left(\overline{\mathrm{dim}} S + \varepsilon\right) a\rho\right),$$

for all $x \in M, \varepsilon > 0$, and some constant C' that does not depend on x or ρ.

3. We construct a bounded subharmonic function $\Phi : M \to \mathbb{R}$ with

$$\Delta\Phi \geq e^{-a\mathrm{dist}(\cdot,K)},$$

under the assumption that $\overline{\mathrm{dim}} S < n-1$. The construction essentially consists of two steps.

We first construct bounded subharmonic functions $\Phi_d(x)$, that take the form $g(\mathrm{dist}(x, \tilde{r}(x)))$ for some function $g : \mathbb{R} \to \mathbb{R}$, such that $\Delta\Phi_d \geq 1$ on $N_{d+1}(K) \setminus N_d(K)$, for all $d \geq D$, for some large enough D. The Φ_d are constructed so that $\sup_d \|\Phi_d\|_\infty < \infty$. This is done in Lemma 3.23.

We then use the estimate from Lemma 3.15 to construct a bounded subharmonic function of the form $\Phi_0(x) = \sum_{y\in V} \hat{g}(\mathrm{dist}(x, y))$, for some appropriate function $\hat{g} : \mathbb{R} \to \mathbb{R}$ and some discrete set V, with $\Delta\Phi_0 \geq 1$ on $N_C(K)$, for some $C > 0$ large. This is done in Lemma 3.24.

Finally, since $\sup_d \|\Phi_d\|_\infty < \infty$, we can take Φ to be an appropriately rescaled $\Phi_0 + \sum_{d\geq D} e^{-ad}\Phi_d$.

4. Suppose we are given a function $f : M \to N$ between pinched Hadamard manifolds with $\|\tau(f)\| \lesssim e^{-a\mathrm{dist}(\cdot,K)}$. On any ball $B_R(x)$, we can construct a harmonic map $h_R : B_R(x) \to N$ with $h_R = f$ on $\partial B_R(x)$. An estimate by Schoen and Yau from [60] shows that

$$\Delta\mathrm{dist}(f, g) \geq -\|\tau(f)\| - \|\tau(g)\|,$$

for any smooth functions $f, g : M \to N$. In fact we use a general formula from which Schoen and Yau derive this estimate in §3.4.1 to bound the Laplacian of $\delta(x) = \mathrm{dist}(x, \tilde{r}(x))$.

Using this formula we see that $\Delta(\mathrm{dist}(h_R, f) + C\Phi) > 0$ for some large constant C. Therefore $\mathrm{dist}(h_R, f) + C\Phi \leq C \sup|\Phi|$ by the maximum principle. It follows that $\mathrm{dist}(h_R, f)$ is bounded uniformly in R. A classical argument combining Cheng's lemma, Schauder elliptic estimates and Arzela-Ascoli theorem, outlined in Proposition 2.10, then shows that we can take the limit of such harmonic maps as $R \to \infty$, to get a harmonic map $h_\infty : M \to N$ at a bounded distance from f. The details of the argument described in this paragraph are in §3.4.4.

This idea to use Φ to bound $\sup \operatorname{dist}(h_R, f)$ also appears in the work of Donnelly [17]. Specifically, it is shown in [17, Lemma 3.1] that given a function $f : M \to N$ and a bounded non-negative function $\Phi : M \to \mathbb{R}$ with $\Delta\Phi > \left\|\tau(f)\right\|$, there exists a harmonic map $h : M \to N$ such that $\operatorname{dist}(h(x), f(x)) \leq \sup_x \Phi(x)$, in a way completely analogous to our proof here. In [17], Φ is constructed using an assumption on the integral of the Green's function. By contrast, here we are able to construct Φ directly using the distance functions $\operatorname{dist}(\cdot, K)$ and $\operatorname{dist}(\cdot, x)$ for various $x \in M$.

Theorem 3.1 then follows immediately from Theorem 3.7 and Step 1 in the outline above. To conclude Theorem 3.6, we need to show that the invariant upper Minkowski dimension of a quasicircle is less than $n - 1$, which we do in §3.5.

3.2 Deforming the nearest-point retraction to a smooth map

In this section we deform the nearest-point retraction $r : M \to K$ to a convex set K to a smooth map $\tilde{r}$ with $\sup_{x \in M} \operatorname{dist}(r(x), \tilde{r}(x)) < \infty$, so that

$$\begin{aligned} \|D\tilde{r}\| &\lesssim e^{-a\operatorname{dist}(\cdot,K)}, \\ \left\|D^2\tilde{r}(X,X)\right\| &\lesssim e^{-a\operatorname{dist}(\cdot,K)}\|X\|^2 . \end{aligned}$$

In §3.2.1, we describe how to modify any Lipschitz map $f : X \to Y$ to a smooth map where the first two derivatives at $x \in X$ are controlled by the local Lipschitz constant of f near x. This is essentially a restatement of the results of [5, §2.2] suitable for our purposes. In §3.2.2 we show that the local Lipschitz constant of the nearest-point retraction decays exponentially with the distance from the convex set.

3.2.1 Deforming Lipschitz maps to smooth maps

We do this by the methods of Benoist and Hulin in [5, §2.2]. We collect their results as Lemma 3.8.

Lemma 3.8. *Let $f : X \to Y$ be a Lipschitz map between pinched Hadamard manifolds X and Y. Then there exists a smooth map $\tilde{f} : X \to Y$ at a bounded distance from f*

and a polynomial P with non-negative coefficients and $P(0) = 0$ such that whenever $x \in X$ and f is L-Lipschitz in a neighbourhood of x, we have

$$\left\| D\tilde{f} \right\|_\infty \leq P(L) \text{ and } \left\| D^2\tilde{f} \right\|_\infty \leq P(L).$$

Proof. This is the result of [5, Lemma 2.8] and a slight strengthening of [5, Lemma 2.7]. We state these results as propositions below, and combine them as in [5, Proof of Proposition 2.4. Second step].

Proposition 3.9 (Lemma 2.8 in [5])**.** *There exist constants $r_0 > 0$ and $N_0 \in \mathbb{Z}_{>0}$ such that for each $r < r_0$, any r-separated subset of X can be decomposed into at most N_0 disjoint subsets, each of which is $4r$-separated.*

Proposition 3.10 (Strengthening of Lemma 2.7 in [5])**.** *Let $g : X \to Y$ be a map between pinched Hadamard manifolds. Then for all $r > 0$ small enough, there exists a map $g_{r,x} : X \to Y$ such that*

$$g_{r,x}(z) = \begin{cases} g(x) & \text{if } \operatorname{dist}(x,z) \leq \frac{r}{2}, \\ g(z) & \text{if } \operatorname{dist}(x,z) \geq r. \end{cases}$$

Moreover, if g is L-Lipschitz on $B_r(x)$, so is $g_{r,x}$. If we further assume that g is C^2 on some neighbourhood U of x, then

$$\left\| D^2 g_{r,x} \right\|_{L^\infty(U)} \lesssim L + L^2 + \left\| D^2 g \right\|_{L^\infty(U)},$$

where the implied constants depend on r.

Proof. To construct $g_{r,x}$, we use harmonic coordinates, that we already used in the proof of Theorem 2.7. We state in the proposition below the properties of these coordinates that we will use.

Proposition 3.11 (Lemma 2.6 in [5])**.** *There exist constants $r_0 > 0$ and $c_0 > 1$ such that for any $y \in Y$, there exists a chart $\Phi_y : B_{r_0}(y) \to U_y \subseteq \mathbb{R}^{\dim Y}$ such that $\Phi_y(y) = 0$ and*

$$\left\| D\Phi_y \right\|, \left\| D\left(\Phi_y^{-1}\right) \right\|, \left\| D^2\Phi_y \right\|, \left\| D^2\left(\Phi_y^{-1}\right) \right\| \leq c_0.$$

Here $U_y \subseteq \mathbb{R}^{\dim Y}$ is given the standard Euclidean metric. In particular, we have for $r < r_0$,

$$B_{c_0^{-1}r}(0) \subseteq \Phi_y(B_r(y)) \subseteq B_{c_0 r}(0).$$

Let $\chi : \mathbb{R} \to [0,1]$ be a smooth function with $\chi|_{[-\frac{1}{2},\frac{1}{2}]} = 0$, and $\chi|_{\mathbb{R}\setminus[-1,1]} = 1$. Write $\chi_r(x) = \chi(\frac{x}{r})$, and let $\Phi.$ be the coordinates given by Proposition 3.11. We define

$$g_{r,x}(z) = \begin{cases} g(x) & \text{if } \operatorname{dist}(x,z) \le \frac{r}{2} \\ \Phi_x^{-1}\left(\chi_r(\operatorname{dist}(x,z))\Phi_x(g(z))\right) & \text{if } \frac{r}{2} \le \operatorname{dist}(x,z) \le r \\ g(z) & \text{otherwise} \end{cases}.$$

Note that this is the exact same function as in [5, Lemma 2.7], and is well-defined when $c_0^2 r \operatorname{Lip}(g) < r_0$. In this proof, denote by $d(z) = \operatorname{dist}(x,z)$, for ease of notation. We have for $\frac{r}{2} \le d(z) \le r$,

$$D_z g_{r,x} = D_{\chi_r(d(z))\Phi_x(g(z))}\Phi_x^{-1}\left(\chi_r'(d(z))\left(D_z d\right)\Phi_x(g(z)) + \chi_r(d(z))D_{g(z)}\Phi_x D_z g\right),$$

so

$$\left\|D_z g_{r,x}\right\| \lesssim \left\|\Phi_x(g(z))\right\| + \left\|D_z g\right\| \lesssim \left\|D_z g\right\|.$$

Taking one more derivative, we see that

$$\left\|D_z^2 g_{r,x}\right\| \lesssim \left\|D_z g\right\|^2 + \left\|D_z g\right\| + \left\|D_z^2 g\right\|.$$

The result now follows from $\|D_z g\| \lesssim \operatorname{Lip}(g|_{B_r(x)})$. □

The rest of the proof is completely analogous to [5, Proof of Proposition 2.4. Second step]. Let $r > 0$ be small enough to be chosen later. Let X_0 be a maximal $\frac{r}{2}$-separated subset of X. By Proposition 3.9 we can write $X_0 = X_1 \cup X_2 \cup \cdots \cup X_{N_0}$ where each X_i is $2r$-separated. Define $f_0 = f$, and set

$$f_i(z) = \begin{cases} (f_{i-1})_{r,x}(z) & \text{if } z \in B_r(x) \text{ for some } x \in X_i \\ f_{i-1}(z) & \text{otherwise} \end{cases}$$

Set $\tilde{f} = f_{N_0}$. Then since each point $z \in X$ is in $B_{r/2}(x)$ for some $x \in X_0$ (by maximality of X_0), some f_i is locally constant near x and hence f_{N_0} is smooth at x.

By Proposition 3.10, whenever f has Lipschitz constant at most L near $x \in X$, we have

$$\left\|D_x \tilde{f}\right\| \lesssim L,$$
$$\left\|D_x^2 \tilde{f}\right\| \le \sum_{i=0}^{N_0} P^{\circ i}(L),$$

where $P(z) = \Lambda(z + z^2)$ for some large enough constant $\Lambda > 0$. The result follows. □

3.2.2 Local Lipschitz constant of the nearest-point retraction

We now show that the local Lipschitz constant of the nearest-point retraction to K decays exponentially with the distance from K. This is a basic result in CAT$(-a^2)$ geometry, and is probably not new.

Proposition 3.12. *Let K be a closed convex subset of M, and let $r : M \to K$ be the nearest-point retraction. Then its r has local Lipschitz constant at $x \in M$ at most $Ce^{-a\mathrm{dist}(x,K)}$, for some constant C depending only on M.*

Proof. Near K the claim follows from the fact that r is 1-Lipschitz. Assume therefore that $s = \mathrm{dist}(x, K) \geq 10a^{-1}$, and let $y \in M \backslash N_{s-a^{-1}}(K)$ be such that $\mathrm{dist}(r(x), r(y)) \leq a^{-1}$.

We note that since K is convex, we have $[r(x), r(y)] \subseteq K$. Since r is the nearest-point retraction, we have $\measuredangle_{r(x)}(x, r(y)), \measuredangle_{r(y)}(y, r(x)) \geq \pi/2$.

Pick comparison triangles $\overline{r(x)xy}$ for $r(x)xy$ and $\overline{r(x)yr(y)}$ for $r(x)yr(y)$ as in Theorem 2.3. We glue them appropriately to get a hyperbolic polygon in $\mathbb{H}^2$. We work in the disk model, and we can suppose without loss of generality that $\overline{r(x)}$ corresponds to the origin. Suppose that $\overline{x}$, $\overline{y}$ and $\overline{r(y)}$ correspond to $A, B, C \in \mathbb{D}$, respectively. We know that

1. the angle between $[0, A]$ and $[0, C]$ is at least $\pi/2$,
2. the angle between $[C, 0]$ and $[C, B]$ is at least $\pi/2$,
3. $\mathrm{dist}(A, B) = a\mathrm{dist}(x, y) \leq 1$, and
4. $\mathrm{dist}(0, C) = a\mathrm{dist}(r(x), r(y))$, $\mathrm{dist}(0, [A, B]) \geq as - 1$.

Claim 3.13. We have $\mathrm{dist}(A, B) \gtrsim e^{as}\mathrm{dist}(0, C)$.

Proof. We can suppose without loss of generality that $\measuredangle_0(A, C) = \measuredangle_C(0, B) = \frac{\pi}{2}$.

From hyperbolic law of cosines applied twice to the right-angled triangle $0BC$, we see that

$$\begin{aligned}\cos \measuredangle_0(B, C) &= \frac{\cosh \mathrm{dist}(0, C) \cosh \mathrm{dist}(0, B) - \cosh \mathrm{dist}(B, C)}{2 \sinh \mathrm{dist}(0, B) \sinh \mathrm{dist}(0, C)} \\ &= \cosh \mathrm{dist}(B, C) \frac{\cosh^2 \mathrm{dist}(0, C) - 1}{2 \sinh \mathrm{dist}(0, B) \sinh \mathrm{dist}(0, C)} \\ &= \frac{\cosh \mathrm{dist}(B, C)}{2 \sinh \mathrm{dist}(0, B)} \sinh \mathrm{dist}(0, C) \approx \mathrm{dist}(0, C).\end{aligned}$$

Thus $\sin \measuredangle_0(A,B) \approx \mathrm{dist}(0,C)$, and consequently $\measuredangle_0(A,B) \approx \mathrm{dist}(0,C)$.

Since the metric on $\mathbb{D}$ is $d\rho^2 + \sinh^2 \rho d\theta^2$ in geodesic polar coordinates and $\mathrm{dist}(0,[A,B]) \geq as-1$, we have

$$\mathrm{dist}(A,B) \geq \measuredangle_0(A,B)\sinh(as-1) \approx e^{as}\mathrm{dist}(0,C).$$

□

From the claim we have

$$\mathrm{dist}(x,y) \gtrsim \mathrm{dist}(A,B) \gtrsim e^{as}\mathrm{dist}(0,C) \approx e^{as}\mathrm{dist}(r(x),r(y)),$$

which concludes the proof. □

Applying Lemma 3.8 to the nearest-point retraction gives the following corollary.

Corollary 3.14. *For any convex subset $K \subseteq M$ there exists a map $\tilde{r}: M \to K$ with $\sup_{x\in M} \mathrm{dist}(r(x),\tilde{r}(x)) < \infty$ so that*

$$\|D\tilde{r}\| \lesssim e^{-a\mathrm{dist}(\cdot,K)},$$
$$\left\|D^2\tilde{r}\right\| \lesssim e^{-a\mathrm{dist}(\cdot,K)}.$$

3.3 Upper bound on the volume of the convex hull within a large ball

In this section we show that given an upper bound on the invariant upper Minkowski dimension of a set $S \subseteq \partial_\infty M$, we get an upper bound on $\mathrm{vol}(B_\rho(x) \cap N_d(\mathrm{CH}(S)))$ for any $d > 0$.

Recall that the invariant upper Minkowski dimension is defined using the visual metric $\mathrm{dist}_x^{\mathrm{vis}}(\cdot,\cdot)$, where $x \in M$ is some fixed basepoint, satisfying

$$A^{-1}e^{-a\mathrm{dist}(x,[y,z])} \leq \mathrm{dist}_x^{\mathrm{vis}}(y,z) \leq Ae^{-a\mathrm{dist}(x,[y,z])},$$

for all $y,z \in \partial_\infty M$.

Lemma 3.15. *Let S be a set in the boundary $\partial_\infty M$. Then for all $x \in M$, we have*

$$\mathrm{vol}(B_\rho(x) \cap N_d(\mathrm{CH}(S))) \lesssim e^{a\rho\beta},$$

for any $\beta > \overline{\dim}S, d > 0$, where the implicit constant depends only on M, d and β.

We first outline the proof of Lemma 3.15. We will estimate the volume of the intersection of $N_d(\mathrm{CH}(S))$ with spherical shells $\mathrm{An}(R) = B_{R+1}(x) \setminus B_R(x)$ (here An stands for annulus). To achieve this, we cover the set S with balls $B_1, B_2, ..., B_N$ of radius e^{-aR} in the visual metric d_x^{vis}. The proof has three ingredients, sketched below.

1. We first show that $\mathrm{CH}(S) \subseteq N_C(\mathrm{Cone}(x, S))$ for some absolute constant $C > 0$, where $\mathrm{Cone}(x, S) = \bigcup_{y \in S}[x, y]$.
2. We next show that $N_C(\mathrm{Cone}(x, B_i)) \setminus B_R(x) \subseteq \mathrm{Cone}(x, \tilde{B}_i)$, where $\tilde{B}_i$ is the ball with the same center as B_i, but has radius larger by a bounded factor. This result explicitly uses that B_i has radius e^{-aR}.
3. Finally, we show that $\mathrm{Cone}(x, \tilde{B}_i) \cap \mathrm{An}(R)$ has bounded diameter independent of R, and hence bounded volume.

Combining these three ingredients, we see that $\mathrm{vol}(B_\rho(x) \cap N_d(\mathrm{CH}(S))) \lesssim N = N(R) \lesssim e^{a\rho\overline{\dim} S}$ by assumption. The implicit constant in this inequality does not depend on $x \in M$ since in Definition 3.2 we require the estimate (3.1.1) to be uniform in $x \in M$. The rest of this section is devoted to proving Lemma 3.15.

3.3.1 Notation

For $S \subseteq \partial_\infty M$, denote by $\mathrm{Cone}(x, S)$ the union of geodesic rays with one endpoint x and the other endpoint (at infinity) in S. We denote by $\pi_x : M \setminus \{x\} \to \partial_\infty M$ the projection that maps $y \in M \setminus \{x\}$ to the unique point $z \in \partial_\infty M$ so that $y \in [x, z]$. We also write, for the duration of this proof $\mathrm{An}(R) = B_{R+1}(x) \setminus B_R(x)$. We also remind the reader that $[a, b]$ denotes the geodesic segment connecting $a, b \in \overline{M}$ (potentially infinite on one or both sides).

3.3.2 Estimating the convex hull with the cone

The purpose of this subsection is to show the proposition below.

Proposition 3.16. *There exists a constant C such that for all $S \subseteq \partial_\infty M$ and $x \in M$, we have $\mathrm{CH}(S) \subseteq N_C(\mathrm{Cone}(x, S))$.*

Proof. Denote by $\mathrm{GH}(S)$ the union of all geodesics with both endpoints in S. Clearly $\mathrm{GH}(S) \subseteq \mathrm{CH}(S)$. The following claim is a basic result due to Bowditch [10, Lemma 2.6].

Claim 3.17. For some constant C' depending only on M, we have $\mathrm{CH}(S) \subseteq N_{C'}(\mathrm{GH}(S))$.

Note that for any $a, b \in S$, since M is δ-hyperbolic (as a metric space) for some $\delta > 0$, we have $[a, b] \subseteq N_\delta([x, a] \cup [x, b])$, and therefore

$$\mathrm{GH}(S) \subseteq N_\delta(\mathrm{Cone}(x, S)),$$

for all $x \in M$ and $S \subseteq \partial_\infty M$. In particular $\mathrm{CH}(S) \subseteq N_{C'+\delta}(\mathrm{Cone}(x, S))$, so we set $C = C' + \delta$. □

3.3.3 Estimating the neighbourhood of a cone

In this subsection, we show the following proposition.

Proposition 3.18. *For any $R > 0$ and constant $C > 0$, there exists a constant $\tilde{C} = \tilde{C}(C, M)$, such that for all sets $S \subseteq \partial_\infty M$, we have*

$$N_C(\mathrm{Cone}(x, S)) \setminus B_R(x) \subseteq \mathrm{Cone}(x, N_{\tilde{C}e^{-aR}}(S)).$$

Proof. This is essentially equivalent to the following claim. We remind the reader that $\pi_x(y)$ is the unique point of intersection of the half-ray xy with the boundary at infinity $\partial_\infty M$.

Claim 3.19. For any $C > 0$, there exists $D = D(C)$ such that for all $x, y, z \in M$ with $\mathrm{dist}(y, z) \leq C$, we have

$$\mathrm{dist}(x, y) - \mathrm{dist}(x, [\pi_x(y), \pi_x(z)]) \leq D(C).$$

Proof. Since M is $\mathrm{CAT}(-a^2)$, it is also δ-Gromov hyperbolic, for some $\delta = \delta(a) > 0$. Standard facts about Gromov hyperbolic spaces [38, pp. 4-5] imply the claim below.

Claim 3.20. Let X be a $\mathrm{CAT}(-1)$ metric space, and let $a, b, c \in X$. Let $\hat{a} \in [b, c], \hat{b} \in [a, c], \hat{c} \in [a, b]$ be such that

$$\begin{aligned}
\mathrm{dist}(a, \hat{b}) = \mathrm{dist}(a, \hat{c}) &= \frac{\mathrm{dist}(a, b) + \mathrm{dist}(a, c) - \mathrm{dist}(b, c)}{2}, \\
\mathrm{dist}(b, \hat{a}) = \mathrm{dist}(b, \hat{c}) &= \frac{\mathrm{dist}(b, a) + \mathrm{dist}(b, c) - \mathrm{dist}(a, c)}{2}, \\
\mathrm{dist}(c, \hat{a}) = \mathrm{dist}(c, \hat{b}) &= \frac{\mathrm{dist}(c, a) + \mathrm{dist}(c, b) - \mathrm{dist}(a, b)}{2}.
\end{aligned}$$

Then there exists an absolute constant $\delta > 0$ (independent of X, a, b, c) such that

$$\begin{gathered}
\left|\mathrm{dist}(a, \hat{b}) - \mathrm{dist}(a, [b, c])\right| \leq \delta, \\
\text{dist-H}([a, \hat{b}], [a, \hat{c}]) \leq \delta, \\
\text{dist-H}([b, \hat{a}], [b, \hat{c}]) \leq \delta, \\
\text{dist-H}([c, \hat{a}], [c, \hat{b}]) \leq \delta, \\
\mathrm{dist}(\hat{a}, \hat{b}), \mathrm{dist}(\hat{a}, \hat{c}) \leq \delta,
\end{gathered}$$

where dist-H denotes the Hausdorff distance.

Proof. Note that [38, pp. 4-5] states all but the last estimate for Gromov hyperbolic spaces, while the last estimate easily follows from comparison with the hyperbolic plane. ☐

Considering the triangle $x\pi_x(y)\pi_x(z)$, it follows that there exist points $\bar{y} \in [x, \pi_x(y)]$, $\bar{z} \in [x, \pi_x(z)]$, and $w \in [\pi_x(y), \pi_x(z)]$, with the following properties

$$\operatorname{dist}(x, \bar{y}) = \operatorname{dist}(x, \bar{z}) = \operatorname{dist}(x, [\pi_x(y), \pi_x(z)]) + O(\delta), \tag{3.3.1}$$

$$\text{dist-H}\left([x, \bar{y}], [x, \bar{z}]\right) \leq \delta, \tag{3.3.2}$$

$$\text{dist-H}\left([\pi_x(y), \bar{y}], [\pi_x(y), w]\right) \leq \delta, \tag{3.3.3}$$

$$\text{dist-H}\left([\pi_x(z), \bar{z}], [\pi_x(z), w]\right) \leq \delta, \tag{3.3.4}$$

$$\operatorname{dist}(w, \bar{y}) \leq \delta. \tag{3.3.5}$$

By (3.3.1), it suffices to show that $\operatorname{dist}(x, y) - \operatorname{dist}(x, \bar{y}) \leq C + O(\delta)$.

We assume that $\operatorname{dist}(x, y) - \operatorname{dist}(x, \bar{y}) \geq C + 10\delta$. It follows that $\operatorname{dist}(y, [x, \bar{y}]) \geq C + 10\delta$, and hence $\operatorname{dist}(y, [x, \bar{z}]) \geq C + 9\delta$. Therefore $\operatorname{dist}(z, [x, \bar{z}]) \geq 8\delta$, and hence $\operatorname{dist}(z, [\pi_x(z), w]) \leq \delta$.

It now follows that $\operatorname{dist}(z, [\pi_x(y), w]) \leq \delta$ and $\operatorname{dist}(z, [\pi_x(z), w]) \leq \delta$ by (3.3.3) and (3.3.4), respectively. Therefore there exist point $\hat{z} \in [\pi_x(z), w]$ and $\hat{y} \in [\pi_x(y), w]$, such that

$$\operatorname{dist}(y, \hat{y}) \leq \delta,$$
$$\operatorname{dist}(z, \hat{z}) \leq \delta.$$

In particular, $\operatorname{dist}(\hat{y}, \hat{z}) \leq C + 12\delta$, and hence $\operatorname{dist}(\hat{y}, w) \leq C + 12\delta$. By (3.3.5), we have $\operatorname{dist}(\hat{y}, \bar{y}) \leq C + 13\delta$, and thus $\operatorname{dist}(y, \bar{y}) \leq C + 14\delta$. But $\operatorname{dist}(y, \bar{y}) = \operatorname{dist}(x, y) - \operatorname{dist}(x, \bar{y})$, which concludes the proof. ☐

Let $z \in N_C(\operatorname{Cone}(x, S)) \setminus B_R(x)$, so that for some $y \in \operatorname{Cone}(x, S)$ we have $\operatorname{dist}(y, z) \leq C$. By Claim 3.19,

$$\operatorname{dist}(x, [\pi_x(y), \pi_x(z)]) \geq \operatorname{dist}(x, y) - D,$$

so that $\operatorname{dist}_x^{\text{vis}}(\pi_x(y), \pi_x(z)) \leq Ae^{aD}e^{-a\operatorname{dist}(x,y)}$. But $\pi_x(y) \in S$ since $y \in \operatorname{Cone}(x, S)$, and hence $\operatorname{dist}_x^{\text{vis}}(\pi_x(z), S) \leq Ae^{aD}e^{-a\operatorname{dist}(x,y)}$. Since $\operatorname{dist}(x, z) \geq R$, we have $\operatorname{dist}(x, y) \geq R - \operatorname{dist}(y, z) \geq R - C$ and hence $\operatorname{dist}_x^{\text{vis}}(\pi_x(z), S) \leq Ae^{a(C+D)}e^{-aR}$. We thus let $\tilde{C} = Ae^{a(C+D(C))}$, and see that $\pi_x(z) \in N_{\tilde{C}}(S)$, and hence $z \in \operatorname{Cone}(x, N_{\tilde{C}}(S))$. Since $\tilde{C}$ or this argument do not depend on the choice of z, we are done. ☐

3.3.4 Decomposition

In this and the next two subsections, we show Lemma 3.15 using Propositions 3.16 and 3.18.

Set $\varepsilon = e^{-aR}$, and cover the set S by $N = N(\varepsilon)$ balls of radius ε, centered at $y_1, y_2, ..., y_N$. Since $\mathrm{CH}(S) \subseteq N_C(\mathrm{Cone}(x, S))$ by Proposition 3.16, and $S \subseteq \bigcup_{i=1}^N B_\varepsilon(y_i)$, we have

$$\begin{aligned} N_d(\mathrm{CH}(S)) \cap \mathrm{An}(R) &\subseteq \mathrm{An}(R) \cap \bigcup_{i=1}^{N} N_{C+d}(\mathrm{Cone}(x, B_\varepsilon(y_i))) \setminus B_R(x) \\ &\subseteq \mathrm{An}(R) \cap \bigcup_{i=1}^{N} \mathrm{Cone}(x, B_{\varepsilon+\tilde{C}e^{-aR}}(y_i)), \end{aligned} \tag{3.3.6}$$

where we used Proposition 3.18 in going from the first to the second line.

3.3.5 Volume bound on cones over visual balls

In this subsection, we show that each piece in the decomposition (3.3.6) has bounded volume.

Claim 3.21. Fix a constant C. Then for points $x \in M, y \in \partial_\infty M$, define the set

$$S_{R,C}(x,y) = \mathrm{Cone}\left(x, \{z \in \partial_\infty M : \mathrm{dist}(x, [y,z]) \geq R\}\right) \cap B_{R+C}(x) \setminus B_R(x).$$

Then the diameter of $S_{R,C}$ is bounded by a constant $D(C)$ depending only on C and M.

Proof. Let $b \in [x,y]$ be such that $\mathrm{dist}(x,b) = R$. We will show that $S_{R,C}(x,y)$ is contained in a ball around b of bounded radius.

Let $a \in S_{R,C}(x,y)$ and let $z \in \partial_\infty M$ be such that $a \in [x,z]$. Applying Claim 3.20 to the triangle xyz, we see that there exist points $u \in [x,y], v \in [x,z]$ such that $\mathrm{dist}(x,u) = \mathrm{dist}(x,v) = \mathrm{dist}(x,[y,z]) + O(1) \geq R - O(1)$ and such that

$$\text{dist-H}([x,u],[x,v]) \leq O(1).$$

Note that since $R \leq \mathrm{dist}(x,a), \mathrm{dist}(x,b) \leq R+C$, we have $\mathrm{dist}(a,[x,v]) \leq C+O(1)$ and $\mathrm{dist}(b,[x,u]) \leq C+O(1)$. It follows that

$$\mathrm{dist}(a,[x,u]), \mathrm{dist}(b,[x,u]) \leq C + O(1).$$

Thus there exist points $\hat{a}, \hat{b} \in [x,u]$ such that $\mathrm{dist}(a,\hat{a}), \mathrm{dist}(b,\hat{b}) \leq C + O(1)$. Thus $\mathrm{dist}(x,\hat{a}), \mathrm{dist}(x,\hat{b}) = R + O(C+1)$, and hence $\mathrm{dist}(\hat{a},\hat{b}) = O(C+1)$. It now follows that $\mathrm{dist}(a,b) \leq O(C+1)$, concluding the proof. □

Denote by $V(C)$ the maximal volume of a ball of radius $D(C)$. Note that for $z \in B_{\varepsilon+\tilde{C}e^{-aR}}(y_i)$, we have

$$\varepsilon + \tilde{C}e^{-aR} = e^{-aR}(1+\tilde{C}) \geq \mathrm{dist}_x^{\mathrm{vis}}(y_i, z) \geq A^{-1}e^{-a\mathrm{dist}(x,[y_i,z])},$$

and hence $\mathrm{dist}(x,[y_i,z]) \geq R - \frac{1}{a}\log A(1+\tilde{C})$. Therefore

$$B_{R+1}(x) \cap \mathrm{Cone}(x, B_{\varepsilon(1+AC')}(y_i)) \subseteq S_{R-\frac{1}{a}\log A(1+\tilde{C}), 1+\frac{1}{a}\log A(1+\tilde{C})}(x, y_i),$$

and hence

$$\mathrm{vol}(\mathrm{An}(R) \cap N_{C+d}(\mathrm{Cone}(x, B_\varepsilon(y_i)))) \leq V(1 + a^{-1}\log A(1+\tilde{C})) =: V_0(d).$$

3.3.6 Finishing the proof

We now combine previous results to show the main volume estimate. We have

$$\mathrm{vol}(N_d(\mathrm{CH}(S)) \cap \mathrm{An}(R)) \leq N(e^{-aR})V_0 \lesssim e^{aR\beta}V_0.$$

Hence we have

$$\mathrm{vol}(N_d(\mathrm{CH}(S)) \cap B_\rho(x)) \leq \sum_{r=0}^{\lfloor \rho \rfloor} \mathrm{vol}(\mathrm{CH}(K) \cap \mathrm{An}(r)) \lesssim e^{a\rho\beta}V_0(d),$$

where the implicit constant depends only on M and β.

