

Holonomy of Parabolic Geometries Near Isolated Higher-Order Fixed Points

Jacob W. ERICKSON

School of Mathematics, Georgia Institute of Technology, Atlanta, USA

E-mail: jerickson41@gatech.edu

URL: <https://erickson.math.gatech.edu/>

Received December 01, 2025, in final form August 12, 2026; Published online August 28, 2026

<https://doi.org/10.3842/SIGMA.2026.084>

Abstract. For Cartan geometries admitting an automorphism with isotropy satisfying a particular, loosely dynamical property on their model geometries, we demonstrate the existence of an open subset of the geometry with trivial holonomy. This property, which generalizes characteristics of isotropies corresponding to isolated higher-order fixed points in parabolic geometries that are known to require a nearby open subset to have vanishing curvature, only relies upon the behavior of the isotropy in the model geometry, and therefore applies regardless of initial curvature assumptions, such as regularity or normality. Along the way to proving our main results, we also derive a couple of results for working with holonomy, relating to limits of sequences of developments and the existence of antidevelopments, that are useful in their own right. To showcase the effectiveness of the techniques developed, we use them to completely characterize all almost c-projective and almost quaternionic structures that admit a nontrivial automorphism with a higher-order fixed point, as well as all nondegenerate partially integrable almost CR structures that admit a higher-order fixed point with non-null isotropy.

Key words: Cartan geometries; parabolic geometries; essential automorphisms; holonomy

2020 Mathematics Subject Classification: 53C29; 58D19; 57S25; 53C30

1 Introduction

In each dimension greater than two, a celebrated theorem of Ferrand and Obata tells us that, up to conformal isomorphism, there are only two examples of conformal structures of Riemannian signature admitting automorphism groups that act non-properly: the standard conformal sphere and Euclidean space. A similar result, due to Schoen and Webster, holds for strictly pseudoconvex CR structures: if the CR automorphisms of such a structure act non-properly, then it must be isomorphic to either the standard CR structure on the sphere or a left-invariant CR structure on the Heisenberg group. Eventually, it was realized that Cartan geometries—a class of geometric structures that extend the spirit of Klein’s Erlangen program somewhat analogously to the way that Riemannian manifolds extend Euclidean geometry—are the natural setting for this type of result, and in [7], Frances gave a generalization of both results, extending the idea to all parabolic Cartan geometries of real rank one, assuming the curvature satisfies a standard regularity condition. A convenient and concise history of this result can be found in [13].

These Ferrand–Obata type theorems provide striking examples where the global structure of a given Cartan geometry is almost completely determined by the behavior of its automorphisms. Naturally, though, we are led to ask: how far can we take this? How much can we say about the global structure of a Cartan geometry just from the behavior of its automorphisms, or even from the behavior of a *single* automorphism?

In the vein of this latter question, we specifically want to look at automorphisms of parabolic geometries with *higher-order fixed points*. These are automorphisms with isotropies, over some

fixed point of the geometry, lying in the unipotent radical of the parabolic subgroup used as the stabilizer subgroup for the corresponding model geometry. For many of the most common types of parabolic geometries, such as projective and conformal geometry, these just amount to automorphisms whose derivatives on the base manifold are the identity transformation on the tangent space over some fixed point.

The existence of a higher-order fixed point often guarantees that the curvature of the geometry vanishes over some open subset containing the higher-order fixed point in its closure. Results in this direction essentially go all the way back to [15] in the projective case and [9] in the conformal case, though a more unified approach was introduced in [2] and expanded upon in [11, 14], among other places. Notably, however, these prior results are all local in nature, unless we make quite strong assumptions on the geometry, such as real-analyticity or—as explored in [12, 15]—metrizable.

In this paper, we will provide a way to extend these existing local results to semi-global ones when the higher-order fixed point is isolated. Our main result, Theorem 4.4, proves that if an automorphism has isotropy satisfying a kind of dynamical property in the model geometry that applies to the cases with isolated higher-order fixed points where local results are known, then there is a certain invariant open subset of the geometry on which the *holonomy* must be trivial. If we know a bit more about the topology of this open subset, then we can even prove genuine global results along the lines of Ferrand–Obata type theorems. To showcase this, we specifically prove such results for c-projective structures, quaternionic structures, and CR structures; as a courtesy to those unfamiliar with Cartan geometries, we state these theorems in their “native” languages, though in the rest of the paper, we will almost exclusively work from the unified perspective provided by the Cartan machinery.

Theorem A. *Suppose that $(M, J, [\nabla])$ is an almost c-projective structure. If $(M, J, [\nabla])$ admits a nontrivial automorphism α with a higher-order fixed point $x \in M$, so that $\alpha(x) = x$ and $\alpha_{*x} = \text{id}_{T_x M}$, then $(M, J, [\nabla])$ embeds geometrically onto a dense open subspace of the standard flat c-projective structure on $\mathbb{CP}^m$. In particular, if M is compact, then such an automorphism and fixed point exist if and only if $(M, J, [\nabla])$ is isomorphic to the flat c-projective structure on $\mathbb{CP}^m$.*

Theorem B. *Suppose that M is an almost quaternionic manifold. If M admits a nontrivial automorphism α with a higher-order fixed point $x \in M$, so that $\alpha(x) = x$ and $\alpha_{*x} = \text{id}_{T_x M}$, then M embeds geometrically onto a dense open subspace of $\mathbb{HP}^m$ equipped with its standard quaternionic structure. In particular, if M is compact, then such an automorphism and fixed point exist if and only if M is isomorphic to $\mathbb{HP}^m$.*

Theorem C. *Suppose that (M, J, D) is a partially integrable almost CR structure with Levi form of signature (p, q) , where J is an almost complex structure on the contact distribution D . If (M, J, D) admits a nontrivial automorphism α with a non-null higher-order fixed point $x \in M$, so that $\alpha(x) = x$ and either $\alpha_{*x} = \text{id}_{T_x M}$ or the image of the difference $\alpha_{*x} - \text{id}_{T_x M}$ consists of non-null vectors for the Levi form on the contact distribution D , then (M, J, D) embeds geometrically onto a dense open subspace of $\text{Null}(h_{p,q}) := \{\mathbb{C}^\times u \in \mathbb{CP}^{p+q+1} \mid h_{p,q}(u, u) = 0\}$ equipped with the standard flat CR structure, where the Hermitian form $h_{p,q}$ on $\mathbb{C}^{p+q+2}$ has corresponding quadratic form given by*

$$\begin{bmatrix} z_0 \\ \vdots \\ z_{p+q+1} \end{bmatrix} \mapsto 2\text{Re}(\bar{z}_0 z_{p+q+1}) + \sum_{j=1}^p |z_j|^2 - \sum_{j=p+1}^{p+q} |z_j|^2.$$

In particular, if M is compact, then such an automorphism and fixed point exist if and only if (M, J, D) is isomorphic to the flat CR structure on $\text{Null}(h_{p,q})$.

These results will follow from Theorem 5.2, which is itself a corollary of our main result, Theorem 4.4. To get to these results, however, we will need to prove a couple of intermediate facts that are useful for working with development and holonomy in this setting, which we provide in Section 3 after reviewing terminology and notation in Section 2. Then, in Section 4, we will describe the dynamical property for isotropies alluded to above and use it to prove our main result, before finishing the paper with the applications in Section 5.

We should remark that Theorems A, B, and C all appeared in the original preprint for [6]. This current paper and its methods grew, at least in part, out of efforts to simplify the proofs and assumptions in that preprint, to avoid having to use the somewhat heavy-handed sprawl construction. It should be considered as separate from—though obviously related to—the current version of [6], which is substantially different and does not include these results.

2 Preliminaries

This section is meant as a way of establishing terminology and notation; it is probably not a good introduction to the topic. We recommend [16] for an actual introduction to Cartan geometries, [3] for an introduction to parabolic geometries, and [5] for an overview of the results on holonomy.

2.1 Relevant model geometries

Cartan geometries are modeled on homogeneous geometries in the sense of Klein, though we present these geometries in a way that emphasizes the role of the Lie group as a principal bundle over the homogeneous space.

Definition 2.1. A *model* (or *model geometry*) is a pair (G, H) , where G is a Lie group and H is a closed subgroup of G such that G/H is connected. In this case, G is called the *model group* and H is called the *isotropy* or *stabilizer subgroup*.

Writing $\text{Euc}(m) := \mathbb{R}^m \rtimes \text{O}(m)$ for the group of Euclidean isometries on $\mathbb{R}^m$ (with the dot product metric structure) and $\text{Aff}(m) := \mathbb{R}^m \rtimes \text{GL}_m \mathbb{R}$ for the group of affine transformations on $\mathbb{R}^m$, two standard examples of model geometries are $(\text{Euc}(m), \text{O}(m))$, which corresponds to Euclidean geometry on $\mathbb{R}^m \cong \text{Euc}(m)/\text{O}(m)$, and $(\text{Aff}(m), \text{GL}_m \mathbb{R})$, which corresponds to affine geometry on $\mathbb{R}^m \cong \text{Aff}(m)/\text{GL}_m \mathbb{R}$. Cartan geometries that are modeled on these two examples correspond to Riemannian structures (modulo torsion considerations) and affine connections, respectively.

The models relevant to Theorems A, B, and C are all *parabolic*.

Definition 2.2. A model (G, P) is *parabolic* whenever the model group G is semisimple and the isotropy¹ P is a parabolic subgroup.

For parabolic models, we get an Ad_P -invariant filtration of $\mathfrak{g}$ given by

$$\mathfrak{g} = \mathfrak{g}^{-k} \supset \mathfrak{g}^{-k+1} \supset \cdots \supset \mathfrak{g}^k \supset \{0\},$$

with $\mathfrak{g}^0 := \mathfrak{p}$ and $\mathfrak{g}^1 := \mathfrak{p}_+$, where $\mathfrak{p}_+$ is the nilradical of $\mathfrak{p}$. From this filtration, we can also get a grading $\sum_{i=-k}^k \mathfrak{g}_i$ on $\mathfrak{g}$ with $\mathfrak{g}_i \approx \mathfrak{g}^i/\mathfrak{g}^{i+1}$; this grading corresponds to a choice of Cartan involution, but we will keep this choice implicit, since we will not need the Cartan involution here. Notably, the filtration and grading satisfy $[\mathfrak{g}^i, \mathfrak{g}^j] \subseteq \mathfrak{g}^{i+j}$ and $[\mathfrak{g}_i, \mathfrak{g}_j] \subseteq \mathfrak{g}_{i+j}$ for all i and j ,

¹When we assume that the model is parabolic, we will always denote the isotropy subgroup by P . This will help us to visually distinguish between results that apply to more general Cartan geometries and results that apply to particular parabolic geometries.

so $\mathfrak{g}_- := \sum_{i < 0} \mathfrak{g}_i$ and $\mathfrak{g}_0$ are subalgebras. We denote by G_- and P_+ the connected nilpotent subgroups generated by $\mathfrak{g}_-$ and $\mathfrak{p}_+$, respectively, and by G_0 the closed subgroup of P such that $\text{Ad}_{G_0}(\mathfrak{g}_i) = \mathfrak{g}_i$ for every grading component $\mathfrak{g}_i$.

In Theorems A, we will focus on the model $(\text{PGL}_{m+1} \mathbb{C}, P)$, where $\text{PGL}_{m+1} \mathbb{C}$ is the quotient of $\text{GL}_{m+1} \mathbb{C}$ by its center $\mathbb{C}^\times \mathbf{1}$ and

$$P := \left\{ \begin{pmatrix} r & p \\ 0 & A \end{pmatrix} \in \text{PGL}_{m+1} \mathbb{C} \mid r \in \mathbb{C}^\times, p^\top \in \mathbb{C}^m, A \in \text{GL}_m \mathbb{C} \right\}.$$

This model geometry is parabolic, with grading

$$\begin{aligned} \mathfrak{g}_- = \mathfrak{g}_{-1} &:= \left\{ \begin{pmatrix} 0 & 0 \\ v & 0 \end{pmatrix} \in \mathfrak{pgl}_{m+1} \mathbb{C} \mid v \in \mathbb{C}^m \right\}, \\ \mathfrak{g}_0 &:= \left\{ \begin{pmatrix} r & 0 \\ 0 & R \end{pmatrix} \in \mathfrak{pgl}_{m+1} \mathbb{C} \mid r \in \mathbb{C}, R \in \mathfrak{gl}_m \mathbb{C} \right\}, \end{aligned}$$

and

$$\mathfrak{p}_+ = \mathfrak{g}_1 := \left\{ \begin{pmatrix} 0 & p \\ 0 & 0 \end{pmatrix} \in \mathfrak{pgl}_{m+1} \mathbb{C} \mid p^\top \in \mathbb{C}^m \right\},$$

and associated subgroups

$$\begin{aligned} G_- &:= \left\{ \begin{pmatrix} 1 & 0 \\ v & \mathbf{1} \end{pmatrix} \in \text{PGL}_{m+1} \mathbb{C} \mid v \in \mathbb{C}^m \right\}, \\ G_0 &:= \left\{ \begin{pmatrix} r & 0 \\ 0 & A \end{pmatrix} \in \text{PGL}_{m+1} \mathbb{C} \mid r \in \mathbb{C}^\times, A \in \text{GL}_m \mathbb{C} \right\}, \end{aligned}$$

and

$$P_+ := \left\{ \begin{pmatrix} 1 & p \\ 0 & \mathbf{1} \end{pmatrix} \in \text{PGL}_{m+1} \mathbb{C} \mid p^\top \in \mathbb{C}^m \right\}.$$

In this case, the model encodes complex projective geometry over $\text{PGL}_{m+1} \mathbb{C}/P \cong \mathbb{CP}^m$, and Cartan geometries of type $(\text{PGL}_{m+1} \mathbb{C}, P)$ (with certain curvature restrictions) correspond to almost c-projective structures. An overview of such structures can be found in [1].

