UNSTABLE MINIMAL SURFACES IN SYMMETRIC SPACES OF NON-COMPACT TYPE

NATHANIEL SAGMAN AND PETER SMILLIE

Abstract. We prove that if Σ is a closed surface of genus at least 3 and G is a split real semisimple Lie group of rank at least 3 acting faithfully by isometries on a symmetric space N, then there exists a Hitchin representation $\rho : \pi_1(\Sigma) \to G$ and a ρ-equivariant unstable minimal map from the universal cover of Σ to N. This follows from a new lower bound on the index of high energy minimal maps into an arbitrary symmetric space of non-compact type. Taking $G = \mathrm{PSL}(n, \mathbb{R})$, $n \geq 4$, this disproves the Labourie conjecture.

1. Introduction

1.1. **Hitchin representations.** Let Σ_g be a closed orientable surface of genus $g \geq 2$, G the adjoint form of a split real simple Lie group, and $\mathrm{Rep}(\pi_1(\Sigma_g), G)$ the representation variety. There is a unique irreducible embedding $i_G : \mathrm{PSL}(2, \mathbb{R}) \to G$. Given a Fuchsian representation ρ from $\pi_1(\Sigma_g)$ to $\mathrm{PSL}(2, \mathbb{R})$, $i_G \circ \rho$ defines a representation to G. The Hitchin component, consisting of the Hitchin representations, is the connected component of $\mathrm{Rep}(\pi_1(\Sigma_g), G)$ containing representations of the form $i_G \circ \rho$. Understood to generalize Teichmüller space, Hitchin components are the first and most important examples of higher Teichmüller spaces.

In his seminal paper [20], Hitchin found using Higgs bundles that the Hitchin component is homeomorphic to a cell, just like Teichmüller space. Around the same time, Choi-Goldman showed that the Hitchin component for $G = \mathrm{PSL}(3, \mathbb{R})$ parametrizes geometric structures on surfaces: it is the space of marked convex $\mathbb{RP}^2$ structures on Σ_g [6]. About a decade later, Labourie [23, Theorem 1.5] and Fock-Goncharov [15] independently proved that Hitchin representations are discrete and faithful. Now known to share many dynamical properties with Fuchsian representations, Hitchin representations—and more generally Anosov representations into non-compact real Lie groups—have become fruitful objects of study. For more information we recommend the survey [41].

Let K be a maximal compact subgroup of G, so that G/K with its Killing metric is a Riemannian symmetric space. Given an irreducible representation $\rho : \pi_1(\Sigma_g) \to G$, Donaldson [11] and Corlette [9] proved that for every Riemann surface structure S on Σ_g with universal cover $\tilde{S}$, there exists a unique ρ-equivariant harmonic map $h : \tilde{S} \to G/K$. Labourie showed that for Hitchin representations, there is a choice of Riemann surface structure such that the harmonic map h is minimal [24]. Labourie's minimal surfaces are stable minima for the equivariant area functional.

Labourie conjectured that for $G = \mathrm{PSL}(n, \mathbb{R})$, every Hitchin representation $\rho : \pi_1(\Sigma_g) \to G$ determines a unique ρ-equivariant minimal map into the symmetric space. It is natural to extend this conjecture to all Hitchin representations into split real simple Lie groups. Labourie's existence theory applies even more generally to all representations in higher Teichmuüller spaces. These include maximal and Θ-positive representations (see [3] and [18]).

The uniqueness question has been posed for these representations, and is also referred to in some places as the Labourie Conjecture. To distinguish, we call the uniqueness conjecture for all higher Teichmüller spaces the Generalized Labourie Conjecture.

The uniqueness of the minimal map for $\mathrm{PSL}(2,\mathbb{R})$ follows more or less from the definition of a Fuchsian representation. Labourie noted that for $\mathrm{PSL}(3,\mathbb{R})$, uniqueness had been established by Loftin [26]. Labourie then proved the generalized conjecture for all split real simple G of rank 2 [25] (this adds $\mathrm{PSp}(4,\mathbb{R})$ and G_2), and Collier-Tholozan-Toulisse proved it for maximal representations into Hermitian Lie groups of rank 2 [8]. See also the work of Collier [7], Alessandrini-Collier [1], and Nie [31].

The Labourie Conjecture and its generalization have attracted significant interest in higher Teichmüller theory. Whenever the uniqueness conjecture holds for Hitchin representations into a split real Lie group G, so in particular for all such G of rank 2, the Hitchin component admits a natural complex structure. More precisely, there is a mapping class group equivariant parametrization as the total space of an explicit holomorphic vector bundle over Teichmüller space (see [24, section 2] and [25, section 1]). This is preferable to Hitchin's original parametrization, which depends on a choice of complex structure on Σ_g and doesn't enjoy mapping class group symmetry. For spaces of maximal representations in rank 2, similar equivariant parametrizations as complex manifolds have been obtained in [7] and [1], although the complex structures are more complicated.

As far as we're aware, the Generalized Labourie Conjecture was believed to be true until March 2021, when Marković demonstrated the existence of a product of Fuchsian representations into $\mathrm{PSL}(2,\mathbb{R})^3$ with multiple minimal surfaces in the corresponding product of closed hyperbolic surfaces [28] (see also [27]). This gives a counterexample to the Generalized Labourie Conjecture for maximal representations into Hermitian Lie groups of rank at least 3. In recent joint work with Marković, we prove that every unstable equivariant minimal surface in $\mathbb{R}^n$ gives rise to an unstable minimal surface in a product of n closed hyperbolic surfaces [29].

In this paper we settle Labourie's original conjecture. We first prove that the index of a high energy minimal map into a non-compact symmetric space of rank at least 3 is bounded below by the index of a related equivariant minimal map into $\mathbb{R}^n$. Specializing to split real Lie groups, we use the new lower bound to deduce that the generalized conjecture fails for Hitchin representations into every split real group of rank at least 3.

For split real G splitting into simple factors as $G = G_1 \times \cdots \times G_m$, we use Hitchin to mean that each component representation to G_i is Hitchin. In the semisimple case the associated symmetric space can have different invariant metrics leading to different meanings of "minimal," and our theorem is true for any invariant metric.

Theorem A. *Let G be the adjoint form of a split real semisimple Lie group with rank at least 3 acting faithfully by isometries on a symmetric space N. For every $g \geq 3$, there exists a Hitchin representation $\rho : \pi_1(\Sigma_g) \to G$ and an unstable ρ-equivariant minimal map into N. In particular, there are at least two equivariant minimal surfaces in N.*

We restrict to $g \geq 3$ because there are no unstable equivariant minimal surfaces in $\mathbb{R}^n$ of genus 2; we expect the result to extend to $g = 2$ (see [29, Theorem D] for $G = \mathrm{PSL}(2,\mathbb{R})^3$). By general principles, we can use Theorem A to prove the following corollary.

Corollary A. *For every $g \geq 3$ and pair G, N as in Theorem A, there exists a Hitchin representation $\rho : \pi_1(\Sigma_g) \to G$ with at least two area minimizing ρ-equivariant minimal maps into N.*

We remark that there are other candidates for an equivariant and holomorphic parametrization of the Hitchin components. See [16] and [32].

1.2. **The case $\mathbf{PSL}(2,\mathbb{R})^n$.** Before describing the proof of Theorem A, we explain our previous work with Marković [29]. On a fixed Riemann surface structure S over Σ_g with universal cover $\tilde{S}$, every holomorphic quadratic differential ϕ on S corresponds to a unique Fuchsian representation to $\mathrm{PSL}(2,\mathbb{R})$ and an equivariant harmonic map from $\tilde{S}$ to the hyperbolic space $\mathbb{H}^2$ (see [20], [40], and [42]). The quadratic differential ϕ also corresponds to a unique equivariant harmonic maps to an $\mathbb{R}$-tree: in one direction, the leaf space (T,d) for the vertical foliation of ϕ is an $\mathbb{R}$-tree, and the projection map from $\tilde{S}\to(T,d)$ is harmonic. Given ϕ, we consider the ray of representations with harmonic maps determined by $R\phi$, $R>0$. After suitably rescaling the metric on the target $\mathbb{H}^2$, the maps converge in some sense as $R\to\infty$ to a harmonic map to the $\mathbb{R}$-tree for ϕ (see [43] for details).

If quadratic differentials $\phi_1,\dots,\phi_n$ sum to 0, the equivariant map to $(\mathbb{H}^2)^n$ determined by $(R\phi_1,\dots,R\phi_n)$ is minimal and converges as $R\to\infty$ to a minimal map in a product of $\mathbb{R}$-trees. We prove that every unstable minimal surface for $\mathrm{PSL}(2,\mathbb{R})^n$ converges along the ray to an unstable minimal surface in a product of trees, and conversely every unstable minimal surface in a product of trees is approximated by unstable $\mathrm{PSL}(2,\mathbb{R})^n$-minimal surfaces [29, Theorem B2].

To produce unstable minimal surfaces in products of trees, we observed that when all the $\phi_i's$ are squares, each $\mathbb{R}$-tree folds onto a copy of $\mathbb{R}$, and the minimal surface in the product maps onto an equivariant minimal surface in $\mathbb{R}^n$. We prove that instability in these two contexts is equivalent [29, Theorem C]. We then noted that unstable equivariant minimal surfaces in $\mathbb{R}^n$ are abundant for $n\geq 3$: the most natural example is the lift of an unstable minimal surface in the n-torus to $\mathbb{R}^n$, and, in fact, any non-planar equivariant minimal surface in dimension 3 is unstable (see [29, section 5.3]).

Finally, we found that minimal surfaces in $\mathbb{R}^n$ are at the heart of the problem: in the general setting where we start with n quadratic differentials summing to 0, giving an equivariant minimal map into a product of trees, there is a branched covering on which the differentials lift to squares. If the corresponding equivariant minimal map from the branched cover into $\mathbb{R}^n$ can be destabilized through equivariant variations that are invariant under the Deck group of the branched covering, then the minimal surface in the original product of $\mathbb{R}$-trees is unstable.

1.3. **The idea of the proof.** Let G be a semisimple real Lie group of real rank r with finite center and no compact factors. Parreau has compactified $\mathrm{Rep}(\pi_1(\Sigma_g),G)$, with boundary objects corresponding to $\pi_1(\Sigma_g)$-actions on rank r buildings [33]. We approach the boundary along Hitchin rays, which generalize our construction for $\mathrm{PSL}(2,\mathbb{R})^n$ by way of G-Higgs bundles. It is conjectured in [21] that, with generic hypothesis, along these Hitchin rays the harmonic maps converge in an appropriate sense to a harmonic map to a building. Theorem B2 from [29] suggests that any counterexample to the Labourie conjecture should come from an unstable minimal surface in a building. It's then feasible to imagine that these minimal surfaces behave like minimal surfaces in $\mathbb{R}^r$. Perhaps after passing to a branched covering, the apartments of the buildings can fold onto each other in a way that produces equivariant minimal surfaces in $\mathbb{R}^r$. This geometric picture gives an intuitive reason to expect stability in rank 2 and instability in rank at least 3.

We now describe the Hitchin rays in more detail for G isogenous to $\mathrm{PGL}(n,\mathbb{R})$, in which case the G-Higgs bundles are essentially classical Higgs bundles as defined by Hitchin. Note that the rank is $n-1$. A Higgs bundle $(E,\overline{\partial}_E,\phi)$ is a holomorphic vector bundle $(E,\overline{\partial}_E)$ with a holomorphic section $\phi \in H^0(S, \mathrm{End}(E) \otimes \mathcal{K})$ called the Higgs field. By the non-abelian Hodge correspondence, every harmonic map to a $\mathrm{PGL}(n, \mathbb{C})$-symmetric space (including harmonic maps to the $\mathrm{PGL}(n, \mathbb{R})$ subspace) arises from a unique polystable Higgs bundle. Although the Higgs bundle is a tractable holomorphic object, the harmonic map is not easy to understand from the Higgs bundle data. Indeed, the harmonic map is related to the Higgs bundle through Hitchin's self-duality equations, which form a system of n^2 coupled non-linear PDE's. Accordingly, the precise analysis of convergence to harmonic maps to buildings isn't fully fleshed out.

A Higgs bundle $(E, \overline{\partial}_E, \phi)$ is generically semisimple if the Higgs field is semisimple on the complement of a finite set. Fortunately, for such Higgs bundles it is possible to extract from Hitchin's equations certain scalar subsolutions, and then to use the maximum principle to understand the behavior of the harmonic map as $R \to \infty$. This is carried out by Mochizuki, using estimates of Simpson, who calls the resulting phenomenon asymptotic decoupling [30]. He essentially proves and convergence to a map to a building in open sets that should map into a single apartment.

Every generically semisimple Higgs bundle has an associated branched covering space called the cameral cover ([21, Section 3], [10]) and an equivariant minimal map $\tilde{f}$ from a covering space of the cameral cover to $\mathbb{R}^{n-1}$. We extrapolate from Mochizuki's analysis that off the branch locus, the intrinsic data (the induced metric and part of the second fundamental form) of the minimal map to the symmetric space associated to the Higgs bundles $(E, \overline{\partial}_E, R\phi)$ converges to the intrinsic data of the map $\tilde{f}$ to $\mathbb{R}^{n-1}$. Unfortunately, the precise convergence to buildings at points where the Higgs field is not semisimple is yet to be understood. The local picture at these points has been worked out only for minimal surfaces associated to $\mathrm{PSL}(3, \mathbb{R})$ and $\mathrm{PSp}(4, \mathbb{R})$ Hitchin representations [12], [37].

Happily, the difficulties at non-semisimple points are readily overcome: by the log cut-off trick, any destabilizing variation of a minimal surface can be perturbed to a destabilizing variation supported on the complement of any discrete set. The upshot of Mochizuki's work and this last observation is that we can maneuver around the technicalities of buildings, and treat high energy minimal maps as if they approximate minimal surfaces in $\mathbb{R}^{n-1}$. With this guiding principle, we prove in Theorem B that if the limiting minimal surface in $\mathbb{R}^{n-1}$ is unstable under variations equivariant by a certain action, then so are the high energy minimal surfaces in the symmetric space. This instability result is the technical step, although the computation becomes simple and natural if we work in the setting of harmonic bundles and vary by well-chosen gauge transformations.

Finally, we show that for any equivariant minimal map $\tilde{f}$ into $\mathbb{R}^{n-1}$, one can find a Higgs bundle for a Hitchin representation to $\mathrm{PSL}(n, \mathbb{R})$ whose limiting minimal map is isometric to $\tilde{f}$. Briefly, one can easily turn the Weierstrass data for any minimal surface into a Cartan-valued holomorphic 1-form for $\mathfrak{sl}(n, \mathbb{C})$, and evaluating the Hitchin fibration on this form defines a point in the Hitchin base. Applying a Hitchin section to this point returns the desired Higgs bundle. To conclude the argument, one picks $\tilde{f}$ to be any unstable equivariant minimal surface in $\mathbb{R}^{n-1}$, which is possible only for $n \geq 4$. The Hitchin section is defined for all split real G, and we use this to prove Theorem A in general.

1.4. **The index of minimal maps.** We now give a careful statement of Theorem B. To keep things classical in the introduction, we focus mainly on the $\mathrm{PGL}(n, \mathbb{C})$ case.

Let $(E, \overline{\partial}_E, \phi)$ be a generically semisimple stable Higgs bundle. There exists $m \leq n$ and a finite subset B such that for every $p \in S - B$, there is a neighbourhood $U \subset S - B$ of p on which the Higgs bundle decomposes holomorphically as $E|_U = \oplus_{i=1}^m E_i$, and ϕ acts on each E_i by multiplication by a holomorphic 1-form ϕ_i. Let $\mathrm{End}^0(E) \subset \mathrm{End}(E)$ be the subbundle of traceless endomorphisms. Ranging over a covering of $S-B$, the locally defined projections to the E_i span a flat holomorphic subbundle $F^0_\phi \subset \mathrm{End}^0(E)|_{S-B}$ with parallel metric and real structure, which we call the toral bundle of ϕ. ϕ itself is a holomorphic F^0_ϕ-valued 1-form. Using the real structure on F^0_ϕ, ϕ further defines a flat Riemannian affine bundle M_ϕ over S, which we call the apartment bundle and comes equipped with a harmonic section f whose holomorphic derivative is ϕ. This is constructed in section 3.2.