3.4 Constructing bounded subharmonic functions

In this section we show that assuming $\overline{\dim} S < n-1$, there exists a bounded subharmonic map $\Phi : M \to \mathbb{R}$ with

$$\Delta\Phi \geq e^{-a\mathrm{dist}(\cdot,K)}.$$

For some $C > 0$ large enough, on $M \setminus N_C(K)$, we construct Φ as some function of the distance $\delta(x) = \mathrm{dist}(x, \tilde{r}(x))$. Hence to bound $\Delta\Phi$, we need bounds on the Laplacian and derivative of δ. This is already known (see e.g. Benoist and Hulin [6, Remark 4.6]), but we include a different proof in §3.4.1 as Proposition 3.22 for completeness. We construct Φ on $M \setminus N_C(K)$ in §3.4.2.

Cover $N_C(K)$ by balls of radius 1 with centers in a suitable discrete set V. On $N_C(K)$, we construct Φ as a sum $\sum_{v\in V} f(\mathrm{dist}(v,\cdot))$ for some smooth function $f : [0,\infty) \to \mathbb{R}$. We choose f to ensure that $\Delta\Phi \geq 1$ on $N_C(K)$. The assumption that $\overline{\dim} S < n-1$ gives us bounds on $\left|V \cap B_\rho(x)\right|$, which we use to ensure that the sum $\sum_{v\in V} f(\mathrm{dist}(v,\cdot))$ converges to a smooth function. This is done in §3.4.3.

Finally we use Φ to show Theorem 3.7 in §3.4.4.

3.4.1 Properties of the distance function

We set $\delta(x) = \operatorname{dist}(x, \tilde{r}(x))$. We will consider δ on $M \setminus N_C(K)$, for some C large to be chosen later. In particular, we let C be large enough so that $x \neq \tilde{r}(x)$ for all $x \in M \setminus N_C(K)$.

Proposition 3.22. *For some $C > 0$ large enough, the distance function δ is smooth on $M \setminus N_C(K)$ and has*

$$\Delta\delta \gtrsim 1 \text{ and } \|D\delta\| \lesssim 1.$$

Proof. Smoothness follows from the fact that dist : $M \times M \to \mathbb{R}$ is smooth away from the diagonal. The second estimate follows from the bounds

$$\|Dx\| \lesssim 1 \text{ and } \|D\tilde{r}\| \lesssim 1.$$

To show the first estimate, we compare $\Delta\delta$ near x to the function $\Delta\operatorname{dist}(\cdot, \tilde{r}(x))$, using formulae from [60]. Let e_α be an orthonormal frame near x, and let $p = \tilde{r}(x)$. We let e_i be an orthonormal frame near p, and set

$$\tilde{r}_* e_\alpha = \sum_i \tilde{r}^i_\alpha e_i.$$

We let $X_\alpha = e_\alpha + \tilde{r}_* e_\alpha \in T_x M \oplus T_{\tilde{r}(x)} M = T_{(x,\tilde{r}(x))} M \times M$. Note that since $\|D\tilde{r}\| \lesssim e^{-aC}$, we have $\left|\tilde{r}^i_\alpha\right| \lesssim e^{-aC}$. Then by a suitable modification of [60, equation (2.11)] we have

$$\Delta_x\left(\delta^2\right) = 2(D_{e_\alpha}\operatorname{dist})^2 + 2\sum_{\alpha\neq i}(\tilde{r}^i_\alpha D_{e_i}\operatorname{dist})^2 + 2\delta\sum_\alpha D^2_x(\operatorname{dist})(X_\alpha, X_\alpha) + 2\delta\operatorname{dist}_*\tau(\tilde{r}),$$

$$\Delta\operatorname{dist}(x,p)^2 = 2(D_{e_\alpha}\operatorname{dist})^2 + 2\operatorname{dist}(x,p)\sum_\alpha D^2_x(\operatorname{dist})(e_\alpha, e_\alpha).$$

Hence we have $\Delta\delta - \Delta\operatorname{dist}(\cdot, p) \lesssim e^{-aC}$. It is well-known (see e.g. [5, Lemma 3.1]) that $\Delta\operatorname{dist}(\cdot, p) \gtrsim 1$, so this concludes the proof. □

3.4.2 Constructing bounded subharmonic functions far from the convex set

We use the following lemma to construct Φ on $N_{d+1}(K) \setminus N_d(K)$.

Lemma 3.23. *There exists a constant $D > 0$, such that for all $d \geq D$, there exists a bounded subharmonic function $\Phi_d : M \to \mathbb{R}$ such that*

$$\Delta\Phi_d \geq 1 \text{ on } N_{d+1}(K) \setminus N_d(K),$$

so that $\sup_x\left|\Phi_d(x)\right|$ does not depend on d.

Proof. Let $f : [0, \infty) \to [0, \infty)$ be a C^2 function. Then

$$\Delta(f \circ \delta) = f'(\delta)\Delta\delta + f''(\delta)\|D\delta\|^2 .$$

By Proposition 3.22 we can suppose $\Delta\delta \geq A$ and $\|D\delta\|^2 \leq B$, for some positive constants A, B. Let $u : \mathbb{R} \to [0, \infty)$ be a C^1 function with the following properties:

1. We have $u = 0$ on $\left(-\infty, -\frac{1}{2}\right]$ and $u = 1$ on $[0, 1]$.

2. On $(-\infty, 1]$, u is non-decreasing, and on $[1, \infty)$, u is decreasing.

3. We have the differential inequality $Au + Bu' \geq 0$.

4. We have $u \approx e^{-\varepsilon x}$ for x large enough, for some $\varepsilon > 0$.

We set $f(x) = \int_0^x u(t - d)dt$. By exponential decay of u at infinity, f is bounded. When $d \geq C$ from Proposition 3.22, we have

$$\Delta(f \circ \delta) \geq Au(\delta - d) + B \min(u'(\delta - d), 0) \geq 0,$$

so $f \circ \delta$ is subharmonic, and moreover $\Delta(f \circ \delta) = \Delta\delta \gtrsim 1$ whenever $d \leq \delta(x) \leq d + 1$, or equivalently $x \in N_{d+1}(K) \setminus N_d(K)$. We rescale $f \circ \delta$ to get Φ_d. ☐

3.4.3 Constructing bounded subharmonic functions in a large neighbourhood of the convex set

In the following Lemma we construct Φ on $N_C(K)$ for arbitrarily large $C > 0$.

Lemma 3.24. *Suppose that $\overline{\dim} S < n - 1$. Then for any $C > 0$, there exists a bounded subharmonic function $\Phi_0 : M \to \mathbb{R}$ with*

$$\Delta\Phi_0 \geq 1 \text{ on } N_C(K).$$

Proof. We first show the Lemma for $K = \{v\}$ and $C = 1$ for some $v \in M$. We construct Φ_0 as $\Phi_0^v(x) = f(\operatorname{dist}(v, x))$ for some smooth non-decreasing function $f : [0, \infty) \to \mathbb{R}$. Then

$$\begin{aligned}\Delta\Phi_0 &= f'(\operatorname{dist}(v, \cdot))\Delta\operatorname{dist}(v, \cdot) + f''(\operatorname{dist}(v, \cdot))\left\|D\operatorname{dist}(v, \cdot)\right\|^2 \\ &\geq (n-1)af'(\operatorname{dist}(v, \cdot)) + f''(\operatorname{dist}(v, \cdot)),\end{aligned}$$

where we used the fact that $\Delta\operatorname{dist}(v, \cdot) \geq (n-1)a$, and $\left\|D\operatorname{dist}(v, \cdot)\right\| = 1$. The equality follows from the fact that $\operatorname{dist}(v, \cdot)$ is 1-Lipschitz and linear with slope 1 along any geodesic through v. The inequality follows from the well-known (see e.g. [5, Lemma 2.3])

$$D^2(\mathrm{dist}(v,\cdot)) \geq a\coth(a\mathrm{dist}(v,\cdot))(g - d\mathrm{dist}(v,\cdot)^{\otimes 2}),$$

where g is the Riemannian metric on M. We construct the function f in the following claim.

Claim 3.25. There exists a C^2 function $f : [0,\infty) \to \mathbb{R}$ with the following properties,

1. In a neighbourhood of 0, we have $f(x) = f(0) + x^2$,
2. We have $(n-1)af'(x) + f''(x) \geq 0$ for all $x \geq 0$,
3. For $x \in [0,1]$, we have $(n-1)af'(x) + f''(x) \geq 2$.
4. We have $\left|f^{(k)}(x)\right| \lesssim e^{-a(n-1)x}$ for all $k \geq 0$.

Proof. Note that the second property is equivalent to the function $x \to f'(x)e^{(n-1)ax}$ being non-decreasing. Since $2xe^{(n-1)ax}$ is increasing near 0, we can construct a non-decreasing smooth function $g : [0,\infty) \to \mathbb{R}$ that agrees with $2xe^{(n-1)ax}$ on $[0,1]$, and that is constant on $[2,\infty)$. We can then set

$$f(x) = -\int_x^\infty g(t)e^{-(n-1)at}dt.$$

By construction, properties 1) and 2) hold for such f. Property 4) holds since $g(t)e^{-(n-1)at}$ is proportional to $e^{-(n-1)at}$ for t large enough.

To verify property 3), note that for $0 \leq x \leq 1$ we have

$$(n-1)af'(x) + f''(x) = e^{-(n-1)ax}g'(x) = 2 + 2x(n-1)a \geq 2.$$

□

We now have for $d = \mathrm{dist}(v,\cdot)$,

$$\Delta\Phi_0^v \geq (n-1)af'(d) + f''(d) \geq \begin{cases} 1 & \text{when } d \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

We also record the following properties of Φ_0^v,

1. since $f(x) = f(0) + x^2$ in a neighbourhood of 0 and $\mathrm{dist}(v,\cdot)^2$ is C^2, it follows that Φ_0^v is C^2, and

2. since $\left|f^{(k)}(x)\right| \lesssim e^{-(n-1)ax}$ for $0 \leq k \leq 2$, and by boundedness of $\Delta \mathrm{dist}(v, \cdot)$ and $\left\|D\mathrm{dist}(v, \cdot)\right\|$, we have

$$|\Phi_0|, \|D\Phi_0\|, |\Delta\Phi_0| \lesssim e^{-a\mathrm{dist}(v,\cdot)}.$$

We now turn to the general case of $K = \mathrm{CH}(S)$. Let V be a maximal 1-separated subset of $N_C(K)$. Set

$$\Phi_0(x) = \sum_{v \in V} \Phi_0^v(x). \tag{3.4.1}$$

We now show that the sum above converges, and that Φ_0 is bounded. Set $V_\rho(x) = \{v \in V : \rho \leq \mathrm{dist}(x, v) < \rho + 1\}$. Note that since V is 1-separated, the balls $\{B_{1/2}(v) : v \in V\}$ are disjoint, hence

$$\left|V_\rho(x)\right| \leq \frac{\mathrm{vol}(B_{\rho+3/2}(x) \cap N_C(K))}{\inf_{x\in M} \mathrm{vol}(B_{1/2}(x))} \lesssim e^{a\rho(\overline{\dim}S+\varepsilon)},$$

for any $\varepsilon > 0$ by Lemma 3.15. Therefore

$$\sum_{v\in V_\rho(x)} \left|\Phi_0^v(x)\right| \lesssim e^{-a\rho(n-1-\overline{\dim}S-\varepsilon)},$$
$$\sum_{v\in V_\rho(x)} \left\|D\Phi_0^v(x)\right\| \lesssim e^{-a\rho(n-1-\overline{\dim}S-\varepsilon)},$$
$$\sum_{v\in V_\rho(x)} \left|\Delta\Phi_0^v(x)\right| \lesssim e^{-a\rho(n-1-\overline{\dim}S-\varepsilon)}.$$

It follows from the first equation that the sum from (3.4.1) converges absolutely and locally uniformly. Since the constants are independent of x, it follows that Φ_0 is bounded. The second and third equations imply that

$$\Delta\Phi_0(x) = \sum_{v\in V} \Delta\Phi_0^v(x) \geq \left|\{v \in V : \mathrm{dist}(x, v) \leq 1\}\right| \geq \begin{cases} 1 & \text{on } N_C(K), \\ 0 & \text{otherwise,} \end{cases}$$

by maximality of V. □

3.4.4 Proof of Theorems 3.7 and 3.1

Proof of Theorem 3.7. Let $\Phi_d : M \to \mathbb{R}$ for $d = 0$ or $d \geq D$ be the functions from Lemmas 3.23 and 3.24, and let

$$\Phi(x) = \Phi_0(x) + \sum_{d\geq D} e^{-ad}\Phi_d(x).$$

Then Φ is bounded and subharmonic, with

$$\Delta\Phi \gtrsim e^{-a\mathrm{dist}(\cdot,K)}.$$

Let $h_d : B_d(x) \to N$ be the harmonic map such that $h_d = f$ on $\partial B_d(x)$. Then by [60], we have

$$\Delta\mathrm{dist}(h_d(x), f(x)) \geq -\big\|\tau(f)\big\| \gtrsim -e^{-a\mathrm{dist}(\cdot,K)}.$$

Hence for a suitable constant $C' > 0$ that does not depend on d, we have

$$\Delta\left(\mathrm{dist}(h_d(x), f(x)) + C'\Phi(x)\right) > 0.$$

By the maximum principle, we have

$$\mathrm{dist}(h_d(x), f(x)) \leq 2C' \sup_{x\in M}\big|\Phi(x)\big| = C''.$$

This concludes the proof by Proposition 2.10. □

Proof of Theorem 3.1. By Corollary 3.14, there exists a smooth map $\tilde{r} : M \to M$ with

$$\sup_{x\in M} \mathrm{dist}(\tilde{r}(x), r(x)) < \infty,$$
$$\big\|\tau(\tilde{r})\big\| \lesssim e^{-a\mathrm{dist}(\cdot,K)}.$$

Then by Theorem 3.7 there exists a harmonic map $h : M \to M$ with

$$\sup_{x\in M} \mathrm{dist}(h(x), \tilde{r}(x)) < \infty,$$

and hence $\sup_{x\in M} \mathrm{dist}(h(x), r(x)) < \infty$. □

3.5 Proof of Theorem 3.6

Let S be the image of $\partial\iota : S^1 \to \partial\mathbb{H}^n$, and let K be the convex hull of S.

Claim 3.26. For some d large enough, there exists a Lipschitz map $f : N_d(K) \to \mathbb{H}^2$ such that

$$\sup_{x\in\mathbb{H}^2} \mathrm{dist}(f \circ \iota(x), x) < \infty.$$

Proof. Note that $\mathbb{H}^2 = \bigcup_{z \in S^1} [-z, z]$, so that

$$\iota(\mathbb{H}^2) = \bigcup_{z \in S^1} \iota([-z, z]).$$

Each $\iota([-z, z])$ is a quasigeodesic with the same constants, so Morse lemma implies that $\iota([-z, z]) \subseteq N_C([\iota(-z), \iota(z)])$ for some $C > 0$ that depends only on quasi-isometry constants of ι. Therefore $\iota(\mathbb{H}^2) \subseteq N_C(K)$. Similarly we have $[\iota(z), \iota(w)] \subseteq N_C([z, w])$, so $\mathrm{GH}(S) \subseteq N_C(\iota(\mathbb{H}^2))$. We have already shown as part of the first Claim of Lemma 3.15 that $K \subseteq N_{C'}(\mathrm{GH}(S))$ for some constant $C' > 0$ that depends only on n, so that $K \subseteq N_{C+C'}(\iota(\mathbb{H}^2))$, and therefore the quasi-isometric embedding $\iota : \mathbb{H}^2 \to N_C(K)$ is quasisurjective.

Therefore there exists a quasi-inverse $\tilde{f} : N_C(K) \to \mathbb{H}^2$, for all C large enough, meaning $\sup_{x \in \mathbb{H}^2} \mathrm{dist}(x, \tilde{f} \circ \iota(x)) < \infty$. We in fact construct a quasi-inverse on a larger set $\tilde{f} : N_{C+1}(K) \to \mathbb{H}^2$, so that by the construction of Benoist and Hulin from [5, Proposition 2.4] we can construct a Lipschitz map $f : N_C(K) \to \mathbb{H}^2$ so that $\sup_x \mathrm{dist}(\tilde{f}(x), f(x)) < \infty$. Then f is a Lipschitz quasi-inverse of ι, as claimed. □

Now let $r : \mathbb{H}^n \to K$ be the nearest-point retraction, and write $\hat{r} = f \circ r$, for some f as in the Claim. Then $\hat{r} : \mathbb{H}^n \to \mathbb{H}^2$ is Lipschitz with

$$\mathrm{Lip}\left(\hat{r}|_{\mathbb{H}^n \setminus N_d(K)}\right) \lesssim e^{-ad},$$

for all $d > 0$, by Proposition 3.12. By Lemma 3.8 there exists a smooth map $\tilde{r} : \mathbb{H}^n \to \mathbb{H}^2$ so that

$$\begin{aligned} \sup_{x \in \mathbb{H}^n} \mathrm{dist}(\tilde{r}(x), \hat{r}(x)) &< \infty, \\ \|D_x \tilde{r}\| &\lesssim e^{-a\mathrm{dist}(x,K)}, \\ \left\|D_x^2 \tilde{r}\right\| &\lesssim e^{-a\mathrm{dist}(x,K)} \end{aligned} \tag{3.5.1}$$

for all $x \in \mathbb{H}^n$. Applying Theorem 3.7 to $\tilde{r}$ shows Theorem A.1(1). It follows from the first inequality and the Claim that

$$\sup_{x \in \mathbb{H}^2} \mathrm{dist}(x, \tilde{r} \circ \iota(x)) < \infty. \tag{3.5.2}$$

Note that each set in $\{\gamma S : \gamma \in \mathrm{Isom}(\mathbb{H}^n)\}$ is a quasicircle with the same quasisymmetry constants as S, so by [26, Theorem 18, Lemma 16] we see that

$$\overline{\dim} S < n - 1.$$

Therefore by Theorem 3.7 and (3.5.1), there exists a harmonic map $h : \mathbb{H}^n \to \mathbb{H}^2$ such that $\sup_{\mathbb{H}^n} \operatorname{dist}(h, \tilde{r}) < \infty$. Thus from (3.5.2), we get

$$\sup_{x\in\mathbb{H}^2} \operatorname{dist}(x, h \circ \iota(x)) \leq \sup_{\mathbb{H}^n} \operatorname{dist}(h, \tilde{r}) + \sup_{x\in\mathbb{H}^2} \operatorname{dist}(x, \tilde{r} \circ \iota(x)) < \infty,$$

concluding the proof of Theorem 3.6.

Chapter 4

Harmonic projections to large convex sets

4.1 Introduction

We continue our study from Chapter 3 of harmonic maps that are at a finite distance from a nearest-point projection to a convex set in a pinched Hadamard manifold. Here we show an analogue of Theorem 3.1 for convex sets that are sufficiently large.

Definition 4.1. A closed convex subset K of a pinched Hadamard manifold M is called admissible if there exists an angle θ and a distance R_0 with the following property. For any $x \in K, R > R_0$, there exists a point $\xi \in \partial_\infty M$ such that

$$\partial B_R(x) \cap \mathrm{Cone}(x\xi, \theta) \subseteq \partial B_R(x) \cap K,$$

where $\mathrm{Cone}(x\xi, \theta) = \{y \in M : \measuredangle_x(y, \xi) < \theta\}$.

Our main result is that nearest point retractions to admissible convex sets can be deformed to harmonic maps.

Theorem 4.2. *Let K be an admissible closed convex subset of a pinched Hadamard manifold M. There exists a harmonic map $h : M \to M$ that is a finite distance away from the nearest-point retraction $r : M \to K$.*

In hyperbolic spaces $\mathbb{H}^n$, a rich class of admissible convex sets is provided by convex hulls of open sets in $\partial_\infty \mathbb{H}^n \cong \mathbb{S}^{n-1}$ with sufficiently regular boundary.

Theorem 4.3. *Let $U \subseteq \partial_\infty \mathbb{H}^n = \mathbb{S}^{n-1}$ be an open set with quasiconformal boundary. Then the convex hull of U is admissible.*

Here by quasiconformal boundary we mean that near any point $x \in \partial U$, there exists a local quasiconformal map that sends U to $\mathbb{R}_+ \times \mathbb{R}^{n-2}$ and x to the origin. Note that Theorem 4.3 immediately shows Theorem A.1(2).

The proof of Theorem 4.2 uses a generalization of the "interior estimate" of Benoist–Hulin [5, §4]. Recall that the main result of [5] states that given any quasi-isometry f between pinched Hadamard manifolds, there exists a harmonic map at a bounded distance from f. As a separate application of our generalized interior estimate, we weaken the quasi-isometry requirement on f in their proof. Recall that in a pinched Hadamard manifold, we denote by $\sigma_{x,R}$ the harmonic measure on $\partial B_R(x)$ as seen from x.

Definition 4.4. A Lipschitz map $f : M \to N$ between pinched Hadamard manifolds is non-collapsing if the following two conditions hold

1. there exist constants $c, R_0 > 0$, such that for any $x \in M, R > R_0$, we have

$$\int_{\partial B_R(x)} \operatorname{dist}(f(x), f(y)) d\sigma_{x,R}(y) \geq cR,$$

   and

2. for any $\varepsilon > 0$, there exist $\theta, R_0 > 0$ such that for any $x \in M, R > R_0$ and $\xi \in \partial_\infty N$, we have

$$\sigma_{x,R}\left(\{y \in \partial B_R(x) : f(y) \neq f(x) \text{ and } \measuredangle_{f(x)}(\xi, f(y)) < \theta\}\right) < \varepsilon,$$

   where $\measuredangle_a(b, c)$ denotes the angle at a between the geodesics $[a, b]$, joining a and b, and $[a, c]$, joining a and c.

Theorem 4.5. *For any non-collapsing Lipschitz map $f : M \to N$ between pinched Hadamard manifolds, there exists a harmonic map $h : M \to N$ such that* $\sup \operatorname{dist}(h, f) < \infty$.

Note that nearest point projections are in general not non-collapsing, so Theorem 4.2 can not be derived directly from Theorem 4.5. In §4.1.1 below, we state our generalized interior estimate as Theorem 4.10, as well as a slightly more general version of Theorem 4.5 that we will in fact prove. In §4.1.2 we then describe how the rest of this chapter is organized.

4.1.1 More precise results

We will in fact prove a slightly stronger version of Theorem 4.5.

Definition 4.6. Let $\omega : \mathbb{R}_+ \to \mathbb{R}_+$ be a function such that $\omega(x) \to \infty$ and $\frac{\omega(x)}{x} \to 0$ as $x \to \infty$. Then a Lipschitz map $f : M \to N$ is called ω-weakly non-collapsing (weakly non-collapsing map with size function ω) if the following two conditions hold

1. there exist constants $c, R_0 > 0$, such that for any $x \in M, R > R_0$, we have

$$\int_{\partial B_R(x)} \mathrm{dist}(f(x), f(y)) d\sigma_{x,R}(y) \geq cR,$$

   and

2. for any $\varepsilon > 0$, there exist $\theta, R_0 > 0$ such that for any $x \in M, R > R_0$ and $\xi \in \partial_\infty N$, we have

$$\sigma_{x,R}\left(\{y \in \partial B_R(x) : \mathrm{dist}(f(x), f(y)) \geq \omega(R) \text{ and } \measuredangle_{f(x)}(\xi, f(y)) < \theta\}\right) < \varepsilon.$$

We call an $\omega : \mathbb{R}_+ \to \mathbb{R}_+$ with $\omega(x) \to \infty$ and $\frac{\omega(x)}{x} \to 0$ as $x \to \infty$ a sub-linear size function. A Lipschitz map is weakly non-collapsing if it is ω-weakly non-collapsing for some sub-linear size function ω.

Theorem 4.7. *For any weakly non-collapsing Lipschitz map $f : M \to N$, there exists a harmonic map $h : M \to N$ such that* $\sup \mathrm{dist}(h, f) < \infty$.

Remark 4.8.
1. Note that a non-collapsing map as in Definition 4.4 is a weakly non-collapsing map with size function 0, so Theorem 4.5 follows immediately from Theorem 4.7.

2. We will show below that, if f is a weakly non-collapsing map, and $\tilde{f}$ is a Lipschitz map such that $\sup \mathrm{dist}(f, \tilde{f}) < \infty$, then $\tilde{f}$ is also weakly non-collapsing (albeit with a different size function). In particular, the harmonic map obtained either from Theorem 4.5 or Theorem 4.7 is weakly non-collapsing, but not necessarily with size function 0.

3. If f is an ω-weakly non-collapsing map, and $\tilde{\omega} \geq \omega$ is a sub-linear size function, then f is also an $\tilde{\omega}$-weakly non-collapsing. Thus the condition $\omega(x) \to \infty$ as $x \to \infty$ in Definition 4.6 is superfluous, and is there merely for convenience.

Both Theorem 4.2 and 4.7 follow from our generalized interior estimate, stated below. Recall that a pointed topological space is a pair (X, x), where X is a topological space and $x \in X$ is a distinguished point, and that a map between pointed spaces $f : (X, x) \to (Y, y)$ is a map $f : X \to Y$ such that $f(x) = y$.

Definition 4.9. Let $\mathcal{F}$ be a family of smooth maps between pointed pinched Hadamard manifolds. Then $\mathcal{F}$ is uniformly non-collapsing if it is uniformly Lipschitz, if the domain and range of any function in $\mathcal{F}$ have uniformly bounded pinching constants, and if the following two conditions hold

1. There exist constants $c, R_0 > 0$, such that for any $f : (M, x) \to (N, y)$ in $\mathcal{F}$ and any $R > R_0$, we have
$$\int_{\partial B_R(x)} \operatorname{dist}(f(x), f(y)) d\sigma_{x,R}(y) \geq cR,$$
and
2. There exists a sub-linear size function $\omega : \mathbb{R}_+ \to \mathbb{R}_+$ such that for any $\varepsilon > 0$, there exist $\theta > 0, R_0 > 0$ such that, for any $f : (M, x) \to (N, y)$ in $\mathcal{F}$ and $R > R_0$, and any $\xi \in \partial_\infty N$, we have
$$\sigma_{x,R}\left(\{y \in \partial B_R(x) : \measuredangle_{f(x)}(\xi, f(y)) < \theta \text{ and } \operatorname{dist}(f(x), f(y)) \geq \omega(R)\}\right) < \varepsilon.$$

Theorem 4.10 (Generalized interior estimate)**.** *Let $\mathcal{F} = \{f_n : (M_n, x_n) \to (N_n, y_n) : n = 1, 2, ...\}$ be a uniformly non-collapsing family. Suppose R_n is a sequence of positive real numbers with $R_n \to \infty$, and let $h_n : B_{R_n}(x_n) \to N_n$ be a sequence of harmonic maps, such that the maximum of* $\operatorname{dist}(h_n, f_n)$ *is achieved at $x_n \in M_n$. Then* $\sup_n \sup_{B_{R_n}(x_n)} \operatorname{dist}(f_n, h_n) < \infty$.

4.1.2 Organization and a brief outline

Here we briefly describe the contents of each section in the chapter.

In §4.2, we show that any weakly non-collapsing Lipschitz map can be deformed to a smooth weakly non-collapsing map with bounds on the first two derivatives. This is achieved by using the same argument as in §3.2, that is in turn a slight generalization of the argument of Benoist–Hulin [5, §2]. In particular, here we merely verify that the property of being weakly non-collapsing is preserved under finite distance deformations (although the size function is not preserved). This is an important step, as the proofs of both Theorem 4.7 and Theorem 4.2 depend on computations of the

Laplacian of the distance function, using the classical computation of Schoen–Yau [60]. For this we need the underlying maps to be at least C^2, and moreover we need control on the tension field of the map that we are trying to deform to a harmonic map.

In §4.3 we prove Theorem 4.10. The main technical result in this section is Lemma 4.14, that easily implies Theorem 4.10, and that we believe is of independent interest. Lemma 4.14 is a more precise quantitative version of the "interior estimate" of [5, §4]. The proof of Theorem 4.10 relies on the observation that since $\operatorname{dist}(f_n(x_n), h_n(\cdot))$ is a subharmonic function, we have

$$\int_{\partial B_{R_n}(x_n)} \big(\operatorname{dist}(f_n(x_n), h_n(y)) - \operatorname{dist}(f_n(x_n), h_n(x_n))\big)\, d\sigma_{x_n,R_n}(y) \geq 0.$$

In the proof of Theorem 4.10, we then estimate of the integrand on the left-hand side in the regime where $\operatorname{dist}(f_n(x_n), h_n(x_n)) \to \infty$ as $n \to \infty$, along the lines of [5, §4]. This section is the heart of the chapter, and a more detailed outline can be found at the start of §4.3. In §4.4 we derive Theorem 4.7 from Theorem 4.10, using the ideas of [5] and the limiting argument in Proposition 2.10.

In §4.5 we show Theorem 4.2. The overall strategy is similar to the proof of Theorem 4.7, and is still based on Proposition 2.10. One notable difference from Theorem 4.7 is in the fact that getting the analogue of the "boundary estimate" ([5, §3.4]) requires the ideas of §3.4.2.

Finally in §4.6 we show Theorem 4.3. We give here a brief outline of the proof. Firstly, it is easy to see that the only way admissibility can fail is along a sequence of points converging to the boundary at infinity $\partial_\infty \mathbb{H}^n$. If this sequence converges to a point in U, admissibility holds. Assume therefore that the sequence converges to a point ξ in ∂U. We then use quasiconformal maps to straighten out ∂U near ξ. Note that in the model case where $U = \mathbb{R}_+ \times \mathbb{R}^{n-2} \subseteq \mathbb{R}^{n-1} \cup \{\infty\} = \mathbb{S}^{n-1}$, it is easy to show Theorem 4.3 by hand. Therefore it suffices to show that applying this quasiconformal map preserves the condition in Definition 4.1. This follows from the fact, due to Tukia–Väisälä [71], that local quasiconformal maps in $\mathbb{S}^{n-1}$ can be extended to local bi-Lipschitz maps in $\mathbb{H}^n$.

4.2 Deforming to smooth maps

Our aim here is to show that any weakly non-collapsing map can be deformed to a smooth weakly non-collapsing map, with control on the first two derivatives. Note that from Lemma 3.8, any Lipschitz map can be deformed to a smooth map with first two derivatives bounded. The following proposition is thus the aim of this section.