In Theorem B, we will focus on the model $(\text{PGL}_{m+1} \mathbb{H}, P)$. Here, $\text{PGL}_{m+1} \mathbb{H}$ is the quotient of the quaternionic general linear group $\text{GL}_{m+1} \mathbb{H}$ of right $\mathbb{H}$ -module automorphisms of $\mathbb{H}^{m+1}$ by its center, which corresponds to those automorphisms which left-multiply every element of $\mathbb{H}^{m+1}$ by a nonzero real number, and

$$P := \left\{ \begin{pmatrix} r & p \\ 0 & A \end{pmatrix} \in \text{PGL}_{m+1} \mathbb{H} \mid r \in \mathbb{H}^\times, p^\top \in \mathbb{H}^m, A \in \text{GL}_m \mathbb{H} \right\}.$$

This model geometry is also parabolic, with grading, subgroups, and subalgebras the same as those for $(\text{PGL}_{m+1} \mathbb{C}, P)$ but with $\mathbb{C}$ replaced by $\mathbb{H}$ everywhere. It encodes the quaternionic analogue of projective geometry over $\text{PGL}_{m+1} \mathbb{H}/P \cong \mathbb{HP}^m$, and Cartan geometries of type $(\text{PGL}_{m+1} \mathbb{H}, P)$ (again with certain curvature restrictions) correspond to almost quaternionic structures, as described in [3, Proposition 4.1.8].

Finally, in Theorem C, we will focus on the model $(\text{PU}(\mathfrak{h}_{p,q}), P)$, where $\mathfrak{h}_{p,q}$ is the Hermitian form on $\mathbb{C}^{p+q+2}$ with quadratic form given by

$$\begin{bmatrix} z_0 \\ \vdots \\ z_{p+q+1} \end{bmatrix} \mapsto 2\text{Re}(\bar{z}_0 z_{p+q+1}) + \sum_{j=1}^p |z_j|^2 - \sum_{j=p+1}^{p+q} |z_j|^2,$$

and $\text{PU}(\mathfrak{h}_{p,q})$ is the quotient of the group of unitary transformations for $\mathfrak{h}_{p,q}$ by its center, consisting of multiples of the identity matrix by elements of $\text{U}(1)$. Denoting by $I_{p,q}$ the $(p+q) \times (p+q)$ diagonal matrix with the first p diagonal entries equal to $+1$ and the last q diagonal entries equal to -1 , the Lie group $\text{PU}(\mathfrak{h}_{p,q})$ has Lie algebra of the form

$$\mathfrak{pu}(\mathfrak{h}_{p,q}) := \left\{ \begin{pmatrix} r & \beta & \text{is} \\ v & R & -I_{p,q}\bar{\beta}^\top \\ it & -\bar{v}^\top I_{p,q} & -\bar{r} \end{pmatrix} \mid \begin{array}{l} s, t \in \mathbb{R}, v, \beta^\top \in \mathbb{C}^{p+q}, \\ r \in \mathbb{C}, \text{ and } R \in \mathfrak{u}(p, q) \end{array} \right\},$$

where elements of the Lie algebra are considered equivalent if their difference is an imaginary multiple of the identity matrix. The parabolic subgroup P , then, is

$$P := \left\{ \begin{pmatrix} r & r\beta & r(\text{is} - \frac{1}{2}\beta I_{p,q}\bar{\beta}^\top) \\ 0 & A & -AI_{p,q}\bar{\beta}^\top \\ 0 & 0 & \bar{r}^{-1} \end{pmatrix} \in \text{PU}(\mathfrak{h}_{p,q}) \mid \begin{array}{l} s \in \mathbb{R}, \beta^\top \in \mathbb{C}^{p+q}, \\ r \in \mathbb{C}^\times, \text{ and } A \in \text{U}(p, q) \end{array} \right\},$$

with grading given by

$$\begin{aligned} \mathfrak{g}_{-2} &:= \left\langle \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \text{i} & 0 & 0 \end{pmatrix} \right\rangle, \\ \mathfrak{g}_{-1} &:= \left\{ \begin{pmatrix} 0 & 0 & 0 \\ v & 0 & 0 \\ 0 & -\bar{v}^\top I_{p,q} & 0 \end{pmatrix} \in \mathfrak{pu}(\mathfrak{h}_{p,q}) \mid v \in \mathbb{C}^{p+q} \right\}, \\ \mathfrak{g}_0 &:= \left\{ \begin{pmatrix} r & 0 & 0 \\ 0 & R & 0 \\ 0 & 0 & -\bar{r} \end{pmatrix} \in \mathfrak{pu}(\mathfrak{h}_{p,q}) \mid r \in \mathbb{C}, R \in \mathfrak{u}(p, q) \right\}, \\ \mathfrak{g}_1 &:= \left\{ \begin{pmatrix} 0 & \beta & 0 \\ 0 & 0 & -I_{p,q}\bar{\beta}^\top \\ 0 & 0 & 0 \end{pmatrix} \in \mathfrak{pu}(\mathfrak{h}_{p,q}) \mid \beta^\top \in \mathbb{C}^{p+q} \right\}, \end{aligned}$$

and

$$\mathfrak{g}_2 := \left\langle \begin{pmatrix} 0 & 0 & \text{i} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \right\rangle,$$

with corresponding subgroups

$$\begin{aligned} G_- &:= \left\{ \begin{pmatrix} 1 & 0 & 0 \\ v & \mathbf{1} & 0 \\ it - \frac{1}{2}\bar{v}^\top I_{p,q}v & -\bar{v}^\top I_{p,q} & 1 \end{pmatrix} \in \text{PU}(\mathfrak{h}_{p,q}) \mid v \in \mathbb{C}^{p+q}, t \in \mathbb{R} \right\}, \\ G_0 &:= \left\{ \begin{pmatrix} r & 0 & 0 \\ 0 & A & 0 \\ 0 & 0 & \bar{r}^{-1} \end{pmatrix} \in \text{PU}(\mathfrak{h}_{p,q}) \mid r \in \mathbb{C}^\times, A \in \text{U}(p, q) \right\}, \end{aligned}$$

and

$$P_+ := \left\{ \begin{pmatrix} 1 & \beta & \text{is} - \frac{1}{2}\beta I_{p,q}\bar{\beta}^\top \\ 0 & \mathbf{1} & -I_{p,q}\bar{\beta}^\top \\ 0 & 0 & 1 \end{pmatrix} \in \text{PU}(\mathfrak{h}_{p,q}) \mid \beta^\top \in \mathbb{C}^{p+q}, s \in \mathbb{R} \right\}.$$

Within P_+ , we will also say that

$$a = \begin{pmatrix} 1 & \beta & \text{is } -\frac{1}{2}\beta I_{p,q}\bar{\beta}^\top \\ 0 & \mathbb{1} & -I_{p,q}\bar{\beta}^\top \\ 0 & 0 & 1 \end{pmatrix} \in P_+$$

is *non-null* if and only if it is nontrivial and $\beta I_{p,q}\bar{\beta}^\top \neq 0$ if $\beta \neq 0$.

The homogeneous space $\text{PU}(\mathfrak{h}_{p,q})/P$ is naturally diffeomorphic to the null-cone

$$\text{Null}(\mathfrak{h}_{p,q}) := \{\mathbb{C}^\times u \in \mathbb{CP}^{p+q+1} \mid \mathfrak{h}_{p,q}(u, u) = 0\}$$

for $\mathfrak{h}_{p,q}$ in $\mathbb{CP}^{p+q+1}$. This null-cone is a compact and simply connected smooth manifold of (real) dimension $2(p+q)+1$; for $q=0$ or $p=0$, $\text{Null}(\mathfrak{h}_{p,q})$ is diffeomorphic to the sphere. Cartan geometries of type $(\text{PU}(\mathfrak{h}_{p,q}), P)$ (again, with certain curvature restrictions) correspond to nondegenerate partially integrable almost CR structures with Levi form of signature (p, q) , as detailed in [3, Section 4.2.4].

2.2 Generalities for Cartan geometries

In essence, the idea of a Cartan geometry of type (G, H) is to specify a $\mathfrak{g}$ -valued one-form ω that behaves like the Maurer-Cartan form $\omega_G: X_g \in T_g G \mapsto L_{g^{-1}*} X_g \in \mathfrak{g}$ does on G , where $L_a: g \mapsto ag$ denotes left-translation by $a \in G$.

Definition 2.3. Let (G, H) be a model. A *Cartan geometry of type (G, H) over a (smooth) manifold M* is a pair $(\mathcal{G}, \omega)$, where $\mathcal{G}$ is a principal H -bundle over M with quotient map² $q_H: \mathcal{G} \rightarrow M$ and ω is a $\mathfrak{g}$ -valued one-form on $\mathcal{G}$ such that

- for every $\mathcal{G} \in \mathcal{G}$, $\omega_{\mathcal{G}}: T_{\mathcal{G}}\mathcal{G} \rightarrow \mathfrak{g}$ is a linear isomorphism;
- for every $h \in H$, $R_h^* \omega = \text{Ad}_{h^{-1}} \omega$;
- for every $Y \in \mathfrak{h}$, the flow of the vector field $\omega^{-1}(Y)$ is given by $\exp(t\omega^{-1}(Y)) = R_{\exp(tY)}$ for all $t \in \mathbb{R}$.

A natural example of a Cartan geometry of type (G, H) is always the *Klein geometry* of that type, which encodes the geometric structure of the model geometry as a Cartan geometry.

Definition 2.4. For a model (G, H) , the *Klein geometry of type (G, H)* is the Cartan geometry (G, ω_G) over G/H , where G is the model group and ω_G is the Maurer-Cartan form on G .

Throughout, we will want to compare different Cartan geometries of the same type. To do this, we will use local diffeomorphisms that preserve the Cartan-geometric structure, called *geometric maps*.

Definition 2.5. Given two Cartan geometries $(\mathcal{G}, \omega)$ and $(\mathcal{Q}, \nu)$ of type (G, H) , a *geometric map* $\varphi: (\mathcal{G}, \omega) \rightarrow (\mathcal{Q}, \nu)$ is an H -equivariant smooth map $\varphi: \mathcal{G} \rightarrow \mathcal{Q}$ such that $\varphi^* \nu = \omega$.

A geometric map $\varphi: (\mathcal{G}, \omega) \rightarrow (\mathcal{Q}, \nu)$ always induces a corresponding local diffeomorphism between the base manifolds of $\mathcal{G}$ and $\mathcal{Q}$, given by $q_H(\mathcal{G}) \mapsto q_H(\varphi(\mathcal{G}))$ for each $\mathcal{G} \in \mathcal{G}$. We find it convenient and natural to not waste symbols to distinguish between these maps; throughout, whenever we have a geometric map φ , we will also denote its induced map on the base manifolds by the same symbol φ . The meaning should always be clear from context.

²In an effort to declutter notation, we will always denote the quotient map of a principal H -bundle by q_H , even if there are multiple relevant principal bundles. The meaning of the map should always be clear from context. Similarly, we will always denote the right-action of $h \in H$ on a principal bundle by $R_h: \mathcal{G} \mapsto \mathcal{G}h$.

Of course, some geometric maps tell us more than others. We say that a geometric map $\varphi: (\mathcal{G}, \omega) \rightarrow (\mathcal{Q}, \nu)$ is a *geometric embedding* when φ is injective, and when φ is bijective, we further say that it is a *(geometric) isomorphism*. A geometric isomorphism from $(\mathcal{G}, \omega)$ to itself is then called a *(geometric) automorphism*.

Geometric maps are quite rigid: if the base manifold of $(\mathcal{G}, \omega)$ is connected, then a geometric map $\varphi: (\mathcal{G}, \omega) \rightarrow (\mathcal{Q}, \nu)$ is uniquely determined by the image $\varphi(\mathfrak{g})$ of an arbitrary single element $\mathfrak{g} \in \mathcal{G}$. As such, the *automorphism group* $\text{Aut}(\mathcal{G}, \omega)$, consisting of all automorphisms of $(\mathcal{G}, \omega)$, will act freely on $\mathcal{G}$.