The cameral cover is a minimal degree branched covering $\tau : C \to S$ on which the monodromy of F^0_ϕ trivializes. On $C - \tau^{-1}(B)$, ϕ admits globally defined eigen-1-forms $\phi_1, \ldots, \phi_m$, which then extend to all of C by the removable singularities theorem. Lifting to a covering space $\tilde{C}$ and integrating the real parts gives a harmonic map to $\mathbb{R}^m$

$$\tilde{f} = (\tilde{f}_1, \ldots, \tilde{f}_m),\ \tilde{f}_i(z) = \sqrt{r_i}\mathrm{Re}\int_{z_0}^{z} \phi_i,$$

where $r_i = \mathrm{rank}(E_i)$, unique up to translation and equivariant with respect to a representation to $\mathbb{R}^m$. It's tidier to work with the equivalent harmonic section of the apartment bundle M_ϕ, but it's helpful to keep this geometric picture in mind.

Using Hitchin and Simpson's side of the non-abelian Hodge correspondence, the Higgs bundle produces a surface group representation $\rho : \pi_1(\Sigma_g) \to \mathrm{PGL}(n, \mathbb{C})$ and a ρ-equivariant harmonic map h from $\tilde{S}$ to the symmetric space of $\mathrm{PGL}(n, \mathbb{C})$. This is equivalent to the data of a bundle E with projectively flat connection D and harmonic metric H. If we rescale the Higgs field via the $\mathbb{R}^+$-action $R \cdot (E, \overline{\partial}_E, \phi) = (E, \overline{\partial}_E, R\phi)$, then we get new pairs of representations and maps (ρ_R, h_R), or equivalently triples (E_R, D_R, H_R). While our results concern the maps h_R, for the problem at hand it's easier to compute with the metrics H_R.

For each R, the map h_R is minimal if and only if f is. Assuming minimality of f, let $\mathrm{Ind}(H_R)$ be the equivariant index for the area of h_R, and $\mathrm{Ind}(f)$ the index of f as a section of M_ϕ, which is the same as the equivariant index for the area of $\tilde{f}$. The indices are rigorously defined in section 3.2.

Theorem B for PGL$(n, \mathbb{C})$. $\liminf_{R\to\infty} \mathit{Ind}(H_R) \geq \mathit{Ind}(f)$.

Inspired by the situation for $\mathrm{PSL}(2, \mathbb{R})^n$, we conjecture that $\mathrm{Ind}(H_R)$ is non-decreasing with R and $\lim_{R\to\infty} \mathrm{Ind}(H_R) = \mathrm{Ind}(f)$.

We turn now to the most general case of a real semisimple Lie group G acting on a symmetric space N, and a generically semisimple stable G-Higgs bundle $(P, A^{0,1}, \phi)$. Here $(P, A^{0,1})$ is a holomorphic $K^{\mathbb{C}}$ bundle and ϕ is a holomorphic 1-form valued in $P \times_{K^{\mathbb{C}}} \mathfrak{p}^{\mathbb{C}}$. As in the classical case, there is an $\mathbb{R}^+$-family of representations and harmonic maps, or G-harmonic bundles (M_R, h_R), in which M_R is the bundle with fiber N assoicated to P, which inherits a flat connection, and h_R is a harmonic section.

The right generalization of the bundle F^0_ϕ to this setting is what we call the G-toral bundle of ϕ, whose fiber is the center of the centralizer of ϕ in $P \times_{K^{\mathbb{C}}} \mathfrak{p}^{\mathbb{C}}$. As in the $\mathrm{PGL}(n, \mathbb{C})$ case, this is really a vector bundle away from some bad set B, and its rank is at most the rank of G. From the toral bundle, we build the G-apartment bundle M^G_ϕ, together with its minimal section f (section 3.4.1). Then we prove

Theorem B. $\liminf_{R\to\infty} \mathit{Ind}(h_R) \geq \mathit{Ind}(f)$.

The proof almost follows directly from the case of $\mathrm{PGL}(n, \mathbb{C})$ by using the adjoint representation of G. The only minor complication arises from our insistence, in the semisimple case, of allowing an arbitrary invariant metric on the symmetric space N.

Remark 1.1. By the work of Toledo [39], the associated energy functional on Teichmüller space is plurisubharmonic. Since the energy index is equal to the area index [14, Theorem 3.4], $\mathrm{Ind}(h_R)$ is at most $3g-3$.

1.5. **Acknowledgements.** We would like to thank Nicolas Tholozan, Eugen Rogozinnikov, and Max Riestenberg for helpful conversations, and Mike Wolf in particular for pointing out the connection between [28] and minimal surfaces in $\mathbb{R}^n$.

2. Preliminaries

2.1. **Hermitian metrics.** For a complex vector space V, we write $V^\vee$ for its dual and $\overline{V}$ for its complex conjugate. Similarly, for a linear map of complex vector spaces $f : V \to W$, we write $f^\vee : W^\vee \to V^\vee$ for the transpose and $\overline{f} : \overline{V} \to \overline{W}$ for the conjugate. A Hermitian metric H on V is a linear isomorphism $H : V \to \overline{V}^\vee$ such that $H = \overline{H}^\vee$ and the associated sesquilinear form is positive definite. If $f \in \mathrm{End}(V)$ and H is a Hermitian metric on V, then the adjoint with respect to H is $f^{*_H} = H^{-1}\overline{f}^\vee H : V \to V$.

Let $\mathbb{P}\mathrm{Met}(V)$ be the real projectivization of the space of Hermitian metrics on V. We will use $[H]$ to denote the equivalence class of a metric H in $\mathbb{P}\mathrm{Met}(V)$. The group $G = \mathrm{PGL}(V)$ acts transitively on $\mathbb{P}\mathrm{Met}(V)$ via

$$g \cdot [H] = [\overline{\hat{g}^{-1}}^\vee H \hat{g}^{-1}] \tag{1}$$

where $\hat{g}$ is any lift of g to $\mathrm{GL}(V)$. The stabilizer of a class $[H]$ is a maximal compact subgroup $K_H \subset G$, which exhibits $\mathbb{P}\mathrm{Met}(V)$ as a symmetric space of non-compact type.

Let $\mathfrak{g}$ be the Lie algebra of G, which we identify with the traceless endomorphisms $\mathrm{End}^0(V)$. Let $\mathrm{Hom}_\mathbb{R}(V, \overline{V}^\vee)$ be the real subspace of $\mathrm{Hom}(V, \overline{V}^\vee)$ consisting of maps equal to their conjugate transpose. For any $[H] \in \mathbb{P}\mathrm{Met}(V)$, the tangent space $T_{[H]}\mathbb{P}\mathrm{Met}(V)$ is the quotient of $\mathrm{Hom}_\mathbb{R}(V, \overline{V}^\vee)$ by the line $[H]$. If $X \in \mathfrak{g}$, then the derivative of the action (1) is given by

$$X \cdot [H] = [-\overline{X}^\vee H - HX] \in T_{[H]}\mathbb{P}\mathrm{Met}(V).$$

If $Y \in \mathrm{End}(V)$, write Y^0 for its traceless part. There is a $\mathfrak{g}$-valued 1-form φ on $\mathbb{P}\mathrm{Met}(V)$ defined by

$$\varphi_{[H]} = -\frac{1}{2}(H^{-1}dH)^0. \tag{2}$$

for any lift H to $\mathrm{Hom}_\mathbb{R}(V, \overline{V}^\vee)$. Combining the last two equations, we find

$$\varphi_{[H]}(X \cdot [H]) = \frac{1}{2}(X^{*_H} + X). \tag{3}$$

Note that the operator $*_H$ depends only on the projective equivalence class $[H]$.

Let $\mathrm{End}^0_H \subset \mathfrak{g}$ be the space of H-self-adjoint endomorphisms of V. Then $\varphi_{[H]}(X \cdot [H])$ is the projection of X to End^0_H, and it follows that φ defines an isomorphism from $T\mathbb{P}\mathrm{Met}(V)$ to the bundle over $\mathbb{P}\mathrm{Met}(V)$ whose fiber over a metric H is End^0_H.

The Killing metric on $\mathbb{P}\mathrm{Met}(V)$ is defined to be

$$2n\,\mathrm{tr}(\varphi \otimes \varphi). \tag{4}$$

In the sequel, we will omit the tensor product $\otimes$ from our notation, writing this as $2n\,\mathrm{tr}(\varphi^2)$, which is to be read as a bilinear form. This metric is G-invariant, and every G-invariant metric on $\mathbb{P}\mathrm{Met}(V)$ is a constant multiple of the Killing one.

To simplify the statement of the following proposition, we use φ to identify $T\mathbb{P}\mathrm{Met}(V)$ with the subbundle End^0_H of the trivial $\mathfrak{g}$ bundle over $\mathbb{P}\mathrm{Met}(V)$ whose fiber over $[H]$ is the subspace End^0_H. The Lie bracket and differential d are defined on this trivial $\mathfrak{g}$ bundle.

Proposition 2.1 (Theorem X.2.6 in [22])**.** *For sections X, Y, Z of End^0_H, implicitly identified with vector fields on the symmetric space, the Levi-Civita connection ∇^{lc} is given by*

$$\nabla^{lc}_X Y = d_X Y - [X, Y]. \tag{5}$$

and curvature of ∇^{lc} is given by

$$R(X, Y)Z = -[[X, Y], Z]$$

Remark 2.2. Although the metric is canonical only up to a constant, the Levi-Civita connection and curvature are independent of this choice.

2.2. **Harmonic maps, bundles, and minimal surfaces.** Let S be a possibly non-compact Riemann surface and $h : S \to N$ a smooth map from S to a Riemannian manifold N. Let $\mathcal{K}$ be the canonical bundle of S, and let ∂h be the $(1, 0)$-part of the derivative of h, which is a section of $h^*TN \otimes_{\mathbb{R}} \mathcal{K}$. Let $\overline{\partial}$ be the $(0, 1)$-part of the connection on $h^*TN \otimes_{\mathbb{R}} \mathbb{C}$ induced from the Levi-Civita connection of N, as well as its natural extension to $h^*TN \otimes_{\mathbb{R}} \mathbb{C}$-valued forms.

Definition 2.3. The map h is harmonic if $\overline{\partial}\partial h = 0$.

If $\nu = \langle \cdot, \cdot \rangle$ is the Riemannian metric on N, the Hopf differential is the section of $\mathcal{K}^2$ defined by $\langle \partial h, \partial h \rangle$. We call h conformal if $\langle \partial h, \partial h \rangle = 0$. We eschew words like almost conformal and branched conformal, instead calling a classically conformal map a conformal immersion.

Definition 2.4. A harmonic and conformal map is called a minimal map.

The area density of h is the $(1, 1)$-form on S given by

$$a_h = \sqrt{\langle \partial h, \overline{\partial} h \rangle^2 - |\langle \partial h, \partial h \rangle|^2} \tag{6}$$

where the expression inside the square-root is interpreted as a real section of $\mathcal{K}^2 \otimes \overline{\mathcal{K}}^2$. If S is compact, we write the area of $h(S)$ as

$$A(h) = \int_S a_h. \tag{7}$$

Analogously, we define the energy density $e_h = \langle \partial h, \overline{\partial}\, h \rangle$ and total energy $E(h) = \int_S e_h$. For arbitrary S, h is harmonic if and only if it is a critical point of E with respect to compactly supported variations on all compact subsurfaces, and if h is harmonic, then it is minimal if and only if it is a critical point of A in this sense. In particular, a minimal map defines a minimal surface in the classical sense if it is an immersion and an embedding.

Remark 2.5. A minimal map h is allowed to be constant. If it is non-constant, then since ∂h is a holomorphic section of a holomorphic vector bundle, its vanishing locus is discrete. By conformality, h is an immersion wherever ∂h is nonzero. Therefore, a minimal map is an immersion away from a discrete set.

2.2.1. *Flat Riemannian bundles.* Most of the time, for us S will be a compact Riemann surface. We are not interested in minimal maps from S per se, but rather in minimal sections of flat Riemannian bundles over S, which we now define. Let $M \to S$ be a fiber bundle. A flat connection on M is a foliation transverse to the fibers of M of complementary dimension, which induces isomorphisms between nearby fibers, and hence local flat trivializations $M|_U = U \times M_z$, where $z \in U \subset S$ and M_z is the fiber over z. A Riemannian structure on M is a Riemannian metric on the fibers. M is a flat Riemannian bundle if it has both a flat connection and a Riemannian structure, and the induced identifications between the fibers are isometries.

Let $M \to S$ be a flat Riemannian bundle, and let h be a section. We say that h is minimal (or harmonic or conformal respectively) if the corresponding map to M_z is minimal (resp. harmonic, conformal) for every local flat trivialization $M|_U = U \times M_z$. Since these notions are invariant under isometry, it suffices to check them in a set of trivializations which covers S.

If we don't want to choose trivializations, we can also give a more abstract definition of minimality. Let $T^{\text{vert}}M$ be the vertical tangent bundle of M. The metric on M induces a metric on $T^{\text{vert}}M$, and the flat connection on M lets us lift the Levi-Civita connection on the fibers of M to a connection on $T^{\text{vert}}M$, which we also call the Levi-Civita connection. If h is a section of M, we simply interpret ∂h as a 1-form valued in $h^*T^{\text{vert}}M$ and $\overline{\partial}$ in terms of the Levi-Civita connection on $T^{\text{vert}}M$. Having done so, the original definitions of harmonic and conformal still make sense. Naturally, the formulas (6) and (7) still define the area density and area for sections of flat Riemannian bundles, and minimal sections are still critical points of area. The condition that h is an immersion is potentially ambiguous; the correct condition is that the the area form a_h (defined with respect to the vertical tangent bundle) is nonvanishing. It remains true that a minimal section is an immersion away from a discrete set.

2.2.2. *Bundles and equivariance.* If $\tilde{S}$ is a universal covering space of S, and $\tilde{M}$ is the pullback of M to $\tilde{S}$, the flat connection gives canonical isomorphisms between each fiber of $\tilde{M}$, each of which is furthermore identified with the space of flat sections of $\tilde{M}$. Let $\Gamma^{\text{flat}}(\tilde{M})$ be this space. Viewing it as the space of flat sections, the Deck group of $\tilde{S}$ acts on $\Gamma^{\text{flat}}(\tilde{M})$ by pullback. Call this action ρ.

If h is a section of M over S, we can pull it back to a section $\tilde{h}$ of $\tilde{M}$. By identifying each fiber of $\tilde{M}$ with the space of flat sections, we can view $\tilde{h}$ as a map $\tilde{h} : \tilde{S} \to \Gamma^{\text{flat}}(\tilde{M})$. This map will be equivariant by the action ρ. Conversely, ρ-equivariant maps from $\tilde{S}$ to $\Gamma^{\text{flat}}(\tilde{M})$ descend to sections of M on S. Clearly h is minimal if and only if $\tilde{h}$ is.

2.2.3. *Harmonic and minimal bundles.* Let E be a complex vector bundle over S. A connection D on E is called projectively flat if its curvature is equal to ωId, where ω is a smooth 2-form. Since the identity is in the center of $\text{End}(E)$, a projectively flat connection on E induces a flat connection on $\text{End}(E)$.

If (E, D) is a projectively flat bundle, we define the projectivized bundle of metrics $\mathbb{P}\text{Met}(E)$, whose fiber at a point z is $\mathbb{P}\text{Met}(E(z))$. A choice of invariant metric ν on $\mathbb{P}\text{Met}(E)$ makes it into a flat Riemannian bundle. Unless we specify ν, we will always take the Killing metric (4). The projectively connection D defines a flat connection on $\mathbb{P}\text{Met}(E)$. In this way, $\mathbb{P}\text{Met}(E)$ has the structure of a flat Riemannain bundle.

The main object of study in this paper will be triples (E, D, H), where E is a vector bundle over S, D a projectively flat connection, and H a smooth Hermitian metric. Given such a triple, let $[H]$ be the corresponding section of $\mathbb{P}\mathrm{Met}(E)$. The first part of the following definition is standard, but the second part is not.