Proposition 4.11. *Let $\omega : \mathbb{R}_+ \to \mathbb{R}_+$ be a sub-linear size function, and let $f : M \to N$ be an ω-weakly non-collapsing map between pinched Hadamard manifolds. Suppose $\tilde{f} : M \to N$ is a Lipschitz map such that $D = \sup \operatorname{dist}(f, \tilde{f}) < \infty$. Then $\tilde{f}$ is a weakly non-collapsing map with size function $\omega(x) + 2D$.*

Proof. Suppose that $\sup_M \operatorname{dist}(f, \tilde{f}) = D$.

To check Definition 4.6(1), we write

$$\begin{aligned}\int_{\partial B_R(x)} \operatorname{dist}(\tilde{f}(x), \tilde{f}(y)) d\sigma_{x,R}(y) &\geq \int_{\partial B_R(x)} \big(\operatorname{dist}(f(x), f(y)) - 2D\big)\, d\sigma_{x,R}(y) \\ &\geq cR - 2D \geq \frac{c}{2}R\end{aligned}$$

for $R > \max(R_0, 4c^{-1}D)$, where R_0, c are constants from Definition 4.6(1) for f.

It remains to show Definition 4.6(2). We first observe that for any set $S \subseteq N$, we have

$$\tilde{f}^{-1}(S) \subseteq f^{-1}\left(N_D(S)\right). \tag{4.2.1}$$

The proof relies on the following proposition, that we show in the next subsection.

Proposition 4.12. *For any $D, \theta > 0$, there exist $\hat{D}, \hat{\theta} > 0$ such that for any two points $x, y \in M$ at a distance at most D and any $\xi \in \partial_\infty M$, we have*

$$N_D\left(\operatorname{Cone}(x\xi, \hat{\theta}) \setminus B_{\hat{D}}(x)\right) \subseteq \operatorname{Cone}(y\xi, \theta).$$

Now let $\varepsilon > 0$ be arbitrary. Let θ_0, R_0 be such that

$$\sigma_{x,R}\left(f^{-1}\left(\operatorname{Cone}(f(x)\xi, \theta_0) \setminus B_{\omega(R)}(f(x))\right) \cap \partial B_R(x)\right) < \varepsilon. \tag{4.2.2}$$

for any $x \in M, \xi \in \partial_\infty N, R > R_0$. Then choose $\theta, \hat{D} > 0$ as in Proposition 4.12 such that

$$N_D\left(\operatorname{Cone}(\tilde{f}(x)\xi, \theta) \setminus B_{\hat{D}}(\tilde{f}(x))\right) \subseteq \operatorname{Cone}(f(x)\xi, \theta_0).$$

In particular, for R large enough, we have $\omega(R) > \hat{D} - 2D$, and then we have

$$\begin{aligned}N_D\left(\operatorname{Cone}(\tilde{f}(x)\xi, \theta) \setminus B_{\omega(R)+2D}(\tilde{f}(x))\right) &\subseteq N_D\left(\operatorname{Cone}(\tilde{f}(x)\xi, \theta) \setminus B_{\hat{D}}(\tilde{f}(x))\right) \\ &\subseteq \operatorname{Cone}(f(x)\xi, \theta_0),\end{aligned}$$

and hence

$$N_D\left(\mathrm{Cone}(\tilde f(x)\xi,\theta)\setminus B_{\omega(R)+2D}(\tilde f(x))\right)\subseteq \mathrm{Cone}(f(x)\xi,\theta_0)\setminus B_{\omega(R)+D}(\tilde f(x))$$
$$\subseteq \mathrm{Cone}(f(x)\xi,\theta_0)\setminus B_{\omega(R)}(f(x)). \quad (4.2.3)$$

Combining (4.2.1), (4.2.3) and (4.2.2), we see that for R large enough, we have

$$\sigma_{x,R}\left(\tilde f^{-1}\left(\mathrm{Cone}(\tilde f(x)\xi,\theta_0)\setminus B_{\omega(R)+2D}(\tilde f(x))\right)\cap \partial B_R(x)\right)<\varepsilon,$$

for all $x\in M,\xi\in\partial_\infty N$. Thus $\tilde f$ satisfies Definition 4.6(2) with size function $\tilde\omega(x)=\omega(x)+2D$. □

4.2.1 Moving the apex of a cone

Here we show Proposition 4.12. We fix $D,\theta>0$. Let $\hat D$ (resp. $\hat\theta$) be an arbitrary positive constant, that we will freely increase (resp. decrease) over the course of the proof. By Proposition 3.18, it suffices to show

$$\mathrm{Cone}(x\xi,\hat\theta)\setminus B_{\hat D}(x)\subseteq \mathrm{Cone}(y\xi,\theta). \quad (4.2.4)$$

Remark 4.13. Note that in Chapter 3, we work with the visual metric on $\partial_\infty M$, whereas here we are interested in the angle metric. It is classical that the two are Hölder equivalent, and the direction we need follows readily from Claim 4.16 and [9, §2.5].

Let $z\in\mathrm{Cone}(x\xi,\hat\theta)\setminus B_{\hat D}(x)$ and let w be the point on $x\xi$ closest to z. Our first assertion is that

$$\mathrm{dist}(x,w)\ge\min\left(\hat D,a^{-1}\log\frac{1}{\hat\theta}\right)+O(1). \quad (4.2.5)$$

By comparison with the hyperbolic plane for the triangle xzw, we see that

$$\sinh\left(a\mathrm{dist}(z,w)\right)\le\sin\measuredangle_x(z,w)\sinh\left(a\mathrm{dist}(x,z)\right). \quad (4.2.6)$$

This in particular shows that

$$\mathrm{dist}(z,w)\le\max\left(0,\mathrm{dist}(x,z)+a^{-1}\log\measuredangle_x(z,w)\right)+O(1). \quad (4.2.7)$$

Therefore by the triangle inequality

$$\begin{aligned}\mathrm{dist}(x,w)&\ge\mathrm{dist}(x,z)-\mathrm{dist}(z,w)\\ &\ge\min\left(\mathrm{dist}(x,z),a^{-1}\log\frac{1}{\measuredangle_x(z,w)}\right)+O(1)\\ &\ge\min\left(\hat D,a^{-1}\log\frac{1}{\hat\theta}\right)+O(1),\end{aligned}$$

thus showing (4.2.5).

Let δ be the Gromov constant of M as a hyperbolic metric space. By (4.2.5), since $\operatorname{dist}(x, y) \le D$, by choosing $\hat{D}$ large enough and $\hat{\theta}$ small enough, we can arrange it so that $\operatorname{dist}(w, xy) > 10\delta$. Thus, by considering the ideal triangle $x\xi y$, we see that $\operatorname{dist}(w, y\xi) \le \delta$. Therefore

$$\operatorname{dist}(z, y\xi) \le \operatorname{dist}(z, x\xi) + \delta. \tag{4.2.8}$$

Similarly to (4.2.6), by comparison to the hyperbolic plane, we see that

$$\sinh(b\operatorname{dist}(z, y\xi)) \ge \sinh(b\operatorname{dist}(z, y)) \sin \measuredangle_y(z, \xi) \gtrsim e^{b(\operatorname{dist}(x,z)-D)} \measuredangle_y(z, \xi).$$

It follows from (4.2.8) that

$$\measuredangle_y(z, \xi) \lesssim e^{b(\operatorname{dist}(z,x\xi)-\operatorname{dist}(x,z))},$$

where we absorbed $e^{b(D+\delta)}$ into the implicit constant. Applying (4.2.7), we get

$$\begin{aligned} \measuredangle_y(z, \xi) &\lesssim \exp\Big(\max\big(-b\operatorname{dist}(x, z), ba^{-1} \log \measuredangle_x(z, w)\big)\Big) \\ &\lesssim \exp\left(-\min\left(b\hat{D}, ba^{-1} \log \frac{1}{\hat{\theta}}\right)\right). \end{aligned}$$

By increasing $\hat{D}$ and decreasing $\hat{\theta}$ further, we can ensure that $\measuredangle_y(z, \xi) < \theta$. Since z was arbitrary, and none of our constants or choices of $\hat{D}, \hat{\theta}$ depended on z, this concludes the proof of (4.2.4).

4.3 Generalized interior estimate

This section is devoted to proving Theorem 4.10, which follows from the technical Lemma 4.14 below.

Lemma 4.14. *Suppose M, N are pinched Hadamard manifolds of dimension at most n with pinching constants $-b^2 \le -a^2 < 0$, and let $x \in M$. Let $f : B_R(x) \to N$ and $h : B_R(x) \to N$ be a smooth and harmonic map, respectively. Suppose that* $\operatorname{dist}(h, f)$ *achieves its maximum at x. Then for any $0 < r < R$, there exist positive constants $A = A(r, R, a, b, n), B = B(r, R, a, b, n)$, such that either*

$$\sup \operatorname{dist}(h, f) \le A\operatorname{diam}(f(B_r(x))) + B, \tag{4.3.1}$$

or

$$\int_{\partial B_r(x)} \min\left(a\rho(y), \log\frac{\pi}{\theta(y)}\right) d\sigma_{x,r}(y) \geq \frac{1}{2}\int_{\partial B_r(x)} a\rho(y) d\sigma_{x,r}(y) - 2,$$

where

$$\rho(y) = \operatorname{dist}(f(x), f(y)) \text{ and } \theta(y) = \measuredangle_{f(x)}(h(x), f(y)).$$

The proof of Lemma 4.14 is a quantitative version of the proof of the "interior estimate" [5, §4]. We first outline the proof of Lemma 4.14 briefly. We divide the outline into three steps.

1. We first observe that $\operatorname{dist}(f(x), h(\cdot))$ is a subharmonic function, so in particular

$$\int_{\partial B_r(x)} \big(\operatorname{dist}(f(x), h(y)) - \operatorname{dist}(f(x), h(x))\big)\, d\sigma_{x,r}(y) \geq 0. \tag{4.3.2}$$

   The entirety of the proof of Lemma 4.14 is estimating the integrand on the left-hand side under the assumption that $\operatorname{dist}(h(x), f(x))$ is very large.

2. If $\operatorname{dist}(h(x), f(x)) =: D$ is large enough, we have

$$\inf_{y\in B_r(x)} \operatorname{dist}(f(x), h(y)) \geq \frac{D}{2}, \tag{4.3.3}$$

$$\sup_{y\in B_r(x)} \measuredangle_{f(x)}(h(x), h(y)) \leq C\exp\left(-\frac{a}{2}D\right). \tag{4.3.4}$$

   Inequality (4.3.3) follows from the fact that $\operatorname{dist}(f(x), h(\cdot))$ is a positive subharmonic function defined on $B_r(x)$ that takes the value D at the center x, and is bounded above by $D+2\operatorname{diam}(f(B_r(x)))$. For D large enough, $D^{-1}\operatorname{diam}(f(B_r(x)))$ is very small, which forces $\inf_{y\in B_r(x)} \operatorname{dist}(f(x), h(y))$ to be comparable to D. Inequality (4.3.4) then follows from (4.3.3) and Cheng's lemma.

3. We then have the chain of inequalities

$$\begin{aligned}\operatorname{dist}(f(x), h(y)) - \operatorname{dist}(f(x), h(x)) &\leq \operatorname{dist}(f(x), h(y)) - \operatorname{dist}(f(y), h(y))\\ &\leq 2a^{-1}\log\frac{1}{\measuredangle_{f(x)}(f(y), h(y))} - \operatorname{dist}(f(x), f(y)) + O(1).\end{aligned}$$

   The inequality in the first line follows from the fact that

$$\operatorname{dist}(f(x), h(x)) = \sup_{B_r(x)} \operatorname{dist}(h, f),$$

   and the inequality in the second line follows from the comparison of the triangle with vertices $f(x), f(y), h(y)$ with the hyperbolic plane. Plugging the final inequality into (4.3.2) along with the bound (4.3.4) yields Lemma 4.14.

We first show Theorem 4.10 assuming Lemma 4.14 below, and then we show Lemma 4.14 in §4.3.1.

Proof of Theorem 4.10. We assume that $\sup_{B_{R_n}(x_n)} \operatorname{dist}(f_n, h_n) \to \infty$, possibly after passing to a subsequence. Fix a large constant $R \geq 1$, that we will choose later, and pass to a subsequence such that $R_n > R$ for all n. Our proof strategy is to apply Lemma 4.14 to $B_R(x_n)$.

Since $\sup_{B_R(x_n)} \operatorname{dist}(f_n, h_n) \to \infty$, we eventually have violation of (4.3.1). Thus for large n, we have

$$\int_{\partial B_R(x_n)} \min\left(a\rho_n(y), \log\frac{\pi}{\theta_n(y)}\right) d\sigma_{x_n,R}(y) \gtrsim R,$$

where

$$\rho_n(y) = \operatorname{dist}(f_n(x_n), f_n(y)) \text{ and } \theta_n(y) = \measuredangle_{f_n(x_n)}(h_n(x_n), f_n(y)).$$

Note that in this proof, we suppress the dependence of implicit constants on the constants of $\mathcal{F}$ coming from Definition 4.9. We observe that since f_n are uniformly Lipschitz, we have $\rho_n(y) \lesssim R$.

Let ω be the size function of $\mathcal{F}$. We now let

$$S_n = \{y \in \partial B_R(x_n) : \theta_n(y) \leq \pi e^{-a\omega(R)} \text{ and } \rho_n(y) \geq \omega(R)\}.$$

Then

$$\begin{aligned} R &\lesssim \int_{\partial B_R(x_n)} \min\left(a\rho_n(y), \log\frac{\pi}{\theta_n(y)}\right) d\sigma_{x_n,R}(y) \\ &\leq \int_{S_n} R d\sigma_{x_n,R} + \int_{\partial B_R(x_n)\setminus S_n} a\omega(R) d\sigma_{x_n,R} \leq R\sigma_{x_n,R}(S_n) + a\omega(R). \end{aligned}$$

By sub-linearity of ω, we have $\sigma_{x_n,R}(S_n) \gtrsim 1$. However, observe that

$$S_n = f_n^{-1}\left(\operatorname{Cone}\left(f_n(x_n)h_n(x_n), e^{-a\omega(R)}\right) \setminus B_{\omega(R)}(f_n(x_n))\right) \cap \partial B_R(x_n).$$

Thus for R large enough, depending on the constants of $\mathcal{F}$, we reach a contradiction with Definition 4.9(2) for $\mathcal{F}$. $\square$

4.3.1 Proof of Lemma 4.14

For clarity of notation, we introduce the notation

$$\rho_f(y) = \operatorname{dist}(f(x), f(y)),$$
$$\rho_h(y) = \operatorname{dist}(f(x), h(y)).$$

The proof will follow from the following two inequalities. The first is

$$\measuredangle_{f(x)}(h(x), h(y)) \leq C(x,r)a\frac{\rho_h(x) + \|\rho_f\|_\infty}{\sinh\left(a\rho_h(x) - C(x,r)a\|\rho_f\|_\infty\right)}, \tag{4.3.5}$$

where $C(x,r) > 0$ is a constant depending only on x, r. The second is

$$\int_{\partial B_r} \min\left(a\rho_f(y), \log\frac{\pi}{\measuredangle_{f(x)}(h(y), f(y))}\right) \geq \frac{1}{2}\int_{\partial B_r}(a\rho_f - 2), \tag{4.3.6}$$

provided $\rho_h(x) \geq (C(x,r)+2)\|\rho_f\|_\infty$. Here we are integrating against $\sigma_{x,r}$, but we drop the $d\sigma_{x,r}(y)$ in formulas for brevity. We first prove Lemma 4.14 assuming inequalities (4.3.5) and (4.3.6), that we show in the next two subsections.

Let $\varepsilon = \frac{\pi}{4}\exp\left(-a\|\rho_f\|_\infty\right)$, and

$$\mathcal{C}_\varepsilon = \{y \in \partial B_r : \measuredangle_{f(x)}(h(x), f(y)) \leq \varepsilon\}.$$

By (4.3.5), if $\rho_h(x) \geq A\|\rho_f\|_\infty + B$ for some suitable constants A and B depending only on M, N, we have

$$\sup_y \measuredangle_{f(x)}(h(x), h(y)) \leq \frac{1}{2}\varepsilon.$$

We observe that for $y \in \partial B_r \setminus \mathcal{C}_\varepsilon$, we have

$$\measuredangle_{f(x)}(h(y), f(y)) \geq \measuredangle_{f(x)}(h(x), f(y)) - \measuredangle_{f(x)}(h(x), h(y)) \geq \frac{1}{2}\measuredangle_{f(x)}(h(x), f(y)).$$

Thus for $y \in \partial B_r \setminus \mathcal{C}_\varepsilon$, we have

$$\log\frac{\pi}{\measuredangle_{f(x)}(h(y), f(y))} \leq 1 + \log\frac{\pi}{\measuredangle_{f(x)}(h(x), f(y))}.$$

Therefore

$$\int_{\partial B_r} \min\left(a\rho_f(y), \log\frac{\pi}{\measuredangle_{f(x)}(h(y), f(y))}\right) \leq A + B, \tag{4.3.7}$$

where

$$A = \int_{\partial B_r \setminus \mathcal{C}_\varepsilon} \min\left(a\rho_f(y), 1 + \log \frac{\pi}{\measuredangle_{f(x)}(h(x), f(y))} \right),$$

$$B = \int_{\mathcal{C}_\varepsilon} a\rho_f.$$

We note that by choice of ε, for $y \in \mathcal{C}_\varepsilon$,

$$1 + \log \frac{\pi}{\measuredangle_{f(x)}(h(x), f(y))} \geq a\|\rho_f\|_\infty \geq a\rho_f(y).$$

By (4.3.7) and (4.3.6), we have

$$\frac{1}{2}\int_{\partial B_r} (a\rho_f - 2) \leq \int_{\partial B_r} \min\left(a\rho_f(y), 1 + \log \frac{\pi}{\measuredangle_{f(x)}(h(x), f(y))} \right).$$

4.3.1.1 Proof of (4.3.5).

The proof of (4.3.5) depends on the following estimate.

Claim 4.15. Let $f : B_r(x) \to \mathbb{R}$ be a subharmonic function. Then there exists $\lambda = \lambda(r, a, b, n) > 0$, such that

$$f(x) \leq \lambda \inf_{B_r(x)} f + (1 - \lambda) \sup_{B_r(x)} f.$$

Proof. If f is constant, there is nothing to prove. Therefore we post-compose f with a linear function so that $\inf_{B_r(x)} f = 0$ and $\sup_{B_r(x)} f = 1$. Let $y \in B_r(x)$ be such that $f(y) = 0$, and let $\rho = \operatorname{dist}(x, y)$. If $\rho = 0$, then $x = y$ and there is nothing to prove. Otherwise, note that

$$f(x) \leq \int_{\partial B_\rho(x)} f(z) d\sigma_{x,\rho}(z). \tag{4.3.8}$$

Note that by Cheng's lemma, $\sup_{B_r(x)}\|Df\| \leq C = C(r)$. For some $\theta = \theta(r)$, the following holds by comparison to the hyperbolic plane: given $a, b \in \partial B_\rho(x)$ for $\rho \leq r$ with $\measuredangle_x(a, b) < \theta$, we have $\operatorname{dist}(a, b) < \frac{1}{2C}$.

From (4.3.8), we get

$$\begin{aligned} f(x) &\leq \frac{1}{2}\sigma_{x,\rho}\left(\partial B_\rho(x) \cap \operatorname{Cone}(xy, \theta)\right) + 1 - \sigma_{x,\rho}\left(\partial B_\rho(x) \cap \operatorname{Cone}(xy, \theta)\right) \\ &= 1 - \frac{1}{2}\sigma_{x,\rho}\left(\partial B_\rho(x) \cap \operatorname{Cone}(xy, \theta)\right). \end{aligned}$$

By [4], there is some absolute $\mu = \mu(r, n)$ such that

$$\sigma_{x,\rho}\left(\partial B_\rho(x) \cap \operatorname{Cone}(xy, \theta)\right) \geq \mu.$$

Therefore $f(x) \leq 1 - \frac{1}{2}\mu$, and the claim is shown with $\lambda = \frac{1}{2}\mu$. ∎

Since $\mathrm{dist}(h(y), f(y)) \le \rho_h(x)$, we have

$$\rho_h(y) \le \rho_f(y) + \mathrm{dist}(h(y), f(y)) \le \rho_f(y) + \rho_h(x),$$

and hence $\|\rho_h\|_\infty \le \rho_h(x) + \|\rho_f\|_\infty$. From Claim 4.15 and the fact that ρ_h is subharmonic, it follows that

$$\rho_h(x) \le \lambda \inf_{B_r} \rho_h + (1-\lambda)\left(\|\rho_f\|_\infty + \rho_h(x)\right).$$

In particular, we have

$$\inf_{B_r} \rho_h \ge \rho_h(x) - \frac{1-\lambda}{\lambda}\|\rho_f\|_\infty \ge \rho_h(x) - C\|\rho_f\|_\infty. \tag{4.3.9}$$

By comparison to the hyperbolic plane, we have

$$a\mathrm{length}(h([x,y])) \ge \sinh(a \inf_{B_r} \rho_h) \measuredangle_{f(x)}(h(x), h(y)).$$

By Cheng's lemma,

$$\mathrm{length}(h([x,y])) \le r\|Dh\|_\infty \le C\|\rho_h\|_\infty \le C\left(\rho_h(x) + \|\rho_f\|_\infty\right).$$

Therefore,

$$\measuredangle_{f(x)}(h(x), h(y)) \le Ca\frac{\rho_h(x) + \|\rho_f\|_\infty}{\sinh\left(a\rho_h(x) - Ca\|\rho_f\|_\infty\right)}.$$

4.3.1.2 Proof of (4.3.6).

We first relate the deficiency (i.e. the slack in the triangle inequality) of the triangle with vertices $h(x), f(y), h(y)$ and the angle $\measuredangle_{f(x)}(h(y), f(y))$, that we denote by $\theta(y)$ for this subsection only, by abuse of notation.

Claim 4.16. Let $D(y) = \rho_f(y) + \rho_h(y) - \mathrm{dist}(f(y), h(y))$. Assuming $\rho_h \ge 2\rho_f$ and $\rho_f \ge a^{-1}$, we have

$$\log\frac{\pi}{\theta(y)} \ge \frac{a}{2}D(y) - 1.$$

Proof. This follows from comparison with the hyperbolic plane. By the hyperbolic law of cosines, we have

$$\begin{aligned}\cosh(a\mathrm{dist}(f(y), h(y))) &\ge \cosh(a\rho_f(y))\cosh(a\rho_h(y)) - \sinh(a\rho_f(y))\sinh(a\rho_h(y))\cos\theta(y)\\ &= \cosh(a(\rho_f(y) - \rho_h(y))) + 2\sin^2\frac{\theta(y)}{2}\sinh(a\rho_f(y))\sinh(a\rho_h(y))\end{aligned}$$

Therefore

$$\sin^2\frac{\theta}{2} \le \frac{\sinh\left(\frac{a}{2}(\operatorname{dist}(f,h)+\rho_f-\rho_h)\right)\sinh\left(\frac{a}{2}(\operatorname{dist}(f,h)+\rho_h-\rho_f)\right)}{\sinh(a\rho_f)\sinh(a\rho_h)}.$$

Since $a\rho_h \ge 2a\rho_f \ge 2$, we have $\min\left(a\rho_f, a\rho_h\right) \ge 1$, so

$$\sin^2\frac{\theta}{2} \le \frac{e^{-a\left(\rho_f+\rho_h-\operatorname{dist}(f,h)\right)}}{(1-e^{-2a\rho_f})(1-e^{-2a\rho_h})} \le e^{-aD(y)}(1-e^{-2})^{-2}.$$

Since $\sin\frac{\theta}{2} \ge \frac{\theta}{\pi}$, we see that

$$\frac{\pi}{\theta} \ge (1-e^{-2})e^{\frac{a}{2}D(y)},$$

so $\log\frac{\pi}{\theta} \ge \frac{a}{2}D(y) - \log\frac{1}{1-e^{-2}} \ge \frac{a}{2}D(y) - 1$. □

By (4.3.9), and the assumption that $\rho_h(x) \ge (C+2)\|\rho_f\|_\infty$, we have $\inf_{B_r}(a\rho_h) \ge 2\|\rho_f\|_\infty$. We let $G = \{y \in \partial B_r : a\rho_f(y) \ge 1\}$. We observe that for $y \in G$, we have by Claim 4.16

$$\log\frac{\pi}{\theta(y)} \ge \frac{a}{2}D(y) - 1.$$

Note that $D(y) = \rho_f(y) + \rho_h(y) - \operatorname{dist}(f(y), h(y)) \le 2\rho_f(y)$ by the triangle inequality, and hence $\frac{a}{2}D(y) - 1 \le a\rho_f(y)$. Therefore for $y \in G$,

$$\min\left(a\rho_f(y), \log\frac{\pi}{\theta(y)}\right) \ge \frac{a}{2}D(y) - 1. \tag{4.3.10}$$

On the other hand, for $y \notin G$, we have

$$\min\left(a\rho_f(y), \log\frac{\pi}{\theta(y)}\right) \ge 0 \ge a\rho_f(y) - 1 \ge \frac{a}{2}D(y) - 1.$$

Hence (4.3.10) holds for all $y \in \partial B_r$. Integrating (4.3.10) over ∂B_r, we get

$$\begin{aligned}\int_{\partial B_r} \min\left(a\rho_f(y), \log\frac{\pi}{\theta(y)}\right) &\ge \frac{1}{2}\int_{\partial B_r}\big(aD(y) - 2\big)\\ &\ge \frac{1}{2}\int_{\partial B_r}\big(a\rho_h(x) + a\rho_f(y) - a\operatorname{dist}(f(y), h(y)) - 2\big)\\ &\ge \frac{1}{2}\int_{\partial B_r}(a\rho_f - 2),\end{aligned}$$

where we used subharmonicity of ρ_h in going from the first to the second line.

4.4 Weakly non-collapsing maps

In this section we prove Theorem 4.7, that in turn immediately implies Theorem 4.5, as explained in Remark 4.8. The proof of Theorem 4.7 has four steps, following the outline in §2.1.3.

1. By combining Lemma 3.8 and Proposition 4.11, it follows immediately that the map f is at a finite distance from a map that is weakly non-collapsing with first two derivatives bounded.

2. We then construct harmonic maps h_n on larger and larger balls B_n that agree with f on ∂B_n.

3. The boundary estimate of Benoist–Hulin [5, Proposition 3.7], stated below as Proposition 4.17, then shows that in any finite distance neighbourhood of the boundary ∂B_n, the distance between f and h_n remains bounded. The generalized interior estimate Theorem 4.10 shows that the distance between f and h_n remains bounded far from the boundary of B_n.

4. The limiting argument Proposition 2.10 then shows the existence of the desired harmonic map.

Proposition 4.17 ([5, Proposition 3.7]). *Let $f : M \to N$ be a smooth map between two pinched Hadamard manifolds with first two derivatives bounded. Let $x_0 \in M, R > 0$, and let $h : B_R(x_0) \to N$ be a harmonic map that agrees with f on $\partial B_R(x_0)$. Then there exists a constant C that depends only on $\|Df\|_\infty, \|D^2 f\|_\infty$, and the pinching constants and dimensions of M, N, such that*

$$\mathrm{dist}(h(x), f(x)) \leq C\mathrm{dist}(x, \partial B_R(x_0))$$

4.4.1 Proof of Theorem 4.7

Let $f : M \to N$ be an ω-weakly non-collapsing map between pinched Hadamard manifolds. By Lemma 3.8 there exists a smooth $\tilde{f} : M \to N$ such that $D\tilde{f}, D^2\tilde{f}$ are bounded and $\sup \mathrm{dist}(f, \tilde{f}) < \infty$. Proposition 4.11 then guarantees that $\tilde{f}$ is a weakly non-collapsing map, possibly with a different size function.

Fix an arbitrary point $x \in M$. Then let $h_n : B_n(x) \to N$ be the harmonic map that agrees with $\tilde{f}$ on $\partial B_n(x)$. If $\sup \mathrm{dist}(\tilde{f}, h_n)$ is a bounded sequence, by Proposition 2.10, we are done. Assume therefore, after passing to a subsequence, that $\sup \mathrm{dist}(\tilde{f}, h_n) \to \infty$.

Let $x_n \in B_n(x)$ be a sequence of points such that the maximum of $\mathrm{dist}(\tilde{f}, h_n)$ is achieved at x_n. By Proposition 4.17, we have

$$R_n = \mathrm{dist}(x_n, \partial B_n(x)) \to \infty.$$

We observe that the family of maps $\{\tilde{f} : (M, x_n) \to (N, f(x_n)) \text{ for } n = 1, 2, ...\}$ is uniformly non-collapsing by definition. Applying Theorem 4.10 to the harmonic maps $h_n : B_{R_n}(x_n) \to N$, we get that

$$\sup_n \sup_{B_n(x)} \mathrm{dist}(\tilde{f}, h_n) = \sup_n \mathrm{dist}(\tilde{f}(x_n), h_n(x_n)) < \infty,$$

which is a contradiction.

4.5 Nearest-point projections to admissible convex sets

This section is devoted to showing Theorem 4.2. We first give a rough outline of the proof. As in the proof of Theorem 4.7, we construct harmonic maps h_n defined on larger and larger balls $B_n(o)$ for some fixed $o \in M$, agreeing with r on the boundaries $\partial B_n(o)$. The goal is to use the limiting argument in Proposition 2.10 to get a harmonic map defined on all of M. It therefore suffices to show that $\sup_M \mathrm{dist}(h_n, r)$ is a bounded sequence. We do this in two steps.

1. We first show that for some fixed $D > 0$, we have

$$\sup_{M \setminus N_D(K)} \mathrm{dist}(h_n, r) \leq \sup_{N_D(K)} \mathrm{dist}(h_n, r) + O(1).$$

   This inequality is derived analogously to the proof of the "boundary estimate" of Benoist–Hulin [5, Proposition 3.7], and it follows from the existence of a bounded subharmonic function Φ such that $\Delta\Phi \gtrsim e^{-a\mathrm{dist}(\cdot,K)}$ on $M \setminus N_D(K)$, for some fixed $D > 0$, that we showed in §3.4.2, and from the classical inequality of Schoen–Yau [60] on the Laplacian of the distance between two functions.

2. It therefore remains to show that $\sup_{N_D(K)} \mathrm{dist}(h_n, r)$ is bounded. This follows from our generalized interior estimate Theorem 4.10, since the map r is non-collapsing near the convex set K. This bound is contained in Proposition 4.18 below.

To state Proposition 4.18, we first note that, given any admissible convex set K and a nearest-point projection map $r : M \to K$, by Corollary 3.14, there exists a smooth map $\tilde{r} : M \to M$ such that $\mathcal{D} := \sup_M \operatorname{dist}(r, \tilde{r}) < \infty$, and

$$\|D\tilde{r}\| \lesssim e^{-a\operatorname{dist}(\cdot,K)},$$
$$\|\tau(\tilde{r})\| \lesssim e^{-a\operatorname{dist}(\cdot,K)}.$$

Proposition 4.18. *Let $D > 0$. There exist constants $R_0 = R_0(D) > 0$ and $C = C(D) > 0$, such that, for any $x \in N_D(K)$ and $R > R_0$ we have the following property. Given a harmonic map $h : B_R(x) \to M$ such that $\operatorname{dist}(h, \tilde{r})$ achieves its maximum at x, we have $\operatorname{dist}(h, \tilde{r}) < C$.*

We first show Theorem 4.2 assuming Proposition 4.18 in §4.5.1. We then show Proposition 4.18 in §4.5.2.