Following our convention of not distinguishing between geometric maps and the corresponding induced maps on the base manifolds, when we talk about fixed points of an automorphism $\alpha \in \text{Aut}(\mathcal{G}, \omega)$, we will mean fixed points of the induced map on the base manifold; since $\text{Aut}(\mathcal{G}, \omega)$ acts freely, automorphisms will not have actual fixed points in the bundle $\mathcal{G}$. We denote by

$$\text{Fix}_M(\alpha) := \{q_H(\mathfrak{g}) \in M \mid \alpha(q_H(\mathfrak{g})) = q_H(\mathfrak{g})\}$$

the set of fixed points for $\alpha \in \text{Aut}(\mathcal{G}, \omega)$.

For each $\mathfrak{g} \in \mathcal{G}$ such that $q_H(\mathfrak{g}) \in \text{Fix}_M(\alpha)$, there is some $a \in H$ such that $\alpha(\mathfrak{g}) = \mathfrak{g}a$. This element $a \in H$ is called the *isotropy* of α at $\mathfrak{g}$.

Definition 2.6. For $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ and $e \in \mathcal{G}$ such that $\alpha(e) \in eH$, the *isotropy* of α at e is the unique element $a \in H$ such that $\alpha(e) = ea$.

For the applications considered in the paper, we will primarily focus on automorphisms $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ of parabolic geometries $(\mathcal{G}, \omega)$ of type (G, P) with isotropy $a \in P_+ < P$ at some $e \in \mathcal{G}$. In that case, we say that α has a *higher-order fixed point* at $q_P(e)$.

Another core idea for Cartan geometries is that of *curvature*, which tells us when the geometry locally differs from the Klein geometry.

Definition 2.7. Given a Cartan geometry $(\mathcal{G}, \omega)$ of type (G, H) , its *curvature* is the $\mathfrak{g}$ -valued two-form given by $\Omega := d\omega + \frac{1}{2}[\omega, \omega]$.

When the curvature vanishes in a neighborhood of a point, then the geometry is locally equivalent to that of the Klein geometry near that point. In other words, when Ω vanishes on some neighborhood of an element $e \in \mathcal{G}$, we can find a geometric embedding

$$\psi: (q_H^{-1}(U), \omega_G) \hookrightarrow (\mathcal{G}, \omega)$$

from a neighborhood $q_H^{-1}(U)$ of the identity $e \in G$ in the Klein geometry such that $\psi(e) = e$. When the curvature vanishes everywhere, we say that the geometry is *flat*; such geometries are equivalent to locally homogeneous geometric structures in the sense of [10].

2.3 Development and holonomy

Again, we recommend [5] for an overview of our techniques involving holonomy, as well as [16, Chapter 3, Section 7] for a review of basic results on developments of paths.

We would like to pretend that Cartan geometries $(\mathcal{G}, \omega)$ of type (G, H) “are” their model geometries. The notions of *development* and *holonomy* allow us to do this somewhat judiciously.

Definition 2.8. Given a (piecewise smooth)³ path $\gamma: [0, 1] \rightarrow \mathcal{G}$ in a Cartan geometry $(\mathcal{G}, \omega)$ of type (G, H) , the *development* γ_G of γ is the unique (piecewise smooth) path $\gamma_G: [0, 1] \rightarrow G$ such that $\gamma_G(0) = e$ and $\gamma^*\omega = \gamma_G^*\omega_G$.

³Throughout, whenever we refer to a “path”, we will always mean a piecewise smooth path.

Development allows us to identify paths in a Cartan geometry with paths in the model, and if we fix a pretend “identity element” $e \in \mathcal{G}$, we can even give a kind of correspondence between elements of $\mathcal{G}$ and elements of G .

Definition 2.9. For a Cartan geometry $(\mathcal{G}, \omega)$ of type (G, H) and points $e, g \in \mathcal{G}$, we say that $g \in G$ is a *development of g from e* if and only if there exists a path $\gamma: [0, 1] \rightarrow \mathcal{G}$ and $h \in H$ such that $\gamma(0) = e$, $\gamma(1)h = g$, and $\gamma_G(1)h = g$.

Developments of elements in $\mathcal{G}$ are usually not unique. Thankfully, the *holonomy group* $\text{Hol}_e(\mathcal{G}, \omega)$ tells us precisely how this happens: if g is a development of g from e , then the set of all possible developments from e to g is precisely $\text{Hol}_e(\mathcal{G}, \omega)g$.

Definition 2.10. For a Cartan geometry $(\mathcal{G}, \omega)$ of type (G, H) , the *holonomy group* of $(\mathcal{G}, \omega)$ at $e \in \mathcal{G}$ is the subgroup of G given by

$$\text{Hol}_e(\mathcal{G}, \omega) := \left\{ \gamma_G(1)h_\gamma^{-1} \in G \mid \begin{array}{l} \gamma: [0, 1] \rightarrow \mathcal{G} \text{ is a path such that,} \\ \text{for } h_\gamma \in H, \gamma(0) = e = \gamma(1)h_\gamma^{-1} \end{array} \right\}.$$

Because $\text{Hol}_e(\mathcal{G}, \omega)$ completely encapsulates the ambiguity in taking developments from e , we can assign unique developments to elements of $\mathcal{G}$ when $\text{Hol}_e(\mathcal{G}, \omega) = \{e\}$. In other words, whenever $\text{Hol}_e(\mathcal{G}, \omega) = \{e\}$, we get a geometric map $\text{dev}_e: (\mathcal{G}, \omega) \rightarrow (G, \omega_G)$, which we could reasonably call the *developing map based at e* , given by $\gamma(1)h \mapsto \gamma_G(1)h$ for (piecewise smooth) paths $\gamma: [0, 1] \rightarrow \mathcal{G}$ starting at $\gamma(0) = e$ and $h \in H$. When the geometry is flat and the base manifold is simply connected, the holonomy group is always trivial, and the induced map of dev_e on the base manifold is precisely the developing map in the usual sense for locally homogeneous geometric structures.

Crucially, it follows that the curvature vanishes over an open set if and only if loops that are null-homotopic in that open set have trivial holonomy.

3 Some useful tools for working with holonomy

To make appropriate use of holonomy and development for our present purposes, we will need a way to tell when a sequence of holonomies converges, as well as a way to reverse the process of development for sufficiently small paths. In this section, we will provide a couple of results that will allow us to do this, which are useful in their own right and will almost certainly see significant use in subsequent works.

3.1 Convergence of developments

For a Cartan geometry $(\mathcal{G}, \omega)$ of type (G, H) , an inner product g on the Lie algebra $\mathfrak{g}$ determines a Riemannian metric g_ω on $\mathcal{G}$ given by $g_\omega(\xi_1, \xi_2) := g(\omega(\xi_1), \omega(\xi_2))$; we call such a Riemannian metric on $\mathcal{G}$ an ω -*constant Riemannian metric*. Given an element $e \in \mathcal{G}$, we can use an arbitrary ω -constant Riemannian metric g_ω to imbue the set $\text{Path}_e(\mathcal{G})$ of (piecewise smooth) paths $\gamma: [0, 1] \rightarrow \mathcal{G}$ in $\mathcal{G}$ starting at $\gamma(0) = e$ with a useful topology.

Definition 3.1. The *first-order compact-open topology* on

$$\text{Path}_e(\mathcal{G}) := \{ \gamma \in C^0([0, 1]; \mathcal{G}) \mid \gamma \text{ is a piecewise smooth path and } \gamma(0) = e \},$$

the space of piecewise smooth paths $\gamma: [0, 1] \rightarrow \mathcal{G}$ starting at $\gamma(0) = e$, is the topology induced by the metric

$$d_{g_\omega}(\gamma_1, \gamma_2) := \sup_{t \in [0, 1]} (\text{dist}_{g_\omega}(\gamma_1(t), \gamma_2(t))) + \sup_{t \in [0, 1]} \|\omega(\dot{\gamma}_1(t)) - \omega(\dot{\gamma}_2(t))\|_g,$$

where the second supremum is taken over all $t \in [0, 1]$ such that both $\dot{\gamma}_1(t)$ and $\dot{\gamma}_2(t)$ are well defined.

This is essentially just an extension of the usual C^1 compact-open topology on smooth paths so that it is defined on piecewise smooth paths instead.

When we say that a sequence of paths converges in $\text{Path}_e(\mathcal{G})$, we will mean that the sequence converges with respect to the first-order compact-open topology.

Lemma 3.2 (continuous dependence on parameters). *Suppose (f_λ) is a net of functions from $[0, 1]$ to $\mathfrak{g}$ that converges uniformly to some function $f_\infty: [0, 1] \rightarrow \mathfrak{g}$. If there exists a solution $\gamma_\infty \in \text{Path}_e(\mathcal{G})$ to the ODE given by $\omega(\dot{\gamma}_\infty) = f_\infty$, and solutions $\gamma_\lambda \in \text{Path}_e(\mathcal{G})$ to $\omega(\dot{\gamma}_\lambda) = f_\lambda$ for each λ , then (γ_λ) converges to γ_∞ in $\text{Path}_e(\mathcal{G})$.*

Proof. As the labeling implies, this result is effectively just a version of the usual “continuous dependence on parameters” lemma that one typically encounters for ODEs, except that it is adapted to our current setting. The main tool we will need is the “fundamental lemma” 1.5.1 from [4, Chapter 2], which tells us that, with the norm $\|\cdot\|$ on a Banach space E , a subset $U \subseteq E$, and an interval I , if a function $f: I \times U \rightarrow E$ is k -Lipschitz, $\gamma: I \rightarrow U$ solves the ODE given by $\dot{\gamma}(t) = f(t, \gamma(t))$, and $\bar{\gamma}: I \rightarrow U$ satisfies $\|\dot{\gamma}(t) - f(t, \bar{\gamma}(t))\| \leq \varepsilon$ for all $t \in I$, then for $\tau \in I$, the distance $\|\gamma(t) - \bar{\gamma}(t)\|$ satisfies the bound

$$\|\gamma(t) - \bar{\gamma}(t)\| \leq \|\gamma(\tau) - \bar{\gamma}(\tau)\| e^{k|t-\tau|} + \varepsilon \frac{e^{k|t-\tau|} - 1}{k}.$$

The strategy, in essence, is just to apply this “fundamental lemma” in finitely many coordinate balls inside charts along γ_∞ . This will show that γ_λ converges uniformly to γ_∞ in each of the finitely many balls, hence it converges to γ_∞ uniformly, and since the derivatives converge uniformly by assumption, it follows that γ_λ converges to γ_∞ in the first-order compact-open topology.

To be more explicit, let us cover $\gamma_\infty([0, 1])$ in finitely many coordinate balls $B_0, \dots, B_n$, and cover $[0, 1]$ by closed subintervals $I_0, \dots, I_n$ with $\max(I_i) \in I_{i+1}$ and $\gamma_\infty(I_i) \subseteq B_i$ for each i . Since $f_\infty(t)$ is constant in $\mathcal{G} \in \mathcal{G}$, we can find $k_i > 0$ for which the coordinate expressions for the smooth vector fields $\omega^{-1}(f_\infty(t))$ are k_i -Lipschitz on $\bar{B}_i$ with respect to the Euclidean norm $\|\cdot\|_{\text{coord}}$ induced by the coordinates; for clarity, note that this norm only gives a metric on the chart, not all of $\mathcal{G}$, but since $\bar{B}_i$ is compact, the restriction of the coordinate metric to $\bar{B}_i$ is bi-Lipschitz equivalent to the restriction of an ω -constant metric, so uniform convergence on $\bar{B}_i$ with respect to this metric implies uniform convergence with respect to an ω -constant one and vice versa. By compactness of $[0, 1]$ and $\bar{B}_i$,

$$\varepsilon_i(\lambda) := \max_{\substack{t \in [0, 1] \\ \mathcal{G} \in \bar{B}_i}} \|\omega_\mathcal{G}^{-1}(f_\lambda(t)) - \omega_\mathcal{G}^{-1}(f_\infty(t))\|_{\text{coord}},$$

exists for each i and λ , and it converges to 0 as $\lambda \rightarrow \infty$ because f_λ converges uniformly to f_∞ by assumption. Starting with $t_0 := 0$, so that $\|\gamma_\lambda(t_0) - \gamma_\infty(t_0)\|_{\text{coord}} = 0$, the “fundamental lemma” of [4] tells us that, for λ sufficiently far along toward ∞ ,

$$\|\gamma_\lambda(t) - \gamma_\infty(t)\|_{\text{coord}} \leq \varepsilon_0(\lambda) \frac{e^{k_0 t} - 1}{k_0} \leq \varepsilon_0(\lambda) \frac{e^{k_0} - 1}{k_0}$$

for all $t \in I_0$, so as $\lambda \rightarrow \infty$, $\gamma_\lambda|_{I_0}$ converges uniformly to $\gamma_\infty|_{I_0}$. Next, choose $t_1 \in I_0 \cap I_1$; by taking λ to be sufficiently far along toward ∞ , we have just seen (now using the coordinate metric in B_1 rather than B_0) that we can get $\|\gamma_\lambda(t_1) - \gamma_\infty(t_1)\|_{\text{coord}}$ as close to 0 as we would like, so using the “fundamental lemma” of [4], we get that

$$\|\gamma_\lambda(t) - \gamma_\infty(t)\|_{\text{coord}} \leq \|\gamma_\lambda(t_1) - \gamma_\infty(t_1)\|_{\text{coord}} e^{k_1} + \varepsilon_1(\lambda) \frac{e^{k_1} - 1}{k_1}$$

for all $t \in I_1$, hence $\gamma_\lambda|_{I_1}$ converges uniformly to $\gamma_\infty|_{I_1}$ as $\lambda \rightarrow \infty$. Iteratively, we repeat this process, picking $t_i \in I_{i-1} \cap I_i$ and then using the uniform convergence over I_{i-1} and the “fundamental lemma” from [4] to show that, for λ sufficiently far along toward ∞ , we get

$$\|\gamma_\lambda(t) - \gamma_\infty(t)\|_{\text{coord}} \leq \|\gamma_\lambda(t_i) - \gamma_\infty(t_i)\|_{\text{coord}} e^{k_i} + \varepsilon_i(\lambda) \frac{e^{k_i} - 1}{k_i}$$

for all $t \in I_i$, hence $\gamma_\lambda|_{I_i}$ converges uniformly to $\gamma_\infty|_{I_i}$ for each of the finitely many i . As noted above, this implies that γ_λ converges uniformly to γ_∞ on all of $[0, 1]$, and since, by assumption, the derivatives converge uniformly as well, γ_λ converges to γ_∞ in $\text{Path}_e(\mathcal{G})$. ■

This naturally leads to the following corollary, which seems implicit in much of the existing literature.