Definition 2.6. (E, D, H) is a harmonic bundle if the associated section $[H]$ of $\mathbb{P}\mathrm{Met}(E)$ is harmonic. If $[H]$ is moreover minimal, we call (E, D, H) a minimal harmonic bundle.

Note that the definition of minimal would be equivalent had we chosen any other scaling of the metric on $\mathbb{P}\mathrm{Met}(E)$, and that the definition of harmonic does not reference the metric at all, only the Levi-Civita connection.

In the introduction, we mentioned the existence theorem for harmonic maps of Donaldson [11] and Corlette [9]. Rephrasing in terms of harmonic metrics, this gives one side of the non-abelian Hodge correspondence. A representation $\rho : \pi_1(S_g) \to \mathrm{PGL}(n, \mathbb{C})$ is called reductive if the Zariski closure of the image subgroup is a reductive subgroup of $\mathrm{PGL}(n, \mathbb{C})$. It is irreducible if it is not contained in a non-trivial parabolic subgroup.

Theorem 2.7 (Non-abelian Hodge correspondence I)**.** *Let (E, D) be a projectively flat vector bundle of rank n over a closed surface S, and suppose that the holonomy representation is reductive. Then there exists a harmonic metric H on (E, D). If the holonomy is irreducible, then $[H]$ is unique.*

We will often write that (E, D, H) is determined by NAH I, even though only $[H]$ is uniquely defined; since we only ever use the operator $*_H$, which depends only on $[H]$, we hope the reader will excuse this laziness.

2.3. **Higgs bundles.** We owe an interpretation of the definitions 2.3 and 2.4 of harmonic and minimal in the context of triples (E, D, H). First we need a good description of $[H]^*T^{\mathrm{vert}}\mathbb{P}\mathrm{Met}(E)$. This is provided by the endomorphism-valued 1-form φ defined in (2). On each fiber $E(z)$, φ is an $\mathrm{End}^0(E(z))$-valued 1-form on $\mathbb{P}\mathrm{Met}(E(z))$, which for each $[H'] \in \mathbb{P}\mathrm{Met}(E(z))$, identifies $T_{[H']}\mathbb{P}\mathrm{Met}(E(z))$ with $\mathrm{End}^0_{H'}(E(z))$, the space of H'-self-adjoint endomoprhisms of $E(z)$. Note the complexification of $\mathrm{End}^0_{H'}(E)$ is just $\mathrm{End}^0(E)$, so φ identifies the complexification of the tangent space at each point $[H']$ with $\mathrm{End}^0(E)$.

In particular, if (E, D, H) is a flat bundle with a metric, and $[H]$ is the corresponding section of $\mathbb{P}\mathrm{Met}(E)$, then the pullback $\varphi_H := [H]^*\varphi$ of φ by the section $[H]$ identifies $[H]^*T^{\mathrm{vert}}\mathbb{P}\mathrm{Met}(E)$ with $\mathrm{End}^0_H(E)$, and therefore identifies the space of variations of the section $[H]$ with the space of sections of $\mathrm{End}^0_H(E)$ over S. Since φ_H is a an $\mathrm{End}(E)$-valued 1-form, we can define a connection on E by

$$\nabla^H = D - \varphi_H$$

This induces the connection $D - [\varphi_H, \cdot]$ on $\mathrm{End}(E)$, which we also write ∇^H. If X and Y are sections of $\mathrm{End}^0_H(E)$ (which we implicitly identify with $[H]^*T^{\mathrm{vert}}(M)$), then using equation (3) and the fact that X is already equal to its H-self-adjoint part, we have

$$\nabla^H_X Y = D_X Y - [X, Y].$$

Comparing this with the formula (5) in Proposition 2.1, we conclude

Proposition 2.8. *The subbundle $End^0_H(E)$ of traceless H-self-adjoint endomorphisms is invariant by ∇^H, and the restriction of ∇^H to this bundle is the pullback by the section $[H]$ of the Levi-Civita connection on $T^{\mathrm{vert}}\mathbb{P}\mathrm{Met}(E)$.*

Proposition 2.8 allows us to express harmonicity of (E, D, H) in terms of the connection ∇^H. Using the Riemann surface structure on S, let ϕ be the $(1, 0)$-part of φ_H, and let $\overline{\partial}_E$ be the $(0, 1)$-part of the connection ∇^H. Recall also that the Killing metric on $\mathrm{End}^0_H(E)$ is given by $\langle X, Y\rangle = 2n\,\mathrm{tr}(XY)$. Definitions 2.3 and 2.4 translate to the result below.

Proposition 2.9. *(E, D, H) is a harmonic bundle if $\overline{\partial}_E \phi = 0$. It is a minimal bundle if in addition $tr(\phi^2) = 0$.*

Since we're working on a Riemann surface, the Koszul-Malgrange theorem guarantees the existence of a complex structure on E with del-bar operator $\overline{\partial}_E$. This motivates the introduction of Higgs bundles.

Definition 2.10. A Higgs bundle on S is the data $(E, \overline{\partial}_E, \phi)$, where $(E, \overline{\partial}_E)$ is a holomorphic vector bundle of rank n and ϕ is a holomorphic section of $\text{End}(E) \otimes \mathcal{K}$ called the Higgs field.

In this context, Theorem 2.7 says that every projectively flat bundle (E, D) with reductive holonomy gives the data of a Higgs bundle. Actually, this passage from projectively flat bundles to Higgs bundles is generically reversible. Assume now that S is a closed Riemann surface. If $\deg(E)$ is the degree of E and $\text{rank}(E)$ the rank, the slope of a complex vector bundle E over S is the quantity

$$\mu(E) = \frac{\deg(E)}{\text{rank}(E)}.$$

Definition 2.11. A Higgs bundle $(E, \overline{\partial}_E, \phi)$ is stable if for all Φ-invariant holomorphic subbundles $E' \subset E$ of positive rank, we have $\mu(E') < \mu(E)$. It is polystable if it is a direct sum of stable Higgs bundles all of the same slope.

The proposition below is clear from the definitions, and we record it for future use.

Proposition 2.12. *If $(E, \overline{\partial}_E, \phi)$ is stable, then for $\gamma \in \mathbb{C}^*$, $(E, \overline{\partial}_E, \gamma\phi)$ is stable.*

The converse to Theorem 2.7 is due to Hitchin [20] and Simpson [35]. Given a Hermitian metric H on a holomorphic bundle $(E, \overline{\partial}_E)$, let ∂^H be the unique $(1,0)$-connection with the property that $\overline{\partial}_E + \partial^H$ is H-unitary.

Theorem 2.13 (Non-abelian Hodge correspondence II)**.** *If a Higgs bundle $(E, \overline{\partial}_E, \phi)$ is stable, then there exists a Hermitian metric H on E such that $(E, \overline{\partial}_E + \partial^H + \phi + \phi^{*_H})$ is a projectively flat bundle. The harmonic section $[H]$ of $\mathbb{P}\text{Met}(E)$ is unique.*

Throughout the paper, we use the abbreviation NAH for the non-abelian Hodge correspondence.

Remark 2.14. If $(E, \overline{\partial}_E, \phi)$ is a direct sum of stable Higgs bundles $(E_i, \overline{\partial}_{E_i}, \phi_i)$ (i.e., is polystable), then NAH II gives a harmonic bundle structure to each E_i, and hence a harmonic map $[H] = \prod_i [H_i]$ into $M = \prod_i \mathbb{P}\text{Met}(E_i)$. There is an $n-1$ dimensional space of natural embeddings of $\prod_i \mathbb{P}\text{Met}(E_i)$ into $\mathbb{P}\text{Met}(E)$ corresponding to different relative scalings of the n factors. By choosing an embedding, we can construct a harmonic bundle (E, D, H).

It is also sometimes useful to not embed M into $\mathbb{P}\text{Met}(E)$, but rather to view it as a flat Riemannian bundle of symmetric spaces in its own right, with a harmonic section $[H]$. Note that M itself has an n-dimensional family of invariant metrics ν corresponding to choosing a rescaling of the Killing metric on each factor. This is related to but not the same as the family of embeddings; for instance, there is a single metric in this n-dimensional family which makes all of the embeddings of M into $\mathbb{P}\text{Met}(E)$ isometric. Note that the definition of a minimal section of M depends on a choice ν of invariant metric on its fibers.

Definition 2.15. We will say that a polystable Higgs bundle is minimal with respect to an invariant metric ν if the associated harmonic map to M is minimal with respect to ν.

2.4. **Symmetric spaces and G-Higgs bundles.** This subsection can be skipped if the reader is only interested in Theorem B for $\mathrm{PGL}(n, \mathbb{C})$ or Theorem A for $\mathrm{PSL}(n, \mathbb{R})$.

2.4.1. *Symmetric spaces of non-compact type.* To us, a symmetric space of non-compact type is a connected and simply connected Riemannian manifold with an inversion symmetry about each point, whose de Rham decomposition contains only non-compact symmetric spaces, and no factors of $\mathbb{R}$. The isometry group of a symmetric space of non-compact type is always semi-simple with no compact factors and trivial center. Such groups are our primary interest, but in order to make it easy to apply the results to other standard settings, we allow certain subgroups and non-faithful actions as well.

Definition 2.16. A Lie group G is admissible if it is semisimple with finite center and no non-compact factors. An admissible pair (G, K) is an admissible Lie group G together with a maximal compact subgroup K.

Definition 2.17. An action of an admissible Lie group G is essentially faithful if its kernel is discrete.

If (G, K) is an admissible pair, we will always write $\mathfrak{g} = \mathfrak{k} + \mathfrak{p}$ for the corresponding Lie algebra decomopsition. Let N be the pointed space G/K. Using a generalization of the $\mathfrak{g}$-valued 1-form φ, one can show that the G-invariant metric on N are in bijection with the K-invariant positive definite bilinear forms on $\mathfrak{p}$. Any choice of metric makes N into a symmetric space of non-compact type with an essentially faithful transitive G-action.

Definition 2.18. An admissible triple (G, K, ν) is an admissible pair (G, K) together with a K-invariant positive definite bilinear form ν on $\mathfrak{p}$.

If N is a symmetric space of non-compact type, and h a point of N, this determines canonically a Lie algebra $\mathfrak{g}$ of Killing fields and a decomposition $\mathfrak{g} = \mathfrak{k}_h + \mathfrak{p}_h$, with $\mathfrak{k}_h$ the Killing fields vanishing at h, as well as the metric ν. The pair $(G = \mathrm{Aut}(N), K = \mathrm{Stab}_G(h))$ is admissible with these Lie algebras, but N does not uniquely determine the Lie group G.

Remark 2.19. If G is simple, then $\mathfrak{g}$ has a unique invariant bilinear form up to scale. For the purpose of studying minimal surfaces in N, scaling the metric is unimportant. However, if N is reducible, then different scalings of its different factors determine geniunely different notions of minimal surfaces. Thus, in the reducible case N contains an important additional piece of information that G does not.

Theorem X.2.6 of [22] proves that the formulas for the curvature and connection of $\mathbb{P}\mathrm{Met}(E)$ in Proposition 2.1 continue to hold for N. Observe that these formulas only depend on G, and not on the choice invariant bilinear form. Alternatively, we can deduce them by isometrically embedding N as a totally geodesic subspace of a product of symmetric spaces of the type $\mathbb{P}\mathrm{Met}$.

Proposition 2.20. *If (G, K, ν) is an admissible triple, there is a representation $\sigma = \prod_i \sigma_i : G \to \prod_i \mathrm{SL}_{n_i}(\mathbb{C})$ sending K to $\prod_i \mathrm{SU}(n_i)$, and constants a_i, such that the induced map $(G/K, \nu) \to \prod_i (\mathbb{P}\mathrm{Met}(\mathbb{C}^{n_i}), a_i\, tr)$ is a totally geodesic isometry.*

Proof. This is essentially Theorem 2.6.5 of [13]. The K-invariant form ν on $\mathfrak{p}$ extends uniquely to a G-invariant form on $\mathfrak{g}$ which we also call ν. It also determines a positive definite form ν_K on $\mathfrak{g}$, which is equal to ν on $\mathfrak{p}$ and $-\nu$ on $\mathfrak{k}$. Take an orthonormal basis of $\mathfrak{g}$ with respect to the form ν_K, and use the adjoint representation to map G to to $\mathrm{SL}(\dim(\mathfrak{g}), \mathbb{C})$. Since the adjoint representation preserves the splitting of $\mathfrak{g}$ into simple pieces $\mathfrak{g}_i$, it actually maps to $\prod_i \mathrm{SL}(n_i, \mathbb{R})$, where $n_i = \dim(\mathfrak{g}_i)$. Including $\mathrm{SL}(n_i, \mathbb{R})$ into $\mathrm{SL}(n_i, \mathbb{C})$ gives the required representation. □

The most important invariant of a symmetric space N of non-compact type is its rank, which is the largest dimensional subspace of its tangent space at any point on which the sectional curvature vanishes. Equivalently, this is the dimension of a maximal abelian subalgebra of $\mathfrak{g}$ contained in $\mathfrak{p}$. We call such an algebra a maximal toral subalgebra of $\mathfrak{p}$, using the word toral to emphasize that it necessarily consists of semisimple elements.

Let us now additionally fix a maximal toral subalgebra $\mathfrak{a}$ of $\mathfrak{p}$. The Weyl group of $\mathfrak{a}$ (sometimes called the restricted Weyl group of $\mathfrak{a}$) is the normalizer of $\mathfrak{a}$ in K modulo its centralizer. The roots of $\mathfrak{a}$ (sometimes called restricted roots) are the nonzero characters of $\mathfrak{a}$ that arise in its adjoint action on $\mathfrak{g}$. We reserve the word Cartan subalgebra for a (complex) maximal toral subalgebra of the complexification $\mathfrak{g}^{\mathbb{C}}$. The real Lie algebra $\mathfrak{g}$ is called split if the complexification of a maximal toral subalgebra of $\mathfrak{p}$ is a Cartan subalgebra of $\mathfrak{g}^{\mathbb{C}}$. A Lie group G is called split if its Lie algebra is split.

Let $K^{\mathbb{C}} \subset G^{\mathbb{C}}$ and $\mathfrak{a}^{\mathbb{C}} \subset \mathfrak{p}^{\mathbb{C}} \subset \mathfrak{g}^{\mathbb{C}}$ be complexifications of K, $\mathfrak{a}$, and $\mathfrak{p}$ respectively. Let $\mathcal{O}(\mathfrak{p})^K$ the ring of K-invariant complex polynomials on $\mathfrak{p}$, and $\mathcal{O}(\mathfrak{a})^W$ the ring of W-invariant polynomials on $\mathfrak{a}$. These are the same as the ring of $K^{\mathbb{C}}$ invariant polynomials on $\mathfrak{p}^{\mathbb{C}}$ and W-invariant poynomials on $\mathfrak{a}^{\mathbb{C}}$ respectively. The Chevalley restriction theorem gives an isomorphism $\mathcal{O}(\mathfrak{p})^K = \mathcal{O}(\mathfrak{a})^W$, and says that $\mathcal{O}(\mathfrak{p})^K$ is free on r generators, where r is the dimension of $\mathfrak{a}$, i.e., the rank of N. For example, if $G = \mathrm{PGL}(n, \mathbb{R})$, so that $\mathfrak{p}$ consists of symmetric matrices, then the elementary symmetric polynomials p_i, $i = 2, \dots, n$ applied to the eigenvalues is a minimal set of generators for $\mathrm{O}(\mathfrak{p})^K$.

An element X of $\mathfrak{p}^{\mathbb{C}}$ is called regular semisimple if it is semisimple and contained in a unique Cartan subalgebra $\mathfrak{a}^{\mathbb{C}}_X$ of $\mathfrak{p}^{\mathbb{C}}$. For $X \in \mathfrak{a}^{\mathbb{C}}$, X is regular semisimple if and only if no root of $\mathfrak{a}$ vanishes on X. We will use the following standard lemma in section 4.