4.5.1 Proof of Theorem 4.2

From Lemma 3.23, for some $D > 0$ large enough, there exist bounded subharmonic functions $\phi_n : M \to \mathbb{R}$ for $n > D - 1$, such that

$$\Delta\phi_n \geq 1 \text{ on } N_{n+1}(K) \setminus N_n(K),$$

and $\sup_n \|\phi_n\|_\infty < \infty$. We now construct the function

$$\Phi = \sum_{n=\lfloor D \rfloor}^{\infty} e^{-an}\phi_n,$$

such that Φ is a bounded subharmonic function, with the property that $\Delta\Phi \gtrsim e^{-a\operatorname{dist}(\cdot,K)}$ on $M \setminus N_D(K)$.

We now fix an arbitrary point $o \in M$, and let $h_N : B_N(o) \to M$ be the harmonic map that agrees with $\tilde{r}$ on $\partial B_N(o)$. By Proposition 2.10, it suffices to show the following claim.

Claim 4.19. The sequence $\sup_{B_N(o)} \operatorname{dist}(h_N, \tilde{r})$ is bounded.

Proof. Assume that, possibly after passing to a subsequence, we have

$$\sup_{B_N(o)} \operatorname{dist}(h_N, \tilde{r}) \to \infty.$$

Note that from [60], we have

$$\Delta\operatorname{dist}(h_N, \tilde{r}) \gtrsim -\|\tau(\tilde{r})\| \gtrsim -e^{-a\operatorname{dist}(\cdot,C)}.$$

Therefore, for a suitably chosen constant c, the function

$$\mathrm{dist}(h_N, \tilde{r}) + c\Phi$$

is subharmonic on $M \setminus N_D(K)$.

Let $x_N \in B_N(o)$ be the point where the maximum of $\mathrm{dist}(h_N, \tilde{r})$ is achieved. If $x_N \in M \setminus N_D(K)$ for infinitely many N, then

$$\mathrm{dist}(h_N(x_N), \tilde{r}(x_N)) + c\Phi(x_N) \leq \sup_{\partial B_N(o)} (\mathrm{dist}(h_N, \tilde{r}) + c\Phi) \lesssim \|\Phi\|_\infty \lesssim 1,$$

which is a contradiction since Φ is bounded. Thus, for infinitely many N, we have $x_N \in N_D(K)$. Proposition 4.17 shows that $\mathrm{dist}(x_N, \partial B_N(o)) \to \infty$ as $N \to \infty$. In particular, for N large enough, we may apply Proposition 4.18, to get that $\sup_N \sup_{B_N(o)} \mathrm{dist}(h_N, \tilde{r}) < \infty$. This is a contradiction. □

4.5.2 Proof of Proposition 4.18

This follows immediately from Theorem 4.10, once we show that the family

$$\{\tilde{r} : (M, x) \to (M, \tilde{r}(x)) \text{ for } x \in N_D(K)\}$$

is uniformly non-collapsing. Note that a proof identical to that of Proposition 4.11 shows the following.

Proposition 4.20. *Let $\mathcal{F}$ be a uniformly non-collapsing family, let $D > 0$, and let $\tilde{\mathcal{F}}$ be a uniformly Lipschitz family of maps between pointed pinched Hadamard manifolds. Assume that for any $\tilde{f} : (M, x) \to (N, y)$ in $\tilde{\mathcal{F}}$, there exists a map $f : M \to N$ in $\mathcal{F}$, such that $\sup_M \mathrm{dist}(f, \tilde{f}) < D$. Then $\tilde{\mathcal{F}}$ is uniformly non-collapsing.*

In particular, it suffices to show that the family

$$\{r : (M, x) \to (M, r(x)) \text{ for } x \in N_D(K)\}$$

is uniformly non-collapsing. The rest of this subsection is devoted to showing this.

We first check Definition 4.9(1). Fix some $x \in N_D(K)$, and set $\rho(y) = \mathrm{dist}(r(x), r(y))$. Let $\hat{x} \in K$ be such that $\mathrm{dist}(x, \hat{x}) \leq 2D$. Then from Definition 4.1, there exist θ, R_0, such that for some $\xi \in \partial_\infty M$ we have $\partial B_R(\hat{x}) \cap \mathrm{Cone}(\hat{x}\xi, \theta) \subseteq \partial B_R(x) \cap K$ for all $R > R_0$. By Proposition 4.12, there exist absolute constants $\hat{\theta}, \hat{D}$ such that

$$\mathrm{Cone}(x\xi, \hat{\theta}) \cap \partial B_R(x) \subset \mathrm{Cone}(\hat{x}\xi, \theta) \setminus B_{R-2D}(\hat{x}) \subseteq K$$

for $R > \max(\hat{D}, R_0 + D)$. Then we have

$$\begin{aligned}\int_{\partial B_R(x)} \rho(y) \geq \int_{\partial B_R(x)\cap K} \rho(y) &= \int_{\partial B_R(x)\cap K} \mathrm{dist}(r(x), y) \\ &\geq \sigma_{x,R}(\mathrm{Cone}(x\xi, \hat{\theta}) \cap \partial B_R(x))(R - \mathrm{dist}(x, r(x))) \\ &\gtrsim R - D \approx R,\end{aligned}$$

where we used the fundamental estimate of Benoist–Hulin [4] that $\sigma_{x,R}(\mathrm{Cone}(x\xi, \hat{\theta}) \cap \partial B_R(x)) \gtrsim 1$.

We now turn to Definition 4.9(2).

Claim 4.21. Let $x, y \in K$ be such that $\mathrm{dist}(x, y) = R$. Then for any $z \in r^{-1}(y)$, we have $\measuredangle_x(y, z) \leq \pi e^{-aR}$.

Proof. Let $\bar{x}\bar{y}\bar{z}$ be a comparison triangle for xyz in the hyperbolic plane of curvature $-a^2$. Then

$$\measuredangle_x(y, z) \leq \measuredangle_{\bar{x}}(\bar{y}, \bar{z}),$$

so it suffices to estimate $\measuredangle_{\bar{x}}(\bar{y}, \bar{z})$. Note that $\measuredangle_{\bar{y}}(\bar{x}, \bar{z}) \geq \measuredangle_y(x, z) \geq \frac{\pi}{2}$, and we can assume without loss of generality that $\measuredangle_{\bar{y}}(\bar{x}, \bar{z}) = \frac{\pi}{2}$. Then the dual hyperbolic law of cosines shows

$$\cos \measuredangle_{\bar{z}}(\bar{x}, \bar{y}) = \sin \measuredangle_{\bar{x}}(\bar{y}, \bar{z}) \cosh(aR) \geq \frac{2\measuredangle_{\bar{x}}(\bar{y}, \bar{z})}{\pi} \frac{e^{aR}}{2},$$

and hence $\measuredangle_{\bar{x}}(\bar{y}, \bar{z}) \leq \pi e^{-aR}$, which concludes the proof. □

We now have for $x \in N_D(K)$ and any $C > 0, \xi \in \partial_\infty M$,

$$r^{-1}\left(\mathrm{Cone}(r(x)\xi, \theta) \setminus B_C(r(x))\right) \subseteq \mathrm{Cone}(r(x)\xi, \theta + \pi e^{-aC}),$$

and hence in particular

$$\begin{aligned}\partial B_R(x) \cap r^{-1}\left(\mathrm{Cone}(r(x)\xi, \theta) \setminus B_{\sqrt{R}}(r(x))\right) &\subseteq \partial B_R(x) \cap \mathrm{Cone}(r(x)\xi, \theta + \pi e^{-a\sqrt{R}}) \\ &\subseteq \partial B_R(x) \cap \mathrm{Cone}(x\xi, \tilde{\theta}(R, \theta)),\end{aligned}$$

where $\tilde{\theta}(R, \theta) \to 0$ as $R \to \infty, \theta \to 0$. Here in going from the first to the second line, we used Proposition 4.12. From the work of Benoist–Hulin [4], we see that $\sigma_{x,R}(\partial B_R(x) \cap \mathrm{Cone}(x\xi, \tilde{\theta})) \to 0$ as $\theta \to 0, R \to \infty$. Therefore Definition 4.9(2) holds, and Proposition 4.18 is shown.

4.6 Admissible convex sets in hyperbolic spaces

In this section we prove Theorem 4.3, that readily follows from the lemma below. For a set $S \subseteq \partial_\infty \mathbb{H}^n$, denote by $\mathrm{CH}(S)$ the closed convex hull of S.

Lemma 4.22. *Let $S \subseteq \mathbb{S}^{n-1}$ be an open set with quasiconformal boundary. Then there exists an angle $\theta > 0$, such that for any $x \in \mathrm{CH}(S)$, there exists $\xi \in S$ such that $\mathrm{Cone}(x\xi, \theta) \cap \mathbb{S}^{n-1} \subseteq S$.*

We prove Lemma 4.22 by contradiction. Assuming there is a sequence $x_i \in \mathrm{CH}(S)$ that provides a contradiction to the claim in Lemma 4.22, it is easy to see that it has to converge to some point in the boundary at infinity $\xi \in S \subseteq \partial_\infty \mathbb{H}^n = \mathbb{S}^{n-1}$. Then by assumption we can map S in a neighbourhood of ξ to some standard model using a quasiconformal map. By a classical result of Tukia–Väisälä [71], this quasiconformal map can be extended to a bi-Lipschitz self-map F of $\mathbb{H}^n$. We then prove the claim of Lemma 4.22 for this standard model, and transport it back to S using the map F. All of this is done in §4.6.1.

Theorem 4.3 then follows by simple hyperbolic geometry, that we explain in §4.6.2.

4.6.1 Boundary analysis: Proof of Lemma 4.22

Suppose that the conclusion of Lemma 4.22 fails. Then there exists a sequence $x_i \in \mathbb{H}^n$ such that $x_i \in \mathrm{CH}(S)$, and such that

$$\sup\{\theta : \mathrm{Cone}(x_i\xi, \theta) \cap \mathbb{S}^{n-1} \subseteq S \text{ for some } \xi \in S\} \to 0 \tag{4.6.1}$$

as $i \to \infty$. Note that if x_i remain in some compact set, after passing to a subsequence, we may assume that $x_i \to x_\infty \in \mathbb{H}^n$. Since S is an open set, this is a contradiction with (4.6.1). Assume therefore that $x_i \to s \in \partial_\infty \mathbb{H}^n$, possibly after passing to a subsequence. Since $x_i \in \mathrm{CH}(S)$, we have $s \in \bar{S}$.

Before continuing with the proof, we define a different version of the cone that will be more convenient for us to work with. We define for $x \in \mathbb{H}^n, \xi \in \mathbb{S}^{n-1}$, and $D > 0$, the set

$$C^D_{x\xi} = \{\eta \in \mathbb{S}^{n-1} : \mathrm{dist}(x, [\xi, \eta]) \geq D\}.$$

It is classical that for some absolute constant $C > 1$, we have

$$\mathrm{Cone}(x\xi, C^{-1}e^{-D}) \cap \partial_\infty \mathbb{H}^n \subseteq C^D_{x\xi} \subseteq \mathrm{Cone}(x\xi, Ce^{-D}) \cap \partial_\infty \mathbb{H}^n.$$

so it suffices to show that there exists $D > 0$ and $\xi_i \in S$ such that $C^D_{x_i\xi_i} \subseteq S$ for all i large enough. The rest of the proof is devoted to showing this.

If s lies in the interior of S, we may set $\xi_i = s$ and pick an arbitrary $D > 0$. Therefore assume $s \in \partial S$. Let U be an open set containing s, and $f : U \to V \subseteq \mathbb{R}^{n-1}$ be a quasiconformal homeomorphism, such that $f(s) = 0$ and

$$f(S \cap U) = V \cap (\mathbb{R}_+ \times \mathbb{R}^{n-2}).$$

By [71, Theorem 3.2], we can extend f to a map

$$F : \mathcal{U} \to \mathcal{V},$$

where $\mathcal{U}$ and $\mathcal{V}$ are neighbourhoods of U, V in $\mathbb{H}^n$, respectively, such that F is L-bi-Lipschitz for the hyperbolic metric, meaning

$$L^{-1}\mathrm{dist}(a, b) \leq \mathrm{dist}(F(a), F(b)) \leq L\mathrm{dist}(a, b).$$

We let $y_i = F(x_i)$. Note that since F is bi-Lipschitz, by the Morse lemma we have

$$\sup_i \mathrm{dist}(y_i, \mathrm{CH}(f(S \cap U))) < \infty.$$

Claim 4.23. There exists a sequence $\eta_i \in V \cap (\mathbb{R}_+ \times \mathbb{R}^{n-2})$ and $D < \infty$, such that $C^D_{y_i\eta_i} \subseteq V \cap (\mathbb{R}_+ \times \mathbb{R}^{n-2})$.

We first show how to complete the proof assuming Claim 4.23. Suppose therefore that η_i, D are as in Claim 4.23. We set $\xi_i = F^{-1}(\eta_i)$. Since F is L-Lipschitz,

$$f\left(C^{L(D+M)}_{x_i\xi_i}\right) \subseteq C^D_{y_i\eta_i}.$$

Here $M = M(L)$ is a constant with the property that

$$\mathrm{dist}(F([a, b]), [F(a), F(b)]) \leq M.$$

The existence of such a constant is the well-known Morse lemma. From the conclusion of Claim 4.23, we now have

$$C^{L(D+M)}_{x_i\xi_i} \subseteq f^{-1}\left(V \cap (\mathbb{R}_+ \times \mathbb{R}^{n-2})\right) \subseteq S,$$

as desired.

Proof of Claim 4.23. Suppose the conclusion of the claim does not hold, and pass to a subsequence such that $\sup\{D : C^D_{y_i\eta_i} \subseteq V\} \to \infty$ as $i \to \infty$.

Let A_i be an isometry of $\mathbb{H}^n$ such that $A_i(y_i) = y_0$, and $A(0) = 0$, where $0 \in \mathbb{R}^{n-1} = \partial\mathbb{H}^n$. Then $A_i(\mathbb{R}_+ \times \mathbb{R}^{n-2})$ is the boundary at infinity of a halfspace in a totally geodesic copy G_i of $\mathbb{H}^{n-1}$ in $\mathbb{H}^n$. But $\sup_i \operatorname{dist}(y_0, G_i) < \infty$ and $0 \in G_i \cap \partial_\infty\mathbb{H}^n$, so we can pass to a subsequence so that $G_i \to G$, and thus $A_i(\mathbb{R}_+ \times \mathbb{R}^{n-2}) \to G \cap \partial_\infty\mathbb{H}^n$, where G is some halfspace in a totally geodesic copy of $\mathbb{H}^{n-1}$ lying in $\mathbb{H}^n$. Then there exist some $\eta \in G \cap \partial_\infty\mathbb{H}^n, D < \infty$ such that

$$C^D_{y_0\eta} \cap \partial G \cap \partial_\infty\mathbb{H}^n = \emptyset.$$

Taking $\eta_i = A_i^{-1}(\eta)$, we see that for large enough i,

$$C^{2D}_{y_i\eta_i} \subseteq V \cap (\mathbb{R}_+ \times \mathbb{R}^{n-2}),$$

which is a contradiction. □

4.6.2 Proof of Theorem 4.3

Let $K = \mathrm{CH}(U)$ of an open set U with quasiconformal boundary. For any $x \in N_D(K)$, there exists by Lemma 4.22 an angle $\theta = \theta(D) > 0$ and $\xi \in \partial_\infty\mathbb{H}^n$ such that

$$\mathrm{Cone}(x\xi, \theta) \cap \partial_\infty\mathbb{H}^n \subseteq U.$$

We claim that for all $R > R_0 = R_0(D)$,

$$\mathrm{Cone}\left(x\xi, \frac{\theta}{12}\right) \cap \partial B_R(x) \subseteq \mathrm{CH}\left(\mathrm{Cone}(x\xi, \theta) \cap \partial_\infty M\right). \tag{4.6.2}$$

Note that (4.6.2) immediately shows admissibility of $\mathrm{CH}(U)$, so the rest of this subsection is devoted to showing (4.6.2).

Let $y \in \mathrm{Cone}\left(x\xi, \frac{\theta}{12}\right) \cap \partial B_R(x)$ be arbitrary. Pick any point $\eta_1 \in \partial_\infty\mathbb{H}^n$ such that $\frac{\theta}{6} < \measuredangle_x(\eta_1, \xi) < \frac{\theta}{3}$, and let $\eta_2 \in \partial_\infty\mathbb{H}^n$ be such that $y \in [\eta_1, \eta_2]$. Then in particular we have

$$\frac{\theta}{12} < \measuredangle_x(\eta_1, y) < \frac{5}{12}\theta.$$

Claim 4.24 below shows that, for R large enough depending on θ, we have

$$\measuredangle_x(y, \eta_2) < \frac{\theta}{12}.$$

Thus $\measuredangle_x(\eta_1, \eta_2) < \frac{\theta}{2}$, and hence $\eta_2 \in \mathrm{Cone}(x\xi, \theta) \cap \partial_\infty\mathbb{H}^n$. Then the set $\mathrm{CH}\left(\mathrm{Cone}(x\xi, \theta) \cap \partial_\infty\mathbb{H}^n\right)$ contains the entire geodesic $[\eta_1, \eta_2]$, and hence also contains y.

Claim 4.24. Let $x, y \in \mathbb{H}^n$ and $\xi, \eta \in \partial_\infty \mathbb{H}^n$ be such that $y \in [\xi, \eta]$. If $\measuredangle_x(\xi, y) = \alpha$ and $\operatorname{dist}(x, y) = R$, we have

$$\measuredangle_x(y, \eta) \lesssim e^{-2R},$$

where the implicit constant depends on α.

Proof. By the dual hyperbolic law of cosines applied to $xy\xi$ and to $xy\eta$, we see that

$$1 = -\cos\alpha \cos\measuredangle_y(\xi, x) + \sin\alpha \sin\measuredangle_y(\xi, x)\cosh(R), \tag{4.6.3}$$

$$1 = \cos\beta \cos\measuredangle_y(\xi, x) + \sin\beta \sin\measuredangle_y(\xi, x)\cosh(R). \tag{4.6.4}$$

It follows from (4.6.3) that for large R, we have $\measuredangle_y(x, \xi) \lesssim e^{-R}$. Straightforward analysis of (4.6.4) then implies $\beta \lesssim \measuredangle_y(\xi, x)^2 \lesssim e^{-2R}$. □

Chapter 5

Putman–Wieland conjecture

In this chapter, we show our results Theorems B.1 and B.2 on the Putman–Wieland conjecture. Both of these results rely essentially on the following theorem.

Theorem 5.1. *Let μ be a Beltrami form on a marked Riemann surface $S \in \mathcal{T}_g$. Let ϕ be the holomorphic 1-form on S such that $[\mathrm{Re}(\phi)] = \chi$. Then the Laplacian of the Hodge norm of $\chi \in H^1(\Sigma_g, \mathbb{R})$ in the complex direction defined by μ is*

$$\Delta_\mu \mathrm{E}_\chi = 4i \int_S \theta \wedge \bar{\theta},$$

where θ is a holomorphic 1-form such that $\mu\phi - \bar{\theta}$ is $\bar{\partial}$-exact.

For the definitions of Hodge norm and plurisubharmonicity, the reader should go back to §2.2.1.2 and §2.2.3, respectively. Note that θ in Theorem 5.1 always exists by Hodge theory.

Theorem 5.1 has implications for Teichmüller dynamics, and can be obtained from the computations of Forni [24], a good exposition of which can also be found in [25]. When χ is an integral cohomology class, Theorem 5.1 can also be shown by analyzing harmonic maps $\omega : S \to S^1$ that induce $\chi : \pi_1(S) \to \mathbb{Z}$, as was done by Marković and the author in [50].

In Chapter 6, we will derive a version of Theorem 5.1 for non-abelian cohomology (for the most general version, see Theorem 6.9 and Theorem 6.8). We include a direct proof of Theorem 5.1 in §5.1 for completeness, as the abelian case has a significantly simpler proof, some ideas of which will also appear in Chapter 6. We now outline how to derive Theorems B.1 and B.2 from Theorem 5.1.

Recall that Theorem B.1 states that a covering map $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ and a cohomology class $\chi \in H^1(\Sigma_h, \mathbb{R})$ give a counterexample to the Putman–Wieland conjecture if and only if the energy E_χ is constant on the subspace $\sigma_\pi(\mathcal{T}_{g,n})$ of $\mathcal{T}_h$.

Both of the directions of this result rely on the fact, shown in §5.2, that $\sigma_\pi : \mathcal{T}_{g,n} \to \mathcal{T}_h$ is λ-equivariant, for some homomorphism $\lambda : \Theta_\pi \to \mathrm{Mod}_h$, where $\Theta_\pi \leq \mathrm{Mod}_{g,n}$ is a finite index subgroup. We now outline the two directions of Theorem B.1 separately.

"If" Suppose that the energy E_χ is constant on $\sigma_\pi(\mathcal{T}_{g,n})$. Then E_χ is in particular constant on the $(\Gamma_\pi \cap \Theta_\pi)$-orbit of any Riemann surface $X \in \mathcal{T}_{g,n}$. However, by equivariance of σ_π, it follows that χ has a bounded Hodge norm on the $\lambda(\Gamma_\pi)$ orbit of $\sigma_\pi(X)$. In particular, the Γ_π orbit of χ consists of forms of bounded Hodge norm on $\sigma_\pi(X)$. Since there are only finitely many integral cohomology classes with Hodge norm at most C, for any $C > 0$, it follows that $\Gamma_\pi \cdot \chi$ is finite.

"Only if" If $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ is a finite cover, and if $\chi \in H^1(\Sigma_h, \mathbb{Z})$ has finite Γ_π-orbit, then there exists a finite index subgroup $\Gamma \leq \Gamma_\pi$ such that χ is Γ-invariant. It follows that the Hodge norm of χ descends to a map $\mathrm{E}_\chi \circ \sigma_\pi : \mathcal{M}_{g,n}^\Gamma \to \mathbb{R}$. Since plurisubharmonicity is a local condition, this map is also plurisubharmonic. We then show that $\mathrm{E}_\chi \circ \sigma_\pi$ is bounded, so since the complex manifold $\mathcal{M}_{g,n}^\Gamma$ is quasiprojective, the energy $\mathrm{E}_\chi \circ \sigma_\pi$ is constant by Lemma 2.19.

We give the details of both of these directions in §5.3.

From our computation in §5.1, we get an explicit criterion for when E_χ is not strictly plurisubharmonic in some complex direction defined by a Beltrami form μ.

Corollary 5.2. *The map* E_χ *is plurisubharmonic for any cohomology class* $\chi \in H^1(\Sigma_h, \mathbb{R})$. *Moreover,* E_χ *is not strictly plurisubharmonic in the direction defined by a Beltrami form* μ *on a Riemann surface* S *if and only if* $\mu\phi$ *is* $\bar{\partial}$*-exact, where* ϕ *is the holomorphic 1-form on* S *with* $[\mathrm{Re}(\phi)] = \chi$.

We derive Theorem B.2 in §5.4 from analyzing this criterion. Our main tools here are the Riemann–Roch theorem and the Li–Yau result comparing gonality and the bottom of the spectrum of the Laplacian [45].

5.1 Plurisubharmonicity of the Hodge norm

This section is devoted to proving Theorem 5.1. We first set up some notation in §5.1.1. We then compute the first two derivatives of ϕ in this setup in §5.1.2, then the second derivative of the Hodge norm in §5.1.3, and finally the Laplacian of the Hodge norm in §5.1.4.

5.1.1 Notation and setup

Throughout this section, we fix a Beltrami form μ on a Riemann surface S and denote by ϕ the holomorphic 1-form on S such that $[\mathrm{Re}(\phi)] = \chi$. We denote by $S^{t\mu}$ the Riemann surface obtained by solving the Beltrami equation for $f^{t\mu} : S \to S^{t\mu}$,

$$\bar{\partial} f^{t\mu} = t\mu \partial f^{t\mu}.$$

We let ϕ_t be the holomorphic 1-form on $S^{t\mu}$ such that $[\mathrm{Re}(\phi_t)] = \chi$. We abuse notation to denote $(f^{t\mu})^*\phi_t$ by ϕ_t and think of it as a form on S. With these conventions,

$$\mathrm{E}_\chi(S^{t\mu}) = \frac{i}{2}\int_S \phi_t \wedge \bar{\phi}_t,$$

and hence

$$\left(\frac{d^2}{dt^2}\mathrm{E}_\chi(S^{t\mu})\right)_{t=0} = \frac{i}{2}\int_S \ddot{\phi} \wedge \bar{\phi} + 2\dot{\phi} \wedge \dot{\bar{\phi}} + \phi \wedge \ddot{\bar{\phi}}, \tag{5.1.1}$$

where $\phi_t = \phi + t\dot{\phi} + \frac{t^2}{2}\ddot{\phi} + O(t^3)$.

5.1.2 Computing the derivatives of ϕ

The following operator will play a key role in the rest of this section. For a Riemann surface S, we denote by $H^{1,0}(S)$ the (complex) vector space of holomorphic 1-forms, and by $H^{0,1}(S)$ the (complex) vector space of antiholomorphic 1-forms.

Definition 5.3. We define the second fundamental form $A_\mu : H^{1,0}(S) \to H^{0,1}(S)$ by requiring

$$\int_S \theta \wedge A_\mu(\phi) = \int_S \theta \wedge \mu\phi, \tag{5.1.2}$$

for any $\theta, \phi \in H^{1,0}(S)$. In other words, $A_\mu(\phi)$ is the antiholomorphic part of the Hodge decomposition of the $(0,1)$-form $\mu\phi$.

Remark 5.4. The tensor $A \in T\mathcal{M}_g \otimes \mathrm{Hom}(H^{1,0}(S), H^{0,1}(S))$ is the second fundamental form of the bundle $H^{1,0}$ over $\mathcal{M}_g$ with fiber $H^{1,0}(S)$ over S, as a subbundle of the flat bundle $H^1(\Sigma_g, \mathbb{C}) \times \mathcal{M}_g$. This is elaborated in [25, §2.3, §2.5].

Claim 5.5. We have

$$\dot{\phi} = A_\mu(\phi) - \overline{A_\mu(\phi)} + dg, \tag{5.1.3}$$

for the smooth function $g : S \to \mathbb{C}$ such that $\mu\phi = A_\mu(\phi) + \bar{\partial} g$. We also have

$$\ddot{\phi}^{1,0} + \overline{\ddot{\phi}^{0,1}} \text{ is } \partial\text{-exact,} \tag{5.1.4}$$

and

$$\ddot{\phi}^{0,1} = 2\mu\dot{\phi}^{1,0}. \tag{5.1.5}$$

Proof. We first give two preliminary observations.

- Since ϕ_t is closed with an exact real part for all t, differentiating we obtain that $\dot{\phi}$ and $\ddot{\phi}$ are closed with exact real parts.
- The form ϕ_t has type $(1,0)$ on $S^{t\mu}$, so is proportional to

  $$df^{t\mu} = (1 + t\mu)\partial f^{t\mu},$$

  and hence $\phi_t^{0,1} = t\mu\phi_t^{1,0}$, where $(\cdot)^{0,1}$ and $(\cdot)^{1,0}$ refer to the complex structure on S. Differentiating once and twice, we obtain

  $$\dot{\phi}^{0,1} = \mu\phi^{1,0} = \mu\phi,$$
  $$\ddot{\phi}^{0,1} = 2\mu\dot{\phi}^{1,0}.$$

  Note that this already shows (5.1.5).

We now let $\dot{\phi} = \alpha + \bar{\beta} + dh$ for some smooth function $h : S \to \mathbb{C}$ and $\alpha, \beta \in H^{1,0}(S)$ be the Hodge decomposition of $\dot{\phi}$. Since $\dot{\phi}$ has an exact real part, we see that $\mathrm{Re}(\alpha + \bar{\beta}) = \mathrm{Re}(\alpha + \beta)$ is exact. It follows from standard Hodge theory that $\alpha = -\beta$, so that $\dot{\phi} = \bar{\beta} - \beta + dh$. Therefore

$$A_\mu(\phi) + \bar{\partial} g = \mu\phi = \dot{\phi}^{0,1} = \bar{\beta} + \bar{\partial} h.$$

Therefore $\bar{\beta} = A_\mu(\phi)$ and $h - g$ is constant. Hence

$$\dot{\phi} = A_\mu(\phi) - \overline{A_\mu(\phi)} + dg,$$

which is (5.1.3).

The 1-form

$$\ddot{\phi}^{1,0} + \overline{\ddot{\phi}^{0,1}} = \left(\ddot{\phi} + \overline{\ddot{\phi}}\right)^{1,0} = \left(2\mathrm{Re}(\ddot{\phi})\right)^{1,0}$$

is ∂-exact since $2\mathrm{Re}(\ddot{\phi})$ is exact, which is (5.1.4). □

5.1.3 Computing the second derivative of the Hodge norm

Lemma 5.6. *Let $\mu\phi = A_\mu(\phi) + \bar{\partial} g$ be the Hodge decomposition of $\mu\phi$, where $g : S \to \mathbb{C}$. Then*

$$\left(\frac{d^2}{dt^2}\mathrm{E}_\chi(S^{t\mu})\right)_{t=0} = 2i\int_S \overline{A_\mu(\phi)} \wedge A_\mu(\phi) + 2\mathrm{Im}\int_S \partial g \wedge \bar{\partial} g.$$

Proof. From (5.1.1), (5.1.4), (5.1.5) and Stokes' theorem we have

$$\begin{aligned}\left(\frac{d^2}{dt^2}\mathrm{E}_\chi(S^{t\mu})\right)_{t=0} &= \frac{i}{2}\int_S \ddot{\phi}^{1,0} \wedge \bar{\phi} + 2\dot{\phi} \wedge \bar{\dot{\phi}} + \phi \wedge \overline{\ddot{\phi}^{1,0}} \\ &= -\frac{i}{2}\int_S \overline{\ddot{\phi}^{0,1}} \wedge \bar{\phi} - 2\dot{\phi} \wedge \bar{\dot{\phi}} + \phi \wedge \ddot{\phi}^{0,1} \\ &= i\int_S \dot{\phi} \wedge \bar{\dot{\phi}} - i\int_S \phi \wedge \mu\dot{\phi}^{1,0} - \overline{\phi \wedge \mu\dot{\phi}^{1,0}} \\ &= i\int_S \dot{\phi} \wedge \bar{\dot{\phi}} + 2\mathrm{Im}\int_S \phi \wedge \mu\dot{\phi}^{1,0}.\end{aligned}$$

We compute the two terms separately.