Proposition 3.3. *The development map $(\cdot)_G: \text{Path}_e(\mathcal{G}) \rightarrow \text{Path}_e(G)$ given by $\gamma \mapsto \gamma_G$ is continuous.*

Proof. Since $\text{Path}_e(\mathcal{G})$ is a metric space, it suffices to prove that if $\gamma_k \rightarrow \gamma_\infty$ in $\text{Path}_e(\mathcal{G})$, then $(\gamma_k)_G \rightarrow (\gamma_\infty)_G$ in $\text{Path}_e(G)$. But $\gamma_k \rightarrow \gamma_\infty$ in $\text{Path}_e(\mathcal{G})$ implies that $\omega(\dot{\gamma}_k) = \omega_G((\dot{\gamma}_k)_G)$ converges uniformly to $\omega(\dot{\gamma}_\infty) = \omega_G((\dot{\gamma}_\infty)_G)$, so by Lemma 3.2, $(\gamma_k)_G \rightarrow (\gamma_\infty)_G$ in $\text{Path}_e(G)$. ■

In particular, it follows from Proposition 3.3 that a sequence of paths converging to e in $\text{Path}_e(\mathcal{G})$ must have developments converging to the identity in G .

Lemma 3.4. *Suppose (γ_k) is a sequence in $\text{Path}_e(\mathcal{G})$. If (γ_k) converges in $\text{Path}_e(\mathcal{G})$ to the constant path at e , then the developments $(\gamma_k)_G(1)$ converge to the identity element $e \in G$.*

For the purposes of using this lemma, it will be helpful to keep two particular facts in mind.

First, the endpoint of the development of a path does not depend on parametrization: if τ is a (piecewise smooth) reparametrization of $[0, 1]$ with $\tau(0) = 0$ and $\tau(1) = 1$, then $\gamma^*\omega = \gamma_G^*\omega_G$ implies

$$(\gamma \circ \tau)^*\omega = \tau^*\gamma^*\omega = \tau^*\gamma_G^*\omega_G = (\gamma_G \circ \tau)^*\omega_G,$$

so $(\gamma \circ \tau)_G = \gamma_G \circ \tau$ and $(\gamma \circ \tau)_G(1) = \gamma_G(\tau(1)) = \gamma_G(1)$. Therefore, even if a sequence of paths in $\text{Path}_e(\mathcal{G})$ converges to the constant path *modulo parametrization*, the endpoints of their developments will still converge to the identity in G .

Second, the lemma is applicable even when the converging paths we are considering are on the base manifold: if (γ_k) is a sequence of paths in $\mathcal{G}$ such that $\gamma_k(0) \in eH$ and the paths $q_H(\gamma_k)$ in the base manifold converge to the constant path at $q_H(e)$ to first order, then we can find paths $\beta_k: [0, 1] \rightarrow H$ such that $\gamma_k(0)\beta_k(0)^{-1} = e$ and $\gamma_k\beta_k^{-1} \rightarrow e$ in $\text{Path}_e(\mathcal{G})$. In this situation, we should take particular note of how the sequence (β_k) of paths in H plays a similar role to what we call *ballast sequences*—the original term “holonomy sequence”, as used in [2, 7, 14], among several other places, is clearly inappropriate given the present context, since they are not sequences of holonomies for the geometry—in Section 4 below: while the paths γ_k might not converge to anything, we can use the paths β_k to counterbalance them so that the end result $\gamma_k\beta_k^{-1}$ still converges.

3.2 Existence of antidevelopments

It would be convenient to be able to reverse the process of development: given a piecewise smooth path $\gamma: [0, 1] \rightarrow G$ with $\gamma(0) = e$, we would like to be able to construct a corresponding path $\gamma_{\mathcal{G}}$ in $\mathcal{G}$ such that γ is the development of $\gamma_{\mathcal{G}}$. We can reasonably call such a path—if it exists—an *antidevelopment* of γ .

Definition 3.5. For a Cartan geometry $(\mathcal{G}, \omega)$ of type (G, H) and a path $\gamma \in \text{Path}_e(G)$, an *antidevelopment* for γ starting at $e \in \mathcal{G}$ is a path $\gamma_{\mathcal{G}} \in \text{Path}_e(\mathcal{G})$ such that $(\gamma_{\mathcal{G}})_G = \gamma$.

Unlike with developments, antidevelopments do not always exist. Conceptually, we can regard this asymmetry as a consequence of the fact that ω -constant metrics g_{ω} for a Cartan geometry $(\mathcal{G}, \omega)$ are not necessarily complete, whereas ω_G -constant metrics on G are always complete. In other words, there will often be paths that “escape” $\mathcal{G}$ despite having finite length with respect to some g_{ω} , but this can never happen for finite-length paths in G because ω_G -constant metrics on G must be complete.

However, if we fix a particular element $e \in \mathcal{G}$ and an ω -constant metric g_{ω} , then we can always find a sufficiently small radius $r_e > 0$ for which the metric ball $B_{r_e}(e)$ has compact closure. This radius will usually depend on the choice of element e unless g_{ω} is complete, but it will always be greater than 0. Paths starting at e of length less than r_e will be trapped inside the compact subspace $\overline{B_{r_e}(e)}$, so if $\gamma: [0, 1] \rightarrow G$ is a path such that $\gamma(0) = e$ whose length with respect to g_{ω_G} is less than r_e , then its antidevelopment $\gamma_{\mathcal{G}}$ will not be able to escape $\overline{B_{r_e}(e)}$. In particular, we have the following.

Lemma 3.6. *Suppose $(\mathcal{G}, \omega)$ is a Cartan geometry of type (G, H) and g is an inner product on $\mathfrak{g}$. For each $e \in \mathcal{G}$, there exists an $r_e > 0$ such that every path $\gamma \in \text{Path}_e(G)$ whose length with respect to g_{ω_G} is less than r_e has a well-defined antidevelopment $\gamma_{\mathcal{G}}$ starting at e .*

In summary, sufficiently short paths $\gamma \in \text{Path}_e(G)$ will always have antidevelopments at a given point $e \in \mathcal{G}$.

4 Isotropies that “ensnare” a region of the model

In [2, 14], the overall strategy for demonstrating local flatness near a given higher-order fixed point was to find an open region of the geometry whose points were pulled toward that higher-order fixed point in a particular way. Here, we will establish two criteria to capture the behavior underlying these results when the higher-order fixed point is *isolated*, meaning that some neighborhood of it contains no other higher-order fixed points with the same isotropy up to conjugation. These criteria—codified in Definition 4.2—are entirely determined by the behavior of the isotropy in the relevant model geometry, which we will not require to be parabolic in general. Because this behavior only depends on the isotropy, this allows us to transfer this desired behavior to every Cartan geometry admitting an automorphism with such isotropy.

4.1 Criteria for the isotropy

To understand the expected behavior of automorphisms near an isolated higher-order fixed point, which we will take as inspiration for the ideas below, it is useful to look at some examples on low-dimensional model geometries. We have sketched such examples for real and complex projective geometry in Figure 1.

Given an automorphism α with a higher-order fixed point, the results in [2, 14] suggest we should look for particular types of *ballast sequences* to find an open set on which the curvature vanishes.

Definition 4.1. Let $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ be an automorphism for a Cartan geometry $(\mathcal{G}, \omega)$ of type (G, H) . A sequence (b_k) in H is called a *ballast sequence* for α at $\mathfrak{g} \in \mathcal{G}$ with attractor $e \in \mathcal{G}$ if and only if there exists a sequence $(\mathfrak{g}_k)$ in $\mathcal{G}$ such that $\mathfrak{g}_k \rightarrow \mathfrak{g}$ and $\alpha^k(\mathfrak{g}_k)b_k^{-1} \rightarrow e$.

In particular, for isolated higher-order fixed points, the techniques in [2, 14] suggest looking for ballast sequences (b_k) contained in the Levi subgroup $G_0 < P$ of our parabolic isotropy P such




Figure 1. Sketches of automorphisms with an isolated higher-order fixed point for 2-dimensional real projective geometry ($\mathrm{PGL}_3 \mathbb{R}, P$) (left) and 1-dimensional complex projective geometry ($\mathrm{PGL}_2 \mathbb{C}, P$) (right).

that the restriction $\mathrm{Ad}_{b_k}|_{\mathfrak{g}_-}$ of the adjoint action to $\mathfrak{g}_-$ converges to 0. We should think of this as an assumption on the behavior of the derivatives of the automorphism on the tangent bundle over the base manifold M . To understand why, suppose that (b_k) is such a ballast sequence at $\mathcal{G}$, so that $\alpha^k(\mathcal{G}_k)b_k^{-1} \rightarrow e$ for some sequence $(\mathcal{G}_k)$ in $\mathcal{G}$ converging to $\mathcal{G}$ and $\mathrm{Ad}_{b_k}(v) \rightarrow 0$ for all $v \in \mathfrak{g}_-$. For sufficiently small $v \in \mathfrak{g}_-$, we therefore get a well-defined sequence of elements $\exp(\omega^{-1}(v))\mathcal{G}_k$ converging to $\exp(\omega^{-1}(v))\mathcal{G}$, where $\exp(\omega^{-1}(v))$ denotes the time-1 flow along the ω -constant vector field $\omega^{-1}(v)$, so that

$$\alpha^k(\exp(\omega^{-1}(v))\mathcal{G}_k)b_k^{-1} = \exp(\omega^{-1}(\mathrm{Ad}_{b_k}v))(\alpha^k(\mathcal{G}_k)b_k^{-1}) \rightarrow e.$$

Taking such v to be in a bounded neighborhood V of $0 \in \mathfrak{g}_-$, we can make $\mathrm{Ad}_{b_k}(V)$ as close to 0 as we would like for k sufficiently large, so for every neighborhood of $q_P(e)$, we may find a large enough k that $\alpha^k q_P(\exp(\omega^{-1}(V))\mathcal{G}_k)$ is contained in that neighborhood. In other words, the existence of such a ballast sequence for each $\mathcal{G} \in q_P^{-1}(U)$ over some open set $U \subseteq M$ guarantees that the iterates α^k , viewed as maps on the base manifold, converge *locally uniformly* to the constant map at $q_P(e)$ on U . Moreover, for whatever $w \in \mathfrak{g}_-$ we choose, we can define $\exp(t\omega^{-1}(w))\exp(\omega^{-1}(v))\mathcal{G}_k$ for sufficiently small $t \in \mathbb{R}$, and

$$\begin{aligned} \alpha^k(\exp(t\omega^{-1}(w))\exp(\omega^{-1}(v))\mathcal{G}_k)b_k^{-1} \\ = \exp(t\omega^{-1}(\mathrm{Ad}_{b_k}w))\exp(\omega^{-1}(\mathrm{Ad}_{b_k}v))(\alpha^k(\mathcal{G}_k)b_k^{-1}) \rightarrow e; \end{aligned}$$

restricting our possible choices of w to a bounded set in $\mathfrak{g}_-$, we can again get all $\mathrm{Ad}_{b_k}(w)$ as close to 0 as we would like, so the corresponding curves $t \mapsto \alpha^k(q_P(\exp(t\omega^{-1}(w))\exp(\omega^{-1}(v))\mathcal{G}_k))$ converge to first order to the constant curve at $q_P(e)$. Thus, we are looking for automorphisms α with a region U over which iterates of the derivatives of α converge locally uniformly to the zero vector over some point of the base manifold.