Lemma 2.21. *If $X \in \mathfrak{p}^{\mathbb{C}}$ is regular semisimple, $Y \in \mathfrak{p}^{\mathbb{C}}$ is arbitrary, and every invariant polynomial $p \in \mathcal{O}(\mathfrak{p})^K$ has $p(X) = p(Y)$, then $X = \mathrm{Ad}_k Y$ for some k in K. If k and k' are two elements of K conjugating Y to X, and $Z \in \mathfrak{p}^{\mathbb{C}}$ commutes with Y, then $\mathrm{Ad}_k Z = \mathrm{Ad}_{k'} Z$.*

2.4.2. *G-Higgs bundles.* Let N be a symmetric space of non-compact type, and let $G = \mathrm{Isom}(N)$. Let M be a flat Riemannian bundle over S, each of whose fibers is isometric to N. If $\tilde{S}$ is a universal cover of S, then M defines, up to conjugacy, a holonomy representation from the Deck group of $\tilde{S}$ to G. Similar to the $\mathrm{PGL}(n, \mathbb{C})$ case, the bundle is reductive if the Zariski closure of any representative of the conjugacy class gives a reductive subgroup of G. It is irreducible if no representative is contained in a non-trivial parabolic subgroup of G. Corlette's existence result on harmonic maps is formulated in this setting [9].

Theorem 2.22 (G-NAH I)**.** *There exists a harmonic section $h : S \to M$ if and only if the holonomy of M is reductive. If the holonomy is irreducible, then h is unique.*

Let Q be the bundle whose fiber at $z \in S$ is the G-torsor of isometries from N to M_z. Then Q is a G-principle bundle which inherits a flat connection A (an equivariant $\mathfrak{g}$-valued 1-form) from the flat connection on M. If we fix once and for all a point h_0 in N with stabilizer K and corresponding decomposition $\mathfrak{g} = \mathfrak{k} + \mathfrak{p}$, then a section h of M gives a reduction Q^h of the structure group of Q to K, whose fiber at z is the subset of isometries sending h_0 to $h(z)$. The pullback of the vertical tangent bundle of M by the section h is therefore identified with $Q^h \times_K \mathfrak{p}$, with K acting on $\mathfrak{p}$ by conjugation. Write ∇^h for the pullback of the Levi-Civita connection by h to $Q^h \times_K \mathfrak{p}$, given by Proposition 2.1. The harmonicity of h is equivalent to the holomorphicity of ∂h with respect to the $(0,1)$-part of ∇^h. Similarly, the conformality of h with respect to the metric ν on N is equivalent to the condition $\nu((\partial h)^2) := \nu(\partial h \otimes \partial h) = 0$.

Definition 2.23. We say that (M, h) is a G-harmonic bundle if h is harmonic. It is a G-minimal bundle with respect to a metric ν if h is minimal with respect to ν.

This leads us to a definition of a G-Higgs bundle. Let (G, K) be any admissible pair.

Definition 2.24. A G-Higgs bundle is a pair $(P, A^{0,1}, \phi)$, where P is a holomorphic principal $K^{\mathbb{C}}$ bundle on S with compatible $(0,1)$-connection $A^{0,1}$, and ϕ is a holomorphic section of $\mathrm{ad}\mathfrak{p}^{\mathbb{C}} \otimes \mathcal{K}$ called the Higgs field.

Here we abbreviate $\mathrm{ad}\mathfrak{p}^{\mathbb{C}} = P \times_{K^{\mathbb{C}}} \mathfrak{p}^{\mathbb{C}}$.

The other direction of NAH is also true in this context. A reduction P^c of the structure group of P to $K \subset K^{\mathbb{C}}$ defines a real structure on the Lie algebra $\mathrm{ad}\mathfrak{k}^{\mathbb{C}}$, which we write as $*_c$. The Chern connection of $A^{0,1}$ with respect to this reduction is $A^{0,1} + (A^{0,1})^{*_c}$, and its curvature F^c is a real $\mathrm{ad}\mathfrak{k}^{\mathbb{C}}$-valued 2-form.

In addition, P^c defines a real structure on $\mathrm{ad}\mathfrak{p}^{\mathbb{C}}$. The connection $\theta = A^{0,1} + (A^{0,1})^{*_c} + \phi + \phi^{*_c}$ is flat if and only if the curvature F^c satisfies

$$F^c - [\phi, \phi^{*_c}] = 0$$

Stability and polystability for G-Higgs bundles are defined analogously to the $\mathrm{PGL}(n, \mathbb{C})$ case (see [17]). The analog of Proposition 2.12 remains true for G-Higgs bundles. The following is proved in [5].

Theorem 2.25 (G-NAH II)**.** *Let $(P, A^{0,1}, \phi)$ be a stable G-Higgs bundle. There exists a unique reduction of the structure group P^c such that $(P^c \times_K G, A^{0,1} + (A^{0,1})^{*_c} + \phi + \phi^{*_c})$ is a flat G-principal bundle (h therefore defines a harmonic section of the associated flat bundle).*

Definition 2.26. A stable G-Higgs bundle $(P, A^{0,1}, \phi)$ is called minimal with respect to an invariant bilinear form ν on the symmetric space G/K if $\nu(\phi^2) = 0$.

Let's give a word on the relation to ordinary Higgs bundles. First, suppose that H is a complex Lie group. Since we only define G-Higgs bundles for real Lie groups G, to make sense of an H-Higgs bundle, we must first restrict scalars from $\mathbb{C}$ to $\mathbb{R}$. The complexification of $G = \mathrm{Res}_{\mathbb{C}/\mathbb{R}} H$ is isomorphic to $H \times \overline{H}$. With repsect to any Cartan involution, both $\mathfrak{k}^{\mathbb{C}}$ and $\mathfrak{p}^{\mathbb{C}}$ are isomorphic to the Lie algebra $\mathfrak{h}$. In this way, we recover the usual notion of an H-Higgs bundle for a complex group H.

Next, if $(P, A^{0,1}, \phi)$ is a $\mathrm{PGL}(n, \mathbb{C})$-Higgs bundle, the obstruction to holomorphically extending the structure group of P to $\mathrm{GL}(n, \mathbb{C})$ is a class in $H^2(S, \mathcal{O}^*)$, where $\mathcal{O}^*$ is the sheaf of invertible holomorphic functions. This space is trivial. We have to choose an extension, but having done so, we get a classical Higgs bundle with E the associated vector bundle. A reduction of the structure group from $\mathrm{PGL}(n, \mathbb{C})$ to $\mathrm{PU}(n)$ determines a Hermitian metric on E up to scaling by a real-valued function on S.

Finally, if G and G' are two real semisimple groups, a homomorphism $\sigma : G \to G'$ turns a G-Higgs bundles into a G'-Higgs bundle. If $(P, A^{0,1}, \phi)$ is stable, then $(\sigma(P), \sigma(\phi))$ will be polystable, but may not be stable (see [4]). By the discussion on polystability in 2.3, the failure of stability is not a big concern.

Putting together the previous three paragraphs, this gives a way of associating a classical Higgs bundle to a G-Higgs bundle provided we are given a representation of G into $\mathrm{PGL}_n(\mathbb{C})$. We will use this association in order to apply the asymptotic decoupling estimates to G-Higgs bundles, under the assumption that the representation is essentially faithful.

3. Convergence of the second variation along the $\mathbb{R}^+$-flow

The purpose of this section is to prove Theorem B.

3.1. **The second variation of area.** Before studying the second variation of area for maps to symmetric spaces, we give a general formula for the second variation of area for a minimal immersion to a Riemannian manifold. Let $h : S\times(-\epsilon,\epsilon)\to N$ be a one-parameter variation of a minimal map. We write h for the initial map and hope there is no confusion. Let ∇ be the connection on $S\times(-\epsilon,\epsilon)$ obtained by pulling back the Levi-Civita connection from N. We define sections $\dot{h}$ and $\ddot{h}$ of $h^*TN|_{S\times\{0\}}$ by

$$\dot{h} = dh\Big(\frac{\partial}{\partial t}\Big)|_{t=0},\ \ddot{h} = \Big(\nabla_{\frac{\partial}{\partial t}} dh\Big(\frac{\partial}{\partial t}\Big)\Big)|_{t=0}.$$

Similarly, we have 1-forms with values in $h^*TM\otimes_{\mathbb{R}}\mathbb{C}$,

$$\partial\dot{h} = \Big(\nabla^{1,0} dh\Big(\frac{\partial}{\partial t}\Big)\Big)|_{t=0},\ \overline{\partial}\dot{h} = \Big(\nabla^{0,1} dh\Big(\frac{\partial}{\partial t}\Big)\Big)|_{t=0}.$$

Let $a(t)$ be the area form induced on S from the map h at time t, and let $\ddot{a} = \frac{d^2}{dt^2}\big|_{t=0} a(t)$. In this context, the standard second variation of area formula takes the following form. We remark that we make no assumption about $\dot{h}$ being a normal variation.

Proposition 3.1 (Second variation of area)**.** *If h is a minimal immersion at time 0, then the second derivative of the induced area form at time 0 is given by*

$$\frac{1}{2}\ddot{a} = \langle\partial\dot{h},\overline{\partial}\dot{h}\rangle + \langle R(\dot{h},\partial h)\dot{h},\overline{\partial}h\rangle - \frac{2|\langle\partial\dot{h},\partial h\rangle|^2}{\langle\partial h,\overline{\partial}h\rangle} + \frac{1}{2}d(\langle\ddot{h},\overline{\partial}h\rangle - \langle\partial h,\ddot{h}\rangle) \tag{8}$$

This should be a standard formula. We give a proof for the reader's convenience.

Proof. From equation (6) we have $a(t) = \sqrt{\langle\partial h_t,\overline{\partial}h_t\rangle^2 - |\langle\partial h_t,\partial h_t\rangle|^2}$. Since h is conformal, $\langle\partial h_t,\partial h_t\rangle$ is $O(t)$, so to find the second derivative in t we can Taylor expand:

$$a(t) = \langle\partial h_t,\overline{\partial}h_t\rangle - \frac{|\langle\partial h_t,\partial h_t\rangle|^2}{2\langle\partial h,\overline{\partial}h\rangle^2} + O(t^3) \tag{9}$$

Since the Levi-Civita connection is torsion-free, we get that

$$\Big(\nabla_{\frac{\partial}{\partial t}}\partial h_t\Big)|_{t=0} = \partial\dot{h}$$
$$\Big(\nabla_{\frac{\partial}{\partial t}}\nabla_{\frac{\partial}{\partial t}}\partial h_t\Big)|_{t=0} = R(\dot{h},\partial h)\dot{h} + \partial\ddot{h}$$

and similarly for $\overline{\partial}h_t$. We compute each term separately.

$$\begin{aligned}\frac{d^2}{dt^2}\Big|_{t=0}\langle\partial h_t,\overline{\partial}h_t\rangle &= 2\langle\nabla_{\frac{\partial}{\partial t}}\partial h_t,\nabla_{\frac{\partial}{\partial t}}\overline{\partial}\,h_t\rangle + \langle\nabla_{\frac{\partial}{\partial t}}\nabla_{\frac{\partial}{\partial t}}\partial h_t,\overline{\partial}\,h_t\rangle + \langle\partial h_t,\nabla_{\frac{\partial}{\partial t}}\nabla_{\frac{\partial}{\partial t}}\overline{\partial}\,h_t\rangle\Big|_{t=0}\\ &= 2\langle\partial\dot{h},\overline{\partial}\dot{h}\rangle + \langle R(\dot{h},\partial h)\dot{h},\overline{\partial}h\rangle + \langle\partial h,R(\dot{h},\overline{\partial}h)\dot{h}\rangle + \langle\partial\ddot{h},\overline{\partial}h\rangle + \langle\partial h,\overline{\partial}\ddot{h}\rangle\end{aligned}$$

By the symmetries of the Riemannian curvature tensor, the two curvature terms are equal. For the rightmost two terms, we have

$$\begin{aligned}\langle\partial\ddot{h},\overline{\partial}h\rangle + \langle\partial h,\overline{\partial}\ddot{h}\rangle &= \partial\langle\ddot{h},\overline{\partial}h\rangle - \overline{\partial}\langle\partial h,\ddot{h}\rangle\\ &= d(\langle\ddot{h},\overline{\partial}h\rangle - \langle\partial h,\ddot{h}\rangle)\end{aligned}$$

Finally, using that h is conformal, the computation of the second term of (9) simplifies to give

$$\frac{d^2}{dt^2}\Big|_{t=0}\frac{|\langle\partial h_t,\partial h_t\rangle|^2}{2\langle\partial h,\overline{\partial}h\rangle^2} = \frac{4|\langle\partial\dot{h},\partial h\rangle|^2}{\langle\partial h,\overline{\partial}h\rangle}.$$

Putting this all together yields the proposition. □

If $h : S \to N$ is any minimal map, not necessarily an immersion, define $\mathrm{Var}(h)$ to be the space of smooth sections of h^*TN with compact support on S. The stability form Q_h of h is the quadratic form on $\mathrm{Var}(h)$ defined for $X \in \mathrm{Var}(h)$ by

$$Q_h(X) = 2\int_S \langle \partial X, \overline{\partial} X\rangle - \frac{2|\langle \partial X, \partial h\rangle|^2}{\langle \partial h, \overline{\partial} h\rangle} + \langle R(X, \partial h)X, \overline{\partial} h\rangle \tag{10}$$

Proposition 3.2. *If h is a minimal map and $h(t)$ is a variation with $\dot{h} = X \in \mathrm{Var}(h)$, and $A(t)$ the total area of a compact subsurface containing the support of X and of $\ddot{h}$, then*

$$\left.\frac{d^2}{dt^2}\right|_{t=0} A(t) = Q_h(X).$$

Proof. The area form is non-degenerate on the open set $U \subset S$ on which h is an immersion, and so we can apply the formula (8) on U. But it is easily checked that both the area form and the right hand side of the formula (8) smoothly extend to all of S. The result follows from dominated convergence. The final term in formula (8) vanishes upon integration by the divergence theorem. □

Definition 3.3. We call X a destabilizing variation of h if $Q_h(X) < 0$. h is called unstable if it has a destabilizing variation. The index $\mathrm{Ind}(h)$ of h is the maximal dimension of a space of destabilizing variations.

The log cutoff trick (see [29, section 4.4] for details and explanation) implies the following proposition, which is essential to the proof of Theorem B.

Proposition 3.4. *If $U = S - B$, for B a discrete set, then $\mathrm{Ind}(h|_U) = \mathrm{Ind}(h)$.*

Another elementary proposition of a similar flavor is:

Proposition 3.5. *If N has nonpositive sectional curvature, $\iota : N' \to N$ is the inclusion of a totally geodesic subspace, and $h = \iota \circ h'$, then $\mathrm{Ind}(h) = \mathrm{Ind}(h')$.*

If h is a minimal section of a flat Riemannian bundle M over S, then the second variation formula and Proposition 3.2 still hold, as one checks in local flat charts, the only difference being that the bundle h^*TM is replaced everywhere by the bundle $h^*T^{\mathrm{vert}}M$, and the Levi-Civita connection on a fiber M_z is replaced by its natural extension to the whole bundle $T^{\mathrm{vert}}M$. For instance, ∂h is now a $(1,0)$-form valued in $h^*T^{\mathrm{vert}}M$. We extend definitions of variations, stability operator, and index to minimal sections in the obvious way. Propositions 3.4 and 3.5 also generalize immediate to sections of flat Riemannian bundles.

3.1.1. *Minimal harmonic bundles.* We record the stability form for a minimal harmonic bundle (E, D, H). We use the isomorphism φ to identify $\mathrm{Var}(H)$ with the space of smooth sections of $\mathrm{End}^0_H(E)$. One way of thinking about this identification is that if X is a smooth section of $\mathrm{End}^0_H(E)$, then the one-parameter family of metrics

$$H^t = e^{-t\overline{X}^\vee} H e^{-tX}$$

has derivative X at time zero.

If X is a variation, then ∂X should takes values in the complexified tangent bundle, so it becomes an $\mathrm{End}^0(E)$-valued 1-form. Recall that under this isomorphism, the (complexified) Riemannian metric becomes $\langle X, Y\rangle = 2n\,\mathrm{tr}(XY)$. We write the corresponding Hermitian norm as $|X|^2_H = 2n\,\mathrm{tr}(XX^{*_H})$. The operator ∂ on this bundle is the $(1,0)$-part of ∇^H defined in 2.3.