First we note that by (5.1.3) and Stokes' theorem, we have

$$\int_S \dot{\phi} \wedge \bar{\dot{\phi}} = -\int_S (A_\mu(\phi) - \overline{A_\mu(\phi)}) \wedge (A_\mu(\phi) - \overline{A_\mu(\phi)}) = 0.$$

We also have from (5.1.3) that

$$\dot{\phi}^{1,0} = \partial g - \overline{A_\mu(\phi)}.$$

Hence by Stokes' theorem,

$$\begin{aligned}\mathrm{Im}\int_S \phi \wedge \mu\dot{\phi}^{1,0} = \mathrm{Im}\int_S \dot{\phi}^{1,0} \wedge \mu\phi &= -\mathrm{Im}\int_S \overline{A_\mu(\phi)} \wedge \mu\phi + \mathrm{Im}\int_S \partial g \wedge \mu\phi \\ &= -\mathrm{Im}\int_S \overline{A_\mu(\phi)} \wedge A_\mu(\phi) + \mathrm{Im}\int_S \partial g \wedge \bar{\partial} g \\ &= i\int_S \overline{A_\mu(\phi)} \wedge A_\mu(\phi) + \mathrm{Im}\int_S \partial g \wedge \bar{\partial} g,\end{aligned}$$

from which the result follows. □

5.1.4 Computing the Laplacian of the Hodge norm

We have

$$\Delta_\mu \mathrm{E}_\chi = \left(\frac{d^2}{dt^2}\mathrm{E}_\chi(S^{t\mu})\right)_{t=0} + \left(\frac{d^2}{dt^2}\mathrm{E}_\chi(S^{it\mu})\right)_{t=0}.$$

Suppose now $\mu\phi = A_\mu(\phi) + \bar\partial g$. Then $i\mu\phi = iA_\mu(\phi) + \bar\partial(ig)$, which shows in particular that $A_{i\mu}(\phi) = iA_\mu(\phi)$. Then by Lemma 5.6, we have

$$\Delta_\mu \mathrm{E}_\chi = 4i \int_S \overline{A_\mu(\phi)} \wedge A_\mu(\phi) + 2\mathrm{Im} \int_S \partial g \wedge \bar\partial g + \partial(ig) \wedge \bar\partial(ig) = 4i \int_S \overline{A_\mu(\phi)} \wedge A_\mu(\phi).$$

This concludes the proof of Theorem 5.1.

5.2 Equivariance of σ_π

The goal of this brief section is to show the following lemma. Recall that $\Gamma_\pi \leq \mathrm{Mod}_{g,n+1}$ is the subgroup consisting of mapping classes that lift to $\Sigma_{h,m}$, using one of the punctures on $\Sigma_{g,n+1}$ as a basepoint, and that $\Lambda_\pi : \Gamma_\pi \to \mathrm{Mod}_h$ is this lifting homomorphism.

Lemma 5.7. *For any finite cover $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$, there exist finite index subgroups $\hat\Theta_\pi \leq \Gamma_\pi$ and $\Theta_\pi \leq \mathrm{Mod}_{g,n}$, such that the map Λ_π factors as*

$$\hat\Theta_\pi \xrightarrow{\mathcal{F}_{g,n}} \Theta_\pi \xrightarrow{\lambda} \mathrm{Mod}_h,$$

when restricted to $\hat\Theta_\pi$. Moreover, the map $\sigma_\pi : \mathcal{T}_{g,n} \to \mathcal{T}_h$ is λ-equivariant.

Proof. By a slight abuse of notation, we denote the forgetful homomorphisms by $\mathcal{F}_g : \mathrm{Mod}_{g,k} \to \mathrm{Mod}_g$. Recall that the forgetful homomorphism $\mathrm{Mod}_{g,n+1} \to \mathrm{Mod}_{g,n}$ is denoted $\mathcal{F}_{g,n}$.

Claim 5.8. There exists a finite-index subgroup $\hat\Theta \leq \Gamma_\pi$ such that

$$\hat\Theta \cap \pi_1(\Sigma_{g,n}) \leq \ker \Lambda_\pi.$$

Proof. Consider the diagram

$$\begin{array}{ccc} & & \pi_1(\Sigma_{g,n}) \\ & & \big\downarrow {\scriptstyle \iota_{g,n}} \\ \Gamma_\pi & \leq & \mathrm{Mod}_{g,n+1} \\ \big\downarrow {\scriptstyle \Lambda_\pi} & & \big\downarrow {\scriptstyle \mathcal{F}_{g,n}} \\ \mathrm{Mod}_h & & \mathrm{Mod}_{g,n} \end{array} \tag{5.2.1}$$

where the second column is a part of the Birman exact sequence. Let A be a diffeomorphism of $\Sigma_{g,n}$ representing some mapping class in $\Gamma_\pi \cap \pi_1(\Sigma_{g,n})$. Then by definition, A has a lift to $\Sigma_{h,m}$. However since A is homotopic to the identity on $\Sigma_{g,n}$ (by exactness of the Birman sequence at $\mathrm{Mod}_{g,n+1}$), by the homotopy lifting property, the lift of A to $\Sigma_{h,m}$ is homotopic to a homeomorphism covering the identity, i.e. an element of the deck group. Since π is a finite cover, the deck group is finite. In particular, the image of the deck group in Mod_h is finite as well. By a result of Grossman [28], the group Mod_h is residually finite. Therefore, there is some finite index subgroup $\Theta \le \mathrm{Mod}_h$ that has trivial intersection with the deck group. Taking the pre-image of Θ under the map $\lambda : \Gamma \to \mathrm{Mod}_h$ from the diagram (5.2.1), we get a group $\hat{\Theta}$ with the desired property. $\square$

Now let $\Theta = \mathcal{F}_{g,n}(\hat{\Theta})$. Therefore we get a diagram

$$\begin{array}{ccccccccc} 1 & \longrightarrow & \pi_1(\Sigma_{g,n}) \cap \hat{\Theta} & \xrightarrow{\iota_{g,n}} & \hat{\Theta} & \longrightarrow & \Theta & \longrightarrow & 1 \\ & & & & \downarrow{\scriptstyle \Lambda_\pi} & & & & \\ & & & & \mathrm{Mod}_h & & & & \end{array} \tag{5.2.2}$$

Recall that $\iota_{g,n}$ embeds $\pi_1(\Sigma_{g,n})$ as the point-pushing subgroup of $\mathrm{Mod}_{g,n+1}$. It follows that, since $\Lambda_\pi \circ \iota_{g,n}|_{\hat{\Theta}\cap\pi_1(\Sigma_{g,n})} = 1$, that we get an induced homomorphism

$$\lambda : \Theta \to \mathrm{Mod}_h,$$

and we get the desired factorization of Λ_π.

We now show equivariance of σ_π. Let $\Phi : \mathcal{T}_{g,n+1} \to \mathcal{T}_{g,n}$ be the map that forgets a puncture, so that Φ is $\mathcal{F}_{g,n}$-equivariant. Note that the map

$$\sigma_\pi \circ \Phi : \mathcal{T}_{g,n+1} \to \mathcal{T}_h$$

is Λ_π-equivariant. Since Φ and $\mathcal{F}_{g,n} : \hat{\Theta}_\pi \to \Theta_\pi$ are surjective, it follows that σ_π is λ-equivariant. $\square$

5.3 Proof of Theorem B.1

We first recall the statement. Let $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ be a covering map, and let $\chi \in H^1(\Sigma_{h,m})$. We want to show that χ has a finite Γ_π orbit if and only if E_χ is constant along the image of σ_π. We use the notation $\hat{\Theta}_\pi, \Theta_\pi, \lambda$ from Lemma 5.7 in the previous section §5.2.

We first show the "if" direction. Suppose that $\chi \in H^1(\Sigma_h, \mathbb{R})$ is such that E_χ is constant along the image of σ_π. Fix any marked Riemann surface $T \in \mathcal{T}_{g,n}$, and let

$S = \sigma_\pi(T)$. Then for any element $\gamma \in \Theta_\pi$ of the mapping class group $\mathrm{Mod}_{g,n}$ that lifts to Σ_h via λ, we see that

$$\mathrm{E}_{\lambda(\gamma)^*\chi}(S) = \mathrm{E}_\chi(\lambda(\gamma) \cdot S) = \mathrm{E}_\chi(S),$$

where we used the fact that the energy is constant. Thus $\hat{\Theta}_\pi \chi = \{\lambda(\gamma)^*\chi : \gamma \in \Theta_\pi\} \subset H^1(S, \mathbb{Z})$ is a set of integral first cohomology classes with equal energy. Hence $\hat{\Theta}_\pi \chi$ is finite, and since $[\Gamma_\pi : \hat{\Theta}_\pi] < \infty$, the orbit $\Gamma_\pi \chi$ is finite, as desired.

The rest of this section is devoted to the "only if" direction. Suppose that $\chi \in H^1(\Sigma_h, \mathbb{Z})$ is a cohomology class with a finite orbit under Γ_π. We first observe that by Corollary 5.2 and the fact that the map $\sigma_\pi : \mathcal{T}_{g,n} \to \mathcal{T}_h$ is holomorphic, the map

$$\mathrm{E}_\chi \circ \sigma_\pi : \mathcal{T}_{g,n} \to \mathbb{R}$$

is plurisubharmonic. Let $\hat{\Gamma} \leq \Gamma_\pi$ be the stabilizer of χ, which is finite index by assumption. Therefore $\Gamma = \mathcal{F}_{g,n}(\hat{\Theta}_\pi \cap \hat{\Gamma})$ has finite index in Θ_π. The homomorphism $\lambda : \Theta_\pi \to \mathrm{Mod}_h$ induces an action of Γ on $H^1(\Sigma_h, \mathbb{R})$, that leaves χ invariant. In particular, the energy descends to a map

$$\mathrm{E}_\chi \circ \sigma_\pi : \mathcal{M}_{g,n}^{\Gamma} \to \mathbb{R}. \tag{5.3.1}$$

We first briefly conclude the proof of Theorem B.1 assuming that $\mathrm{E}_\chi \circ \sigma_\pi$ is bounded. Note that by [7, Corollary 2.10], there exists a finite index subgroup $\Gamma' \leq \Gamma$ such that $\mathcal{M}_{g,n}^{\Gamma'}$ has a smooth compactification. Thus by Lemma 2.19, the lift of the map (5.3.1) to $\mathcal{M}_{g,n}^{\Gamma'}$ is constant, and hence so is $\mathrm{E}_\chi \circ \sigma_\pi$. In the rest of this section, we show that $\mathrm{E}_\chi \circ \sigma_\pi$ is bounded.

Proof that $\mathrm{E}_\chi \circ \sigma_\pi$ is bounded rests on the following purely topological claim due to Boggi–Putman–Salter [8, Theorem A].

Claim 5.9. Let $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ be a cover, and let $\chi \in H^1(\Sigma_h, \mathbb{R})$ be a cohomology class with finite Γ_π orbit. Then for any lift γ of a simple closed curve on Σ_g, we have $\chi(\gamma) = 0$.

Remark 5.10. Note that in [50], Marković and the author provide a separate more complicated proof of Claim 5.9 using complex analytic techniques and the result of Bridson [11]. The proof of Boggi–Putman–Salter is a much more straightforward topological proof, and we refer the interested reader to [8] for the details.

Suppose that $X_1, X_2, ... \in \mathcal{M}_{g,n}^{\Gamma}$ is a sequence such that $\mathrm{E}_\chi \circ \sigma_\pi(X_i) \to \infty$. Then since $\mathrm{E}_\chi \circ \sigma_\pi$ is continuous, the sequence $(X_i : i \geq 1)$ must leave every compact subset of $\mathcal{M}_{g,n}^{\Gamma}$, and hence also in $\mathcal{M}_{g,n}$.

We suppose now that $X_i \to Y \in \overline{\mathcal{M}_{g,n}}$, so that Y is a noded Riemann surface. Let $Y_1, Y_2, ..., Y_k$ be the irreducible components of Y, thought of as punctured surfaces (where pairs of punctures correspond to a node of Y). We choose lifts of X_i to $\mathcal{T}_g$ (that we also denote X_i) such that $X_i \to Y$ in $\overline{\mathcal{T}_{g,n}}$. Then $\sigma_\pi(X_i) \to Z \in \overline{\mathcal{T}_h}$, where Z is a marked noded Riemann surface, with irreducible components $Z_1, Z_2, ..., Z_l$. Let $\{\gamma_j : j = 1, 2, ..., m\}$ be a set of disjoint simple closed curves on Σ_g that correspond to the nodes of Y. Then the curves corresponding to the nodes on Z are exactly lifts of the curves that correspond to $\{\gamma_j : j = 1, 2, ..., m\}$. In particular, by Claim 5.9, χ vanishes on all curves corresponding to the nodes of Z. Since it also vanishes on all punctures of Z by definition, we can extend χ to a cohomology class

$$\bar{\chi} \in H^1(\overline{Z_1} \sqcup \overline{Z_2} \sqcup \cdots \sqcup \overline{Z_l}, \mathbb{R}),$$

where $\sqcup$ denotes disjoint union.

We now let η be an arbitrary smooth 1-form on $\bigsqcup_{i=1}^{l} \overline{Z_i}$ in the de Rham class $\bar{\chi}$, such that $\eta = 0$ in a neighbourhood of any cusp or node of Z. Denote the pullback of η to Z by ω. After equipping $\sigma_\pi(X_i)$ and $\bigsqcup_{i=1}^{l} Z_i$ with their corresponding hyperbolic metrics, it is classical that there exists a sequence of maps

$$f_i : \sigma_\pi(X_i) \to Z$$

with the following property: for any $\varepsilon > 0$, the Lipschitz constant $L_i(\varepsilon)$ of f_i on the ε-thick part of $\sigma_\pi(X_i)$ (i.e. the part where the injectivity radius is at least ε) has $L_i(\varepsilon) \to 1$ as $n \to \infty$.

It now follows that $f_i^*\omega$ has bounded energy on any thick part, so from the fact that $f_i^*\omega$ vanishes in a neighborhood of $\bigcup_j \gamma_j$, we see that

$$\int_{\sigma_\pi(X_i)} f_i^*\omega \wedge \star(f_i^*\omega)$$

is bounded in $i \geq 1$. However $\mathrm{E}_\chi \circ \sigma_\pi(X_i) \leq \int_{\sigma_\pi(X_i)} f_i^*\omega \wedge \star(f_i^*\omega)$ by Remark 2.14, which is a contradiction. This shows that $\mathrm{E}_\chi \circ \sigma_\pi$ is bounded, which concludes the proof of Theorem B.1.

5.4 Subgroups with the Putman–Wieland property

Using Theorem B.1 we derive the following criterion, which we will use to show Theorem B.2.

Definition 5.11. The gonality of a closed Riemann surface Z is the least degree of a non-constant meromorphic function on Z.

Proposition 5.12. *Let $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$ be a covering map of degree d, and suppose that $\chi \in H^1(\Sigma_h, \mathbb{R})$ has a finite Γ_π-orbit. Then for any marked Riemann surface $X \in \mathcal{T}_{g,n}$, the lift $\sigma_\pi(X)$ has gonality at most d.*

Proof. Fix an arbitrary Riemann surface $X \in \mathcal{T}_{g,n}$, and identify π with the holomorphic map $\sigma_\pi(X) \to X$ in the homotopy class of $\pi : \Sigma_{h,m} \to \Sigma_{g,n}$. Let ϕ be the holomorphic 1-form on $\sigma_\pi(X)$ such that $[\mathrm{Re}(\phi)] = \chi$. By Theorem B.1, we see that E_χ is constant along $\sigma_\pi \mathcal{T}_{g,n} \subseteq \mathcal{T}_h$. Therefore by Corollary 5.2, for any Beltrami form μ on X, the $(0,1)$-form $(p^*\mu)\phi$ is $\bar{\partial}$-exact.

Definition 5.13. Let $p : Y \to X$ be a cover of Riemann surfaces, and Φ be a holomorphic quadratic differential on Y. We define the Θ-projection of Φ to be the holomorphic quadratic differential $\Psi = \Theta_p(\Phi)$ on X given by

$$\Psi_z = \sum_{\tilde{z} \in p^{-1}(z)} (p^*_{\tilde{z}})^{-1}\, \Phi_{\tilde{z}}.$$

For any holomorphic 1-form ψ on $\sigma_\pi(X)$, we have by Stokes' theorem

$$0 = \int_{\sigma_\pi(X)} (p^*\mu)\phi\psi = \int_X \mu \Theta_\pi(\phi\psi).$$

Since μ was arbitrary, we have $\Theta_\pi(\phi\psi) = 0$ for any holomorphic 1-form ψ on $\sigma_\pi(X)$.

Suppose now that $\sigma_\pi(X)$ admits no meromorphic function of degree at most d. Let $x \in X$ be a generic point, and fix some $y \in \pi^{-1}(x)$. Let D be the effective divisor of degree d that corresponds to $p^{-1}(x)$, and $D_{\setminus y} = D - y$. We denote by K the canonical divisor on $\sigma_\pi(X)$. Then by the Riemann-Roch formula, denoting $h^0(D) = \dim\{f \text{ meromorphic with } (f) + D \geq 0\}$, we have

$$h^0(D) - h^0(K - D) = \deg D - h + 1,$$
$$h^0(D_{\setminus y}) - h^0(K - D_{\setminus y}) = \deg D - h.$$

Since $\sigma_\pi(X)$ admits no meromorphic functions of degree at most $d = \deg D > \deg D_{\setminus y}$, we see immediately that $h^0(D) = h^0(D_{\setminus y}) = 1$. Therefore

$$h^0(K - D_{\setminus y}) = h^0(K - D) + 1.$$

But $h^0(K - D)$ is the maximal number of linearly independent holomorphic 1-forms on $\sigma_\pi(X)$ that vanish on $p^{-1}(x)$, and $h^0(K - D_{\setminus y})$ is the maximal number of linearly independent holomorphic 1-forms on $\sigma_\pi(X)$ that vanish on $p^{-1}(x) \setminus \{y\}$. Hence there exists a holomorphic 1-form ψ such that

$$\psi|_{p^{-1}(x)\setminus\{y\}} = 0 \text{ and } \psi(y) \neq 0.$$

Since x was a general point, we can suppose $\phi(y) \neq 0$. Therefore $\Theta_\pi(\phi\psi)$ does not vanish at x. This is a contradiction. □

We now show Theorem B.2. The idea is to use Proposition 5.12, combined with the result of Li–Yau [45] relating the least eigenvalue of the Laplacian and the least degree of a non-constant meromorphic map on a closed Riemann surface.

Proof of Theorem B.2. The following result was originally shown by Li–Yau [45]. Our version was stated, among others, by [20, (11)].

Theorem 5.14. *For any closed Riemann surface X of genus g, the gonality of X is at least $2\lambda_1(X)(g-1)$.*

Now assume $X \in \mathcal{T}_g$ is such that $\lambda_1(\sigma_\pi(X)) > \frac{1}{2(g-1)}$. If we let h be the genus of $\sigma_\pi(X)$, we see that $\sigma_\pi(X)$ has gonality at least

$$2\lambda_1(\sigma_\pi(X))(h-1) > \frac{h-1}{g-1} \geq \deg p,$$

where the final inequality follows from the Riemann–Hurwitz formula. This is a contradiction with Proposition 5.12. □

Chapter 6

Energy of harmonic maps into symmetric spaces

6.1 Introduction

In this chapter, we study the energy functionals associated to equivariant harmonic maps into the symmetric space $X_n = \mathrm{GL}(n, \mathbb{C})/\mathrm{U}(n)$. For any Riemann surface $S \in \mathcal{T}_g$ and any irreducible representation $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$, by one direction of the non-abelian Hodge correspondence (see §2.3) there exists a ρ-equivariant harmonic map $f : \tilde{S} \to X_n$. Using these harmonic maps, we define the energy functional $\mathrm{E}_\rho : \mathcal{T}_g \to \mathbb{R}$ that assigns to $S \in \mathcal{T}_g$ the energy of the map f (i.e. the L^2 norm of the derivative of f over S).

Recall that Toledo [66] showed that the map E_ρ is plurisubharmonic. In this chapter, we investigate directions where E_ρ is not strictly plurisubharmonic, i.e. the kernel of the Levi form of E_ρ, that we denote K_ρ. Note that K_ρ when $\rho : \pi_1(\Sigma_g) \to \mathbb{R}$ is a real cohomology class appeared in our study of the Putman–Wieland conjecture in Chapter 5.

Our main results are Theorems C.2, C.3, and C.4, that are all in turn applications of Theorem C.1. In this introductory section, we give more precise statements of Theorems C.1, C.3 and C.4. We will also state Proposition 6.5, which is an improvement on our results in §5.1 on plurisubharmonicity of the Hodge norm. We first state our results on Higgs bundles (Theorems C.3 and C.4) in §6.1.1, then our results for $n = 1$ (Proposition 6.5) in §6.1.2, and finally our results on Riemannian manifolds with very strongly seminegative curvature (Theorem C.1) in §6.1.3.

6.1.1 Higgs bundles

We first give a more precise statement of Theorem C.3. Recall from §2.3.3.2 that, given a Riemann surface S, the moduli space of polystable rank n degree 0 Higgs bundles over S, denoted $\mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S)$, admits a natural map $H : \mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S) \to \mathcal{B}(S)$ called the Hitchin fibration.

Definition 6.1. The d-th critical locus of H is the set of points in $\mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S)$ where the rank of DH is at most $\dim \mathcal{B}(S) - d$.

These appear in the recent work of Hitchin [35]. He shows that in $\mathrm{SL}(2,\mathbb{C})$, the d-th critical locus is the space of Higgs fields that have at least d zeros (counting multiplicity). The following is a more precise version of Theorem C.3.

Theorem 6.2. *Let $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n,\mathbb{C})$ be an irreducible representation. Then if $d = \dim K_\rho(S)$, the point* $\mathrm{Higgs}(\rho, S)$ *lies in the d-th critical locus of the Hitchin integrable system.*

When ρ is a Hitchin representation in $\mathrm{SL}(n,\mathbb{R})$, Slegers [63] has shown that E_ρ is strictly plurisubharmonic. Our next result uses this to show that for a generic representation ρ, at a generic point in $\mathcal{T}_g$, the energy E_ρ is strictly plurisubharmonic. When S is a marked Riemann surface and $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n,\mathbb{C})$ is a completely reducible representation, denote by $\mathrm{Higgs}(\rho, S)$ the polystable degree 0 Higgs bundle over S that corresponds to ρ by the non-abelian Hodge theorem. The following is a slightly more precise version of Theorem C.4.

Theorem 6.3. *Let $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n,\mathbb{C})$ be an irreducible representation for $g \geq 3$. For any $S \in \mathcal{T}_g$, with the property that* $\mathrm{Higgs}(\rho, S)$ *lies in a smooth fibre of the Hitchin fibration, $K_\rho(S) = \{0\}$. Conversely, for any $g \geq 4, S \in \mathcal{T}_g$, and any $n \geq 2$, there exists a representation $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n,\mathbb{C})$ such that $K_\rho(S) \neq \{0\}$.*

Remark 6.4. Here by a smooth fibre of the Hitchin fibration, we mean a fibre whose corresponding spectral curve is smooth. We refer the reader to §2.3.3.2 for the precise definitions.

6.1.2 The 1-dimensional case

When $n = 1$, we showed in Chapter 5 that E_ρ is plurisubharmonic (Theorem 5.1), and gave a description of K_ρ (Corollary 5.2). Here we continue the study of K_ρ for 1-dimensional representations, and put the results of §5.1 in current context. Before we state our next result, recall that over a Riemann surface S, the space of holomorphic quadratic differentials $\mathrm{QD}(S)$ has dimension $3g-3$, and is naturally isomorphic to the cotangent space to $\mathcal{T}_g$ at S. We denote by $\Omega(S)$ the set of holomorphic 1-forms on S.

Proposition 6.5. *Let $\rho : \pi_1(\Sigma_g) \to \mathbb{C}^*$ be a representation. Given a marked Riemann surface $S \in \mathcal{T}_g$, let ϕ be the holomorphic 1-form whose real part represents the cohomology class $-\frac{1}{2}\log|\rho|$. Then $K_\rho(S)$ is the annihilator of the set $\phi \otimes \Omega(S) \leq \mathrm{QD}(S)$. Moreover, the distribution K_ρ is integrable, and the leaves of the resulting foliation are complex submanifolds of $\mathcal{T}_g$ of codimension g.*

6.1.3 Results on Riemannian manifolds with very strongly seminegative curvature

In this section we give a more precise statement of Theorem C.1, which is Theorem 6.8 below. Let M be a Riemannian manifold with an isometry group $\mathrm{Isom}(M)$. Let $\rho : \pi_1(\Sigma_g) \to \mathrm{Isom}(M)$ be a representation.

Definition 6.6. If $(S_t \in \mathcal{T}_g : t \in \mathbb{D})$ is a holomorphic disk in Teichmüller space, and if $f_t : \tilde{S}_t \to M$ is a smoothly varying family of harmonic ρ-equivariant maps from the universal cover $\tilde{S}_t$ of S_t to M, we say that $((S_t, f_t) : t \in \mathbb{D})$ form a complex disk of ρ-equivariant harmonic maps, based at S_0 with direction $\mu = \left.\frac{\partial S_{x+iy}}{\partial x}\right|_{x=y=0}$.

Definition 6.7. If $f : \tilde{\Sigma}_g \to M$ is a ρ-equivariant smooth map, the pullback bundle f^*TM descends to a bundle on Σ_g, that is naturally equipped with a connection. We call this bundle over Σ_g the equivariant pullback bundle of f.

We remind the reader of the notion of very strongly seminegative curvature introduced by Siu [62] (we will define it in §6.2.2), meaning that the curvature operator is negative semi-definite.

Theorem 6.8. *Suppose that M has very strongly seminegative curvature, and let $((S_t, f_t) : t \in \mathbb{D})$ be a complex disk of equivariant harmonic maps with direction μ. Then $\Delta \mathrm{E}(f_t)(0) = 0$ if and only if there exists a section ξ of the complexified equivariant pullback bundle $E^{\mathbb{C}}$, such that*

$$\mu \partial f_0 = \bar{\partial} \xi \text{ and } R^M(\xi, \partial f_0) = 0. \tag{6.1.1}$$

Moreover, in this case $\xi - \frac{\partial f}{\partial t}$ is a parallel section of $\ker(R^M(-, df_0)) \leq E$.

Given a complex disk of equivariant harmonic maps $((S_t, f_t) : t \in \mathbb{D})$, note that the equivariant pullback bundle E of S_0 has a connection $f_0^*\nabla$ induced from the Levi–Civita connection on TM. Then $(f_0^*\nabla)^{0,1}$ defines a holomorphic structure on $E^{\mathbb{C}} := E \otimes \mathbb{C}$ by the Koszul–Malgrange theorem. In the next result, we rephrase the computation of Toledo in terms of the Hodge theory of this holomorphic bundle.

Theorem 6.9. *Given a complex disk $((S_t, f_t) : t \in \mathbb{D})$ of equivariant harmonic maps with direction μ into some Riemannian manifold M (with no curvature assumptions), let E be as above. Suppose that $\mu\partial f_0 = \bar{\partial}\varphi + \theta$ is the Hodge decomposition of $\mu\partial f_0$, where φ is a section of $E^{\mathbb{C}}$ and θ is a $\Delta_{\bar{\partial}}$-harmonic $E^{\mathbb{C}}$-valued $(0,1)$-form. Then*

$$\Delta\left(\mathrm{E}(f_t)\right)(0) = 8\|\theta\|_{L^2}^2 + 8\left\|\bar{\partial}(w-\varphi)\right\|_{L^2}^2 - 8\int_{\Sigma_g} i\langle R^M(w, \partial f_0) \wedge \bar{\partial} f_0, \overline{w}\rangle,$$

where R^M is the Riemann curvature of M, $\langle\cdot,\cdot\rangle$ is the metric on M, both of which are extended complex linearly, and $w = \left.\frac{\partial f_t}{\partial t}\right|_{t=0} \in \Gamma(S, E^{\mathbb{C}})$.

From Theorem 6.9, it easily follows that whenever M has non-positive Hermitian sectional curvature, the energy $\mathrm{E}(f_t)$ is subharmonic along every complex disk of harmonic maps, which is exactly the result of Toledo [66].

All our results on Higgs bundles are derived from the following corollary of Theorem C.1.

Theorem 6.10. *Let $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ be completely reducible, S be a marked Riemann surface of genus g. If $(E, \phi) = \mathrm{Higgs}(\rho, S)$, then $K_\rho(S)$ is the space of Beltrami forms μ on S such that there exists a section ξ of $\mathrm{End}(E)$ with*

$$\mu\phi = \bar{\partial}\xi \text{ and } [\phi, \xi] = 0. \tag{6.1.2}$$

6.1.4 Outline and organization

In §6.2, we recall some preliminaries about Hodge theory on holomorphic vector bundles over Riemann surfaces, and on Siu's curvature condition and how it applies to symmetric spaces.

6.1.4.1 Riemannian manifolds

In §6.3, we prove Theorems 6.8 and 6.9.

Theorem 6.9 is shown by direct computation, closely following the computation of Toledo [66]. The main differences are that we give a formula for $\Delta\mathrm{E}$, rather than an inequality, and that we introduce the use of Hodge theory on the complexified pullback bundle $f^*TM \otimes \mathbb{C}$, which simplifies some of the expressions.

Theorem 6.8 is derived from Theorem 6.9, one direction being simple: If $\Delta \mathrm{E} = 0$, the derivative f_t of f at $t = 0$ is precisely the ξ from Theorem 6.8. The converse direction is a Bochner argument that depends on the assumption of very strongly seminegative curvature.

6.1.4.2 Higgs bundles

The rest of the paper deals with the specialized situation of completely reducible $\mathrm{GL}(n, \mathbb{C})$ representations.

To the author's knowledge, a proof that E_ρ is smooth does not exist in the literature, although this was shown by Slegers [64] to follow from the classical result of Eells–Lemaire [18] when ρ is Hitchin. We show in §6.4 that when the representation $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ is completely reducible, the harmonic map can be chosen to depend smoothly on the complex structure $S \in \mathcal{T}_g$ and on ρ. It immediately follows that the Higgs field and the harmonic metric also depend smoothly on ρ, S. This allows us to apply Theorem 6.8, and at the same time shows that $\mathrm{E}_\rho, \mathcal{R}_\rho$ are smooth.

Theorem 6.10 then follows immediately from Theorem 6.8. We then prove Theorem C.2 in §6.5. On the one hand, we already have a description of $K_\rho(S)$ in terms of $\mathrm{Higgs}(\rho, S)$ from Theorem 6.10. The description of $\ker d\mathcal{R}_\rho$ in terms of the Higgs bundle $\mathrm{Higgs}(\rho, S)$ follows from a construction of the moduli space of solutions to the Hitchin equation over a varying Riemann surface, that we carry out in §6.4.5. In this case, both directions of the equivalence require a Bochner argument. In §6.6 we show Theorem 6.2.

In §6.7, we analyze the case $n = 1$, proving Proposition 6.5. This is a straightforward corollary of Theorem 6.10, after showing that ϕ from the statement of Proposition 6.5 is the Higgs field associated to ρ.

In §6.8, we show Theorem 6.3. From the general facts about spectral curves, it is easy to show that $K_\rho(S)$ depends only on the image of ρ in the Hitchin fibration associated to S, as long as this fibre is smooth. We then construct in each fibre a representation ρ such that E_ρ is strictly plurisubharmonic at S. This relies on the analysis of Slegers [63], but can also be shown easily from Theorem 6.10.

After that, for an arbitrary $S \in \mathcal{T}_g$ for $g \geq 4$, we construct explicitly Higgs bundles over S in the nilpotent cone for which the system (6.1.2) has many non-zero solutions $\mu \in T_S\mathcal{T}_g$.