Another necessary component in the techniques of [2, 14] is the presence of certain paths that get shrunk to a point of the base manifold by the automorphism. The flat regions constructed in those papers are covered by such shrinking paths, so in particular, there should be at least one of them. We will only need to specify one such path, which we will call ζ_U , in the base manifold of the model geometry, and moreover, we will not need it to be the quotient image of an exponential path, though (for convenience) we will choose for it to be so in all of the applications in Section 5.

Combining these two criteria leads us to the following definition.

Definition 4.2. We say that $a \in H$ *ensnares* an a -invariant connected open subset $U \subseteq G/H$ if and only if the following two criteria hold:

- Over U , the derivative maps $a_*^k: T(G/H) \rightarrow T(G/H)$ of the iterates a^k converge locally uniformly to the constant map at the zero vector $0_{q_H(e)}$ at $q_H(e)$ as $k \rightarrow +\infty$.
- There exists a path $\zeta_U: [0, 1] \rightarrow G/H$ such that $\zeta_U(0) = q_H(e)$, $\zeta_U(1) \in U$, and, for some sequence (τ_k) of orientation-preserving reparametrizations of $[0, 1]$, $a^k(\zeta_U \circ \tau_k)$ converges to the constant path at $q_H(e)$ to first order as $k \rightarrow +\infty$.

Assuming local uniform convergence over U of the derivatives of the iterates of a on G/H provides a convenient way to guarantee that paths inside U are pulled to $q_H(e)$. More precisely, it tells us that whenever we have a path in U , its iterates under a must converge to the constant path at $q_H(e)$ to first order.

For paths γ_1 and γ_2 such that $\gamma_1(1) = \gamma_2(0)$, so that the endpoint of the first is the starting point of the second, we denote by $\gamma_1 \star \gamma_2$ their *concatenation*

$$\gamma_1 \star \gamma_2(t) := \begin{cases} \gamma_1(2t) & \text{if } t \in [0, \frac{1}{2}], \\ \gamma_2(2t - 1) & \text{if } t \in [\frac{1}{2}, 1]. \end{cases}$$

From what Definition 4.2 tells us about ζ_U , together with the aforementioned convergence of paths in U under forward iterates of a , it follows that concatenations of the form $\zeta_U \star q_H(\delta)$, where $\delta: [0, 1] \rightarrow q_H^{-1}(U)$ is a path such that $q_H(\delta(0)) = \zeta_U(1)$, get pulled into $q_H(e)$, so that $a^k((\zeta_U \circ \tau_k) \star q_H(\delta))$ also converges to $q_H(e)$ to first order. As illustrated in Figure 2, the resulting picture can end up looking a bit like a snare being pulled closed, hence the name.




Figure 2. A concatenation of the form $\zeta_U \star q_H(\delta)$ (left) and some of its iterates under a (right).

Before we proceed, it is worth noting again that, while the two criteria of Definition 4.2 were designed to generalize the behavior used by [2, 14] for their results on isolated higher-order fixed points, this parabolic context is not actually required in the definition. In particular, even though our main applications in this paper happen to involve automorphisms of parabolic geometries with isolated higher-order fixed points, we will not need to use this as an explicit assumption for the general results below. For example, in affine geometry $(\text{Aff}(m), \text{GL}_m \mathbb{R})$, the contracting homothety $a = \lambda \mathbf{1} \in \text{GL}_m \mathbb{R}$ (where $|\lambda| < 1$) ensnares the whole model space $U = \text{Aff}(m)/\text{GL}_m \mathbb{R} \cong \mathbb{R}^m$, with ζ_U the constant path at $q_{\text{GL}_m \mathbb{R}}(e)$.

4.2 Existence of ensnared region with trivial holonomy

Since we can pull in concatenations of the form $\zeta_U \star q_H(\delta)$ to be as short as we want, we can show that they always have well-defined antidevelopments when there exists an automorphism with ensnaring isotropy.

Lemma 4.3. *If $(\mathcal{G}, \omega)$ is a Cartan geometry of type (G, H) , $e \in \mathcal{G}$, and $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ such that $\alpha(e) = ea$ for some $a \in H$ that ensnares $U \subseteq G/H$, then every path of the form $\hat{\zeta}_U \star \delta$, where $\hat{\zeta}_U \in \text{Path}_e(G)$ is a lift of ζ_U and $\delta: [0, 1] \rightarrow q_H^{-1}(U)$ such that $\delta(0) = \hat{\zeta}_U(1)$, has a well-defined antidevelopment $(\hat{\zeta}_U \star \delta)_{\mathcal{G}}$ starting at e .*

Proof. Since $a^k((\zeta_U \circ \tau_k) \star q_H(\delta)) = a^k q_H((\hat{\zeta}_U \circ \tau_k) \star \delta)$ converges to the constant path at $q_H(e)$ to first order, there exist $\beta_k: [0, 1] \rightarrow H$ such that $\beta_k(0) = a^k$ and $a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1}$ converges, modulo parametrization, to the constant path at e in $\text{Path}_e(G)$. As a result, for each $r > 0$, there is some $k \geq 0$ for which $a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1}$ has length less than r .

By Lemma 3.6, there exists an $r_e > 0$ for which all paths in $\text{Path}_e(G)$ of length less than r_e have well-defined antidevelopments starting at $e \in \mathcal{G}$. Thus, $a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1}$ has an antidevelopment $(a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1})_{\mathcal{G}}$ starting at e for all sufficiently large k , hence

$$(\hat{\zeta}_U \star \delta)_{\mathcal{G}} := \alpha^{-k}(a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1})_{\mathcal{G}}\beta_k$$

is a well-defined path starting at $\alpha^{-k}(e)\beta_k(0) = ea^{-k}a^k = e$ with development

$$\begin{aligned} ((\hat{\zeta}_U \star \delta)_{\mathcal{G}})_G &= (\alpha^{-k}(a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1})_{\mathcal{G}}\beta_k)_G = \beta_k(0)^{-1}((a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1})_{\mathcal{G}})_G\beta_k \\ &= a^{-k}(a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1})\beta_k = \hat{\zeta}_U \star \delta. \end{aligned} \quad \blacksquare$$

Whenever $(\mathcal{G}, \omega)$ has an automorphism $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ with isotropy $a \in H$ that ensnares $U \subseteq G/H$, we can use Lemma 4.3 to construct an H -invariant subspace

$$q_H^{-1}(U)_{\mathcal{G}} := \{(\hat{\zeta}_U \star \delta)_{\mathcal{G}}(1)h \mid \delta \in \text{Path}_{\hat{\zeta}_U(1)}(q_H^{-1}(U)), h \in H\} \subseteq \mathcal{G}$$

given by the endpoints of antidevelopments for concatenations of the form $\hat{\zeta}_U \star \delta$ together with their right-translations by elements of H . Note that $q_H^{-1}(U)_{\mathcal{G}}$ is necessarily an open subset of $\mathcal{G}$: for each $\delta(1) \in q_H^{-1}(U)$, sufficiently short paths starting at $\delta(1)$ stay in the open set $q_H^{-1}(U)$, so each $(\hat{\zeta}_U \star \delta)_{\mathcal{G}}(1) \in q_H^{-1}(U)_{\mathcal{G}}$ is the center of a small open ball contained in $q_H^{-1}(U)_{\mathcal{G}}$.

We should think of $q_H^{-1}(U)_{\mathcal{G}}$ as the analogue in $\mathcal{G}$ for the open subset $q_H^{-1}(U)$ in G , and we should hope for it to exhibit similar behavior with respect to the automorphism $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ to what $q_H^{-1}(U)$ does for the ensnaring isotropy $a \in H$. Indeed, as we mentioned in the proof of the lemma, each path of the form $\hat{\zeta}_U \star \delta$ admits a sequence of paths $\beta_k: [0, 1] \rightarrow H$ with $\beta_k(0) = a^k$ such that $a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1}$ converges to the constant path e in $\text{Path}_e(G)$ modulo parametrization, and

$$(\alpha^k(\hat{\zeta}_U \star \delta)_{\mathcal{G}}\beta_k^{-1})_G = \beta_k(0)(\hat{\zeta}_U \star \delta)\beta_k^{-1} = a^k(\hat{\zeta}_U \star \delta)\beta_k^{-1},$$

so we must also have that $\alpha^k(\hat{\zeta}_U \star \delta)_{\mathcal{G}}\beta_k^{-1}$ converges to e in $\text{Path}_e(\mathcal{G})$ modulo parametrization. Consequently, the derivatives α_*^k on the base manifold must converge locally uniformly to $0_{q_H(e)}$ over the connected open set $U_{\mathcal{G}} := q_H(q_H^{-1}(U)_{\mathcal{G}})$, just like in the model case. Using this, we can prove the key result for this section: when $a^kga^{-k} \rightarrow e$ implies $g = e$ —as necessarily happens when a is unipotent, for example—the antideveloped region $q_H^{-1}(U)_{\mathcal{G}}$ has trivial holonomy.

Theorem 4.4. *Suppose $(\mathcal{G}, \omega)$ is a Cartan geometry of type (G, H) with an automorphism $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ and an element $e \in \mathcal{G}$ such that $\alpha(e) = ea$ for some $a \in H$ ensnaring $U \subseteq G/H$. If $a^kga^{-k} \rightarrow e$ as $k \rightarrow +\infty$ implies $g = e$, then $\text{Hol}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)}(q_H^{-1}(U)_{\mathcal{G}}, \omega) = \{e\}$.*

Proof. Suppose $\gamma: [0, 1] \rightarrow q_H^{-1}(U)_{\mathcal{G}}$ is a path such that

$$\gamma(0) = \gamma(1)h^{-1} = (\hat{\zeta}_U)_{\mathcal{G}}(1)$$

for some $h \in H$. We want to show that $\gamma_G(1)h^{-1} = e$. To do this, consider the path

$$\gamma_{\zeta_U} := (\hat{\zeta}_U)_{\mathcal{G}} \star \gamma \star \overline{\text{R}_h(\hat{\zeta}_U)_{\mathcal{G}}},$$

where $\bar{\eta}$ is the *reverse* of the path η , given by $\bar{\eta}(t) := \eta(1 - t)$, and R_h denotes right-translation by h .

Since $\alpha_*^k \rightarrow 0_{q_H(e)}$ locally uniformly over $U_{\mathcal{G}}$ as $k \rightarrow +\infty$, the paths $\alpha^k q_H(\gamma_{\zeta_U})$ converge, modulo parametrization, to $q_H(e)$ to first order as before, so there exist paths $\beta_k: [0, 1] \rightarrow H$ with $\beta_k(0) = a^k$ such that, again modulo parametrization, $\alpha^k \gamma_{\zeta_U} \beta_k^{-1}$ converges to e in $\text{Path}_e(\mathcal{G})$. Thus, by Lemma 3.4, the developments

$$(\alpha^k \gamma_{\zeta_U} \beta_k^{-1})_G(1) = \beta_k(0)(\gamma_{\zeta_U})_G(1)\beta_k(1)^{-1} = a^k(\gamma_{\zeta_U})_G(1)\beta_k(1)^{-1}$$

must converge to the identity element $e \in G$. Note, in particular, that the endpoint

$$\alpha^k(\gamma_{\zeta_U}(1))\beta_k(1)^{-1} = \alpha^k(\text{R}_h(\overline{\hat{\zeta}_U})_{\mathcal{G}}(1))\beta_k(1)^{-1} = \alpha^k(eh)\beta_k(1)^{-1} = ea^k h\beta_k(1)^{-1}$$

converges to e , so $a^k h\beta_k(1)^{-1} \rightarrow e$, hence

$$a^k((\gamma_{\zeta_U})_G(1)h^{-1})a^{-k} = (a^k(\gamma_{\zeta_U})_G(1)\beta_k(1)^{-1})(a^k h\beta_k(1)^{-1})^{-1} \rightarrow ee^{-1} = e.$$

But by hypothesis, this implies that $(\gamma_{\zeta_U})_G(1)h^{-1} = e$, and

$$\begin{aligned} (\gamma_{\zeta_U})_G(1) &= ((\hat{\zeta}_U)_{\mathcal{G}} \star \gamma \star \text{R}_h(\overline{\hat{\zeta}_U})_{\mathcal{G}})_G(1) = ((\hat{\zeta}_U)_{\mathcal{G}})_G(1)\gamma_G(1)(\text{R}_h(\overline{\hat{\zeta}_U})_{\mathcal{G}})_G(1) \\ &= \hat{\zeta}_U(1)\gamma_G(1)h^{-1}\hat{\zeta}_U(1)^{-1}h, \end{aligned}$$

so $\gamma_G(1)h^{-1} = \hat{\zeta}_U(1)^{-1}(\gamma_{\zeta_U})_G(1)h^{-1}\hat{\zeta}_U(1) = \hat{\zeta}_U(1)^{-1}\hat{\zeta}_U(1) = e$. ■

It is, perhaps, worth explicitly pointing out that Theorem 4.4 starkly improves previous results that exhibit local flatness near an isolated higher-order fixed point: with considerably weaker assumptions, we have still obtained a large open subset $q_H^{-1}(U)_{\mathcal{G}}$ containing e in its closure, and not only does the curvature vanish on this subset, but its holonomy is trivial, even without additional curvature restrictions or considerations other than the behavior of the isotropy on the model geometry. Moreover, if we also happen to know a bit about the topology of U , then we can even show that this subspace is isomorphic to its analogue in the model.