Proposition 3.6. *Let (E, D, H) be a minimal harmonic bundle on a compact surface S with associated Higgs field ϕ. If X is a smooth section of $\mathrm{End}^0_H(E)$, then the stability form of the section $[H]$ of the associated flat Riemannian bundle $\mathbb{P}\mathrm{Met}(E)$ is*

$$Q_H(X) = 2\int_S |\partial X|_H^2 - \frac{2|\langle \partial X, \phi\rangle|^2}{|\phi|_H^2} + |[X,\phi]|_H^2.$$

Proof. By definition, $\phi = \varphi(\partial H)$, and since the identification φ is implicit, we should replace "∂h" in Equation (10) with ϕ. The curvature term becomes

$$\langle R(X,\phi)X, \phi^{*_H}\rangle = 2n\,\mathrm{tr}(-[[X,\phi],X]\phi^{*_H}) = 2n\,\mathrm{tr}([X,\phi][\phi^{*_H}, X]) = |[X,\phi]|_H^2.$$

□

3.2. **The limiting objects.** Generically regular semisimple Higgs bundles are introduced in [30]. We slightly relax the condition. Let $(E, \overline{\partial}_E, \phi)$ be a Higgs bundle.

Definition 3.7. ϕ is generically semisimple if it is semisimple away from a discrete subset of S.

If ϕ is generically semisimple, then there exists a discrete subset $B \subset S$ and an integer $m \leq n$ such that for every $p \in S - B$, there exists a neighbourhood $U \subset S - B$ of p over which $(E, \overline{\partial}_E, \phi)$ splits in the sense that we have m holomorphic 1-forms $\phi_1, \dots, \phi_m$ such that $\phi_i \neq \phi_j$ at every point in U for $i \neq j$, and a decomposition of the Higgs bundle $E|_U = \oplus_{i=1}^m E_i$, with respect to which ϕ decomposes as $\phi = \sum_{i=1}^m \phi_i \pi_i$, where π_i is the projection onto E_i with respect to the given decomposition of $E|_U$. A generically semisimple Higgs bundle is generically regular semisimple if $m = n$.

Definition 3.8. The critical set B of ϕ is the smallest set such that $(E, \overline{\partial}_E, \phi)$ splits into m distinct eigenspaces in a neighborhood of every point of $S - B$.

Let $(E, \overline{\partial}_E, \phi)$ be a generically semisimple Higgs bundle, and let B be the critical set of ϕ. For any neighborhood U, define ϕ_i and π_i as above. Over $S - B$, these patch together to form an m-dimensional subbundle $F_\phi \subset \mathrm{End}(E)|_{S-B}$ locally spanned by the projections π_i. Let F^0_ϕ be the codimension one subbundle $F^0_\phi = F \cap \mathrm{End}^0(E)$.

Definition 3.9. We call F^0_ϕ the toral bundle of ϕ and F_ϕ the extended toral bundle.

Remark 3.10. Equivalently, the fiber of F^0_ϕ over each point p of $S - B$ is the center of the centralizer of $\phi(p)$ in $\mathrm{End}^0(E)$. If ϕ is regular semisimple, this centralizer is already abelian.

Fix $(E, \overline{\partial}_E, \phi)$, and let F be its extended toral bundle. F has a natural flat connection, flat metric, and real structure defined by the condition that the π_i are real and flat. If X is a section of F, let $dX = \partial X + \overline{\partial} X$ be its derivative with respect to the flat connection, and let $X^\dagger$ be its conjugate with respect to the flat structure. If $X = \sum_i X_i \pi_i$, then

$$X^\dagger = \sum_i \overline{X_i}\pi_i$$

and

$$dX = \sum_i dX_i \pi_i.$$

Since the complex metric on F is just the restriction $\langle X, Y\rangle = 2n\,\mathrm{tr}(XY)$ from $\mathrm{End}^0(E)$, the corresponding Hermitian metric is given by $|X|^2 = 2n\,\mathrm{tr}(XX^\dagger)$. If r_i is the rank of E_i, then

$$\langle \pi_i, \pi_j^\dagger\rangle = 2n\,\mathrm{tr}(\pi_i\pi_j^\dagger) = 2nr_i\delta_{ij}. \tag{11}$$

The toral bundle F^0 of ϕ inherits the flat, metric, and real structures from F. Let $F^0_{\mathbb{R}}$ be its real subbundle, i.e., the $\mathbb{R}$-span of the π_i. Since the π_i's are holomorphic, the flat connection is compatible with the holomorphic structure on $\mathrm{End}(E)$, and hence ϕ can be viewed as a holomorphic F^0-valued 1-form. Let $\psi = \phi + \phi^\dagger$. Since ψ is a closed $F^0_{\mathbb{R}}$-valued 1-form, it defines an affine bundle M_ϕ whose vertical tangent bundle is $F^0_{\mathbb{R}}$ together with a section f of M_ϕ. The construction is as follows. Let $(U_\alpha, w_{\alpha\beta})$ be a presentation of $F^0_{\mathbb{R}}$ by an open cover and locally constant transition functions, and (ψ_α) a representative of ψ, with $\psi_\alpha = w_{\alpha\beta}\psi_\beta$. On each U_α, choose f_α with $df_\alpha = \psi_\alpha$. Since $dw_\alpha = 0$, the functions $t_{\alpha\beta} := f_\alpha - w_{\alpha\beta} f_\beta$ are locally constant. Rearranging this gives

$$f_\alpha = w_{\alpha\beta} f_\beta + t_{\alpha\beta}.$$

The affine transformations $x \mapsto w_{\alpha\beta} x + t_{\alpha\beta}$ define the bundle M_ϕ, and (f_α) defines a section f such that $df = \psi$, and therefore $\partial f = \phi$.

Definition 3.11. We call M_ϕ the appartment bundle associated to $(E, \overline{\partial}_E, \phi)$.

By equation (11), the natural metric on $F^0_{\mathbb{R}}$ is parallel with respect to the flat connection, and therefore the appartment bundle M_ϕ is a flat Riemannian bundle. The section f is harmonic by construction. Assume that $\mathrm{tr}(\phi^2) = 0$, so that the Higgs bundle $(E, \overline{\partial}_E, \phi)$ is minimal. Then the section f is also minimal.

Definition 3.12. The pair (M_ϕ, f) is the limiting object associated to the stable minimal generically semisimple Higgs bundle $(E, \overline{\partial}_E, \phi)$. Its index is by definition the index of f.

A variation $X \in \mathrm{Var}(f)$ is a smooth section of $F^0_{\mathbb{R}}$ with compact support on $S - B$. Since the curvature of the fibers of M_ϕ vanishes, the stability form of f is given by

$$Q_f(X) = 2\int_S |\partial X|^2 - \frac{2|\langle \partial X, \phi\rangle|^2}{|\phi|^2} \tag{12}$$

3.2.1. *Cameral covers.* We now give an equivalent description of the limiting object, which is a slight modification of the equivariant map associated to the section f as in section 1.4. The monodromy of F^0 only permutes the projections π_i (and moreover can only permute any two projections π_i of the same rank), and hence it is contained in a (finite) product of symmetric groups. Therefore there is a finite and minimal degree branched covering $\tau : C \to S$, unbranched over $S - B$, which trivializes the monodromy of F^0. This is the cameral cover of (S, ϕ) (see [21, Section 3] and [10]).

Over $C - \tau^{-1}(B)$, there is a global splitting $\tau^* E|_{S-B} = \oplus_i E_i$ with projections π_i, and $\tau^*\phi|_{S-B}$ acts on each E_i by multiplication by a holomorphic 1-form ϕ_i. We extend the (bounded) 1-forms ϕ_i via the removable singularities theorem to all of C. Each harmonic 1-form $\psi_i = \phi_i + \overline{\phi_i}$ determines a cohomology class, of which it is the unique harmonic representative. Choosing a basepoint z_0 on C, the cohomology class defines a representation from $\pi_1(C)$ to $\mathbb{R}$. Lifting to a harmonic 1-form $\tilde{\psi}_i$ on the universal cover $\tilde{C}$ and integrating from a lift of the basepoint gives a harmonic function

$$\tilde{f}_i(z) = \int_{z_0}^{z} \tilde{\psi}_i$$

which is equivariant by this representation. Scaling each $\tilde{f}_i$ by $\sqrt{r_i}$ and putting them together, we have a representation $\chi : \pi_1(C) \to \mathbb{R}^m$ and a χ-equivariant minimal map $\tilde{f}$ from $\tilde{C}$ to $\mathbb{R}^m$.

Define the equivariant area of $\tilde{f}$ to be the integral of its area form over a fundamental domain of the action of $\pi_1(C)$ on $\tilde{C}$. Note that this is $\deg(\tau)$ times the area of the section

f of M_ϕ. Since the representation χ acts by translations, a variation of the map $\tilde{f}$ that preserves χ-equivariance is the same thing as a χ-invariant map from $\tilde{C}$ to $\mathbb{R}^m$. Of course, this is the same thing as a map from C to $\mathbb{R}^m$. The stability operator for the equivariant area is

$$Q_{\tilde{f}}(X) = 2\deg(\tau)\int_C |\partial X|^2 - \frac{2|\langle \partial X, \partial \tilde{f}\rangle|^2}{|\partial f|^2}.$$

Call a variation X τ-equivariant if it is the pullback of a section of $F^0_{\mathbb{R}}$ on S. Define the (τ,χ) equivariant index of $\tilde{f}$ to be the index of the stability operator restricted to τ-equivariant variations.

The log cut-off trick (Proposition 3.4) shows that we need only consider variations to be supported on $C-\tau^{-1}(B)$. But a τ-equivariant varation of $\tilde{f}$ supported on $C-\tau^{-1}(B)$ is the same as a variation of the section f, and the resulting area forms and stability operators are the same up to a factor of $\deg(\tau)$. We conclude:

Proposition 3.13. *The index of f is the same as the (τ,χ)-equivariant index of $\tilde{f}$.*

3.3. **Convergence under the $\mathbb{R}^+$ action.** There is a $\mathbb{C}^*$-action on the moduli space of Higgs bundles,

$$\gamma\cdot(E,\overline{\partial}_E,\phi) = (E,\overline{\partial}_E,\gamma\phi).$$

By Proposition 2.12, the action restricts to the space of stable Higgs bundles. In this paper, we are only interested in the restricted $\mathbb{R}^+$-action. Since stable Higgs bundles have unique harmonic metrics, the non-abelian Hodge correspondence transfers the $\mathbb{R}^+$-action to an $\mathbb{R}^+$-action on the moduli space of projectively flat bundles with irreducible holonomy. Let $(E,\overline{\partial}_E,\phi)$ be a minimal stable Higgs bundle. By taking $R\mapsto(E,\overline{\partial}_E,R\phi)$, we obtain a family of minimal harmonic bundles $R\mapsto(E,D_R,H_R)$ with holonomy representations ρ_R.

Remark 3.14. The apartment bundle $M_{R\phi}$, and the harmonic section $f_{R\phi}$ depend on R because the translational part is being rescaled by R. However, the special toral bundle $F^0_{R\phi}$ of $(E,R\phi)$ is the same for every R, and hence the space of variations of f_R is well-defined independent of R. The stability operator is independent of R as well.

Recall that the statement of Theorem B is that the liminf as $R\to\infty$ of the indices of the minimal harmonic bundles (E,D_R,H_R) is at least the index of the limiting object (M_ϕ,f). Let Q_{H_R} be the stability operator of (E,D_R,H_R), and let Q_f be the stability operator of (M_ϕ,f). Recall that we can identify $\mathrm{Var}(H_R)$ with the space of smooth sections of $\mathrm{End}^0_{H_R}(E)$, whereas $\mathrm{Var}(f)$ is the space of smooth sections of $F^0_{\mathbb{R}}$ with compact support on $S-B$. In order to relate the operators Q_{H_R} to Q, we need to turn variations of f into variations of H_R. Fortunately, there is a completely natural way to do this. If X is a variation of f, it is in particular a section of $\mathrm{End}^0(E)$, so we can simply project to $\mathrm{End}^0_{H_R}(E)$. Namely, let X_R be the section of $\mathrm{End}^0_{H_R}(E)$ defined by

$$X_R = \frac{X+X^{*_R}}{2} \tag{13}$$

where $*_R$ is an abbreviation of $*_{H_R}$. We will prove

Proposition 3.15. *For any minimal stable Higgs bundle $(E,\overline{\partial}_E,\phi)$ with ϕ generically semisimple, and any $X\in\mathrm{Var}(f)$, we have*

$$\lim_{R\to\infty} Q_{H_R}(X_R) = Q_f(X).$$

Once we have proved this, the proof of Theorem B, at least in the case of $\mathrm{PGL}_n(\mathbb{R})$, follows immediately by applying the following elementary lemma to $Q_R = Q_{H_R}$ and $Q = Q_f$.

Lemma 3.16. *For $R \in \mathbb{R}^+$, let V_R be a family of vector spaces with quadratic forms Q_R. If V is a vector space with quadratic form Q of finite index, and for each R, there is a linear map $X \mapsto X_R$ from V to V_R such that $\lim_{R\to\infty} Q_R(X_R) = Q(X)$ for each R, then $\liminf_{R\to\infty} \operatorname{Ind}(Q_R) \geq \operatorname{Ind}(Q)$.*

Proof. Let W be a k-dimensional subspace of V on which Q is negative definite, where $k = \operatorname{Ind}(Q)$. For each R, $X \mapsto Q_R(X_R)$ is a quadratic form on the finite dimensional space W, so the pointwise convergence to $Q(X)$ implies locally uniform convergence on W. Hence, for large enough R, $Q_R(X_R)$ is also negative definite. For such R, the dimension of $\{X_R | X \in W\}$ must also be k, since otherwise we would have $Q_R(X_R) = 0$ for some vector X, but this is impossible as $Q_R(X_R)$ is negative definite. Hence, for such R, we have $\operatorname{Ind}(Q_R) \geq k$. □

Proving Proposition 3.15 requires an analysis of the behaviour of the minimal harmonic bundles (E, D_R, H_R) as R tends to infinity, and for this we turn to Mochizuki's work [30]. In order to slightly simplify the statements of results below, as well as the proof of Proposition 3.15, we fix a metric σ_0 on $S - B$ to define the norm of 1-forms.

Theorem 3.17 (Proposition 2.3 and 2.10 in [30])**.** *Suppose that over a domain U in S, the bundle E splits as $E = \bigoplus_i E_i$ and ϕ acts on E_i by multiplication by a 1-form ϕ_i. Let π_i be the projection to E_i determined by the splitting, and let $(\pi_i')_R$ be the Hermitian orthogonal projection to E_i. If U' is any relatively compact subdomain of U, then there are constants C_1 and ϵ depending on U, U', the rank of E, and the 1-forms ϕ_i such that*

$$|\pi_i - (\pi_i')_R|_{H_R} < C_1 e^{-\epsilon R} \tag{14}$$

and

$$|\partial_{H_R} \pi_i|_{H_R,\sigma_0} < C_1 e^{-\epsilon R} \tag{15}$$

on U'.

Mochizuki gives a precise description of the constant ϵR, in terms of the gaps between the eigenvalues of $R\phi_i$. All we need is the linear dependence on R. We remark that if E is polystable, the constants C_1 and ϵ do not depend on the choice of harmonic metric H_R.

The second point is stated in Mochizuki with the additional assumption that the E_i are one-dimensional. But the only place he uses it is in the short proof of his Theorem 2.9, and the only property he uses is that ϕ commutes with endomorphisms of E_i. This is true in our case because we specified that ϕ acts on E_i by a multiple of the identity. This slight generalization will be useful for applying Mochizuki's results to groups other than $\mathrm{PGL}(n, \mathbb{C})$ in section 3.4.

We recall that if Z is a section of the toral bundle F^0_ϕ, then $Z^\dagger$ and $dZ = \partial Z + \overline{\partial}\, Z$ are the conjugate and derivative respectively of Z with respect to the real structure and flat connection on F^0_ϕ.