6.2 Preliminaries

6.2.1 Hodge theory on holomorphic vector bundles

Let E be a complex hermitian vector bundle over a Riemann surface S. If E is equipped with a connection ∇, the $(0,1)$-part of this connection defines the structure of a holomorphic vector bundle on E. If ∇ is unitary for the metric on E, then it is equal to the Chern connection on E.

We write $\nabla = \partial + \bar{\partial}$ for the splitting of ∇ into its $(1,0)$ and $(0,1)$ parts. If S is equipped with a volume form, that is automatically Kähler, we may construct the formal adjoint ∇^* of ∇. We split $\nabla^* = \partial^* + \bar{\partial}^*$ into its $(1,0)$ and $(0,1)$ parts. We may now construct holomorphic and antiholomorphic Laplacians on E-valued differential forms,

$$\Delta_{\bar{\partial}} = \bar{\partial}^* \bar{\partial} + \bar{\partial} \bar{\partial}^*,$$
$$\Delta_{\partial} = \partial^* \partial + \partial \partial^*.$$

It is well-known that these are elliptic, and satisfy a Bochner–Kodaira–Nakano identity [15, Theorem (VII.1.2)]

$$\Delta_{\bar{\partial}} = \Delta_{\partial} + i[F_\nabla, \Lambda_\omega],$$

where F_∇ is the curvature of ∇, and Λ_ω is contraction with the Kähler form ω. In particular, Δ_∂ and $\Delta_{\bar{\partial}}$ agree on 1-forms. Standard Hodge theory now shows the following.

Proposition 6.11. *Let ξ be an E-valued $(0,1)$ form. Then there exists a section g of E, and a closed and coclosed E-valued $(0,1)$ form θ, such that $\xi = \bar{\partial} g + \theta$. Moreover, when (E, ∇) comes from the complexification of a real vector bundle, then θ is complex conjugate to a holomorphic E-valued 1-form.*

6.2.2 Curvature of symmetric spaces

Here we recall that symmetric spaces of non-compact type have very strongly seminegative curvature, in the sense of Siu [62]. We first recall the definition of Siu.

Definition 6.12. A manifold X has very strongly seminegative curvature if the sesquilinear form defined on $\bigwedge^2 TX \otimes \mathbb{C}$ by

$$Q(\alpha \wedge \beta, \gamma \wedge \delta) = \langle R(\alpha, \beta)\bar{\delta}, \bar{\gamma} \rangle$$

is negative semi-definite, where R is the Riemann curvature of X.

The following proposition appears in the report of Loustau [46, Corollary 5.5].

Proposition 6.13. *Let X be a locally symmetric space of non-positive curvature. Then X has very strongly seminegative curvature.*

6.3 Levi form of the energy

Let $((S_t, f_t) : t \in \mathbb{D})$ be a complex disk of harmonic maps for some representation $\rho : \pi_1(\Sigma_g) \to \mathrm{Isom}(M)$. We then have

$$\mathrm{E}(f_t) = i \int_{\Sigma_g} \langle \partial f_t \wedge \bar{\partial} f_t \rangle.$$

In this section, we prove Theorems 6.8 and 6.9.

As mentioned in the outline, deriving Theorem 6.8 from Theorem 6.9 consists of two steps: showing that if $\Delta \mathrm{E} = 0$, then the system (6.1.1) has a solution, which follows immediately from Theorem 6.9, and showing the converse, that depends on a Bochner argument that shows that any solution to (6.1.1) has to differ from $\frac{\partial f}{\partial t}$ by a parallel section of $\ker R^M(-, df_0)$. This Bochner argument relies on the assumption of very strongly seminegative curvature for M, and on second order elliptic equations for $\frac{\partial f}{\partial t}$ obtained in §6.3.1 by taking the first-order variation of the harmonic map equation. Using this argument, we prove Theorem 6.8 assuming Theorem 6.9 in §6.3.2.

In §6.3.3, we prove Theorem 6.9, following [66]. There are two major differences between our computation and that in [66]. First, we give an actual equality, rather than an inequality as in [66], and give a slightly more precise analysis of the equality case than [66]. Second, by using Hodge theory on the complexified equivariant pullback bundle f_0^*TM, we simplify some of the expressions.

6.3.1 First-order variation of the harmonic map equation

Suppose first that $((f_t, J_t) : t \in (-1, 1))$ is an open interval of ρ-equivariant harmonic maps. We assemble them into a map

$$F : \Sigma_g \times (-1, 1) \to M.$$

Let $f = f_0$. We now consider the vector bundle $(F^*TM, F^*\nabla)$. Let $\Pi_t : f^*TM \to F(-, t)^*TM$ be the parallel transport in this bundle along vertical paths of the form $\{-\} \times [0, t]$ or $\{-\} \times [t, 0]$. We let $\nabla^t = \Pi_t^*(F^*\nabla)$. We let E be the equivariant pullback bundle of f_0, and $E^{\mathbb{C}} = E \otimes \mathbb{C}$.

Proposition 6.14. *For $\dot f = \frac{\partial F}{\partial t}$, we have $\dot\nabla = R^M(\dot f, df)$, where R^M is the curvature tensor of M, and*

$$d^\nabla\left(\partial^\nabla \dot f + \mu\partial f - \bar\mu\bar\partial f\right) + R^M(\dot f, \bar\partial f)\wedge\partial f = 0,$$
$$d^\nabla\left(\bar\partial^\nabla \dot f - \mu\partial f + \bar\mu\bar\partial f\right) + R^M(\dot f, \partial f)\wedge\bar\partial f = 0.$$

Proof. If X is a vector field on Σ_g, and s is a section of E, we have

$$\begin{aligned}\left.\frac{d}{dt}\right|_{t=0}\nabla^t_X s &= \nabla_{\partial_t}|_{t=0}\nabla_X(\Pi_t s)\\ &= \nabla_X(\nabla_{\partial_t}\Pi_t s) + R^{F^*TM}(\partial_t, X)s\\ &= R^M(\dot f, f_*X)s,\end{aligned}$$

where R^{F^*TM} denotes the curvature of the bundle $(F^*TM, F^*\nabla)$.

We now turn to the second claim. We only show the first equation, since the second one is completely analogous. The harmonic map equation is

$$d^{\nabla^t}\left(\left(\Pi_t^{-1}df_t\right)\circ\frac{\mathrm{id} - iJ_t}{2}\right) = 0.$$

Since $\Pi_t^{-1}df_t = df + td^\nabla\dot f + O(t^2)$, we have by differentiating and using [50, Claim 3.2],

$$R^M(\dot f, df)\wedge\left(df\circ\frac{\mathrm{id}-iJ}{2}\right) + d^\nabla\left(d^\nabla\dot f\circ\frac{\mathrm{id}-iJ}{2}\right) - \frac{i}{2}d^\nabla\left(2i\mu\partial f - 2i\bar\mu\bar\partial f\right) = 0.$$

Therefore

$$d^\nabla\left(\partial^\nabla\dot f + \mu\partial f - \bar\mu\bar\partial f\right) + R^M(\dot f, \bar\partial f)\wedge\partial f = 0.$$

□

Now note that $\nabla^{0,1}$ makes $E^{\mathbb{C}}$ into a holomorphic vector bundle over (Σ_g, J_0). Therefore we can write

$$\mu\partial f = \bar\partial^\nabla\varphi + \overline{\theta},$$

where φ is a section of $E^{\mathbb{C}}$ and θ is a holomorphic $E^{\mathbb{C}}$-valued 1-form. For brevity of notation, from now on we drop the reference to the connection and write $d, \partial, \bar\partial$ for $d^\nabla, \partial^\nabla, \bar\partial^\nabla$.

Corollary 6.15. *In the notation given above,*

$$\bar\partial\partial(\dot f - 2\mathrm{Re}(\varphi)) = -R^M(\dot f, \bar\partial f)\wedge\partial f - f^*R^M\varphi,$$
$$\partial\bar\partial(\dot f - 2\mathrm{Re}(\varphi)) = -R^M(\dot f, \partial f)\wedge\bar\partial f - f^*R^M\bar\varphi.$$

Proof. Again we only show the first equation. Note that

$$\begin{aligned}\bar{\partial}\partial(\dot{f} - \varphi - \bar{\varphi}) &= d^{\nabla}(\partial\dot{f} + \bar{\partial}\varphi - \partial\bar{\varphi}) - (\partial\bar{\partial} + \bar{\partial}\partial)\varphi \\ &= -R^M(\dot{f}, \bar{\partial}f) \wedge \partial f - (d^{\nabla})^2\varphi,\end{aligned}$$

which concludes the proof by definition of the curvature tensor. □

6.3.2 Proof of Theorem 6.8 assuming Theorem 6.9

By Theorem 6.9, $\Delta\mathrm{E}(f_t) = 0$ if and only if for some $\varphi \in \Gamma(E^{\mathbb{C}})$, we have

$$\begin{aligned}\mu\partial f &= \bar{\partial}\varphi, \\ \bar{\partial}(w - \varphi) &= 0, \\ \langle R^M(w, \partial f) \wedge \bar{\partial} f, \bar{w}\rangle &= 0.\end{aligned} \tag{6.3.1}$$

Using elementary linear algebra, we can rewrite (6.3.1) as follows.

Claim 6.16. For any two $X, Y \in TM \otimes \mathbb{C}$, the equality $\langle R^M(X,Y)\bar{Y}, \bar{X}\rangle = 0$ holds if and only if $R^M(X,Y) = 0$.

Proof. We define the sesquilinear form Q on $\bigwedge^2 TM \otimes \mathbb{C}$ by

$$Q(X \wedge Y, Z \wedge W) = \langle R^M(X,Y)\bar{W}, \bar{Z}\rangle.$$

This is well-defined by the standard symmetries of the Riemann curvature R^M. Since M has very strongly seminegative curvature, by definition Q is negative semi-definite. Thus $Q(X \wedge Y, X \wedge Y) = 0$ if and only if $Q(X \wedge Y, -) = 0$, which is in turn equivalent to $R^M(X,Y) = 0$. □

Therefore (6.3.1) is equivalent to

$$R^M(w, \partial f) = 0. \tag{6.3.2}$$

Hence if $\Delta\mathrm{E}(f_t) = 0$, the required solution is $\xi = w$.

Conversely, assume that for a section ξ of $E^{\mathbb{C}}$, we have

$$\mu\partial f = \bar{\partial}\xi \text{ and } R^M(\xi, \partial f) = 0.$$

Note that by Proposition 6.14, we have

$$\begin{aligned}d\left(\bar{\partial}\dot{f}^{\mu} - \mu\partial f + \bar{\mu}\bar{\partial} f\right) + R^M(\dot{f}^{\mu}, \partial f) \wedge \bar{\partial} f = 0, \\ d\left(\bar{\partial}\left(i\dot{f}^{i\mu}\right) + \mu\partial f + \bar{\mu}\bar{\partial} f\right) + R^M(i\dot{f}^{i\mu}, \partial f) \wedge \bar{\partial} f = 0.\end{aligned}$$

Subtracting, and recalling that $w = \frac{1}{2}\left(\dot{f}^{\mu} - i\dot{f}^{i\mu}\right)$, we have

$$\partial\bar{\partial}(w-\xi) + R^M(,\partial f)\wedge\bar{\partial}f = 0.$$

Since $R^M(\xi,\partial f) = 0$, we have for $V = w - \xi$, the equation

$$\partial\bar{\partial}V + R^M(V,\partial f)\wedge\bar{\partial}f = 0.$$

Taking the inner product with $\bar{V}$ and integrating, we have

$$\begin{aligned}\int_{\Sigma_g}\frac{i}{2}\langle R^M(V,\partial f)\wedge\bar{\partial}f,\bar{V}\rangle &= \frac{i}{2}\int_{\Sigma_g} -\langle\bar{V},\partial\bar{\partial}V\rangle \\ &= \frac{i}{2}\int_{\Sigma_g} d\langle\bar{V},\bar{\partial}V\rangle - \langle\bar{V},\partial\bar{\partial}V\rangle \\ &= \frac{i}{2}\int_{\Sigma_g}\langle\partial\bar{V}\wedge\bar{\partial}V\rangle = \left\|\bar{\partial}V\right\|_{L^2}^2.\end{aligned}$$

Since M has very strongly seminegative curvature, it also has non-positive Hermitian sectional curvature. Therefore $\frac{i}{2}\langle R^M(V,\partial f)\wedge\bar{\partial}f,\bar{V}\rangle \le 0$. Therefore we must have

$$\bar{\partial}V = 0 \text{ and } \langle R^M(V,\partial f)\wedge\bar{\partial}f,\bar{V}\rangle = 0. \tag{6.3.3}$$

Thus $\bar{\partial}w = \bar{\partial}\xi$ and $R^M(w,\partial f) = R^M(\xi,\partial f) = 0$, and hence $\Delta\mathrm{E}(f_t) = 0$.

We now turn to the final statement. We want to show that $\partial V = 0$ given $\bar{\partial}V = 0$ and $R^M(V,\partial f) = 0$. This follows from a Bochner-type computation

$$\begin{aligned}\|\partial V\|_{L^2}^2 = \frac{i}{2}\int_{\Sigma_g}\langle\partial V\wedge\bar{\partial}\bar{V}\rangle &= \frac{i}{2}\int_{\Sigma_g} d\langle V,\bar{\partial}\bar{V}\rangle - \langle V,\partial\bar{\partial}\bar{V}\rangle \\ = -\frac{i}{2}\int_{\Sigma_g}\langle V,-\bar{\partial}\partial\bar{V} + f^*R^M\bar{V}\rangle &= -\frac{i}{2}\int_{\Sigma_g}\langle V,f^*R^M\bar{V}\rangle.\end{aligned}$$

We have, in a local holomorphic coordinate z on (Σ_g,J), using the Bianchi identity

$$\begin{aligned}f^*R^M\bar{V} &= R^M(f_z,f_{\bar{z}})\bar{V}dz\wedge d\bar{z} = -(R^M(\bar{V},f_z)f_{\bar{z}} - R^M(f_{\bar{z}},\bar{V})f_z)dz\wedge d\bar{z} \\ &= -R^M(\bar{V},f_z)f_{\bar{z}}dz\wedge d\bar{z}.\end{aligned}$$

Therefore

$$-\frac{i}{2}\langle V,f^*R^M\bar{V}\rangle = \frac{i}{2}dz\wedge d\bar{z}\langle R^M(\bar{V},f_z)f_{\bar{z}},V\rangle \le 0,$$

since M has non-positive Hermitian sectional curvature. Thus $\|\partial V\|_{L^2}^2 \le 0$, so $\partial V = 0$ and $R^M(V,f_{\bar{z}}) = 0$ by Claim 6.16. Along with $\bar{\partial}V = 0$ and $R^M(V,\partial f) = 0$, this implies that $dV = 0$ and $R^M(V,df) = 0$, as desired.

6.3.3 Proof of Theorem 6.9

We first compute some formulas for the second variation of the energy in a direction defined by μ. Thus we are still in the setting where $((J_t, f_t) : t \in (-1,1))$ is an interval of equivariant harmonic maps. As in the previous section set $f = f_0$ and $J = J_0$, and equip $E^{\mathbb{C}}$ with the holomorphic structure coming from $(f^*\nabla)^{0,1}$, and write $\mu\partial f = \bar{\partial}\varphi + \bar{\theta}$.

Note that

$$\mathrm{E}(f_t) = i\int_{\Sigma_g}\left\langle df_t \circ \frac{\mathrm{id} - iJ_t}{2} \wedge df_t \circ \frac{\mathrm{id} + iJ_t}{2}\right\rangle = -\frac{1}{2}\int_{\Sigma_g}\langle df_t \wedge df_t \circ J_t\rangle.$$

We now define $F(s,t) = -\frac{1}{2}\int_{\Sigma_g}\langle df_s \wedge df_s \circ J_t\rangle$. Since f_t is harmonic for J_t, we have $\frac{\partial F}{\partial s}(t,t) = 0$. Therefore

$$\begin{aligned}\left.\frac{d^2}{dt^2}\right|_{t=0} F(t,t) &= \left(\frac{\partial}{\partial s} + \frac{\partial}{\partial t}\right)^2_{t=0} F(t,t) = \left(\frac{\partial}{\partial s} + \frac{\partial}{\partial t}\right)\frac{\partial F}{\partial t}(0,0)\\ &= \frac{\partial^2 F}{\partial t^2}(0,0) + \frac{\partial^2 F}{\partial t \partial s}(0,0).\end{aligned} \tag{6.3.4}$$

We observe that since parallel transport preserves the metric, we have

$$F(s,t) = -\frac{1}{2}\int_{\Sigma_g}\langle \Pi_s^{-1} df_s \wedge \Pi_s^{-1} df_s \circ J_t\rangle.$$

From now on, we transport all derivatives df_s to the bundle E by Π_s, and work exclusively on E.

We will compute the two terms in (6.3.4) separately, but first we recall the first two derivatives of J in the direction μ. This is essentially [50, Claim 3.2].

Claim 6.17. We have $\ddot{J} = 4|\mu|^2 J$, and for any 1-form ω,

$$\omega \circ \dot{J} = 2i\left(\mu\omega^{1,0} - \bar{\mu}\omega^{0,1}\right).$$

6.3.3.1 First term

We have from Claim 6.17,

$$\begin{aligned}\frac{\partial^2 F}{\partial t^2}(0,0) &= -\frac{1}{2}\int_{\Sigma_g} 4|\mu|^2\langle df \wedge df \circ J\rangle = 4i\int_{\Sigma_g}|\mu|^2\langle \partial f \wedge \bar{\partial} f\rangle\\ &= 4i\int_{\Sigma_g}\langle \overline{\mu\partial f} \wedge \mu\partial f\rangle = 4i\int_{\Sigma_g}\langle(\partial\bar{\varphi} + \theta)\wedge(\bar{\partial}\varphi + \bar{\theta})\rangle\\ &= 4\left\|\bar{\partial}\varphi\right\|^2 + 4\|\theta\|^2.\end{aligned} \tag{6.3.5}$$

Here we introduced the L^2 norm for $(1,0)$- or $(0,1)$-forms, by

$$\|\phi\|^2 = i\int_{\Sigma_g} \langle \phi \wedge \bar{\phi}\rangle,$$

for $(1,0)$-forms ϕ, and $\|\bar{\phi}\| = \|\phi\|$. We denote the associated Hermitian inner product by $(\cdot,\cdot)_{L^2}$.

6.3.3.2 Second term

We have

$$\begin{aligned}\frac{1}{2}\frac{\partial^2 F}{\partial s\partial t}(0,0) &= -\frac{1}{2}\int_{\Sigma_g}\langle d^\nabla \dot{f}\wedge df\circ \dot{J}\rangle = -i\int_{\Sigma_g}\langle d^\nabla \dot{f}\wedge \left(\mu\partial f - \bar{\mu}\bar{\partial} f\right)\rangle \\ &= -i\int_{\Sigma_g}\langle d^\nabla \dot{f}\wedge (\bar{\partial}\varphi + \bar{\theta} - \partial\bar{\varphi} - \theta)\rangle = -i\int_{\Sigma_g}\langle \partial\dot{f}\wedge\bar{\partial}\varphi\rangle - \overline{\langle \partial\dot{f}\wedge\bar{\partial}\varphi\rangle} \\ &= 2\mathrm{Im}\int_{\Sigma_g}\langle \partial\dot{f}\wedge\bar{\partial}\varphi\rangle = -(\partial\dot{f},\partial\bar{\varphi})_{L^2} - (\partial\bar{\varphi},\partial\dot{f})_{L^2}. \end{aligned} \tag{6.3.6}$$

We now compute this term in a different way,

$$\frac{1}{2}\frac{\partial^2 F}{\partial s\partial t}(0,0) = -\frac{1}{2}\frac{\partial^2 F}{\partial s^2} = -\frac{1}{2}\int_{\Sigma_g}\langle \dot{f},\mathcal{J}\dot{f}\rangle \mathrm{darea}_{g_0}, \tag{6.3.7}$$

where $\mathcal{J}$ is the Jacobi operator for some background conformal metric g_0 on (Σ_g, J), given by

$$\mathcal{J}V = -\Delta V + \sum_i R^M(\nabla_{e_i} f, V)\nabla_{e_i} f,$$

where e_i form a fibre-wise orthonormal basis of f^*TM. We will review some background on the Jacobi operator in §6.4.1. By Proposition 6.20, we have

$$\mathcal{J}V\,\mathrm{darea}_{g_0} = -2i\left(\partial\bar{\partial}V + R(V,\partial f)\wedge\bar{\partial} f\right).$$

We now return to the setting of a complex disk of equivariant harmonic maps $((J_z, f_z) : z \in \mathbb{D})$, where we have renamed the variable to $z = x + iy$ to avoid confusion. We are interested in $\Delta \mathrm{E}(f_z) = \frac{\partial^2 \mathrm{E}}{\partial x^2} + \frac{\partial^2 \mathrm{E}}{\partial y^2}$. Write $W = 2\frac{\partial f}{\partial z} = \dot{f}^\mu - i\dot{f}^{i\mu}$. We analogously introduce variables t_1, s_1 associated to x, and t_2, s_2 associated to y as above. Then by (6.3.6),

$$\frac{1}{2}\left(\frac{\partial^2 F}{\partial s_1\partial t_1} + \frac{\partial^2 F}{\partial s_2 \partial t_2}\right) = -(\partial\bar{W},\partial\bar{\varphi})_{L^2} - (\partial\bar{\varphi},\partial\bar{W})_{L^2},$$

and by (6.3.7),

$$
\begin{aligned}
\frac{1}{2}\left(\frac{\partial^2 F}{\partial s_1 \partial t_1}+\frac{\partial^2 F}{\partial s_2 \partial t_2}\right) &= -\frac{1}{2}\int_{\Sigma_g}\left(\langle \dot{f}^{\mu}, \mathcal{J}\dot{f}^{\mu}\rangle + \langle \dot{f}^{i\mu}, \mathcal{J}\dot{f}^{i\mu}\rangle\right) d\mathrm{area}_{g_0} \\
&= -\frac{1}{4}\int_{\Sigma_g}\left(\langle W, \mathcal{J}\bar{W}\rangle + \langle \bar{W}, \mathcal{J}W\rangle\right) d\mathrm{area}_{g_0} \\
&= -\frac{1}{2}\mathrm{Re}\int_{\Sigma_g}\langle \bar{W}, \mathcal{J}W\rangle d\mathrm{area}_{g_0} \\
&= \mathrm{Re}\left(i\int_{\Sigma_g}\langle \bar{W}, \partial\bar{\partial}W\rangle + \langle \bar{W}, R(W,\partial f)\wedge\bar{\partial}f\rangle\right) \\
&= -\left\|\bar{\partial}W\right\|^2 + i\int_{\Sigma_g}\langle R(W,\partial f)\wedge\bar{\partial}f, \bar{W}\rangle.
\end{aligned}
$$

Therefore using these two equalities, we have

$$\Delta E(0) = \mathrm{I} + \mathrm{II},$$

where

$$
\begin{aligned}
\mathrm{I} &= 8\left\|\bar{\partial}\varphi\right\|^2 + 8\|\theta\|^2, \\
\mathrm{II} &= -2(\partial\bar{W}, \partial\bar{\varphi})_{L^2} - 2, (\partial\bar{\varphi}, \partial\bar{W})_{L^2} = -2\left\|\bar{\partial}W\right\|^2 + 2\mathcal{R} \\
\mathcal{R} &= i\int_{\Sigma_g}\langle R(W,\partial f)\wedge\bar{\partial}f, \bar{W}\rangle.
\end{aligned}
$$

In particular, we have

$$
\begin{aligned}
\mathrm{II} &= 2\left(-2(\partial\bar{W}, \partial\bar{\varphi})_{L^2} - 2\left(\partial\bar{\varphi}, \partial\bar{W}\right)_{L^2}\right) + 2\left\|\bar{\partial}W\right\|^2 - 2\mathcal{R} \\
&= 2\left\|\bar{\partial}(W-2\varphi)\right\|^2 - 8\left\|\bar{\partial}\varphi\right\|^2 - 2\mathcal{R}.
\end{aligned}
$$

This concludes the proof of Theorem 6.9.

6.4 Smooth dependence on the representation and complex structure

This section is devoted to showing that E_ρ and $\mathcal{R}_\rho$ are smooth, and having sufficient machinery to be able to compute their derivatives.

To show smoothness of E_ρ, we show that the ρ-equivariant harmonic map $f : (\tilde{\Sigma}_g, J) \to X_n$ can be chosen to depend smoothly on ρ and J. We show this by using the Banach manifold implicit function theorem, analogously to the classical result of

Eells–Lemaire [18]. Their argument essentially already shows that when f is unique, it depends smoothly on J. However, since we are interested in showing smooth dependence on the representation as well, we work in the setting of harmonic metrics on flat bundles. Note that here we show smoothness on the representation itself, not on the corresponding element of the character variety. In §6.4.2, we explain how the equation for the harmonic metric on a flat bundle is equivalent to the harmonic map equation, and how its first variation is the Jacobi operator of the associated harmonic map. We will use some standard results on the Jacobi operator in the sequel, so we recall them in §6.4.1. In §6.4.3, we show that f can be chosen to depend smoothly on ρ, J.

It now follows immediately that $\mathcal{R}$ is smooth, since the Higgs bundle (E, ϕ) and the harmonic metric h depend smoothly on ρ, J, so does the flat connection associated to $(E, i\phi)$. However, in proving Theorem C.2, it will be convenient to be able to say that $(\mathcal{R}_\rho)_*\mu = 0$ if and only if the associated solution to the Hitchin equation does not change to first order. For this we need to construct the moduli space of solutions to the Hitchin equation over a varying Riemann surface. We only do this in the locus of stable Higgs bundles, since this simplifies the analysis greatly. Constructing the moduli space of polystable solutions is much harder, even over a single Riemann surface, and was carried out by Fan [21]. Note that our result likely follows from the original paper of Simpson [61], however it seems worthwhile to include an analytic proof. The proof is essentially a repeat of the original proof of Hitchin [33], in a modified setup. This construction is in §6.4.5.

6.4.1 Jacobi operator

Here we collect the definition and some properties of the Jacobi operator that we will use in the sequel, mostly without proofs.

Definition 6.18. If $f : M \to N$ is a harmonic map between Riemannian manifolds, the Jacobi operator $\mathcal{J}_f$, defined on f^*TN, is

$$\mathcal{J}_f(V) = -\Delta V + \sum_{i=1}^{n} R^N(\nabla_{e_i} f, V)\nabla_{e_i} f,$$

where $(e_i : 1 \leq i \leq n)$ form a fibre-wise orthonormal basis of f^*TN, R^N is the curvature tensor of N, and $\Delta V = \mathrm{tr}\nabla^2 V$ is the Laplacian constructed from the Levi-Civita connections on f^*TN, T^*M and the metric on M.

The significance of the Jacobi operator comes from the following well-known fact, the proof of which we will omit.

Proposition 6.19. *Let $f : (M, g) \to (N, h)$ be a harmonic map, and let $f_t : M \times (-1, 1) \to N$ be a variation of $f = f_0$. If we let $\dot{f} = \frac{\partial f_t}{\partial t}\Big|_{M\times\{0\}}$, we have*

$$\frac{d^2}{dt^2}\frac{1}{2}\int_M \mathrm{tr}_g\,(f^*h)\,d\mathrm{vol}_M = \int_M \langle \dot{f}, \mathcal{J}_f \dot{f}\rangle d\mathrm{vol}_M.$$

We will also use the formula of Micaleff–Moore [54, equation (2.3)], stated below as a proposition.

Proposition 6.20. *If M is a Riemann surface, equipped with a Kähler form ω, and $f : M \to N$ is a harmonic map, then*

$$\mathcal{J}_f(V)\omega = -2i\left(\partial\bar{\partial}_A V + R(V, \partial f) \wedge \overline{\partial f}\right).$$

Finally, we state a result of Sunada [65, Lemma 3.4].

Proposition 6.21. *Let (M, g) be a compact connected Riemannian manifold, X be a non-positively curved symmetric space, and $\rho : \pi_1(M) \to \mathrm{Isom}(X)$ be a representation. Let $f : \tilde{M} \to X$ be a ρ-equivariant harmonic map from the universal cover $\tilde{M}$ of M. If $s \in \Gamma(f^*TX)$ satisfies $\mathcal{J}_f s = 0$, then $f_s(x) = \mathrm{Exp}_{f(x)}s(x)$ is harmonic and ρ-equivariant.*

Note that [65, Lemma 3.4] is only stated when the image of ρ acts freely and properly discontinuously on X, but the exact same proof shows the equivariant version in Proposition 6.21.

Corollary 6.22. *Let S be a Riemann surface, $\rho : \pi_1(S) \to \mathrm{GL}(n, \mathbb{C})$ be an irreducible representation, and $f : S \to X_n$ be a harmonic ρ-equivariant map. Then $\ker \mathcal{J}_f$ is generated by the constant vector field whose value at $s \in S$ is $\mathrm{id} \cdot f(s)$, where $\mathrm{id} \in \mathrm{End}(\mathbb{C}^n) \cong \mathfrak{gl}(n, \mathbb{C})$.*

Remark 6.23. Note that since $\mathcal{J}_f$ is elliptic, it does not matter which function space we are referring to, since by elliptic regularity functions in $\ker \mathcal{J}_f$ are automatically smooth.

Proof of Corollary 6.22. Since ρ is irreducible, the flat bundle associated to ρ has a unique harmonic metric up to scaling. It follows that the harmonic map f is unique up to global translation by $\lambda\mathrm{id}$ for $\lambda \in \mathbb{R}^*$. Therefore by Proposition 6.21, if s is in $\ker \mathcal{J}_f$, then it must be a scalar multiple of $\mathrm{id} \cdot f(s)$. Conversely, translating f by $\lambda\mathrm{id}$ gives harmonic maps, so $\mathrm{id} \cdot f$ must be a Jacobi field. □

6.4.2 Dictionary between the harmonic map and the harmonic metric on a flat bundle

In this subsection, we transport the results of Proposition 6.14 and of the previous subsection to equations for the first-order variation of Higgs(ρ, J), as we vary J.

Let E be a vector bundle over Σ_g. Pick a point $x \in \Sigma_g$ and fix one of its lifts $\tilde{x} \in \tilde{\Sigma}_g$, and a frame $e_1, e_2, ..., e_n \in E_x$. Denote by $\mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$ the space of homomorphisms $\pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$. Note that this is not the same as the character variety $\mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$, since we consider homomorphisms differing in a conjugacy as distinct.

We define the extended holonomy map

$$\mathrm{EHol} : \left\{ (D, h) : \begin{array}{c} D \text{ flat connection on } E \\ h \text{ Hermitian metric on } E \end{array} \right\} \longrightarrow \left\{ (\rho, f) : \begin{array}{c} \rho \in \mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) \\ f : \tilde{\Sigma}_g \to X_n \ \rho\text{-equivariant} \end{array} \right\}$$

as follows. Given a Hermitian metric h and a flat connection D on E, let ρ be the holonomy of D with respect to the frame $e_1, e_2, ..., e_n$. Define f as follows: for $\tilde{y} \in \tilde{\Sigma}_g$ lying over $y \in \Sigma_g$,

$$f(\tilde{y}) = (h(\hat{e}_i, \hat{e}_j))_{1 \leq i,j \leq n},$$

where $(\hat{e}_i : 1 \leq i \leq n)$ is the D-parallel transport of the frame $(e_i : 1 \leq i \leq n)$ along the projection to Σ_g of a path from $\tilde{x}$ to $\tilde{y}$. Then we set $\mathrm{EHol}(D, h) = (\rho, f)$. It is easy to see that this is a bijection.