Corollary 4.5. *In the setting of Theorem 4.4, if U is simply connected, then there exists a geometric embedding $\sigma: (q_H^{-1}(U), \omega_G) \hookrightarrow (\mathcal{G}, \omega)$ with $\sigma(q_H^{-1}(U)) = q_H^{-1}(U)_{\mathcal{G}}$.*

Proof. Naturally, we want to try to define $\sigma: q_H^{-1}(U) \rightarrow q_H^{-1}(U)_{\mathcal{G}} \subseteq \mathcal{G}$ by

$$\sigma(\delta(1)h) := (\hat{\zeta}_U \star \delta)_{\mathcal{G}}(1)h$$

for each path $\delta \in \text{Path}_{\hat{\zeta}_U(1)}(q_H^{-1}(U))$ and each $h \in H$. If σ is well defined, then it is H -equivariant and satisfies $\omega(\sigma_*\dot{\delta}(1)) = \omega_G(\dot{\delta}(1))$ by definition, so it is necessarily geometric. It remains, then, to prove that σ is well defined and that σ is injective.

If we have two paths $\delta_1, \delta_2 \in \text{Path}_{\hat{\zeta}_U(1)}(q_H^{-1}(U))$ with the same endpoint $\delta_1(1) = \delta_2(1)$, then we can concatenate δ_1 with the reverse of δ_2 to get a loop $\delta_1 \star \bar{\delta}_2$ based at $\hat{\zeta}_U(1)$. In particular, to prove that $(\hat{\zeta}_U \star \delta_1)_{\mathcal{G}}(1) = (\hat{\zeta}_U \star \delta_2)_{\mathcal{G}}(1)$ for all such paths δ_1 and δ_2 , it suffices to prove that $(\hat{\zeta}_U \star \delta)_{\mathcal{G}}(1) = (\hat{\zeta}_U)_{\mathcal{G}}(1)$ for each loop $\delta \in \text{Path}_{\hat{\zeta}_U(1)}(q_H^{-1}(U))$ based at $\hat{\zeta}_U(1)$.

For each loop $\delta \in \text{Path}_{\hat{\zeta}_U(1)}(q_H^{-1}(U))$, we can find a loop $\beta \in \text{Path}_e(H)$ such that there exists a null-homotopy $c: [0, 1]^2 \rightarrow q_H^{-1}(U)$ from $c_0 = \delta\beta^{-1}$ to the constant loop $c_1 = \hat{\zeta}_U(1)$, since U is simply connected. From Lemma 4.3, we get a well-defined antidevelopment $(\hat{\zeta}_U \star c_s)_{\mathcal{G}}$ for each $s \in [0, 1]$, and we can show that the map $\sigma(c): [0, 1]^2 \rightarrow q_H^{-1}(U)_{\mathcal{G}}$ given by

$$\sigma(c)_s(t) := \sigma(c_s(t)) = (\hat{\zeta}_U \star c_s)_{\mathcal{G}}\left(\frac{1+t}{2}\right)$$

satisfies $\sigma(c)^*\omega = c^*\omega_G$, essentially for the same reason as Theorem 6.1 in [16, Chapter 3]: the distribution $\ker(c^*\omega_G - \omega)$ on $[0, 1] \times q_H^{-1}(U)_{\mathcal{G}}$ is integrable because the curvature $d\omega + \frac{1}{2}[\omega, \omega]$

vanishes on the subspace $q_H^{-1}(U)_{\mathcal{G}}$, so we can build the graph of $\sigma(c)$ locally from the integral submanifolds for the distribution, and $\sigma(c)_s: t \mapsto (\hat{\zeta}_U \star c_s)_{\mathcal{G}}(\frac{1+t}{2})$ is the unique path starting at $\sigma(c)_s(0) = (\hat{\zeta}_U)_{\mathcal{G}}(1)$ such that $\sigma(c)_s^* \omega = c_s^* \omega_G$. Since $\sigma(c)_s(0)$ and $\sigma(c)_s(1)$ are constant in s and $\sigma(c)_1(t)$ is constant in t , it follows that $\sigma(c)_s(0) = \sigma(c)_1(t) = \sigma(c)_s(1) = (\hat{\zeta}_U)_{\mathcal{G}}(1)$, so $\sigma(c)$ is a null-homotopy as well. In particular,

$$\begin{aligned} (\hat{\zeta}_U \star \delta)_{\mathcal{G}}(1) &= (\hat{\zeta}_U \star \delta \beta^{-1})_{\mathcal{G}}(1) \beta(1) = (\hat{\zeta}_U \star \delta \beta^{-1})_{\mathcal{G}}(1) \\ &= (\hat{\zeta}_U \star c_0)_{\mathcal{G}}(1) = \sigma(c)_0(1) = \sigma(c)_0(0) = (\hat{\zeta}_U)_{\mathcal{G}}(1), \end{aligned}$$

so σ is a well-defined geometric map.

Finally, to show that σ is injective, we consider the shifted developing map

$$\hat{\zeta}_U(1) \text{dev}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)}: q_H^{-1}(U)_{\mathcal{G}} \rightarrow q_H^{-1}(U)$$

given by $\gamma(1) \mapsto \hat{\zeta}_U(1) \gamma_G(1)$ for each path $\gamma \in \text{Path}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)}(q_H^{-1}(U)_{\mathcal{G}})$; by Theorem 4.4, we know that $\text{Hol}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)}(q_H^{-1}(U)_{\mathcal{G}}, \omega) = \{e\}$, so the map $\hat{\zeta}_U(1) \text{dev}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)}$ is necessarily well defined on $q_H^{-1}(U)_{\mathcal{G}}$. Because both σ and $\hat{\zeta}_U(1) \text{dev}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)}$ are geometric maps, it follows that the composition $\sigma \circ (\hat{\zeta}_U(1) \text{dev}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)})$ is a geometric map from $q_H^{-1}(U)_{\mathcal{G}}$ to itself such that

$$(\sigma \circ (\hat{\zeta}_U(1) \text{dev}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)}))((\hat{\zeta}_U)_{\mathcal{G}}(1)) = \sigma(\hat{\zeta}_U(1)) = (\hat{\zeta}_U)_{\mathcal{G}}(1),$$

so $\hat{\zeta}_U(1) \text{dev}_{(\hat{\zeta}_U)_{\mathcal{G}}(1)} = \sigma|_{q_H^{-1}(U)_{\mathcal{G}}}^{-1}$, hence σ is injective. ■

In the next section, we will further refine Corollary 4.5 in the case where the open set $U = (G/H) \setminus \text{Fix}_{G/H}(a)$ is simply connected and $\text{Fix}_{G/H}(a)$ has codimension at least two. For such situations, we will see that $q_H^{-1}(U)_{\mathcal{G}}$ already makes up so much of $\mathcal{G}$ that we can essentially just extend σ^{-1} to all of $\mathcal{G}$ by filling in the small gaps where the fixed points would go.

5 Applications

Here, we will prove Theorems A, B, and C from the introduction. To do this, we will prove a more general result, and then we will show that the hypotheses of this result are satisfied in each of the desired cases.

Throughout this section, we will need a notion of codimension; it will be convenient to use the following terminology, which comes from considerations of transversality.

Definition 5.1. For a smooth manifold M and a closed subset $N \subseteq M$, we will say that N has *codimension at least two* if and only if, for each path $\gamma: [0, 1] \rightarrow M$ with $\gamma(0) \notin N$, there exists a homotopy $c: [0, 1]^2 \rightarrow M$, $(s, t) \mapsto c_s(t)$ such that $c_0 = \gamma$, $c_s(0) = \gamma(0)$ for all $s \in [0, 1]$, and $c((0, 1] \times [0, 1]) \cap N = \emptyset$.

Closed submanifolds and subvarieties $N \subseteq M$ will, of course, have codimension at least two in this sense if $\dim(M) - \dim(N) \geq 2$. Moreover, the complement of N must be dense in M if N satisfies this definition: for each point $p \in M$, we can find a sequence of points of the form $c_{s_i}(t_i) \notin N$ with $s_i \rightarrow 0$ such that $c_{s_i}(t_i) \rightarrow p$.

5.1 General result

In this subsection, we will prove the following result, from which Theorems A, B, and C will follow.

Theorem 5.2. *Let $(\mathcal{G}, \omega)$ be a Cartan geometry of type (G, H) over a connected smooth manifold M , $\alpha \in \text{Aut}(\mathcal{G}, \omega)$, and $e \in \mathcal{G}$ such that $\alpha(e) = ea$ for some $a \in H$. Suppose that the isotropy $a \in H$ ensnares $U = (G/H) \setminus \text{Fix}_{G/H}(a)$, that U is simply connected, and that $\text{Fix}_{G/H}(a)$ has codimension at least two in G/H . If $a^k g a^{-k} \rightarrow e$ as $k \rightarrow +\infty$ implies $g = e$, then $(\mathcal{G}, \omega)$ geometrically embeds onto a dense open subset of the Klein geometry (G, ω_G) of type (G, H) .*

By Corollary 4.5, the hypotheses of the theorem tell us that there is a geometric embedding $\sigma: (q_H^{-1}(U), \omega_G) \hookrightarrow (\mathcal{G}, \omega)$. It turns out that this is already sufficient to prove the theorem: U has complement of codimension at least two, so we can use [8, Theorem 1.8] to get a geometric isomorphism between $(\mathcal{G}, \omega)$ and a dense open subset of G . For completeness—and to continue to demonstrate the utility of holonomy for obtaining global results on Cartan geometries—we provide an alternative proof.

Intuitively, the image $\sigma(U) = U_{\mathcal{G}}$ of σ should already make up nearly all of M , and the holonomy is trivial over this subset. The main idea behind our proof of Theorem 5.2, then, is to use the codimension of the complement of U in G/H to homotope loops in M an arbitrarily small amount to get loops in $\sigma(U)$, thus forcing the holonomy to be globally trivial.

Proposition 5.3. *Suppose $(\mathcal{G}, \omega)$ is a Cartan geometry of type (G, H) over a connected smooth manifold M . Let $\mathcal{F} \subseteq G$ be an H -invariant open subset of G such that the complement $G \setminus \mathcal{F}$ has codimension at least two. If there exists a geometric embedding $\sigma: (\mathcal{F}, \omega_G) \hookrightarrow (\mathcal{G}, \omega)$, then $\text{Hol}_e(\mathcal{G}, \omega) = \{e\}$ for any $e \in \mathcal{G}$.*

Proof. Suppose $\gamma: [0, 1] \rightarrow \mathcal{G}$ is the lift of a loop $q_H(\gamma)$ on M with $\gamma(0) = e$, and $h_\gamma \in H$ is such that $\gamma(0)h_\gamma = \gamma(1)$. We want to prove that $\gamma_G(1)h_\gamma^{-1} = e$; since M is connected, we may assume without loss of generality that $e \in \sigma(\mathcal{F})$, and since G acts transitively on itself, we may further assume that $e \in \mathcal{F}$ and $\sigma(e) = e$.

Because $G \setminus \mathcal{F}$ has codimension at least two in G , there exists a homotopy $c: [0, 1]^2 \rightarrow G$ such that $c_0 = \gamma_G$, $c_s(0) = \gamma_G(0) = e$ for all $s \in [0, 1]$, and $c((0, 1] \times [0, 1]) \subseteq \mathcal{F}$; see Figures 3 and 4 for illustrations. Then, since σ is a geometric map and $c_s(0) = e \in \mathcal{F}$ for each $s \in [0, 1]$, $\sigma(c_s)_G = c_s(0)^{-1}c_s = c_s$ for all $s \in (0, 1]$, hence—by Lemma 3.2— $\sigma(c_s)$ converges to $\gamma \in \text{Path}_e(\mathcal{G})$ as $s \rightarrow 0$. In particular,

$$\sigma(c_s(1)h_\gamma^{-1}) = \sigma(c_s(1))h_\gamma^{-1} \rightarrow \gamma(1)h_\gamma^{-1} = \gamma(0) = \sigma(e),$$

so since σ is an embedding, we must have $c_s(1)h_\gamma^{-1} \rightarrow \gamma_G(1)h_\gamma^{-1} = e$, hence $\text{Hol}_e(\mathcal{G}, \omega) = \{e\}$. ■

Recall that, whenever $\text{Hol}_e(\mathcal{G}, \omega) = \{e\}$, the developing map

$$\text{dev}_e: (\mathcal{G}, \omega) \rightarrow (G, \omega_G)$$

based at e is a well-defined geometric map given by $\gamma(1)h \mapsto \gamma_G(1)h$ for every $h \in H$ and piecewise smooth path $\gamma: [0, 1] \rightarrow \mathcal{G}$ starting at $\gamma(0) = e$. Therefore, if we can show that dev_e is injective, then the proof of Theorem 5.2 will be done.