In the following corollary, so as to reduce the number of subscripts, we abbreviate H_R by R in the adjoint, the norm, and the connection.

Corollary 3.18. *Let $(E, \overline{\partial}_E, \phi)$ be a generically semisimple stable Higgs bundle with critical set B and toral bundle $F^0 \subset \operatorname{End}^0(E)|_{S-B}$. Let Z be a smooth section of F^0 on $S - B$ and let U be a subdomain of $S - B$ with compact closure. Then there are constants $C_2, \epsilon > 0$ depending on $(E, \overline{\partial}_E, \phi), U, |Z|$, and $|\overline{\partial}\, Z|_{\sigma_0}$ such that*

(i) $|Z^{*_R} - Z^\dagger|_R$
(ii) $|\partial_R Z^{*_R} - \partial Z^\dagger|_{R,\sigma_0}$

(iii) $|\partial_R Z - \partial Z|_{R,\sigma_0}$, *and*
(iv) $||Z|_R - |Z||$
are all bounded by $C_2 e^{-\epsilon R}$ *on* U.

Proof. First suppose that $(E, \overline{\partial}_E, \phi)$ satisfies the splitting condition of Theorem 3.17 on a neighborhood of U, and write $Z = \sum_i Z_i \pi_i$.

(i) Let $(\rho_i)_R = \pi - (\pi_i')_R$. Since $(\pi_i')_R$ is self-adjoint with respect to H_R, we have $(\rho_i)_R^{*_R} = \pi_i^{*_R} - (\pi_i')_R$ and so
$$|\pi_i - \pi_i^{*_R}|_R = |(\rho_i)_R - (\rho_i)_R^{*_R}|_R \leq 2|(\rho_i)_R|_R \leq 2C_1 e^{-\epsilon R}$$
by equation (14) from Theorem 3.17. So for some constant C_2,
$$|Z^{*_R} - Z^\dagger|_R \leq \sum_i |Z_i||\pi_i^{*_R} - \pi_i|_R \leq C_2 e^{-\epsilon R}.$$

(ii) Writing $Z = \sum_i Z_i \pi_i$ and using $\partial_R \pi_i^{*_R} = 0$, we see
$$|\partial_R Z^{*_R} - \partial Z^\dagger|_{R,\sigma_0} \leq \sum_i |\partial \overline{Z}_i|_{\sigma_0} |\pi_i^{*_R} - \pi_i|_R \leq C_2 e^{-\epsilon R}.$$

(iii) Writing $Z = \sum_i Z_i \pi_i$ yet again,
$$|\partial_R Z - \partial Z|_{R,\sigma_0} \leq \sum_i |Z_i||\partial_R \pi_i|_{R,\sigma_0} \leq C_2 e^{-\epsilon R}$$
where the last inequality is equation (15) from Theorem 3.17.

(iv) We have $|Z|_R^2 - |Z|^2 = \sum_{i,j} Z_i \overline{Z_j}(\mathrm{tr}(\pi_i \pi_j^{*_R}) - \mathrm{tr}(\pi_i \pi_j))$, and
$$|(\mathrm{tr}(\pi_i \pi_j^{*_R}) - \mathrm{tr}(\pi_i \pi_j))|_R \leq |\pi_i|_R |\pi_j^{*_R} - \pi_j|_R.$$
Recall $r_i = \mathrm{rank}(E_i)$. Then expanding the inner product, we find
$$|\pi_i - \pi_i^{*_R}|_R^2 = 2|\pi_i|_R^2 - 2r_i,$$
from which we conclude that $|\pi_i|_R$ is uniformly bounded in R. Thus
$$|\pi_i|_R |\pi_j^{*_R} - \pi_j|_R \leq C_2 e^{-\epsilon R}$$
for some value of C_2.

In general, we can cover U by finitely many sets U_α' such that the conditions of Theorem 3.17 are satisfied on a neighborhood of U_α', and take ϵ (resp. C_2) to be the minimum (resp. maximum) of the value for each U_α'. □

We may now prove Proposition 3.15.

Proof of Proposition 3.15. From Proposition 3.6,
$$Q_{H_R}(X_R) = 2\int_S |\partial_R X_R|_R^2 - \frac{2|\langle \partial_R X_R, \phi\rangle|^2}{|\phi|_R^2} + R^2 |[X_R, \phi]|_R^2.$$
On the other hand from (12) we have
$$Q_f(X) = 2\int_S |\partial X|^2 - \frac{2|\langle \partial X, \phi\rangle|^2}{|\phi|^2}.$$

We prove uniform convergence of each of the first two terms, and uniform convergence of the third term to zero. Starting with the first,
$$\big||\partial_R X_R|_R - |\partial X|_R\big| \leq |\partial_R X_R - \partial X|_R \leq \left|\frac{\partial_R X}{2} - \frac{\partial X}{2}\right|_R + \left|\frac{\partial_R X^{*_R}}{2} - \frac{\partial X}{2}\right|_R$$

which is $O(e^{-\epsilon R})$ by (ii) and (iii) of Corollary 3.18, and so combining with (iv) we see that $|\partial_R X_R|_R - |\partial X|$ is also $O(e^{-\epsilon R})$.

For the second term we have exponential convergence of the denominators by (iv) and for the numerators,

$$\big||\langle \partial_R X_R, \phi\rangle| - |\langle \partial X, \phi\rangle|\big| \leq |\langle \partial_R X_R - \partial X, \phi\rangle| \leq |\partial_R X_R - \partial X|_R |\phi|_R.$$

The term $|\phi|_R$ is uniformly bounded by (iv) and $|\partial_R X_R - \partial X|_R$ is exponentially decreasing by the previous equation.

Finally, since every section of F_0 commutes with ϕ, we have $[X, \phi] = 0$ and so

$$|[X_R, \phi]| = |[\frac{X^{*_R} - X}{2}, \phi]| \leq \frac{1}{2}|X^{*_R} - X|_R |\phi|_R$$

Since $|\phi|_R$ is uniformly bounded by (iv) and $|X^{*_R} - X|_R$ is exponentially small by (i), this term goes to zero (despite the factor of R^2, which is dwarfed by the exponential decay). □

Proof of Theorem B for $\mathrm{PGL}(n, \mathbb{R})$. Apply Proposition 3.15 and Lemma 3.16. □

3.4. **Convergence for G-Higgs bundles.** This subsection is not necessary for the proof of Theorem A for $\mathrm{PSL}(n, \mathbb{R})$.

Fix an admissible triple (G, K, ν). Let $(P, A^{0,1}, \phi)$ be a stable G-Higgs bundle that is minimal with respect to ν. Suppose that ϕ is generically semisimple.

For each $R > 0$, let (M_R, h_R) be the associated minimal G-harmonic bundle. The left hand side of Theorem B is about the index of h_R with respect to the space of smooth sections of $h_R^* T^{\mathrm{vert}} M_R$. Recall that this is identified with $P^c \times_K \mathfrak{p}$. In order to apply the results of the previous section, we use Proposition 2.20 to fix an essentially faithful representation $\sigma = \prod_i \sigma_i : G \to \prod_i \mathrm{SL}_{n_i}(\mathbb{C})$ of G and constants a_i such that the induced map of symmetric spaces is a totally geodesic isometry. By Proposition 3.5, since the product of symmetric spaces $\prod_i \mathbb{P}\mathrm{Met}(\mathbb{C}^{n_i})$ is nonpositively curved, we should be able to express the index of h_R using this isometry. In terms of harmonic bundles, this expectation is realized as follows.

For each i, let $(E_i, \overline{\partial}_{E_i}, \phi_i)$ be the Higgs bundle associated to $(P, A^{0,1}, \phi)$ by σ_i. It is possible that E_i not stable, but only polystable; if so, fix once and for all a relative scaling of its stable factors as in Remark 2.14, which determines for each $R \in \mathbb{R}^+$ a harmonic bundle (E_i, D_i^R, H_i^R). The bundle $\mathbb{P}\mathrm{Met}(E_i)$ inherits the metric $a_i\,\mathrm{tr}$, and for each R the induced map of flat Riemannian bundles from M_R to $\prod_i \mathbb{P}\mathrm{Met}(E_i)$ is a fiberwise totally geodesic isometry, sending h_R to $\prod_i [H_R^i]$. In particular the section $[H_R] := \prod_i [H_R^i]$ is minimal, and its index is the same as the index of h_R.

Let X_R be a variation of $[H_R]$, which we can write as $X_R = \sum_i X_i^R$, where each X_i^R is a section of $\mathrm{End}^0_{H_R^i}(E_i)$. The following two propositions are easy generalizations of Propositions 3.6 and 3.15 respectively.

Proposition 3.19. *The stability operator of H_R applied to X_R is*

$$Q_{H_R}(X_R) = \int_S \Big(\sum_i a_i |\partial_{H_R^i} X_i^R|^2_{H_R^i}\Big) - \frac{2\Big|\sum_i a_i \langle X_i^R, \phi_i\rangle\Big|^2}{\sum_i a_i |\phi_i|^2_{H_R^i}} + R^2\Big(\sum_i a_i |[X_i^R, \phi_i]|^2_{H_R^i}\Big).$$

Proof. This is a consequence of the formula (10), The curvature splits over the direct sum, so we can apply the same manipulations from Proposition 3.6 to get the last term. □

Now let B^σ be the union of the critical sets of each ϕ_i, and define F^σ_ϕ to be the fiber product of the toral bundles $F^0_{\phi_i}$ over $S - B^\sigma$ for each ϕ_i. Each $F^0_{\phi_i}$ inherits the rescaled complex metric $a_i \operatorname{tr}$ from $\operatorname{End}^0(E_i)$, and together these give a parallel metric on F^σ_ϕ. Let M^σ_ϕ be the fiber product of the apartment bundles M_{ϕ_i} for each ϕ_i, and f^σ the product of the canonical sections f_i. Then the minimality of ϕ with respect to ν implies that f^σ is minimal with respect to the metric on M^σ_ϕ. Let Q_{f^σ} be its stability operator.

A variation X of f^σ is a compactly supported section of F^σ_ϕ. Given X, for each R, let X^i_R be the H^i_R-self-adjoint part of X^i as in (13), and set $X_R = \sum_i X^i_R$. Then X_R is a variation of $[H_R]$, and we have:

Proposition 3.20. *For any section $X \in \operatorname{Var}(f^\sigma)$,*

$$\lim_{R\to\infty} Q_{H_R}(X_R) = Q_{f^\sigma}(X).$$

Proof. Exactly parallel to formula (12), the stability operator Q_{f^σ} is given by

$$Q_{f^\sigma}(X) = \int_S \Big(\sum_i a_i |\partial X_i|^2\Big) + \frac{\Big|\sum_i a_i \langle X_i, \phi_i\rangle\Big|^2}{\sum_i a_i|\phi_i|^2}.$$

In the proof of Proposition 3.15, we showed that each term of $Q_{H_R}(X_R)$ converges to the corresponding term of $Q_f(X)$, and even showed the separate convergence of the numerator and denominator of the second term. Hence, the same argument gives the convergence of $Q_{H_R}(X_R)$ to $Q_{f^\sigma}(X)$. □

Proposition 3.20 together with the index Lemma 3.16 therefore implies that

$$\liminf_{R\to\infty} \operatorname{Ind}(H_R) \geq \operatorname{Ind}(f^\sigma).$$

Since $\operatorname{Ind}(h_R) = \operatorname{Ind}(H_R)$, we have almost proved Theorem B. In fact, we have already acheived the main point which is to give a useful lower bound for the index of h_R. To complete the proof, it remains only to define the G-apartment bundle M^G_ϕ and its minimal section f, and show that f has the same index as f^σ.

3.4.1. *G-apartment bundles.* Our definition of M^G_ϕ is a natural generalization of the apartment bundle for a classical Higgs bundle, corresponding to the group $\mathrm{PGL}(n,\mathbb{C})$. We first define the G-toral bundle F^G_ϕ. It is convenient to define it first on a contractible open set $U \subset S$, on which P has been trivialized as $U \times K^{\mathbb{C}}$. This allows us to view the Higgs field ϕ as a $\mathfrak{p}^{\mathbb{C}}$-valued 1-form. Fix a maximal toral subspace $\mathfrak{a}$ of $\mathfrak{p}$, with roots Δ and Weyl group W.

Let $\mathcal{P}$ be the set of linear subspaces V of $\mathfrak{a}^{\mathbb{C}}$ which can be defined by the vanishing of a subset of Δ. If $V \in \mathcal{P}$, let V' be the open subset which is not contained in any smaller subspace in $\mathcal{P}$. Noting that W acts on $\mathcal{P}$, there is a W-invariant stratification of $\mathfrak{a}^{\mathbb{C}}$ indexed by the quotient $\mathcal{P}/W$ whose strata are the W-orbits of V' for $V \in \mathcal{P}$. Since it is W-invariant, it determines a stratification of the locus of semisimple elements of $\mathfrak{p}^{\mathbb{C}}$. Suppose that ϕ is semisimple on all of U. Since ϕ is holomorphic it will lie in some stratum away from a finite set B_U of U. The set B_U does not depend on the trivialization of P over U.

Definition 3.21. The critical set B of ϕ is the union of the set on which ϕ is not semisimple with each set B_U defined above for a covering U.

Let z be a point of $U - B_U$, and let $V \in \mathcal{P}$ be a subspace in the stratum determined by ϕ. That means that there is an element $k \in K^{\mathbb{C}}$ conjugating ϕ at z into V'. Since V'

is open in its stratum of $\mathfrak{a}^{\mathbb{C}}$, we can extend k to a gauge transformation, still called k, in a neighborhood U_z of z conjugating ϕ into V'. Over U_z, we want to define the G-toral bundle F_ϕ^G to be equal to $\mathrm{Ad}_{k^{-1}}V$.

We claim this is independent of the choice of k. Indeed, if k' were any other choice of gauge transformation, then $\mathrm{Ad}k'k^{-1}$ would act on $\mathfrak{a}^{\mathbb{C}}$ by an element of W. Because W acts by reflections in the root hyperplanes, the stabilizer in W of $\mathrm{Ad}_k\phi$ stabilizes the whole subspace V. Hence $\mathrm{Ad}_{(k')^{-1}}V = \mathrm{Ad}_{k^{-1}}V$. It follows that it is also independent of the trivialization of P, since any two trivializations also differ by a gauge transformation.

Since the definition of F_ϕ^G over U_z is independent of choices, these patch together to form a bundle over $S - B$ where $B = \cup_U B_U$ is by definition the critical set of ϕ. This is the G-toral bundle F_ϕ^G.

Remark 3.22. An equivalent definition of the fiber of F_ϕ^G over a point of $S - B$ is the intersection of $\mathrm{ad}\mathfrak{p}^{\mathbb{C}}$ with the center of the centralizer of ϕ in $\mathrm{ad}\mathfrak{p}^{\mathbb{C}}$.

The G-toral bundle F_ϕ^G has a canonical flat connection and real structure with respect to which the roots of $\mathfrak{a}^{\mathbb{C}}$ are real and parallel in any of the local identifications of F_ϕ^G with $V \subset \mathfrak{a}^{\mathbb{C}}$. Since it is a subbundle of $\mathrm{ad}\mathfrak{p}^{\mathbb{C}}$, it inherits the metric ν, which is parallel with respect to the flat connection. As before, ϕ is a holomorphic F_ϕ^G-valued 1-form, so by integrating its real part, we obtain an affine flat Riemannian bundle M_ϕ^G with section f. This is the G-apartment bundle. The proof of Theorem B is completed by

Proposition 3.23. $\mathrm{Ind}(f) = \mathrm{Ind}(f^\sigma)$.

Proof. The representation σ defines an embedding, which we also call σ, of $\mathrm{ad}\mathfrak{p}^{\mathbb{C}}$ to $\prod_i \mathrm{End}^0(E_i)$ which is compatible with the complex linear metrics on each. We need to understand the relationship between F_ϕ^G and $\sigma^{-1}(F_\phi^\sigma)$.