Given a harmonic metric h and a volume form ω on Σ_g, we may consider the tension field $\tau(f)\omega$ of the associated ρ-equivariant map $f : \tilde{\Sigma}_g \to X_n$. This is a f^*TX_n-valued 2-form, so when contracted with the Maurer–Cartan form $\omega^{\mathfrak{p}}$, we get a 2-form taking values in the equivariant pullback of $\mathrm{GL}(n, \mathbb{C}) \times_{U(n)} \mathfrak{p}$, which is precisely the bundle of self-adjoint endomorphisms $\mathfrak{p}(E)$. We denote this 2-form $\tau(h)$ and call it the tension field of the metric h. Note that we have introduced a volume form ω on Σ_g to remove the dependence of the tension field on the background metric on Σ_g. We similarly contract with $\omega^{\mathfrak{p}}$ the operator $\mathcal{J}_f\omega$, to get a second order differential operator $\mathcal{J}_h : C^\infty(\mathfrak{p}(E)) \to \Omega^2(\mathfrak{p}(E))$.

Given any Hermitian metric on E, we can decompose any connection $D = \nabla + \Psi$, where ∇ is a unitary connection on E, and Ψ is self-adjoint $\mathrm{End}(E)$-valued 1-form with respect to h. Explicitly,

$$h(s, \Psi t) = -\frac{1}{2}(Dh)(s, t), \nabla = D - \Psi.$$

Here Ψ will represent the derivative of the associated ρ-equivariant map f, since $\omega^{\mathfrak{p}}(df) = -2\Psi$.

Claim 6.24. When h is a Hermitian metric, we have

$$\tau(h) = 4id^{\nabla}\left(\Psi^{1,0}\right).$$

Moreover, when h is harmonic, we have

$$\mathcal{J}_h = -2i\left(\partial\bar{\partial} - [[\cdot, \Psi^{1,0}], \Psi^{0,1}]\right).$$

Proof. The first equality follows immediately from $\tau(f)\omega = -2i\bar{\partial}\partial f$ and $\omega^{\mathfrak{p}}(df) = -2\Psi$. The second is equivalent to Proposition 6.20. $\square$

6.4.2.1 Variation of the Higgs field

In this subsection, we give equations for the first-order variation of a Higgs field, as we vary the underlying Riemann surface. For convenience, we assume the smoothness result Theorem 6.26.

We first remind the reader that given a harmonic map $f : \left(\tilde{\Sigma}_g, J\right) \to X_n$, the Higgs bundle consists of the bundle $E^{\mathbb{C}}$ which is the equivariant pullback of $\mathrm{GL}(n,\mathbb{C}) \times_{U(n)} \mathfrak{p}^{\mathbb{C}}$ by f, and of the Higgs field

$$\phi = -\frac{1}{2}\omega^{\mathfrak{p}}(\partial f).$$

Note that $\omega^{\mathfrak{p}}(df)$ is by definition self-adjoint, so comparing $(0,1)$-parts of $\omega^{\mathfrak{p}}(df)$ and $\left(\omega^{\mathfrak{p}}(df)\right)^*$, we see that

$$\phi^{*h} = -\frac{1}{2}\omega^{\mathfrak{p}}(\bar{\partial} f).$$

The harmonic metric on $(E^{\mathbb{C}}, \phi)$ is given by the pullback metric on $E^{\mathbb{C}}$, and hence the connection of the flat bundle is given by

$$D = \nabla + \phi + \phi^{*h}.$$

Proposition 6.25. *Let $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n,\mathbb{C})$ be an irreducible representation. Let S_t be a smooth path of Riemann surfaces based at $S_0 = S \in \mathcal{T}_g$, in the direction of $\mu \in \Omega^{-1,1}(S)$. Then there exists a path $(E, \bar{\partial}_t, \phi_t)$ of Higgs bundles over S_t, all with the same harmonic metric h, such that $\nabla_{\bar{\partial}_t,h} + \phi_t + \phi_t^{*h}$ has holonomy ρ, and*

$$\dot{\nabla}_{\bar{\partial},h} = -[V, \phi + \phi^{*h}],$$
$$\dot{\phi} = \partial V + \mu\phi - \bar{\mu}\phi^{*h},$$

where V is the solution to the equation

$$d\left(\partial V + \mu\phi - \bar{\mu}\phi^{*h}\right) = [[V, \phi^{*h}], \phi].$$

Proof. This is just Proposition 6.14 in a different guise. By Theorem 6.26, we get a smooth path of equivariant harmonic maps $f_t : \tilde{S}_t \to X_n$. By Proposition 6.14, we see that after identifying the pullback bundles appropriately, the metric on $f_t^* T X_n$ is constant, and the connection is varying by

$$\dot{\nabla} = R(\dot{f}, df).$$

We also have

$$\frac{d}{dt}(\partial f) = \frac{d}{dt}\left(df \circ \frac{\mathrm{id} - iJ}{2}\right) = \partial \dot{f} + \mu \partial f - \bar{\mu} \bar{\partial} f. \tag{6.4.1}$$

Using the Maurer–Cartan form, we get a smooth path of Higgs bundles with the same harmonic metric $(E, \bar{\partial}_t, \phi_t)$, with the Chern connection ∇_t, such that

$$\dot{\nabla} = -[\dot{f}, \phi + \phi^{*h}].$$

By applying the Maurer–Cartan form to (6.4.1), we see that

$$\dot{\phi} = \partial \dot{f} + \mu \phi - \bar{\mu} \phi^{*h}.$$

Doing the same thing to the first equation in Proposition 6.14, we get

$$d\dot{\phi} - [[\dot{f}, \phi^{*h}], \phi] = 0.$$

□

6.4.3 Main smoothness result

We now show the main result of this section: the smooth dependence of the harmonic map into $X_n = \mathrm{GL}(n, \mathbb{C})/\mathrm{U}(n)$ on the complex structure on Σ_g and on the representation $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$. We remind the reader that by $\mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$ we denote the space of representations $\pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$, and that by $\mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$ we denote the corresponding character variety.

Theorem 6.26. *Given an irreducible representation $\rho_0 : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ and an almost complex structure J_0 on Σ_g, there exists a smooth map*

$$f : U \times \tilde{\Sigma}_g \to X_n$$

where U is a neighbourhood of (ρ_0, J_0) in $\mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) \times \mathcal{T}_g$, such that for $(\rho, J) \in U$, the map $f(\rho, J, -)$ is a ρ-equivariant J-harmonic map $\tilde{\Sigma}_g \to X_n$.

Proof. Let E be a complex vector bundle of rank n over Σ_g. We fix a fibre of E and a frame of that fibre, as in §6.4.2. Let D_0 be the flat connection on E with holonomy ρ_0, as in the definition of EHol, and let h_0 be the corresponding harmonic metric.

We abuse notation slightly and denote by $0 \in \mathcal{T}_g$ the point corresponding to J_0. We extend J_0 to a family of almost complex structures J_t on Σ_g, depending real analytically on $t \in \mathcal{T}_g$, such that (Σ_g, J_t) represents the marked Riemann surface given by t.

We let $\mathrm{Met}(E)$ be the cone of Hermitian metrics on E, and define

$$\tau : \mathcal{F} \times \mathcal{T}_g \times \mathrm{Met}(E) \to \Omega^2(\mathfrak{p}(E)),$$

where $\mathcal{F} = \{B \in \Omega^1(\mathrm{End}(E)) : D_0 B + B \wedge B = 0\}$ is the space of flat connections $D_0 + B$, by

$$\tau(B, t, h) = (D_0 + B - \Psi)\left(\Psi \circ \frac{\mathrm{id} - iJ_t}{2}\right),$$

where

$$h(s, \Psi t) = -\frac{1}{2}((D_0 + B)h)(s, t).$$

From Claim 6.24, we see that $\tau(B, t, h) = 0$ if and only if $\mathrm{EHol}(D_0 + B, h) = (\rho, f)$, where f is a ρ-equivariant harmonic map. It therefore suffices to construct a smooth $h = h(B, t)$ in a neighbourhood of $(0, 0) \in \mathcal{F} \times \mathcal{T}_g$ such that $\tau(B, t, h(B, t)) = 0$.

We extend the definition of τ to the space $\mathrm{Met}(E)_{k+\alpha+2}$ of $C^{k+\alpha+2}$ Hermitian metrics, so that its range is the space $\Omega^2(\mathfrak{p}(E))_{k+\alpha}$ of $C^{k+\alpha}$ forms. These are now Banach manifolds, and we will prove the theorem by appealing to the Banach implicit function theorem.

We first observe that τ is a smooth map between Banach manifolds, and that

$$\frac{\partial \tau}{\partial h} = \frac{\partial^2 \mathrm{E}}{\partial h^2},$$

where E is the energy of the harmonic map associated to h. As can be seen from Proposition 6.19, the derivative $\frac{\partial \tau}{\partial h}$ is the Jacobi operator $\mathcal{J}_{h_0}$ (this is analogous to [18, Lemma 2.6]).

We will apply the Banach manifold implicit function theorem to the following map

$$F : \mathcal{F} \times \mathcal{T}_g \times \mathrm{Met}(E)_{k+\alpha+2} \to \Omega^2(\mathfrak{p}(E))_{k+\alpha}$$

by $F(B,t,h) = \tau(B,t,h) + \log\left|\det(h_x)\right| \frac{\mathrm{id}_E}{n}\omega$. Here ω is a volume form on Σ_g, and by abuse of notation, we refer to the matrix of the metric h_x with respect to $(e_i : 1 \leq i \leq n)$ by h_x as well. Note that F is smooth, and that

$$\frac{\partial F}{\partial h}(\dot{h}) = \mathcal{J}_{h_0}(\dot{h}) + \mathrm{Re}\ \mathrm{tr}\left(h_x^{-1}\dot{h}_x\right)\frac{\mathrm{id}_E}{n}\omega. \tag{6.4.2}$$

Because $\mathcal{J}_{h_0}$ is self-adjoint, its image is the orthogonal complement to its kernel, and hence by Corollary 6.22 (and Claim 6.24),

$$\mathrm{im}(\mathcal{J}_{h_0}) = \left\{\alpha \in \Omega^2(\mathfrak{p}(E))_{k+\alpha} : \mathrm{Re}\int_{\Sigma_g} \mathrm{tr}(\alpha) = 0\right\}. \tag{6.4.3}$$

Therefore if $\frac{\partial F}{\partial h}(\dot{h}) = 0$, integrating and taking the real part of the trace of (6.4.2), we see that $\mathcal{J}_{h_0}(\dot{h}) = 0$ and $\mathrm{Re}\ \mathrm{tr}(h_x^{-1}\dot{h}_x) = 0$. The first condition forces $h^{-1}\dot{h}$ to be a multiple of the identity, and the second to vanish. Thus $\frac{\partial F}{\partial h}$ is injective.

We now show surjectivity of $\frac{\partial F}{\partial h}$. Let $\theta \in \Omega^2(\mathfrak{p}(E))_{k+\alpha}$ be arbitrary, and split

$$\theta = \left(\mathrm{Re}\int_{\Sigma_g} \mathrm{tr}(\theta)\right)\frac{\mathrm{id}_E}{n}\omega + \xi,$$

so that ξ satisfies the condition on the right-hand side of (6.4.3). Thus $\xi = \mathcal{J}_{h_0}(\dot{h})$ for some $\dot{h}$. We now replace $\dot{h}$ by $\dot{h} + \lambda h$ for some $\lambda \in \mathbb{R}$, so that

$$\mathrm{Re}\ \mathrm{tr}\left(h_x^{-1}\dot{h}_x\right) = \mathrm{Re}\int_{\Sigma_g} \mathrm{tr}(\theta).$$

This does not change $\mathcal{J}_{h_0}(\dot{h})$, and hence $\frac{\partial F}{\partial h}(\dot{h}) = \theta$. This concludes the proof that $\frac{\partial F}{\partial h}$ is surjective.

Since $\frac{\partial F}{\partial h}$ is a bijective continuous linear map, by the open mapping theorem it is an isomorphism of Banach spaces. Therefore by the Banach manifold implicit function theorem, we can construct a smooth function $h = h(B,t)$ over U such that $F(B,t,h(B,t)) = 0$.

By Stokes' theorem, as $\tau(h) = d^{\nabla}\Psi^{1,0}$, we have

$$\int_{\Sigma_g} \mathrm{tr}\ \tau(B,t,h) = 0.$$

Therefore $F = 0$ implies, after taking the trace and integrating, that $\mathrm{Re}\ \mathrm{tr}\left(h_x^{-1}\dot{h}\right) = 0$ and that $\tau = 0$. Therefore $h = h(\rho,t)$ is a harmonic metric. $\square$

6.4.4 Consequences of the main smoothness result

6.4.4.1 Proof of Theorem 6.10

Let $\rho = \rho_1 \oplus \rho_2 \oplus ... \oplus \rho_k$ be the decomposition of ρ into its irreducible components where $\rho_i : \pi_1(\Sigma_g) \to \mathrm{GL}(V_i)$ such that $V_1 \oplus V_2 \oplus ... \oplus V_k = \mathbb{C}^n$. We fix a marked Riemann surface $S \in \mathcal{T}_g$. Then each $\rho_i : \pi_1(\Sigma_g) \to \mathrm{GL}(V_i)$ is irreducible, so applying Theorem 6.26, there exists a neighbourhood $U \subset \mathcal{T}_g$ containing S, along with smooth maps

$$f_i : U \times \tilde{\Sigma}_g \to X_{n_i},$$

such that $f_i(J, -)$ is a ρ_i-equivariant J-harmonic map. We combine these maps into a smooth map

$$f = f_1 \times f_2 \times ... \times f_k : U \times \tilde{\Sigma}_g \to \prod_{i=1}^{k} \mathrm{GL}(V_i)/\mathrm{U}(V_i) \subset \mathrm{GL}(n, \mathbb{C})/\mathrm{U}(n) = X_n.$$

The map f has the property that for an arbitrary J, the map $f(J, -)$ is a ρ-equivariant J-harmonic map $\tilde{\Sigma}_g \to X_n$.

Now for any Beltrami form μ on S, we can pick a disk $\iota : \mathbb{D} \to \mathcal{T}_g$ in Teichmüller space with $\iota(0) = S$ and $\iota_z(0) = \mu$, and consider the complex disk of harmonic maps $(f(\iota(t), -) : t \in \mathbb{D})$ based at S with direction μ. The characterization of $K_\rho(S)$ in Theorem 6.10 now follows immediately from Theorem 6.8.

6.4.4.2 Smoothness of $\mathcal{R}$ on the open set of irreducible representations

Note that we will only ever use smoothness of $\mathcal{R}_\rho$ for some fixed irreducible $\rho \in \mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$, but for completeness here we show smoothness of $\mathcal{R}$ over the set of irreducible representations. We first remind the reader that the character variety $\mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$ is smooth near any irreducible representation, so the question of smoothness of $\mathcal{R}$ is well-defined.

Fix an irreducible representation $\rho_0 : \pi_1(\Sigma_g) \to \mathrm{GL}(n, \mathbb{C})$ and a marked Riemann surface $S_0 \in \mathcal{T}_g$. Let $U \subseteq \mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) \times \mathcal{T}_g$ be the open set containing (ρ_0, S_0) from Theorem 6.26. Let $f : U \times \tilde{\Sigma}_g \to X_n$ be the map from Theorem 6.26. Pulling back the tangent bundle of X_n via f, we may construct Higgs bundles $(E, \bar{\partial}_{\rho,S}, \phi_{\rho,S})$ each equipped with a harmonic metric $h_{\rho,S}$, as in §2.3.2, that vary smoothly with $(\rho, S) \in U$, such that $(E, \bar{\partial}_{\rho,S}, \phi_{\rho,S}) = \mathrm{Higgs}(\rho, S)$.

We let $\tilde{\mathcal{R}}(\rho, S) \in \mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$ be the holonomy of $\nabla^{\rho,S} + i\phi_{\rho,S} - i\phi_{\rho,S}^{*h_{\rho,S}}$, where $\nabla^{\rho,S}$ is the Chern connection on $(E, \bar{\partial}_{\rho,S}, h_{\rho,S})$. Here we use the identifications from §6.4.2 to get a representation, and not just an element of the character variety. Note that $\tilde{\mathcal{R}}(\rho, S)$ represents the Higgs field $i\cdot\mathrm{Higgs}(\rho, S)$. Since the Chern connection depends smoothly on both the Hermitian metric and holomorphic structure on E, we see that

$$\tilde{\mathcal{R}} : U \longrightarrow \mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$$

is smooth. Let $p : \mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) \to \mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$ be the natural quotient map. Then in the diagram

$$\begin{array}{ccc} U & \xrightarrow{\quad p\circ\tilde{\mathcal{R}}\quad} & \mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) \\ \big\downarrow{\scriptstyle p\times \mathrm{id}_{\mathcal{T}_g}} & \nearrow{\scriptstyle \mathcal{R}} & \\ \mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) \times \mathcal{T}_g & & \end{array}$$

the vertical map $p \times \mathrm{id}_{\mathcal{T}_g}$ is a surjective submersion onto a neighbourhood of $(\rho_0, S_0) \in \mathrm{Rep}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C})) \times \mathcal{T}_g$ and the map $p \circ \tilde{\mathcal{R}} = \mathcal{R} \circ (p \times \mathrm{id}_{\mathcal{T}_g})$ is smooth. Therefore $\mathcal{R}$ is smooth in some neighbourhood of (ρ_0, S_0).

6.4.5 Moduli of stable Higgs bundles

Here we construct the smooth structure on the moduli space of stable Higgs bundles over a varying Riemann surface. Note that this was carried out by Fan in the general case of polystable Higgs bundles [21] over a single Riemann surface. We follows his analysis, though our case is significantly simpler due to the fact that stable Higgs bundles have trivial stabilizer in the relevant gauge groups.

6.4.5.1 Preliminary definitions

We fix a genus $g \geq 2$, and let J_t be an almost complex structure on Σ_g that represents the point $t \in \mathcal{T}_g$, and such that J_t depends real analytically on t. We also fix a smooth complex vector bundle $E \to \Sigma_g$, with a smooth Hermitian metric h. We define the group of gauge transformations

$$\mathcal{G} = \{T \in \Gamma(\Sigma_g, \mathrm{GL}(E)) : T \text{ is fibre-wise unitary}\}.$$

We also denote by $\mathfrak{u}(E)$ the bundle of skew-adjoint endomorphisms of E, and let $\mathfrak{p}(E)$ be the bundle of self-adjoint endomorphisms of E. Let

$$\mathcal{C} = \{(t, A, \Phi) : t \in \mathcal{T}_g, \Phi \in \Omega^1(\mathfrak{p}(E)) \text{ and } A \text{ is a unitary connection on } E\},$$

$$\mathcal{B} = \left\{ (t, A, \Phi) \in \mathcal{C} : \begin{array}{c} d_A^{0,1}\Phi^{1,0} = 0 \text{ with respect to the complex structure } J_t \\ d_A + \Phi \text{ is a flat connection} \end{array} \right\}.$$

We define $\mathcal{B}^{\mathrm{s}}$ as the set of triples in $\mathcal{B}$ such that $(E, d_A^{0,1}, \Phi^{1,0})$ is a stable Higgs bundle over (Σ_g, J_t). Note that given a stable Higgs bundle $(E, \bar{\partial}^E, \phi)$ over (Σ_g, J_t) for which h is a harmonic metric, the corresponding element in $\mathcal{B}^{\mathrm{s}}$ is given by $(t, \nabla_{\bar{\partial}^E,h}, \phi + \phi^{*h})$, where $\nabla_{\bar{\partial}^E,h}$ is the Chern connection of h. The group of gauge transformations $\mathcal{G}$ admits a natural action on $\mathcal{C}$ that preserves $\mathcal{B}$ and $\mathcal{B}^{\mathrm{s}}$, so we set

$$\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}} = \mathcal{B}^{\mathrm{s}}/\mathcal{G}.$$

6.4.5.2 Deformation complexes and local slices

We are now ready to define local slices for the action of $\mathcal{G}$ on $\mathcal{C}$ near points in $\mathcal{B}^{\mathrm{s}}$. The proof mirrors that of Hitchin's original paper [33].

Define the deformation complex $C_{\mathrm{Hit}}(t, A, \Phi)$

$$\Omega^0(\mathfrak{u}(E)) \xrightarrow{d_1} \Omega^1(\mathfrak{u}(E)) \oplus \Omega^1(\mathfrak{p}(E)) \xrightarrow{d_2} \Omega^2(\mathfrak{u}(E)) \oplus \Omega^2(\mathrm{End}(E)),$$

where

$$d_1\psi = (d_A\psi, [\Phi, \psi]),$$
$$d_2(B, \Psi) = (d_A B + [\Phi, \Psi], d_A^{0,1}\Psi^{1,0} + [B^{0,1}, \Phi]).$$

As shown in [33, pp. 85, 86], this is an elliptic complex that near a stable Higgs bundle has trivial zeroth and second cohomology. We denote by $\mathcal{C}_k$ (resp. $\mathcal{B}^{\mathrm{s}}_k$) the Sobolev space L^2_k of unitary connections A and fields Φ satisfying the same conditions as in the definition of $\mathcal{C}$ (resp. $\mathcal{B}^{\mathrm{s}}$). We similarly denote by $\mathcal{G}_k$ the completion of $\mathcal{G}$ in the L^2_k norm. Then for $k > 2$, by standard theory [55],

$$\mathcal{C}_k/\mathcal{G}_{k+1}$$

is a smooth Banach manifold near a point in $\mathcal{B}^{\mathrm{s}}$. The local slice for this action is precisely $\ker d_1^*$.

We define

$$\mathrm{T} : \mathcal{C} \longrightarrow \Omega^2(\mathfrak{u}(E)) \oplus \Omega^2(\mathrm{End}(E))$$
$$(t, A, \Phi) \longrightarrow (d_A^2 + [\Phi, \Phi], d_A^{0,1}\Phi^{1,0})$$

so that $\mathcal{B} = \mathrm{T}^{-1}(0)$. Note that T is real analytic in t, and as shown by Hitchin [33, pp. 86, 87],

$$D\mathrm{T}_{(t_0,A_0,\Phi_0)}(0, B, \Psi) = d_2(B, \Psi).$$

In particular, $D\mathrm{T}_{(t_0,A_0,\Phi_0)}$ is surjective. By the Banach manifold implicit function theorem, we see that $\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}}$ is a smooth manifold.

We also provide a separate description of $T\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}}$ that will be useful in the rest of the paper. Define $\mathcal{H}^s$ to be the set of triples $(t, \bar{\partial}_A, \Phi)$ where $t \in \mathcal{T}_g, \Phi \in \Omega^1(\mathfrak{p}(E))$, and $\bar{\partial}_A$ is a $(0,1)$-connection (with respect to J_t) on E, such that $(E, \bar{\partial}_A, \Phi^{1,0})$ is a stable Higgs bundle. The set $\mathcal{H}^s$ admits a natural action by the gauge group $\Gamma(\Sigma_g, \mathrm{GL}(E))$, and we define the quotient

$$\mathcal{M}^{\mathrm{s}}_{\mathrm{Higgs}} = \frac{\mathcal{H}^{\mathrm{s}}}{\Gamma(\Sigma_g, \mathrm{GL}(E))}.$$

The non-abelian Hodge correspondence provides a homeomorphism

$$\mathcal{M}^s_{\mathrm{Hit}} \to \mathcal{M}^s_{\mathrm{Higgs}},$$

by taking (t, A, Φ) to $(t, d_A^{0,1}, \Phi)$. We use this homeomorphism to give $\mathcal{M}^s_{\mathrm{Higgs}}$ a smooth structure. We define another deformation complex $C_{\mathrm{Higgs}}(t, \bar{\partial}_A, \Phi)$

$$\Omega^0(\mathrm{End}(E)) \overset{\bar{\partial}_A + \Phi^{1,0}}{\longrightarrow} \Omega^1(\mathrm{End}(E)) \overset{\bar{\partial}_A + \Phi^{1,0}}{\longrightarrow} \Omega^2(\mathrm{End}(E)).$$

It is easy to see that the natural isomorphism

$$\Omega^1(\mathfrak{u}(E)) \oplus \Omega^1(\mathfrak{p}(E)) \to \Omega^1(\mathrm{End}(E))$$

defines an isomorphism $H^1(C_{\mathrm{Hit}}) \cong H^1(C_{\mathrm{Higgs}})$ (the reader can also consult [21, §2]). This isomorphism shows that the natural map

$$\frac{T_{(E,\bar{\partial}_{A_0},\Phi_0)}\mathcal{H}^s}{(\bar{\partial}_{A_0} + \Phi_0^{1,0})\Omega^0(\mathrm{End}(E))} \to T_{(t_0,\bar{\partial}_{A_0},\Phi_0)}\mathcal{M}^{\mathrm{s}}_{\mathrm{Higgs}} \cong T_{(t_0,A_0,\Phi_0)}\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}} \tag{6.4.4}$$

is an isomorphism.

6.5 Proof of Theorem C.2

Here we prove Theorem C.2, that is $K_\rho(S) = \ker D_S\mathcal{R}_\rho$ for some fixed marked Riemann surface S. We have a description of $K_\rho(S)$ from Theorem 6.10, and hence it only remains to describe $\ker D_S\mathcal{R}_\rho$. This will follow from §6.4.5, in particular from the isomorphism (6.4.4).

We first give a preliminary description of $\ker D\mathcal{R}_\rho$ in the following proposition.

Proposition 6.27. *Let $V(\phi)$ be the solution to the equation from Proposition 6.25, and set $V = V(\phi) + iV(i\phi)$. Then $(\mathcal{R}_\rho)_*\mu = 0$ if and only if there exists a section A of* $\mathrm{End}(E)$ *such that*

$$[\phi^{*h}, V] = \bar{\partial} A,$$
$$\partial V - 2\bar{\mu}\phi^{*h} = [\phi, A].$$

Proof. We consider two first order variations of $(d_A, i\phi)$ in $\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}}$, given by

$$-[V(\phi), \phi^{*h} + \phi], i\partial V(\phi) + i\mu\phi - i\bar{\mu}\phi^{*h},$$
$$-[V(i\phi), i\phi - i\phi^{*h}], \partial V(i\phi) + i\mu\phi + i\bar{\mu}\phi^{*h}.$$

By Proposition 6.25, the first is tangent to the path $i \cdot \mathrm{Higgs}(\rho, J_t)$, and the second one is tangent to $\mathrm{Higgs}(\mathcal{R}_\rho(J_0), J_t)$. Subtracting, we get a first order variation

$$-[V, \phi^{*h}] - [V^{*h}, \phi], i\left(\partial V - 2\bar{\mu}\phi^{*h}\right). \qquad (6.5.1)$$

Claim 6.28. We have $(\mathcal{R}_\rho)_*\mu = 0$ if and only if the variation (6.5.1) vanishes as a tangent vector to $\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}}$.

Proof. Suppose first that $(\mathcal{R}_\rho)_*\mu = 0$. Let J_t be a smooth path of almost complex structures on Σ_g, where we identify (Σ_g, J_0) with S, tangent to $\mu \in T_S\mathcal{T}_g$. Let ρ_t be a path in $\mathrm{Hom}(\pi_1(\Sigma_g), \mathrm{GL}(n, \mathbb{C}))$ that is a lift of $\mathcal{R}_\rho(J_t)$ such that $\left(\frac{d}{dt}\right)_{t=0} \rho_t = 0$. By Theorem 6.26, there exists a smooth family of maps $f_{t,s} : \tilde{\Sigma}_g \to X_n$ such that $f_{t,s}$ is a ρ_t-equivariant J_s-harmonic map. Moreover, since $\left(\frac{d}{dt}\right)_{t=0} \rho_t = 0$, we get

$$\left(\frac{\partial}{\partial t}\right)_{t=s=0} f = 0.$$

It follows that the paths $f(t,t)$ and $f(0,t)$ agree to first order. Thus $i \cdot \mathrm{Higgs}(\rho, J_t) = \mathrm{Higgs}(\mathcal{R}_\rho(J_t), J_t)$ and $\mathrm{Higgs}(\mathcal{R}_\rho(J_0), J_t)$ agree to first order (in $\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}}$) at $t = 0$. In particular, the variation (6.5.1) vanishes.

Conversely, assume that (6.5.1) vanishes as an element of $T\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}}$. Then the paths $\mathrm{Higgs}(\mathcal{R}_\rho(J_t), J_t) = i \cdot \mathrm{Higgs}(\rho, J_t)$ and $\mathrm{Higgs}(\mathcal{R}_\rho(J_0), J_t)$ agree to first order at $t = 0$. But the representations $\mathcal{R}_\rho(J_t), \mathcal{R}_\rho(J_0)$ can be recovered as holonomies of $d_{A_t} + \phi_t + \phi_t^{*h}$, and are hence smooth functions on $\mathcal{M}^{\mathrm{s}}_{\mathrm{Hit}}$. Thus $\mathcal{R}_\rho(J_t)$ and $\mathcal{R}_\rho(J_0)$ agree to first order at $t = 0$. This is exactly equivalent to $(\mathcal{R}_\rho)_*\mu = 0$. □

By the isomorphism (6.4.4) from §6.4.5, the variation (6.5.1) vanishes if and only if

$$-[V, \phi^{*h}] = \bar{\partial} A,$$
$$i\left(\partial V - 2\bar{\mu}\phi^{*h}\right) = i[\phi, A],$$

for some section A of $\mathrm{End}(E)$. This concludes the proof of the proposition. □

6.5.1 Forward containment

Suppose that E_ρ has vanishing Laplacian in the direction μ. Let (E, ϕ) be the corresponding Higgs bundle over S, and h be the harmonic metric associated to ρ. By Theorem 6.8, we see that there exists a section ξ of $\mathrm{End}(E)$, such that

$$\mu\phi = \bar{\partial}\xi \text{ and } [\xi, \phi] = 0. \tag{6.5.2}$$

From Proposition 6.25

$$d\left(\partial V(\phi) + \mu\phi - \bar{\mu}\phi^{*h}\right) = [[V(\phi), \phi^{*h}], \phi],$$
$$d\left(\partial V(i\phi) + i\mu\phi + i\bar{\mu}\phi^{*h}\right) = [[V(i\phi), \phi^{*h}], \phi].$$

Combining these two equations, and setting $V = V(\phi) + iV(i\phi)$, we get

$$\bar{\partial}\left(\partial V - 2\bar{\mu}\phi^{*h}\right) = [[V, \phi^{*h}], \phi].$$

Taking the adjoint with respect to h, we get

$$\partial\bar{\partial}\left(V^{*h} - 2\xi\right) = -[[\phi, V^{*h}], \phi^{*h}].$$

Since $[\xi, \phi] = 0$, we see that $\partial\bar{\partial}(V^{*h} - 2\xi) - [[V^{*h} - 2\xi, \phi] \wedge \phi^{*h}] = 0$. Note that this is the exact equation we got for V in the proof of Theorem 6.8. From the Bochner argument in the last paragraph of the proof of Theorem 6.8 and Claim 6.16, we get

$$\bar{\partial}(V^{*h} - 2\xi) = 0 \text{ and } [V^{*h}, \phi] = 0.$$

But $[V^{*h}, \phi] = 0$, so $[V, \phi^{*h}] = 0$, and $\bar{\partial}V^{*h} = 2\bar{\partial}\xi = 2\mu\phi$, so $\partial V^{*h} = 2\bar{\mu}\phi^{*h}$. Thus $(\mathcal{R}_\rho)_*\mu = 0$ by Proposition 6.27.