To do this, we will use the following lemma, which tells us that a geometric map (or more generally any local homeomorphism from a Hausdorff space) is injective when its restriction to the preimage of a dense subset is injective. We imagine that this lemma is well known, but could not find a reference.

Lemma 5.4. *Suppose $\psi: (\mathcal{G}, \omega) \rightarrow (\mathcal{Q}, v)$ is a geometric map and that $V \subseteq \mathcal{Q}$ is a dense subset. If $\psi|_{\psi^{-1}(V)}$ is injective, then ψ is injective.*

Proof. Suppose that $\psi(\mathcal{G}) = \psi(\mathcal{G}')$ for some $\mathcal{G}, \mathcal{G}' \in \mathcal{G}$. Since ψ is a local diffeomorphism, there exists an open neighborhood W of $\mathcal{G}'$ such that $\psi|_W$ is a diffeomorphism onto its image. If $(\mathcal{G}_i)$ is a sequence in $\psi^{-1}(V)$ converging to $\mathcal{G}$, which must exist because density of V in $\mathcal{Q}$ implies density of $\psi^{-1}(V)$ in $\mathcal{G}$, then $\psi(\mathcal{G}_i) \rightarrow \psi(\mathcal{G}) = \psi(\mathcal{G}')$, so $\psi|_W^{-1}(\psi(\mathcal{G}_i)) \rightarrow \psi|_W^{-1}(\psi(\mathcal{G}')) = \mathcal{G}'$. But $\psi|_{\psi^{-1}(V)}$ is injective, so $\psi|_W^{-1}(\psi(\mathcal{G}_i)) = \mathcal{G}_i \rightarrow \mathcal{G}$; since $\mathcal{G}$ is Hausdorff, this means $\mathcal{G} = \mathcal{G}'$. ■




Figure 3. The path γ_G in G , with neighborhoods of its intersection with $G \setminus \mathcal{F}$ highlighted in gray.




Figure 4. A path c_ε in G , which satisfies $c_\varepsilon([0, 1]) \subseteq \mathcal{F}$ and coincides with $\gamma_G = c_0$ outside of the highlighted neighborhoods of the intersections of γ_G with $G \setminus \mathcal{F}$.

Corollary 5.5. *In the setting of Proposition 5.3, the developing map $\text{dev}_e: (\mathcal{G}, \omega) \rightarrow (G, \omega_G)$ is a geometric embedding.*

Proof. As we did in the proof of Proposition 5.3, we may assume that $e \in \mathcal{F}$ and $\sigma(e) = e$. By definition, we have that $\text{dev}_e(\gamma) = \gamma_G$ whenever $\gamma(0) = e$, so $(\text{dev}_e \circ \sigma)(\gamma_G) = \text{dev}_e(\gamma) = \gamma_G$ whenever γ_G is contained in $\mathcal{F}$. Hence, $\text{dev}_e \circ \sigma$ is the identity map on $\mathcal{F}$.

It just remains to show that $\text{dev}_e: \mathcal{G} \rightarrow G$ is injective. To this end, suppose $\text{dev}_e(\mathfrak{g}) \in \mathcal{F}$, and let $\gamma \in \text{Path}_e(\mathcal{G})$ be a path such that $\gamma(1) \in \mathfrak{g}H$. Then, $\text{dev}_e(\gamma)$ is a path in G with $\text{dev}_e(\gamma(0)) = \text{dev}_e(e) = e$ and $\text{dev}_e(\gamma(1)) = \gamma_G(1) \in \text{dev}_e(\mathfrak{g})H$. Since $G \setminus \mathcal{F}$ has codimension at least two, there exists a homotopy $c: [0, 1]^2 \rightarrow G$ such that $c_0 = \text{dev}_e(\gamma) = \gamma_G$, $c_s(0) = e$ for all $s \in [0, 1]$, and $c((0, 1) \times [0, 1]) \subseteq \mathcal{F}$. By continuity, $\sigma(c_s(1)) \rightarrow \sigma(c_0(1)) = \sigma(\gamma_G(1))$ as $s \rightarrow 0$, and by Lemma 3.2, $\sigma(c_s(1)) \rightarrow \gamma(1)$ as $s \rightarrow 0$. Because $\mathcal{G}$ is Hausdorff, this means that $\gamma(1) = \sigma(\gamma_G(1)) \in \sigma(\mathcal{F})$, hence that $\mathfrak{g} \in \sigma(\mathcal{F})$. Thus, $\text{dev}_e^{-1}(\mathcal{F}) = \sigma(\mathcal{F})$, so because $\text{dev}_e|_{\sigma(\mathcal{F})}$ is injective, dev_e is injective by Lemma 5.4. $\blacksquare$

In summary, under the hypotheses of Theorem 5.2, we have shown that the developing map dev_e gives a geometric embedding from $(\mathcal{G}, \omega)$ onto a dense open subset of the Klein geometry (G, ω_G) , so the proof is complete.

5.2 Proof of Theorem A

Recall that, in this case, our model geometry is $(\text{PGL}_{m+1} \mathbb{C}, P)$, which corresponds to complex projective geometry over $\mathbb{CP}^m$. We want to prove that, given a Cartan geometry $(\mathcal{G}, \omega)$ of type $(\text{PGL}_{m+1} \mathbb{C}, P)$ over a connected smooth manifold M and a nontrivial geometric automorphism $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ such that $\alpha(e) \in eP_+$ for some $e \in \mathcal{G}$, there exists a geometric embedding of $(\mathcal{G}, \omega)$ onto a dense open subset of the Klein geometry $(\text{PGL}_{m+1} \mathbb{C}, \omega_{\text{PGL}_{m+1} \mathbb{C}})$.

For $\alpha(e) = ea$, let us write

$$a = \begin{pmatrix} 1 & \beta \\ 0 & 1 \end{pmatrix}$$

for some $\beta^\top \in \mathbb{C}^m$. In view of Theorem 5.2, it suffices to prove that a ensnares the complement $U = \mathbb{CP}^m \setminus \text{Fix}_{\mathbb{CP}^m}(a)$, that U is simply connected, and that $\text{Fix}_{\mathbb{CP}^m}(a)$ has codimension at least two.

For arbitrary

$$q_P(g) = \begin{pmatrix} r \\ x \end{pmatrix} \in \mathbb{CP}^m,$$

with $r \in \mathbb{C}$ and $x \in \mathbb{C}^m$, we have

$$a^k(q_P(g)) = a^k \begin{pmatrix} r \\ x \end{pmatrix} = \begin{pmatrix} 1 & k\beta \\ 0 & 1 \end{pmatrix} \begin{pmatrix} r \\ x \end{pmatrix} = \begin{pmatrix} r + k\beta(x) \\ x \end{pmatrix},$$

which is equal to

$$\begin{pmatrix} 1 \\ \frac{1}{r + k\beta(x)} x \end{pmatrix}$$

whenever $r + k\beta(x) \neq 0$. In particular, fixed points are precisely those points for which $\beta(x) = 0$, and if $\beta(x) \neq 0$, then $a^k(q_P(g))$ must go to $q_P(e)$ as $k \rightarrow +\infty$. Since

$$\text{Fix}_{\mathbb{CP}^m}(a) = \left\{ \begin{pmatrix} r \\ x \end{pmatrix} \in \mathbb{CP}^m \mid \beta(x) = 0 \right\}$$

is a complex projective hyperplane in $\mathbb{CP}^m$, $U = \mathbb{CP}^m \setminus \text{Fix}_{\mathbb{CP}^m}(a)$ is diffeomorphic to $\mathbb{C}^m$, hence it is simply connected. Moreover, since $\beta(x) = 0$ implies $\text{Re}(\beta(x)) = 0$ and $\text{Im}(\beta(x)) = 0$, $\text{Fix}_{\mathbb{CP}^m}(a)$ has (real) codimension 2 in $\mathbb{CP}^m$.

Now, all that is left to do is prove that a ensnares U . For this, note that $\zeta_U: [0, 1] \rightarrow \mathbb{CP}^m$, given by $t \mapsto \begin{pmatrix} 1 \\ ty \end{pmatrix}$ for some $y \in \mathbb{C}^m$ such that $\beta(y) \in (0, +\infty) \subset \mathbb{R}$, shrinks into itself under forward iterates of a , so that $a^k(\zeta_U(t)) = \zeta_U(\frac{t}{1+k\beta(y)})$. Thus, for $\tau_k: t \mapsto \frac{t}{1+k\beta(y)(1-t)}$, we get $a^k(\zeta_U(\tau_k(t))) = \zeta_U(\frac{t}{1+k\beta(y)})$, hence $a^k(\zeta_U \circ \tau_k)$ converges to $q_P(e)$ to first order. Moreover, since $a^k(q_P(g))$ converges to $q_P(e)$ for each $q_P(g) \in U$, we may assume that $r + k\beta(x) \neq 0$ for all $k \geq 0$ after replacing $q_P(g)$ with $a^n(q_P(g))$ for suitably large n . This allows us to work in the coordinate chart $r = 1$, and we know that in those coordinates, a^k is given by $x \mapsto \frac{1}{1+k\beta(x)}x$. Therefore, for $\xi \in T_{q_P(g)}U$ and considering $a_*^k\xi$ as a point derivation on the vector of coordinate functions x ,

$$\begin{aligned} (a_*^k\xi)(x) &= \xi \left(\frac{1}{1+k\beta(x)}x \right) = \frac{1}{1+k\beta(x)} \left(\xi(x) - \frac{k\xi(\beta(x))}{1+k\beta(x)}x \right) \\ &= \frac{1}{(1+k\beta(x))^2} (\xi(x) + k(\beta(x)\xi(x) - \xi(\beta(x))x)), \end{aligned}$$

so since $(1+k\beta(x))^2$ grows quadratically in k while the numerator $\xi(x) + k(\beta(x)\xi(x) - \xi(\beta(x))x)$ only grows linearly in k , $a_*^k\xi \rightarrow 0_{q_P(e)}$ as $k \rightarrow +\infty$. Indeed, for all sufficiently small open subsets $V \subset TU \subseteq T(G/P)$, $\pi: T(G/P) \rightarrow G/P$ the natural projection from the tangent bundle to $G/P \cong \mathbb{CP}^m$, $\|\cdot\|$ a norm on $\mathbb{C}^m$, and $k > (\inf_{x \in \pi(V)} |\beta(x)|)^{-1}$, we have

$$\left\| \frac{1}{1+k\beta(x)}x \right\| = \frac{\|x\|}{|1+k\beta(x)|} \leq \frac{\sup_{x \in \pi(V)} \|x\|}{\inf_{x \in \pi(V)} |\beta(x)|k - 1} \rightarrow 0$$

and

$$\begin{aligned} \|(a_*^k \xi)(x)\| &= \left\| \xi \left(\frac{1}{1 + k\beta(x)} x \right) \right\| \\ &\leq \frac{1}{(\inf_{x \in \pi(V)} |\beta(x)| k - 1)^2} \left(k \left(\sup_{x \in \pi(V)} |\beta(x)| \sup_{\xi \in V} \|\xi(x)\| \right. \right. \\ &\quad \left. \left. + \sup_{\xi \in V} |\xi(\beta(x))| \sup_{x \in \pi(V)} \|x\| \right) + \sup_{\xi \in V} \|\xi(x)\| \right) \rightarrow 0 \end{aligned}$$

for x in $\pi(V)$ and $\xi \in V$ over the chart $r = 1$, so the convergence of a_*^k to $0_{q_P(e)}$ is locally uniform over U . In other words, a ensnares U , so the proof of Theorem A is complete.

5.3 Proof of Theorem B

This proof is largely the same as the one for Theorem A. In this case, the model becomes $(\mathrm{PGL}_{m+1} \mathbb{H}, P)$, with $\mathrm{PGL}_{m+1} \mathbb{H}/P \cong \mathbb{HP}^m$. We want to prove, given a Cartan geometry $(\mathcal{G}, \omega)$ of type $(\mathrm{PGL}_{m+1} \mathbb{H}, P)$ over a connected smooth manifold M and a nontrivial automorphism $\alpha \in \mathrm{Aut}(\mathcal{G}, \omega)$ such that $\alpha(e) \in eP_+$ for some $e \in \mathcal{G}$, that there exists a geometric embedding of $(\mathcal{G}, \omega)$ onto a dense open subset of the Klein geometry $(\mathrm{PGL}_{m+1} \mathbb{H}, \omega_{\mathrm{PGL}_{m+1} \mathbb{H}})$.

As before, for $\alpha(e) = ea$, we write

$$a = \begin{pmatrix} 1 & \beta \\ 0 & \mathbf{1} \end{pmatrix}$$

for some $\beta^\top \in \mathbb{H}^m$. Again, our goal is just to apply Theorem 5.2, so we just need to prove that a ensnares $U = \mathbb{HP}^m \setminus \mathrm{Fix}_{\mathbb{HP}^m}(a)$, that U is simply connected, and that $\mathrm{Fix}_{\mathbb{HP}^m}(a)$ has codimension at least two.