Fix a maximal toral subalgebra $\mathfrak{a}$ of $\mathfrak{p}$. Say that an element of $\mathfrak{a}^\vee$ is a root of σ if it is the difference of two weights of the same factor σ_i. The roots of σ are always invariant by the Weyl group W, and always contain the roots of $\mathfrak{g}$. Hence they define a W-invariant refinement of the stratification of $\mathfrak{a}^{\mathbb{C}}$ by $\mathcal{P}/W$, and in turn a refinement of the stratification of semisimple elemnts of $\mathfrak{p}^{\mathbb{C}}$. It is not hard to see from the definition of toral bundles that $\sigma^{-1}(F_\phi^\sigma)$ is built from this stratification in exactly the same way that F_ϕ^G is built from the original stratifcation; namely, the critical set B^σ is the finite set at which ϕ is in a smaller stratum, and with respect to a local trivialization near $z \in S - B^\sigma$ such that ϕ is contained in $\mathfrak{a}^{\mathbb{C}}$, the fiber of $\sigma^{-1}(F_\phi^\sigma)$ at z is the smallest subspace of $\mathfrak{a}^{\mathbb{C}}$ in the W-invariant stratification containing $\phi(z)$. This implies that $B \subset B^\sigma$, and that over $S - B^\sigma$, $\sigma^{-1}(F_\phi^\sigma) \subset F_\phi^G$.

First suppose that $\sigma^{-1}(F_\phi^\sigma) = F_\phi^G$. Since σ is injective on Lie algebras, it restricts to an isomorphism over $S - B^\sigma$ from F_ϕ^G to F_ϕ^σ. We have already observed that this isomorphism is an isometry with respect to the complex linear metric. Furthermore, since the weights of σ, which determine the real and flat structures of F_ϕ^σ, are constant real linear combinations of the roots of $\mathfrak{g}$, we see that the isomorphism also preserves the real and flat structures. Hence, σ defines an isomorphism of the apartment bundles M_ϕ^G and M_ϕ^σ intertwining the sections f and f^σ. The only difference is that variations of f^σ must vanish near B^σ, whereas variations of f need only vanish near B. By the log cutoff trick (Proposition 3.4), the index is the same.

In the general case where $\sigma^{-1}(F_\phi^\sigma)$ is only a subbundle of F_ϕ^G, it is still a locally constant subbundle since it is cut out by roots of σ, which are locally constant. Therefore, M_ϕ^σ becomes identified with a totally geodesic subbundle of M_ϕ^G. By Proposition 3.5, the index of f and f^σ is still the same.

□

Remark 3.24. Even though we have used Proposition 3.5 to compare indices several times, it was not really necessary. The destabilizing variations X of f^σ in M_ϕ^σ induced from variations of the section f of the G-apartment bundle M_ϕ^G already give varations X_R of H_R which are tangential to the image of M_R inside $\prod_j \mathbb{P}\mathrm{Met}(E_j)$. This is because if $X=\sigma(\xi)$ for a section ξ of $\mathrm{ad}\mathfrak{p}^{\mathbb{C}}$ then $X_R = \sigma(\frac{\xi+\xi^{*c}}{2})$.

4. Unstable minimal surfaces and the Labourie conjecture

After recalling the Hitchin section, we prove Theorem A for $G=\mathrm{PSL}(n,\mathbb{R})$, $n\geq 4$. We then prove Theorem A in general and deduce Corollary A.

4.1. **The Hitchin section.** Let (G,K) be an admissible pair of rank l. Fix homogeneous generators $p_1,\dots,p_l$ of $\mathcal{O}(\mathfrak{p})^K$, and let $m_1,\dots,m_l$ be their degrees. In our terminology, the Hitchin map with respect to the basis $p_1,\dots,p_l$ sends a G-Higgs bundle $(P,A^{0,1},\phi)$ to

$$(p_1(\phi),\dots,p_l(\phi))\in\bigoplus_{i=1}^l H^0(S,\mathcal{K}^{m_i}).$$

The space $\bigoplus_{i=1}^l H^0(S,\mathcal{K}^{m_i})$ is called the Hitchin base.

If G is the adjoint form of a connected split real simple group, Hitchin constructs for each point $(\alpha_1,\dots,\alpha_l)\in\bigoplus_{i=1}^l H^0(S,\mathcal{K}^{m_i})$ a G-Higgs bundle $s(\alpha_1,\dots,\alpha_l)$. He proves:

Theorem 4.1 (Theorem 7.5 in [20])**.** *With respect to the natural complex structure on the moduli space $\mathcal{M}_G(S)$ of G-Higgs bundles over S, s is a holomorphic section of the Hitchin map. Furthermore, $s(\alpha_1,\dots\alpha_l)$ is always stable and s is an isomorphism onto a connected component of $\mathcal{M}_G(S)$.*

By G-NAH II, the image of the Hitchin section s corresponds to a connected component of the representation variety $\mathrm{Rep}(\Sigma_g,G)$. This component of the representation variety is called the Hitchin component, which we write as $\mathrm{Hit}(G)$.

If G is a product of adjoint forms of split real simple groups, we can take the generators p_i to be the union of a basis of generators for each factor. Hitchin's construction gives a G-Higgs bundle for each factor. We extend the Hitchin section to products by simply taking the product of G_i-Higgs bundles for each factor G_i. Theorem 4.1 remains true for such groups G.

Remark 4.2. A metric ν on (G,K) determines a degree two invariant polynomial on $\mathfrak{p}^{\mathbb{C}}$ by $X\mapsto\nu(X^2)$. Since it should be clear in context, we write ν for the polynomial.

Remark 4.3. If $(E,\overline{\partial}_E,\phi)$ lies in the image of the Hitchin section, then ϕ never vanishes [20, Section 5]. It follows that minimal harmonic maps associated to Hitchin representations are immersions.

4.2. **Unstable minimal surfaces for Hitchin representations.** Let (G,K,ν) be an admissible triple with G the adjoint form of a split real semisimple group, and let $N=G/K$. Let $\mathbf{T}_g$ be the Teichmüller space of a surface Σ_g of genus g. Recall that if M is a flat Riemannian N-bundle with irreducible holonomy $\rho:\pi_1(\Sigma_g)\to G$, G-NAH I produces a harmonic section h. Let $\mathbf{E}_\rho(S)$ be the total energy of h with respect to ν, which we view as a function on $\mathbf{T}_g$. The function $\mathbf{E}_\rho$ is smooth [36]. If S is a critical point in $\mathbf{T}_g$, then the section h is minimal with respect to ν, and its index as a minimal map is equal to the index of $\mathbf{E}_\rho$ at S [14, Theorem 3.4]. In the case that ρ is in the Hitchin component, the result below follows from work of Labourie [24, Theorem 1.0.3].

Theorem 4.4 (Labourie)**.** *If ρ is Hitchin, then $\mathbf{E}_\rho$ is proper on $\mathbf{T}_g$.*

Consequently, $\mathbf{E}_\rho$ admits a global minimum, which must in particular be a stable critical point. Therefore, every Hitchin representation has at least one stable equivariant minimal surface. Hence, to prove that there exists a Hitchin representation with more than one equivariant minimal surface, it suffices to produce a Hitchin representation with an unstable minimal surface. This is what we do below. But first, we state a well-known fact which is the key to the construction. Note that a $\mathbb{R}^n$-valued cohomology class of Σ_g defines an action of the Deck group of any covering of Σ_g on R^n by translations.

Proposition 4.5. *Let Σ_g be a closed surface of genus g with $g \geq 3$. For any $n \geq 3$, there is a cohomology class $\beta \in H^1(\Sigma_g, \mathbb{R}^n)$, a conformal structure S on Σ_g, and an unstable minimal map $\tilde{f}$ from a covering space $\tilde{\Sigma}_g$ of Σ_g to $\mathbb{R}^n$, equivariant by the action of the Deck group determined by β.*

Proof. Here is a very direct proof. In the case $g = 3$ and $n = 3$, we take $\tilde{\Sigma}_g$ to be the triply periodic Schwarz P-surface, $\tilde{f}$ the inclusion, S the conformal structure of the quotient by $\mathbb{Z}^3$, and β the coresponding cohomology class. This surface has equivariant index one [34]; to show that the index is at least one, it suffices to consider a normal variation of constant length. If $g > 3$, we can take a branched covering of the (quotient of the) Schwarz P-surface, and if $n > 0$, we can linearly isometrically map $\mathbb{R}^3$ into $\mathbb{R}^n$. □

In fact, there are many such examples, and, at least for large genus, many of very large index. For a discussion in the case $n = 3$, see [29, section 5.3].

By the usual correspondence between equivariant maps and flat bundles, we can also interpret this as a flat Riemannian bundle M_β, whose fibers are $\mathbb{R}^n$ and transition functions are translations given by the cohomology class β, together with a minimal section f. The vertical tangent bundle of M_β is just the trivial $\mathbb{R}^n$ bundle. The equivariant index of $\tilde{f}$ is the same as the index of f. This proposition indicates that we should be looking for instability in Higgs bundles for which the G-toral bundle F_ϕ^G is trivial.

We first prove Theorem A for the main case of interest, $G = \mathrm{PSL}(n, \mathbb{R})$, $n \geq 4$. The proof is simple and distills the main ideas.

Recall from section 2.4 that $\mathfrak{p}^{\mathbb{C}} \subset \mathfrak{g}^{\mathbb{C}} = \mathfrak{sl}(n, \mathbb{C})$ is the subset of complex symmetric matrices. We choose the elementary symmetric polynomials $p_2, \dots, p_n$ as our basis for $\mathcal{O}(\mathfrak{p})^K$ defining the Hitchin base. For $a = \mathrm{diag}(a_1, \dots, a_n)$ in the subalgebra $\mathfrak{a}^C \subset \mathfrak{p}^C$ of diagonal matrices, the p_i are defined

$$p_2(a) = \sum_{1 \leq i < j \leq n} a_i a_j, \; p_3(a) = \sum_{1 \leq i < j < k \leq n} a_i a_j a_k, \; \dots \; p_n(a) = \prod_{i=1}^n a_i.$$

By the Chevalley restriction theorem, this defines p_i on all of $\mathcal{O}(\mathfrak{p})^K$. Note that if q is a polynomial in $\mathbb{C}$ with roots $\lambda_1, \dots, \lambda_n$, then

$$q(z) = \sum_{i=0}^n (-1)^{n-i} p_{n-i}(\lambda) z^i, \tag{16}$$

where $\lambda = \mathrm{diag}(\lambda_1, \dots, \lambda_n)$ and $p_1(\lambda) = \sum_{i=1}^n \lambda_i$ (which vanishes on our Lie algebra).

Proof of Theorem A for PSL$(n,\mathbb{R})$. As above, let $\mathfrak{a}^{\mathbb{C}} \subset \mathfrak{p}^{\mathbb{C}} \subset \mathfrak{g}^{\mathbb{C}}$ be the subalgebra of diagonal matrices,

$$\mathfrak{a}^{\mathbb{C}} = \{\mathrm{diag}(a_1, \dots, a_n) : a_i \in \mathbb{C}, \sum_{i=1}^{n} a_i = 0\}.$$

Let $(\tilde{S}, \beta, \tilde{f})$ be an unstable minimal map to $\mathbb{R}^{n-1}$ as in Proposition 4.5. There is a $\mathrm{SO}(n-1,\mathbb{R})$-torsor of isometries from $\mathbb{R}^{n-1}$ to the subalgebra of real points $\mathfrak{a}$, and we have the freedom to choose one, ι, such that no two components of $\iota(\partial\tilde{f})$ agree at every point of $\tilde{S}$. Let $\mathrm{diag}(\phi_1, \dots, \phi_n) = \iota(\partial\tilde{f})$. This is the holomorphic derivative of the section f of M_β in a particular trivialization, so we denote it by ∂f. From the choice of hyperplane, $p_1(\partial f) = 0$, and since f is minimal and the mapping from $\mathbb{R}^{n-1}$ to $\mathfrak{a}$ is an isometry, $p_2(\partial f) = 0$. By our choice of isometry ι, ∂f is generically regular semisimple.

For each $i = 2, \dots, n$, set $\alpha_i = p_i(\phi_i)$, defining a point $(\alpha_2, \dots, \alpha_n)$ in the Hitchin base, and let $(E, \overline{\partial}_E, \phi)$ be the Higgs bundle associated to $s((\alpha_2, \dots, \alpha_n))$ via the standard representation. Since $\alpha_2 = 0$, $p_2(\phi) = 0$, so the Higgs bundle $(E, \overline{\partial}_E, \phi)$ is minimal. By (16), the characteristic polynomials of ϕ and ∂f agree. Hence, ϕ is not just generically regular semisimple, but it is globally diagonalizable on the complement of its critical set B, with eigen-1-forms $\phi_1, \dots, \phi_n$.

Since ϕ is globally diagonalizable on $S - B$, the extended toral bundle F_ϕ—spanned by globally defined projections to the eigenspaces of ϕ—is trivial. The toral bundle F_ϕ^0 is then trivial as well. Equivalently, S is equal to its cameral cover. After trivializing the vertical tangent bundle of the apartment bundle M_ϕ to $\mathbb{R}^n$, the holomorphic derivative of the minimal section agrees with ∂f up to an isometry. By Theorem B, for R large enough the index of the minimal map h_R associated to $(E, \overline{\partial}_E, \phi)$ is bounded below by the index of the limiting object, which is equal to the equivariant index of $\tilde{f}$, hence positive. This completes the proof. □

The proof of the full Theorem A follows the PSL$(n,\mathbb{R})$ case in an abstract setting.

Proof of Theorem A. Let G be the adjoint form of a split real group of real rank $l \geq 3$, with a maximal compact K and decomposition $\mathfrak{g} = \mathfrak{k} + \mathfrak{p}$. Let $\mathfrak{a}$ be a maximal toral subalgebra of $\mathfrak{p}$ and let ν be an invariant metric on G/K. We set $(\tilde{S}, \beta, \tilde{f})$ to be an unstable minimal map to $\mathbb{R}^l$ from Proposition 4.5. Choose an isometry between $\mathbb{R}^l$ and $\mathfrak{a}$ (with the metric ν) such that no root of $\mathfrak{a}$ vanishes on the derivative $\partial\tilde{f}$ at every point of $\tilde{\Sigma}_g$.

Let M_β be the flat Riemannian bundle with translational holonomy given by β, with section f descended from $\tilde{f}$. Using the chosen isometry, we can identify the vertical tangent bundle of M_β with the trivial $\mathfrak{a}$ bundle, so that ∂f is a $(1,0)$-form valued in $\mathfrak{a}^{\mathbb{C}}$. Since f is minimal, $\nu((\partial f)^2) = 0$.

Let $(p_1, \dots, p_n)$ be a homogeneous generating set for the invariant polynomials on $\mathfrak{p}^{\mathbb{C}}$, and for each $i = 1, \dots, n$, set $\alpha_i = p_i(\partial f)$. Then $(\alpha_1, \dots, \alpha_n)$ is a point in the Hitchin base. Let $(P, A^{0,1}, \phi) = s((\alpha_1, \dots, \alpha_n))$, where s is the Hitchin section. Since s is a section of the Hitchin map, every $K^{\mathbb{C}}$-invariant polynomial on $\mathfrak{p}^{\mathbb{C}}$ takes the same values on ϕ as on ∂f. In particular, since ν is an invariant polynomial, the G-Higgs bundle $(P, A^{0,1}, \phi)$ is minimal with respect to ν. Moreover, since ∂f is generically regular semisimple by the choice of isometry, and s is a section of the Hitchin map, Lemma 2.21 shows that ϕ is conjugate to ∂f in any local trivialization of P, and hence that ϕ is also generically regular semisimple. Let $B \subset S$ be the critical set of ϕ, which is to say the set at which it is not regular semisimple. Let F_ϕ^G be the toral bundle of ϕ over $S - B$ and M_ϕ^G be the G-apartment bundle of ϕ with minimal section f^G.

From Theorem B, for R sufficiently large, the minimal harmonic map associated to $(P, A^{0,1}, R\phi)$ has index bounded below by the index of the minimal section f^G. Since the index of f is positive, we are done if we can produce an isomorphism of flat Riemannian bundles between M_β and M_ϕ^G that intertwines f and f^G. It is even enough to produce an isomorphism between $\mathfrak{a}^{\mathbb{C}}$ and F_ϕ^G which intertwines ∂f and ϕ, since the associated apartment bundles are then constructed by integration in the same way.