6.5.2 Backward containment

Let (E, ϕ) be the stable Higgs bundle that corresponds to ρ, with the harmonic metric h and Chern connection d_A, and let μ be a smooth Beltrami form on S such that $(\mathcal{R}_\rho)_*\mu = 0$. By Proposition 6.27, there exists a section $A \in \Gamma(\mathrm{End}(E))$, with the property that

$$\left(-[V, \phi^{*h}], i\left(\partial V - 2\bar{\mu}\phi^{*h}\right)\right) = \left(\bar{\partial}A, i[\phi, A]\right).$$

Therefore

$$\begin{aligned}\partial\bar{\partial}A &= -[\partial V, \phi^{*h}] = -[[\phi, A] + 2\bar{\mu}\phi^{*h}, \phi^{*h}] \\ &= -[[\phi, A], \phi^{*h}].\end{aligned}$$

Thus $\partial\bar{\partial}A - [[A, \phi], \phi^{*h}] = 0$, so by a Bochner argument and Claim 6.16, we get $\bar{\partial}A = 0$ and $[A, \phi] = 0$. Therefore

$$[\phi, V^{*h}] = 0 \text{ and } \bar{\partial}V^{*h} = 2\mu\phi.$$

By Theorem 6.8, the Laplacian of E_ρ in the direction μ vanishes, so we are done.

6.6 Relationship to the critical points of the Hitchin fibration

In this section, we show Theorem 6.2. This follows easily from Proposition 6.29, that we show below.

We first introduce some notation. Given a Riemann surface S and a Beltrami form μ on S, we define the function

$$\begin{aligned}F_\mu : \mathcal{M}^{\mathrm{ps}}_{\mathrm{Higgs}}(S) &\longrightarrow \mathbb{C} \\ (E, \phi) &\longrightarrow \int_S \mathrm{tr}(\phi \wedge \mu\phi).\end{aligned}$$

Note that F_μ factors through the Hitchin fibration H.

Proposition 6.29. *Let S be a Riemann surface of genus g, and μ be a Beltrami form on S. If there exists a section ξ of $\mathrm{End}(E)$ such that*

$$\mu\phi = \bar{\partial}\xi \text{ and } [\phi, \xi] = 0,$$

then (E, ϕ) is a critical point for F_μ.

Proof. We pick a smooth path $(E, \bar{\partial}_{E,t}, \phi_t) = (E, \bar{\partial}_E + t\dot{A} + O(t^2), \phi + t\dot{\phi} + O(t^2))$ of Higgs bundles starting from $(E, \bar{\partial}_E, \phi)$. Then taking derivatives of $\bar{\partial}_{E,t}\phi_t = 0$, we see that

$$\bar{\partial}_E\dot{\phi} + [\dot{A}, \phi] = 0. \tag{6.6.1}$$

Taking derivatives of $F_\mu(E, \bar{\partial}_{E,t}, \phi_t)$, we get

$$\dot{F}_\mu = 2\int_S \mathrm{tr}(\dot{\phi} \wedge \mu\phi).$$

Assume $\mu\phi = \bar{\partial}\xi$ and $[\phi, \xi] = 0$ for some $\xi \in \Gamma(\mathrm{End}(E))$. Then we have

$$\begin{aligned}\dot{F}_\mu &= 2\int_S \mathrm{tr}(\dot{\phi} \wedge \bar{\partial}\xi) = 2\int_S \mathrm{tr}((\bar{\partial}\dot{\phi})\xi - \bar{\partial}(\dot{\phi}\xi)) \\ &= -2\int_S \mathrm{tr}([\dot{A}, \phi]\xi) = -2\int_S \mathrm{tr}(\dot{A}[\phi, \xi]) = 0.\end{aligned}$$

Here we used Stokes' theorem and (6.6.1) in going from the first to the second line. □

6.6.1 Proof of Theorem 6.2

Suppose now that $d = \dim K_\rho(S)$ and $(E, \phi) = \mathrm{Higgs}(\rho, S)$. Let $\mu_1, ..., \mu_d$ be linearly independent Beltrami forms in $K_\rho(S)$. Then from Proposition 6.29, we see that F_{μ_i} all have critical points at (E, ϕ). Note that

$$(F_{\mu_1}, F_{\mu_2}, ..., F_{\mu_d}) = F \circ H,$$

where $F : \bigoplus_{i=1}^n H^0(S, K_S^i) \to \mathbb{C}^d$ is given by

$$F(\phi_1, \phi_2, ..., \phi_n) = \left(\int_S \mu_i\phi_2 : i = 1, 2, ..., d\right).$$

Note that $\nabla_{(E,\phi)}(F \circ H) = 0$, since all F_{μ_i} have critical points at (E, ϕ). Since the $\{\mu_i : i = 1, 2, ..., d\}$ are linearly independent, F is a linear map of full rank, and thus the rank of $\nabla_{(E,\phi)}H$ is at most $\dim \mathcal{B}(S) - d$.

6.7 The case $n = 1$

Here we show Proposition 6.5. In §6.7.1, we show that ϕ from the statement of Proposition 6.5 corresponds to the Higgs field associated to ρ. In §6.7.2, we derive Proposition 6.5. Characterization of K_ρ follows from Theorem 6.10, while integrability of the distribution K_ρ follows from Theorem C.2.

6.7.1 The Higgs field in rank 1

We remark that a rank 1 Higgs bundle over a Riemann surface S is simply a pair (L, ϕ) consisting of a line bundle L, and a holomorphic 1-form ϕ on S. This follows from the fact that $\mathrm{End}(L) \cong L \otimes L^{-1}$ is canonically isomorphic to the trivial line bundle $\mathcal{O}_S$.

Given a rank 1 Higgs bundle (L, ϕ), the Hitchin equation is

$$F_\nabla = 0,$$

where ∇ is the Chern connection of a Hermitian metric on L. Since ∇ is unitary, its holonomy $\rho_\nabla : \pi_1(\Sigma_g) \to \mathbb{C}^*$ has image in the unitary group $\mathrm{U}(1) = S^1 \leq \mathbb{C}^*$. For $\rho : \pi_1(\Sigma_g) \to \mathbb{C}^*$, the flat connection on $(L, \phi) = \mathrm{Higgs}(\rho, S)$ with holonomy ρ is $\nabla + \phi + \phi^* = \nabla + \phi + \bar{\phi}$, and hence

$$\rho(\gamma) = \rho_\nabla(\gamma) e^{-\int_\gamma \phi + \bar{\phi}} = \rho_\nabla(\gamma) e^{-2\int_\gamma \mathrm{Re}(\phi)}.$$

Therefore $\int_\gamma \mathrm{Re}(\phi) = -\frac{1}{2}\log\big|\rho(\gamma)\big|$.

6.7.2 Proof of Proposition 6.5

Given a rank 1 Higgs bundle $(L, \phi) = \mathrm{Higgs}(\rho, S)$, by Theorem 6.10, the kernel of the Levi form of E_ρ is precisely the space

$$\{\mu \in T_S\mathcal{T}_g : \mu\phi \text{ is } \bar{\partial}\text{-exact}\}.$$

By Serre duality, this is exactly equal to

$$\left\{\mu \in T_S\mathcal{T}_g : \int_S \mu\phi\theta = 0 \text{ for all } \theta \in H^0(K_S)\right\} = \big(\phi \otimes H^0(K_S)\big)^\perp \leq H^0(K_S^2)^\vee.$$

Here $H^0(K_S^2) = \mathrm{QD}(S)$ is the space of holomorphic quadratic differentials on S, and $H^0(K_S^2)^\vee$ denotes its dual. Since $T_S\mathcal{T}_g$ is precisely the dual of $H^0(K_S^2)$, we see that the nullity of the Levi form of E_ρ is equal to $3g - 3 - \dim(\phi \otimes H^0(K_S)) = 2g - 3$.

We now show that the distribution K_ρ on $\mathcal{T}_g$ is integrable. This is a complex distribution of constant dimension. Moreover by Theorem C.2, K_ρ consists of all vectors that are annihilated by $\mathcal{R}_\rho$. Therefore if we write locally in some coordinate system on $\mathrm{Rep}(\pi_1(\Sigma_g), \mathbb{C}^*)$,

$$\mathcal{R}_\rho = (f_1, f_2, ..., f_N)$$

we see that K_ρ is the vanishing locus of $(df_1, df_2, ..., df_N)$. By the Frobenius integrability theorem, since $\dim K_\rho$ is constant, the distribution K_ρ is integrable. Since K_ρ is a complex distribution, the leaves of the resulting foliation are complex submanifolds.

6.8 The case $n \geq 2$

In this section, we show Theorem 6.3.

In §6.8.1, we show that representations in the smooth fibre of the Hitchin fibration over S have $K_\rho(S) = \{0\}$. As described in the outline, from the BNR correspondence [2] and Theorem 6.10, it follows that $K_\rho(S)$ only depends on which Hitchin fibre $\mathrm{Higgs}(\rho, S)$ belongs to. We then construct strictly plurisubharmonic representations in each fibre. This is done by slightly modifying the $\mathrm{SL}(n,\mathbb{R})$ Hitchin section. Note that in the $\mathrm{SL}(n,\mathbb{R})$ Hitchin section, the energy E_ρ is strictly plurisubharmonic by [63], and our result is a slight extension of [63].

In §6.8.2, we construct explicitly stable $\mathrm{SL}(n,\mathbb{C})$-Higgs bundles (E,ϕ) in the nilpotent cone over an arbitrary $S \in \mathcal{T}_g$, for any $g \geq 4, n \geq 2$, such that

$$\dim\{\mu \in T_S\mathcal{T}_g : (6.1.2) \text{ has a solution}\} \geq g - 3.$$

By applying the inverse of the non-abelian Hodge correspondence and Theorem 6.10, this gives representations $\rho : \pi_1(\Sigma_g) \to \mathrm{GL}(n,\mathbb{C})$ such that $\dim K_\rho(S) \geq g - 3$.

6.8.1 Representations in the smooth fibre of the Hitchin fibration

Let $(E,\phi) = \mathrm{Higgs}(\rho, S)$, and S_a be the spectral curve over S associated to $a = H(E,\phi)$. Note that S_a only depends on the fibre $a \in \bigoplus_{i=1}^n H^0(K_S^i)$, and that it is smooth for a generic a, as explained in §6.2.

From the BNR correspondence [2, Proposition 3.6], the bundle E can be written as the pushforward along p of a line bundle L on S_a. Moreover S_a admits a holomorphic 1-form Φ_a, such that ϕ is precisely the pushforward of Φ_a along p. Finally, Φ_a also only depends on a.

At a point $x \in S_a$ that is not a pre-image of a branch point, the characteristic polynomial at $p(x)$ does not have repeated roots. Thus, if $[\xi, \phi] = 0$, then ξ is the pushforward of a function on S_a, in a neighbourhood of any point not lying over a branch point. By continuity of ξ, there is a global function g on S_a such that $\xi = p_* g$.

From Theorem 6.10, it now follows that E_ρ is not strictly plurisubharmonic at S in a direction μ if and only if $\mu\Phi_a$ is $\bar\partial$-exact. However this condition only depends on a, so it suffices to construct a Higgs bundle $\mathrm{Higgs}(\rho, S) = (E,\phi) \in H^{-1}(a)$ such that E_ρ is strictly plurisubharmonic at S. The rest of this subsection is devoted to this construction.

Let $s : \bigoplus_{i=2}^n H^0(S, K_S^i) \to \mathcal{M}_{\mathrm{Higgs}}(S)$ be the Hitchin section. We extend this section to a function $\hat{s} : \bigoplus_{i=1}^n H^0(S, K_S^i) \to \mathcal{M}_{\mathrm{Higgs}}(S)$ by

$$\hat{s}(q_1, q_2, ..., q_n) = q_1 \cdot \mathrm{id}_E + \phi,$$

where $(E, \phi) = s(q_2, q_3..., q_n)$. Note that the characteristic polynomial of $\hat{s}$ is

$$\chi_{\hat{s}}(x) = \det(x\mathrm{id}_E - \hat{s}) = \det((x - q_1)\mathrm{id}_E - s) = \chi_s(x - q_1),$$

Since the coefficients of χ_s run over all possible elements of $\bigoplus_{i=2}^n H^0(S, K_S^i)$, the coefficients of $\chi_{\hat{s}}$ run over all elements of $\bigoplus_{i=1}^n H^0(S, K_S^i)$. The proof is then concluded once the following claim is shown.

Claim 6.30. Let (E, ϕ) be in the image of $\hat{s}$. If $\mu \in \Omega^{-1,1}(S)$ is a Beltrami form such that $\mu\phi = \bar{\partial}\xi$ for some $\xi \in \mathrm{End}(E)$, then μ represents the zero direction in $T_S\mathcal{T}_g$.

Proof. Taking the traceless part of the equation $\mu\phi = \bar{\partial}\xi$, we see that

$$\mu\left(\phi - \mathrm{tr}(\phi)\frac{\mathrm{id}_E}{n}\right) = \bar{\partial}\left(\xi - \mathrm{tr}(\xi)\frac{\mathrm{id}_E}{n}\right).$$

But $\phi - \mathrm{tr}(\phi)\frac{\mathrm{id}_E}{n}$ is in the Hitchin section by construction. In this setting, Slegers [63, Proof of Theorem 1.1, pp. 9] has shown that μ represents the zero direction in $T_S\mathcal{T}_g$. □

6.8.2 Examples of non-strictly plurisubharmonic representations

We assume in this section that $g \geq 4$, and fix an arbitrary Riemann surface $S \in \mathcal{T}_g$. We will construct for all $n \geq 2$ a representation $\rho : \pi_1(\Sigma_g) \to \mathrm{SL}(n, \mathbb{C})$ in the nilpotent cone, such that E_ρ is not strictly plurisubharmonic at S.

We first show the construction when n is odd, and then explain how to modify it in the case of n even.

6.8.2.1 The odd case

We assume first that n is odd. Let $p \in S$ be a generic point, and set $L = \mathcal{O}(p)$. By the geometric Riemann–Roch theorem [1, pp. 12], $h^0(KL^{-1}) = i(p) = g - 2$. Here we denote by h^i the dimension of the i-th cohomology group H^i, and $i(D)$ is the index of specialty of the effective divisor D. Therefore we fix a non-zero element $\psi \in H^0(KL^{-1})$.

We consider the following sequence

$$L_1 = L^{\lfloor \frac{n-1}{2} \rfloor}, L_2 = L^{\lfloor \frac{n-1}{2} \rfloor - 1}, ..., L_n = L^{-\lfloor \frac{n-1}{2} \rfloor}, \tag{6.8.1}$$

so that there are in total n line bundles. Let $E = L_1 \oplus L_2 \oplus ... \oplus L_n$, and $\phi \in H^0(K\mathrm{End}(E))$ be the nilpotent Higgs field that has

$$\phi : L_i \to L_{i+1}K \text{ is multiplication by } \psi,$$

and $\phi(L_n) = 0$.

Claim 6.31. The pair (E, ϕ) is a stable $\mathrm{SL}(n, \mathbb{C})$-Higgs bundle.

Proof. We first observe that $\bigwedge^n E = L_1 L_2 ... L_n \cong \mathcal{O}_S$, and that ϕ is lower triangular in the basis defined by $L_1, L_2, ..., L_n$, and thus traceless. Hence (E, ϕ) is an $\mathrm{SL}(n, \mathbb{C})$-Higgs bundle.

To show stability, suppose that $F \leq E$ is a non-zero ϕ-invariant subbundle. Since ϕ is nilpotent, there exists a maximal $k \geq 0$ such that $\phi^k F \neq 0$. Then $\phi^k F \leq (\ker \phi) K^k = L_n K^k$. Thus $\phi^k F \leq K^k F$ has non-zero intersection with $L_n K^k$, and hence $L_n \leq F$. Repeating the argument with E/L_n in place of E, by induction it follows that for some i, we have

$$F = \bigoplus_{j \geq i} L_j.$$

But unless $F = E$, the degree of F is negative, and hence so is the slope. Since the degree of E is zero, E is stable. □

Let $\mu \in \Omega^{-1,1}(S)$ be an arbitrary Beltrami form, and assume that $\mu\psi = \bar{\partial}\xi$, for some $\xi \in \Gamma(L^{-1})$. Then we can define $\theta \in \Gamma(\mathrm{End}(E))$ by letting $\theta : L_i \to L_{i+1}$ act by multiplication by ξ. Such a θ commutes with ϕ, and has $\bar{\partial}\theta = \mu\phi$.

Therefore if $\mu\psi$ is $\bar{\partial}$-exact, $\mu \in K_\rho(S)$. By Serre duality, $\mu\psi$ is $\bar{\partial}$-exact if and only if

$$\int_S \mu\psi \wedge \eta = 0 \text{ for all } \eta \in H^0(KL). \tag{6.8.2}$$

Since $T_S\mathcal{T}_g$ is dual to the space of holomorphic quadratic differentials on S, the codimension of the space of equivalence classes of Beltrami forms $[\mu] \in T_S\mathcal{T}_g$ for which (6.8.2) holds is

$$\dim\{\psi\eta : \eta \in H^0(KL)\} = h^0(KL).$$

Note that $h^0(KL)$ is the space of meromorphic 1-forms on S with at most a single pole at $p \in S$. But by Stokes' theorem, the residues of the poles of a meromorphic 1-form must sum to zero. Thus forms with a single pole are in fact holomorphic, so $h^0(KL) = h^0(K) = g$. Thus the space of Beltrami forms for which (6.8.2) holds is precisely $3g - 3 - h^0(KL) = 2g - 3$. This concludes the case of odd n.

6.8.2.2 The even case

We now explain how to modify the above construction in the case of even n. We again pick a generic point $p \in S$, set $L = \mathcal{O}(p)$, and let ψ be a non-zero section of $H^0(KL^{-1})$. Set $n = 2m$, and

$$L_1 = L^{m-1}, L_2 = L^{m-2}, ..., L_m = \mathcal{O}_S,$$
$$L_{m+1} = \mathcal{O}_S, L_{m+2} = L^{-1}, ..., L_{2m} = L^{-(m-1)}.$$

Let ω be an arbitrary holomorphic 1-form. Let $E = \bigoplus_{i=1}^{2m} L_i$, and define $\phi \in H^0(K\mathrm{End}(E))$ by

$$\phi : L_i \to L_{i+1}K \text{ is multiplication by } \psi \text{ unless } i = m,$$
$$\phi : L_m \cong \mathcal{O}_S \to K \cong L_{m+1}K \text{ is multiplication by } \omega.$$

The proof of Claim 6.31 applies verbatim to show that (E, ϕ) is a stable $\mathrm{SL}(n, \mathbb{C})$-Higgs bundle.

An analogous argument to the one that led to (6.8.2), shows that $\Delta_\mu \mathrm{E}_\rho = 0$ if both $\mu\psi$ and $\mu\omega$ are $\bar{\partial}$-exact. Again by Serre duality, this is equivalent to

$$\int_S \mu\Phi = 0 \text{ for all } \Phi \in \psi H^0(KL) + \omega H^0(K).$$

The same argument shows that $h^0(KL) = h^0(K) = g$, and hence

$$\dim K_\rho(S) \geq 3g - 3 - h^0(KL) - h^0(K) = g - 3.$$

Bibliography

[1] E. Arbarello, M. Cornalba, P. Griffiths, and J.D. Harris. *Geometry of Algebraic Curves: Volume I*. Grundlehren der mathematischen Wissenschaften. Springer New York, 2013.

[2] Arnaud Beauville, S. Ramanan, and M. S. Narasimhan. Spectral curves and the generalised theta divisor. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1989(398):169–179, July 1989.

[3] Yves Benoist and Dominique Hulin. Harmonic quasi-isometric maps between rank one symmetric spaces. *Annals of Mathematics*, 185(3), May 2017.

[4] Yves Benoist and Dominique Hulin. Harmonic measures on negatively curved manifolds. *Annales de l'Institut Fourier*, 69(7):2951–2971, June 2020.

[5] Yves Benoist and Dominique Hulin. Harmonic quasi-isometric maps II: negatively curved manifolds. *Journal of the European Mathematical Society*, 23(9):2861–2911, April 2021.

[6] Yves Benoist and Dominique Hulin. Harmonic quasi-isometric maps III: quotients of Hadamard manifolds. *Geometriae Dedicata*, 217(3), March 2023.

[7] M. Boggi and M. Pikaart. Galois covers of moduli of curves. *Compositio Mathematica*, 120(2):171–191, 2000.

[8] Marco Boggi, Andrew Putman, and Nick Salter. Generating the homology of covers of surfaces. May 2023. Preprint arXiv:2305.13109.

[9] Marc Bourdon. *Actions quasi-convexes d'un groupe hyperbolique, flot géodésique*. Theses, Université Paris-Sud (1970-2019), February 1993. PhD thesis.

[10] B. H. Bowditch. Some results on the geometry of convex hulls in manifolds of pinched negative curvature. *Commentarii Mathematici Helvetici*, 69(1):49–81, December 1994.

[11] Martin R. Bridson. *Semisimple actions of mapping class groups on CAT(0) spaces*, pages 1–14. Cambridge University Press, February 2010.

[12] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer Berlin Heidelberg, 1999.

[13] S.-Y. Cheng. Liouville theorem for harmonic maps. In Jan Fagerberg, David C. Mowery, and Richard R. Nelson, editors, *Geometry of the Laplace operator*, pages 147–151. American Mathematical Society, 1980.

[14] Kevin Corlette. Flat G-bundles with canonical metrics. *Journal of Differential Geometry*, 28(3), January 1988.

[15] Jean-Pierre Demailly. *Complex analytic and differential geometry*. OpenContent, 2012. Available at `https://www-fourier.ujf-grenoble.fr/~demailly/manuscripts/agbook.pdf`.

[16] S. K. Donaldson. Twisted harmonic maps and the self-duality equations. *Proceedings of the London Mathematical Society*, s3-55(1):127–131, July 1987.

[17] Harold Donnelly. Asymptotic Dirichlet Problem for Harmonic Maps with Bounded Image. *Geometriae Dedicata*, 91(1):1–6, 2002.

[18] J Eells and L Lemaire. Deformations of metrics and associated harmonic maps. *Proceedings of the Indian Academy of Sciences - Section A*, 90(1):33–45, February 1981.

[19] James Eells and J. H. Sampson. Harmonic mappings of riemannian manifolds. *American Journal of Mathematics*, 86(1):109, January 1964.

[20] Jordan S. Ellenberg, Chris Hall, and Emmanuel Kowalski. Expander graphs, gonality, and variation of Galois representations. *Duke Mathematical Journal*, 161(7), May 2012.

[21] Yue Fan. Construction of the moduli space of Higgs bundles using analytic methods. *Mathematical Research Letters*, 29(4):1011–1048, 2022.

[22] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups*. Princeton University Press, December 2011.

[23] Alastair Fletcher and Vladimir Marković. *Quasiconformal maps and Teichmüller theory*. Number 11 in Oxford graduate texts in mathematics. Oxford Univ. Press, Oxford [u.a.], 1. publ. edition, 2007.

[24] Giovanni Forni. Deviation of ergodic averages for area-preserving flows on surfaces of higher genus. *The Annals of Mathematics*, 155(1):1, January 2002.

[25] Giovanni Forni, Carlos Matheus, and Anton Zorich. Lyapunov spectrum of invariant subbundles of the Hodge bundle. *Ergodic Theory and Dynamical Systems*, 34(2):353–408, April 2012.

[26] F. W. Gehring and J. Väisälä. Hausdorff dimension and quasiconformal mappings. *Journal of the London Mathematical Society*, s2-6(3):504–512, May 1973.

[27] William M Goldman. The symplectic nature of fundamental groups of surfaces. *Advances in Mathematics*, 54(2):200–225, November 1984.

[28] Edna K. Grossman. On the residual finiteness of certain mapping class groups. *Journal of the London Mathematical Society*, s2-9(1):160–164, November 1974.

[29] Richard S. Hamilton. *Harmonic Maps of Manifolds with Boundary*. Springer Berlin Heidelberg, 1975.

[30] Joe Harris and Ian Morrison. *Moduli of curves*. Number 187 in Graduate texts in mathematics. Springer, New York, 1998.

[31] P. Hartman. On homotopic harmonic maps. Proc. US-Japan Semin. Diff. Funct. Equat., Minneapolis, Minnesota, 1967.

[32] Sigurdur Helgason. *Differential geometry, Lie groups, and symmetric spaces*. Number volume 34 in Graduate studies in mathematics. American Mathematical Society, Providence, Rhode Island, 2001.

[33] Nigel Hitchin. The self-duality equations on a Riemann surface. *Proceedings of the London Mathematical Society*, s3-55(1):59–126, July 1987.

[34] Nigel Hitchin. Stable bundles and integrable systems. *Duke Mathematical Journal*, 54(1), January 1987.

[35] Nigel Hitchin. Critical Loci for Higgs Bundles. *Communications in Mathematical Physics*, 366(2):841–864, February 2019.

[36] John H. Hubbard. *Teichmüller Theory and Applications to Geometry, Topology, and Dynamics*, volume 1. Ithaca, Matrix editions, 2006.

[37] Nikolai V. Ivanov. Fifteen problems about the mapping class groups. In *Problems on mapping class groups and related topics*, pages 71–80. Providence, RI: American Mathematical Society (AMS), 2006.

[38] Ilya Kapovich and Nadia Benakli. Boundaries of hyperbolic groups, 2002. Preprint arXiv:math/0202286.

[39] S. Kobayashi and K. Nomizu. *Foundations of Differential Geometry, Volume 2*. Foundations of Differential Geometry [by] Shoshichi Kobayashi and Katsumi Nomizu. Wiley, 1963.

[40] François Labourie. Cross ratios, Anosov representations and the energy functional on Teichmüller space. *Annales scientifiques de l'École normale supérieure*, 41(3):439–471, 2008.

[41] François Labourie. Cyclic surfaces and Hitchin components in rank 2. *Annals of Mathematics*, 185(1), January 2017.

[42] Sean Lawton and Adam S. Sikora. Varieties of characters. *Algebras and Representation Theory*, 20(5):1133–1141, March 2017.

[43] Marius Lemm and Vladimir Marković. Heat flows on hyperbolic spaces. *Journal of Differential Geometry*, 108(3), March 2018.

[44] Peter Li and Jiaping Wang. Harmonic rough isometries into Hadamard space. *Asian Journal of Mathematics*, 2(3):419–442, 1998.

[45] Peter Li and Shing-Tung Yau. A new conformal invariant and its applications to the Willmore conjecture and the first eigenvalue of compact surfaces. *Inventiones Mathematicae*, 69(2):269–291, June 1982.

[46] Brice Loustau. Harmonic maps from Kähler manifolds. October 2020. Preprint arXiv:2010.03545.

[47] Eyal Markman. Spectral curves and integrable systems. *Compositio Mathematica*, 93(3):255–290, 1994.

[48] Vladimir Marković. Harmonic maps between 3-dimensional hyperbolic spaces. *Inventiones mathematicae*, 199(3):921–951, July 2014.

[49] Vladimir Marković. Harmonic maps and the Schoen conjecture. *Journal of the American Mathematical Society*, 30(3):799–817, March 2017.

[50] Vladimir Marković and Ognjen Tošić. The second variation of the Hodge norm and higher Prym representations. *To appear in Journal of Topology*, 2021. Preprint on `https://people.maths.ox.ac.uk/~markovic/MT-hodge.pdf`.

[51] Vladimir Marković. Non-uniqueness of minimal surfaces in a product of closed Riemann surfaces. *Geometric and Functional Analysis*, 32(1):31–52, January 2022.

[52] Vladimir Marković. Unramified correspondences and virtual properties of mapping class groups. *Bulletin of the London Mathematical Society*, 54(6):2324–2337, June 2022.

[53] Wolfgang Meyer. Toponogov's Theorem and Applications. 2004. Lecture notes, available at `https://www2.math.upenn.edu/~wziller/math660/TopogonovTheorem-Myer.pdf`.

[54] Mario J. Micallef and John Douglas Moore. Minimal two-spheres and the topology of manifolds with positive curvature on totally isotropic two-planes. *The Annals of Mathematics*, 127(1):199, January 1988.

[55] P. K. Mitter and C. M. Viallet. On the bundle of connections and the gauge orbit manifold in Yang-Mills theory. *Communications in Mathematical Physics*, 79(4):457–472, December 1981.

[56] Peter Petersen. *Riemannian geometry*. Springer New York, 2006.

[57] Andrew Putman and Ben Wieland. Abelian quotients of subgroups of the mapping class group and higher Prym representations. *Journal of the London Mathematical Society*, 88(1):79–96, March 2013.

[58] Nathaniel Sagman and Peter Smillie. Unstable minimal surfaces in symmetric spaces of non-compact type. 2022. Preprint arXiv:2208.04885.

[59] Richard Schoen. The role of harmonic mappings in rigidity and deformation problems, Complex Geometry (Osaka, 1990). *Lecture Notes in Pure Appl. Math.*, 143:179–200, 1993.

[60] Richard Schoen and Shing Tung Yau. Compact group actions and the topology of manifolds with non-positive curvature. *Topology*, 18(4):361–380, 1979.

[61] Carlos T. Simpson. Higgs bundles and local systems. *Publications mathématiques de l'IHÉS*, 75(1):5–95, December 1992.

[62] Yum-Tong Siu. The Complex-Analyticity of Harmonic Maps and the Strong Rigidity of Compact Kahler Manifolds. *The Annals of Mathematics*, 112(1):73, July 1980.

[63] Ivo Slegers. Strict plurisubharmonicity of the energy on Teichmüller space associated to Hitchin representations. 11 2020. Preprint arXiv:2011.03936.

[64] Ivo Slegers. Equivariant harmonic maps depend real analytically on the representation. *manuscripta mathematica*, 169(3–4):633–648, November 2021.

[65] Toshikazu Sunada. Rigidity of certain harmonic mappings. *Inventiones Mathematicae*, 51(3):297–307, October 1979.

[66] Domingo Toledo. Hermitian Curvature and Plurisubharmonicity of Energy on Teichmüller Space. *Geometric and Functional Analysis*, 22:1015–1032, 2011.

[67] Ognjen Tošić. Harmonic projections in negative curvature. *To appear in Communications in Analysis and Geometry*, 2022.

[68] Ognjen Tošić. Non-Strict Plurisubharmonicity of Energy on Teichmüller Space. *International Mathematics Research Notices*, January 2024.

[69] Ognjen Tošić. Harmonic projections in negative curvature II: large convex sets. 2023. Preprint arXiv:2310.05796.

[70] Anthony Tromba. *Teichmüller theory in Riemannian geometry.* Lectures in Mathematics ETH Zürich. Springer, Basel, [Switzerland], 1992.

[71] P. Tukia and J. Vaisala. Quasiconformal extension from dimension n to $n+1$. *The Annals of Mathematics*, 115(2):331, March 1982.

[72] Shing-Tung Yau. Problem section. In Shing-Tung Yau, editor, *Seminar on differential geometry, Annals of Mathematics Studies 102.* Princeton University Press and University of Tokyo Press, 1982.