For arbitrary

$$q_P(g) = \begin{pmatrix} r \\ x \end{pmatrix} \in \mathbb{HP}^m,$$

with $r \in \mathbb{H}$ and $x \in \mathbb{H}^m$, we have

$$a^k(q_P(g)) = a^k \begin{pmatrix} r \\ x \end{pmatrix} = \begin{pmatrix} 1 & k\beta \\ 0 & \mathbf{1} \end{pmatrix} \begin{pmatrix} r \\ x \end{pmatrix} = \begin{pmatrix} r + k\beta(x) \\ x \end{pmatrix},$$

which is equal to

$$\begin{pmatrix} 1 \\ x(r + k\beta(x))^{-1} \end{pmatrix}$$

whenever $r + k\beta(x) \neq 0$. In particular, fixed points are precisely those points for which $\beta(x) = 0$, and if $\beta(x) \neq 0$, then $a^k(q_P(g))$ must go to $q_P(e)$ as $k \rightarrow +\infty$. This time, the set $\mathrm{Fix}_{\mathbb{HP}^m}(a)$ of all points for which $\beta(x) = 0$ is of codimension $\dim(\mathbb{H}) = 4$, since $\beta(x) = 0$ implies its projections to the real subspaces of $\mathbb{H}$ generated by 1, i, j, and ij all vanish as well. Since $\mathbb{HP}^m$ is simply connected and $\mathrm{Fix}_{\mathbb{HP}^m}(a)$ has codimension greater than two, it also follows that $U = \mathbb{HP}^m \setminus \mathrm{Fix}_{\mathbb{HP}^m}(a)$ is simply connected.

We again end the proof by proving that a ensnares U . As before, the path $\zeta_U: [0, 1] \rightarrow \mathbb{HP}^m$ given by $t \mapsto \begin{pmatrix} 1 \\ ty \end{pmatrix}$ for some $y \in \mathbb{H}^m$ such that $\beta(y) \in (0, +\infty) \subset \mathbb{R}$, shrinks into itself under forward iterates of a , so $a^k(\zeta_U)$ must converge to $q_P(e)$ to first order modulo parametrization. After replacing $q_P(g) \in U$ with some forward iterate under a , we can assume $r + k\beta(x) \neq 0$

for all $k \geq 0$, allowing us to work in the coordinate chart $r = 1$, where we know a^k is given by $x \mapsto x(1 + k\beta(x))^{-1}$. Considering $a_*^k \xi$ as a point derivation on the coordinate vector x ,

$$(a_*^k \xi)(x) = \xi(x(1 + k\beta(x))^{-1}) = (\xi(x) - kx(1 + k\beta(x))^{-1}\xi(\beta(x)))(1 + k\beta(x))^{-1},$$

which goes to 0 as $k \rightarrow +\infty$, so $a_*^k(\xi) \rightarrow 0_{q_P(e)}$. Indeed, since the convergence comes from the disparity of degrees in the polynomials in k as before, we again get that a_*^k converges locally uniformly to $0_{q_P(e)}$, hence a ensnares U and Theorem B is proven.

5.4 Proof of Theorem C

This time, our model geometry is given by $(\text{PU}(\mathfrak{h}_{p,q}), P)$, basically corresponding to the natural CR structure on the null-cone for $\mathfrak{h}_{p,q}$ in $\mathbb{CP}^{p+q+1}$. Our goal is to prove that, if $(\mathcal{G}, \omega)$ is a Cartan geometry of type $(\text{PU}(\mathfrak{h}_{p,q}), P)$ over a connected smooth manifold M and $\alpha \in \text{Aut}(\mathcal{G}, \omega)$ is an automorphism such that $\alpha(e) = ea$ for some non-null $a \in P_+$, then there exists a geometric embedding from $(\mathcal{G}, \omega)$ onto a dense open subset of the Klein geometry $(\text{PU}(\mathfrak{h}_{p,q}), \omega_{\text{PU}(\mathfrak{h}_{p,q})})$. Again, this comes down to applying Theorem 5.2.

Let us write

$$a = \begin{pmatrix} 1 & \beta & \text{is} -\frac{1}{2}\beta I_{p,q}\bar{\beta}^\top \\ 0 & \mathbb{1} & -I_{p,q}\bar{\beta}^\top \\ 0 & 0 & 1 \end{pmatrix},$$

where, by hypothesis, $\beta I_{p,q}\bar{\beta}^\top \neq 0$ if $\beta \neq 0$. For arbitrary

$$q_P(g) = \begin{pmatrix} r \\ x \\ c \end{pmatrix} \in \text{Null}(\mathfrak{h}_{p,q}) \subseteq \mathbb{CP}^{p+q+1},$$

with $r, c \in \mathbb{C}$ and $x \in \mathbb{C}^{p+q}$ such that $2\text{Re}(\bar{r}c) + \bar{x}^\top I_{p,q}x = 0$, we get

$$\begin{aligned} a^k(q_P(g)) &= a^k \begin{pmatrix} r \\ x \\ c \end{pmatrix} = \begin{pmatrix} 1 & k\beta & \text{iks} - \frac{k^2}{2}\beta I_{p,q}\bar{\beta}^\top \\ 0 & \mathbb{1} & -kI_{p,q}\bar{\beta}^\top \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} r \\ x \\ c \end{pmatrix} \\ &= \begin{pmatrix} r + k(\beta(x) + \text{ics}) - \frac{ck^2}{2}\beta I_{p,q}\bar{\beta}^\top \\ x - kcI_{p,q}\bar{\beta}^\top \\ c \end{pmatrix}, \end{aligned}$$

so $q_P(g) \in \text{Fix}_{\text{Null}(\mathfrak{h}_{p,q})}(a)$ if and only if $\beta(x) + \text{ics} = 0$ and $cI_{p,q}\bar{\beta}^\top = 0$. If $\beta = 0$ and $s \neq 0$, then this is just the set of points such that $c = 0$, which has codimension $\dim(\mathbb{C}) = 2$, and $U = \text{Null}(\mathfrak{h}_{p,q}) \setminus \text{Fix}_{\text{Null}(\mathfrak{h}_{p,q})}(a)$ is diffeomorphic to $i\mathbb{R} \times \mathbb{C}^{p+q}$ by normalizing to $c = 1$. If $\beta \neq 0$, then we get the set of points such that $\beta(x) = 0$ and $c = 0$, which has codimension 4. Either way, $\text{Fix}_{\text{Null}(\mathfrak{h}_{p,q})}(a)$ has codimension at least two, $U = \text{Null}(\mathfrak{h}_{p,q}) \setminus \text{Fix}_{\text{Null}(\mathfrak{h}_{p,q})}(a)$ is simply connected, and $a^k(q_P(g))$ converges to $q_P(e)$ as $k \rightarrow +\infty$ whenever $q_P(g)$ is not a fixed point.

It remains to prove that a ensnares U . In this case, we can choose $\zeta_U: [0, 1] \rightarrow \text{Null}(\mathfrak{h}_{p,q})$ to be given by

$$\zeta_U(t) := \begin{pmatrix} t(t-s) - \frac{i}{2}\beta I_{p,q}\bar{\beta}^\top \\ -tiI_{p,q}\bar{\beta}^\top \\ t^2i \end{pmatrix},$$

so that $a^k(\zeta_U(t)) = \zeta_U(\frac{t}{1+kt})$ and hence $a^k(\zeta_U(\tau_k(t))) = \zeta_U(\frac{t}{1+k})$ for the reparametrizations $\tau_k: t \mapsto \frac{t}{1+k(1-t)}$ of $[0, 1]$. From this, it follows that the sequence $a^k(\zeta_U \circ \tau_k)$ converges to $q_P(e)$ to first order. After possibly replacing $q_P(g)$ with some forward iterate under a , we can assume that $r + k(\beta(x) + ics) - \frac{ck^2}{2}\beta I_{p,q}\bar{\beta}^\top \neq 0$ for all $k \geq 0$, so we can work in the chart where we normalize to $r = 1$, with coordinates given by

$$\begin{pmatrix} 1 \\ x \\ c \end{pmatrix} = \begin{pmatrix} 1 \\ x \\ i \operatorname{Im}(c) - \frac{\bar{x}^\top I_{p,q} x}{2} \end{pmatrix} \mapsto \begin{bmatrix} x \\ \operatorname{Im}(c) \end{bmatrix} \in \mathbb{C}^{p+q} \times \mathbb{R}.$$

For $r_k := 1 + k(\beta(x) + ics) - \frac{ck^2}{2}\beta I_{p,q}\bar{\beta}^\top$, the expression for a^k in these coordinates is

$$\begin{bmatrix} x \\ \operatorname{Im}(c) \end{bmatrix} \mapsto \begin{bmatrix} \frac{1}{r_k}(x - kcI_{p,q}\bar{\beta}^\top) \\ \operatorname{Im}(\frac{c}{r_k}) \end{bmatrix}.$$

Therefore, since

$$\xi \left(\frac{1}{r_k}(x - kcI_{p,q}\bar{\beta}^\top) \right) = \frac{1}{r_k^2}(r_k(\xi(x) - k\xi(c)I_{p,q}\bar{\beta}^\top) - \xi(r_k)(x - kcI_{p,q}\bar{\beta}^\top))$$

goes to $0 \in \mathbb{C}^{p+q}$ and $\xi(\frac{c}{r_k}) = \frac{\xi(c)r_k - c\xi(r_k)}{r_k^2}$ goes to $0 \in \mathbb{C}$ as $k \rightarrow +\infty$ for each $\xi \in T_{q_P(g)}U$, we have $a_*^k(\xi) \rightarrow 0_{q_P(e)}$. Again, the convergence to $0_{q_P(e)}$ comes from a disparity in the degrees of the polynomials in k , and it follows that a_*^k converges locally uniformly over U , hence a ensnares U , which completes the proof of Theorem C.

Acknowledgements

We would specifically like to thank Karin Melnick for several helpful suggestions and conversations on this topic; she was arguably instrumental in the beginning stages of this work. We are also deeply grateful to Charles Frances, both for inspiring this paper and for tolerating and tempering the author's enthusiasm, and an anonymous referee who reviewed an earlier version of this paper and managed to provide a number of useful improvements. The reviewers at SIGMA also suggested some implemented changes, and we appreciate them for their efforts. Finally, we gratefully acknowledge partial support from the Brin Graduate Fellowship at the University of Maryland and the National Science Foundation under Grant No. 2203493 during the time in which this research was conducted.

References

- [1] Calderbank D.M.J., Eastwood M.G., Matveev V.S., Neusser K., C-projective geometry, *Mem. Amer. Math. Soc.* **267** (2020), v+137 pages, [arXiv:1512.04516](#).
- [2] Čap A., Melnick K., Essential Killing fields of parabolic geometries, *Indiana Univ. Math. J.* **62** (2013), 1917–1953, [arXiv:1208.5510](#).
- [3] Čap A., Slovák J., Parabolic geometries. I. Background and general theory, *Math. Surveys Monogr.*, Vol. 154, *American Mathematical Society*, Providence, RI, 2009.
- [4] Cartan H., Differential calculus, Hermann, Paris, 1971.
- [5] Erickson J.W., Intrinsic holonomy and curved cosets of Cartan geometries, *Eur. J. Math.* **8** (2022), 446–474, [arXiv:2109.05350](#).
- [6] Erickson J.W., A method for determining Cartan geometries from the local behavior of automorphisms, [arXiv:2303.00561](#).

-
- [7] Frances C., Sur le groupe d'automorphismes des géométries paraboliques de rang 1, *Ann. Sci. École Norm. Sup. (4)* **40** (2007), 741–764.
 - [8] Frances C., Rigidity at the boundary for conformal structures and other Cartan geometries, [arXiv:0806.1008](#).
 - [9] Frances C., Local dynamics of conformal vector fields, *Geom. Dedicata* **158** (2012), 35–59, [arXiv:0909.0044](#).
 - [10] Goldman W.M., Geometric structures on manifolds, *Grad. Stud. Math.*, Vol. 227, [American Mathematical Society](#), Providence, RI, 2022.
 - [11] Kruglikov B., The D., Jet-determination of symmetries of parabolic geometries, *Math. Ann.* **371** (2018), 1575–1613, [arXiv:1604.07149](#).
 - [12] Ma T., Geodesic rigidity of Levi-Civita connections admitting essential projective vector fields, *Geom. Dedicata* **205** (2020), 147–166, [arXiv:1705.00460](#).
 - [13] Melnick K., Rigidity of transformation groups in differential geometry, *Notices Amer. Math. Soc.* **68** (2021), 721–732, [arXiv:2009.13937](#).
 - [14] Melnick K., Neusser K., Strongly essential flows on irreducible parabolic geometries, *Trans. Amer. Math. Soc.* **368** (2016), 8079–8110, [arXiv:1410.4647](#).
 - [15] Nagano T., Ochiai T., On compact Riemannian manifolds admitting essential projective transformations, *J. Fac. Sci. Univ. Tokyo Sect. IA Math.* **33** (1986), 233–246.
 - [16] Sharpe R.W., Differential geometry: Cartan's generalization of Klein's Erlangen program, *Grad. Texts in Math.*, Vol. 166, Springer, New York, 1997.