Fix a local trivialization of P on $U \subset S - B$. By the second part of Lemma 2.21, in this trivialization every gauge transformation conjugating ϕ to ∂f defines the same isomorphism from F_ϕ^G to $\mathfrak{a}^{\mathbb{C}}$. Since the isomorphism is unique, these local isomorphisms patch together to give a global flat real isometric isomorphism between the bundle F_ϕ^G and the trivial bundle $\mathfrak{a}^{\mathbb{C}}$ identifying ϕ with ∂f. This completes the proof of Theorem A.

□

4.3. **Non-uniqueness of area minimizers.** To conclude the paper, we prove Corollary A, which states that area minimizers for Hitchin representations need not be unique. First we recall the Labourie map.

Fix a connected split real group G in adjoint form with Cartan decomposition and metric ν on G/K. Given a conformal structure S, let

$$\mathbf{H}_G(S) = \bigoplus_{i=1}^{l} H^0(S, \mathcal{K}^{m_i})$$

be the Hitchin base for G with respect to a homogeneous basis $\{p_i\}$ of invariant polynomials on $\mathfrak{p}$. Let $\mathbf{H}_G$ be the bundle over $\mathbf{T}_g$ whose fiber over S is $\mathbf{H}_G(S)$. Since ν defines a polynomial in $\mathcal{O}(\mathfrak{p})^K$, G-Higgs bundles over S with Higgs field ϕ satisfying $\nu(\phi^2) = 0$ are taken by the Hitchin map onto a subspace $\mathbf{M}_G(S) \subset \mathbf{H}_G(S)$ of complex codimension $3g-3 = \dim(H^0(S, \mathcal{K}^2))$. Varying over $\mathbf{T}_g$, we obtain a codimension $3g-3$ subbundle $\mathbf{M}_G \subset \mathbf{H}_G$ that parametrizes the space of minimal maps associated to Hitchin representations. The Labourie map is the map

$$L_G : \mathbf{M}_G \to \mathrm{Hit}(G)$$

to the Hitchin component of the representation variety, defined by applying the Hitchin section for S on the fiber $\mathbf{M}_G(S)$. Labourie's existence result [24, Theorem 1.0.3] shows that L_G is surjective. Theorem A is equivalent to the following

Theorem 4.6. *For every split real semisimple G of rank at least 3 with maximal compact subgroup K and invariant metric ν, the Labourie map L_G is not injective.*

The theorem implies the following lemma.

Lemma 4.7. *There is no continuous section of the Labourie map.*

Proof. Suppose that $s : \mathrm{Hit}(G) \to \mathbf{M}_G$ is a continuous section of L_G. The spaces $\mathbf{M}_G$ and $\mathrm{Hit}(G)$ are both topologically open balls of the same real dimension $(2g-2)\dim G$. In the case of $\mathrm{Hit}(G)$, this is because the Hitchin section from a fixed Hitchin base is a homoeomorphism onto $\mathrm{Hit}(G)$, and since $\dim \mathbf{T}_g = \dim H^0(\mathcal{K}^2)$ this also shows the dimensions of $\mathbf{M}_G$ and $\mathrm{Hit}(G)$ are the same. Since s is a section, it is both injective and proper. By Brouwer's invariance of domain, s is homeomorphism. But then the Labourie map is also a homeomorphism, which contradicts Theorem A. □

To prove Corollary A, we argue by contradiction, and suppose that for every Hitchin representation ρ, there exists a unique equivariant minimizing minimal surface. If this is the case, we can define a section

$$s_{\min} : \mathrm{Hit}(G) \to \mathbf{M}_G$$

that associates each representation to the conformal structure of the minimal surface and the point in the Hitchin base for the Higgs bundle data. By Lemma 4.7, the proof of Corollary A will follow from the lemma below.

Lemma 4.8. *Assuming its existence, the map s_{min} is continuous.*

The auxiliary result we use is contained in a paper of Tholozan [38].

Definition 4.9. Let X and Y be two metric spaces, and $(F_y)_{y\in Y}$ a family of maps $F_y : X \to \mathbb{R}$ depending continuously on Y in the compact-open topology. We say that $(F_y)_{y\in Y}$ is locally uniformly proper if for every $y_0 \in Y$, there exists a neighbourhood U of y_0 such that for any $C \in \mathbb{R}$, there exists a compact set $K \subset X$ such that for all $y \in U$ and $x \in X\backslash K$, $F_y(x) > C$.

Lemma 4.10 (Proposition 2.6 of [38])**.** *Let X and Y be two metric spaces, and $(F_y)_{y\in Y}$ a locally uniformly proper family of maps $F_y : X \to \mathbb{R}$ depending continuously on Y. Assume that each F_y achieves its minimum at a unique point $x_m(y) \in X$. Then the map from $Y \to X$ defined by*

$$y \mapsto x_m(y)$$

is continuous.

Fix a hyperbolic metric g_0 on Σ_g. For $\gamma \in \pi_1(\Sigma_g)$ let

$$\ell(\rho(\gamma)) = \inf_{x\in G/K} d(\rho(\gamma)x, x)$$

be the translation length for $\rho(\gamma)$ and $\ell_{g_0}(\gamma)$ the g_0-length of the geodesic representative of γ. Hitchin representations are well-displacing, which in [24] means that there exists constants $A, B > 0$ such that

$$\ell(\rho(\gamma)) \geq A\ell_{g_0}(\gamma) - B.$$

More generally, it is proved by Guichard-Wienhard that an Anosov representation defines a quasi-isometric embedding from the Cayley graph of $\pi_1(\Sigma_g)$ to G (with a left-invariant metric), and are hence well-displacing [19].

Proof of Lemma 4.8. We take $X = \mathbf{T}_g$, $Y = \mathrm{Hit}(G)$, and $F_\rho = \mathbf{E}_\rho : \mathbf{T}_g \to \mathbb{R}$. $\rho \mapsto \mathbf{E}_\rho$ is smooth in ρ; this is a consequence of the implicit function theorem [36]. Assuming $(\mathbf{E}_\rho)_{\rho\in\mathrm{Hit}(G)}$ is locally uniformly proper, Lemma 4.10 asserts that the map $m : \mathrm{Hit}(G) \to \mathbf{T}_g$ taking ρ to the conformal structure of the area minimizing surface is continuous.

Toward continuity of $s_{\min}$, it suffices to work locally in $\mathbf{M}_G$, where we can write $s_{min}(\rho) = (m(\rho), \alpha(\rho))$, with $\alpha(\rho) \in \mathbf{M}_G(m(\rho))$. Fixing a basepoint on Σ_g, every Hitchin representation $\rho : \pi_1(\Sigma_g) \to G$ defines a section $\delta_\rho : \mathbf{T}_g \to \mathbf{H}_G$ as follows. For each Riemann surface, we obtain a G-Higgs bundle from ρ using G-NAH I, and then we apply the Hitchin map to get a point in $\mathbf{H}_G$. Assembling these sections into a map $\delta : \mathrm{Hit}(G) \times \mathbf{T}_g \to \mathbf{H}_G$, $\delta(\rho, \cdot) = \delta_\rho(\cdot)$, it is well-understood that δ is continuous (essentially a consequence of the fact that harmonic maps vary smoothly [36]). We express

$$\alpha(\rho) = \delta(\rho, m(\rho))$$

and deduce that $s_{\min}$ is continuous.

It remains to justify that $(\mathbf{E}_\rho)_{\rho\in\mathrm{Hit}(G)}$ is locally uniformly proper, which amounts to going through Labourie's proof that each $\mathbf{E}_\rho$ is proper and observing that the estimates can be made locally uniform in ρ. This requires no new insight, so we only give a brief explanation and point to the relevant references.

Let S_0 be a fixed marked Riemann surface structure on Σ_g. In his proof of [24, Theorem 1.0.3], Labourie shows

$$\mathbf{E}_\rho(S) \geq A_\rho(\mathrm{inter}(S, S_0))^2,$$

where $\mathrm{inter}(\cdot,\cdot) : \mathbf{T}_g^2 \to \mathbb{R}$ is the intersection function (see [2]), and $S \mapsto \mathrm{inter}(S, S_0)$ is known to be proper on $\mathbf{T}_g$ (see [2, Proposition 4] or [24, Proposition 6.2.4]). Stepping into the proof, A_ρ depends only on the minimal well-displacing constant A for ρ. It can be gleaned from the proofs of Propositions 3.3.5 and 3.3.7 of [24] that this constant is uniformly locally controlled with ρ. This establishes the result. □

Proof of Corollary A. Assuming uniqueness of area minimizers, we construct the section $s_{\min} : \mathrm{Hit}(G) \to \mathbf{M}_G$. We apply Lemma 4.7 and Lemma 4.8 to find a contradiction.

□

References

[1] Daniele Alessandrini and Brian Collier. The geometry of maximal components of the $\mathrm{PSp}(4,\mathbb{R})$ character variety. *Geom. Topol.*, 23(3):1251–1337, 2019.

[2] Francis Bonahon. The geometry of Teichmüller space via geodesic currents. *Invent. Math.*, 92(1):139–162, 1988.

[3] Steve Bradlow, Brian Collier, Oscar Garcia-Prada, Peter Gothen, and André Oliveira. A general cayley correspondence and higher teichmüller spaces, 2021.

[4] Steven B. Bradlow, Oscar García-Prada, and Peter B. Gothen. Surface group representations and $\mathrm{U}(p,q)$-Higgs bundles. *J. Differential Geom.*, 64(1):111–170, 2003.

[5] Steven B. Bradlow, Oscar García-Prada, and Ignasi Mundet i Riera. Relative Hitchin-Kobayashi correspondences for principal pairs. *Q. J. Math.*, 54(2):171–208, 2003.

[6] Suhyoung Choi and William M. Goldman. Convex real projective structures on closed surfaces are closed. *Proc. Amer. Math. Soc.*, 118(2):657–661, 1993.

[7] Brian Collier. Maximal $\mathrm{Sp}(4,\mathbb{R})$ surface group representations, minimal immersions and cyclic surfaces. *Geom. Dedicata*, 180:241–285, 2016.

[8] Brian Collier, Nicolas Tholozan, and Jérémy Toulisse. The geometry of maximal representations of surface groups into $\mathrm{SO}_0(2,n)$. *Duke Math. J.*, 168(15):2873–2949, 2019.

[9] Kevin Corlette. Flat G-bundles with canonical metrics. *J. Differential Geom.*, 28(3):361–382, 1988.

[10] Ron Donagi. Decomposition of spectral covers. *Astérisque*, pages 145–175, 1993. Journées de Géométrie Algébrique d'Orsay (Orsay, 1992).

[11] S. K. Donaldson. Twisted harmonic maps and the self-duality equations. *Proc. London Math. Soc. (3)*, 55(1):127–131, 1987.

[12] David Dumas and Michael Wolf. Polynomial cubic differentials and convex polygons in the projective plane. *Geom. Funct. Anal.*, 25(6):1734–1798, 2015.

[13] Patrick B. Eberlein. *Geometry of nonpositively curved manifolds.* Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1996.

[14] Norio Ejiri. A differential-geometric Schottky problem, and minimal surfaces in tori. In *Differential geometry and integrable systems (Tokyo, 2000)*, volume 308 of *Contemp. Math.*, pages 101–144. Amer. Math. Soc., Providence, RI, 2002.

[15] Vladimir Fock and Alexander Goncharov. Moduli spaces of local systems and higher Teichmüller theory. *Publ. Math. Inst. Hautes Études Sci.*, (103):1–211, 2006.

[16] Vladimir Fock and Alexander Thomas. Higher complex structures. *Int. Math. Res. Not. IMRN*, (20):15873–15893, 2021.

[17] Oscar García-Prada. Higgs bundles and surface group representations. In *Moduli spaces and vector bundles*, volume 359 of *London Math. Soc. Lecture Note Ser.*, pages 265–310. Cambridge Univ. Press, Cambridge, 2009.

[18] Olivier Guichard, François Labourie, and Anna Wienhard. Positivity and representations of surface groups, 2021.

[19] Olivier Guichard and Anna Wienhard. Anosov representations: domains of discontinuity and applications. *Invent. Math.*, 190(2):357–438, 2012.

[20] N. J. Hitchin. Lie groups and Teichmüller space. *Topology*, 31(3):449–473, 1992.
[21] Ludmil Katzarkov, Alexander Noll, Pranav Pandit, and Carlos Simpson. Harmonic maps to buildings and singular perturbation theory. *Comm. Math. Phys.*, 336(2):853–903, 2015.
[22] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of differential geometry. Vol. II.* Wiley Classics Library. John Wiley & Sons, Inc., New York, 1996. Reprint of the 1969 original, A Wiley-Interscience Publication.
[23] François Labourie. Anosov flows, surface groups and curves in projective space. *Invent. Math.*, 165(1):51–114, 2006.
[24] François Labourie. Cross ratios, Anosov representations and the energy functional on Teichmüller space. *Ann. Sci. Éc. Norm. Supér. (4)*, 41(3):437–469, 2008.
[25] François Labourie. Cyclic surfaces and Hitchin components in rank 2. *Ann. of Math. (2)*, 185(1):1–58, 2017.
[26] John C. Loftin. Affine spheres and convex $\mathbb{RP}^n$-manifolds. *Amer. J. Math.*, 123(2):255–274, 2001.
[27] Vladimir Marković. Uniqueness of minimal diffeomorphisms between surfaces. *Bull. Lond. Math. Soc.*, 53(4):1196–1204, 2021.
[28] Vladimir Marković. Non-uniqueness of minimal surfaces in a product of closed Riemann surfaces. *Geom. Funct. Anal.*, 32(1):31–52, 2022.
[29] Vladimir Marković, Nathaniel Sagman, and Peter Smillie. Unstable minimal surfaces in $\mathbb{R}^n$ and in products of hyperbolic surfaces, 2022.
[30] Takuro Mochizuki. Asymptotic behaviour of certain families of harmonic bundles on Riemann surfaces. *J. Topol.*, 9(4):1021–1073, 2016.
[31] Xin Nie. Cyclic higgs bundles and minimal surfaces in pseudo-hyperbolic spaces, 2022.
[32] Alexander Nolte. Canonical maps from spaces of higher complex structures to hitchin components, 2022.
[33] Anne Parreau. Compactification d'espaces de représentations de groupes de type fini. *Math. Z.*, 272(1-2):51–86, 2012.
[34] Marty Ross. Schwarz' P and D surfaces are stable. *Differential Geom. Appl.*, 2(2):179–195, 1992.
[35] Carlos T. Simpson. Constructing variations of Hodge structure using Yang-Mills theory and applications to uniformization. *J. Amer. Math. Soc.*, 1(4):867–918, 1988.
[36] Ivo Slegers. Equivariant harmonic maps depend real analytically on the representation. *manuscripta math.*, 2021.
[37] Andrea Tamburelli and Michael Wolf. Planar minimal surfaces with polynomial growth in the $\mathrm{Sp}(4,\mathbb{R})$-symmetric space, 2020.
[38] Nicolas Tholozan. Dominating surface group representations and deforming closed anti-de Sitter 3-manifolds. *Geom. Topol.*, 21(1):193–214, 2017.
[39] Domingo Toledo. Hermitian curvature and plurisubharmonicity of energy on Teichmüller space. *Geom. Funct. Anal.*, 22(4):1015–1032, 2012.
[40] Tom Yau-Heng Wan. Constant mean curvature surface, harmonic maps, and universal Teichmüller space. *J. Differential Geom.*, 35(3):643–657, 1992.
[41] Anna Wienhard. An invitation to higher Teichmüller theory. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures*, pages 1013–1039. World Sci. Publ., Hackensack, NJ, 2018.
[42] Michael Wolf. The Teichmüller theory of harmonic maps. *J. Differential Geom.*, 29(2):449–479, 1989.
[43] Michael Wolf. Harmonic maps from surfaces to **R**-trees. *Math. Z.*, 218(4):577–593, 1995.