Symmetry-Breaking Lie Flows for Optimisation of Discontinuous Functionals

Mikhail Osipov
Independent Researcher, Milan, Italy.

Email: osipov.ma@phystech.edu
Manuscript submitted July 9, 2025; accepted August 20, 2025; published September 23, 2025.
doi: 10.17706/ijapm.2025.15.2.91-101

Abstract: We address the problem of reducing a task cost functional $W(S)$, defined over Sobolev-class signals S, when W is invariant under a global symmetry group $G \subset Diff(M)$ and only accessible as a black-box. Such settings arise in machine learning, imaging, and inverse problems, where performance metrics are non-differentiable and internal to pretrained models. We propose a variational method that leverages symmetry to construct explicit, symmetry-breaking deformations of the input. By minimizing an auxiliary energy, we obtain a gauge field whose induced deformation $h = A_{\phi}[S]$ lies generically transverse to the G-orbit. We show that—even for discontinuous W—a simple double-sign test on h descends to a strictly lower-cost region with positive probability, and almost surely under mild geometric conditions. This method requires no model gradients or labels and operates entirely at test time. It offers a principled mechanism for optimizing invariant cost functionals via Lie-algebraic flows, with applications to black-box systems and symmetry-constrained tasks.

Key words: Symmetry breaking, non-smooth optimisation, machine learning

1. Introduction

Many optimisation problems in applied mathematics, imaging, and machine learning involve cost functionals that are invariant under symmetries of the domain—for example, spatial transformations that preserve classification accuracy or reconstruction quality. Such invariances often arise from physical or geometric considerations and are typically encoded through architectural design or data augmentation.

In this work, we study a distinct setting: test-time optimisation of a cost functional $W(S)$, where:

- $S \in S^s(M)$ lies in a Sobolev space over a smooth compact manifold M,
- $W: S^s(M) \to R$ is invariant under a Lie group $G \subset Diff(M)$,
- W is accessible only as a black box—possibly non-differentiable or discontinuous,
- gradients of W and supervision signals are unavailable at inference time.

The central question is: Can one construct an explicit deformation of S that provably reduces $W(S)$, using only symmetry information and without gradient access?

This problem arises in practical settings where the true evaluation metric—such as classification accuracy, word error rate, or edit distance—is discontinuous and not amenable to classical gradient-based methods. Standard approaches often rely on differentiable surrogates or semi-smooth approximations [1, 2], which may fail near decision boundaries or produce suboptimal solutions due to local trapping.

To address this, we propose a variational framework inspired by field theory. We introduce an auxiliary energy functional defined over scalar fields, whose minimisers induce symmetry-breaking deformations of the form:

$$S \mapsto exp(A_\phi)[S], \quad with \quad A_\phi = \sum_i \phi_i(x)L_i \equiv \phi^i(x)L_i,$$

where $\{L_i\}$ generate the Lie algebra of G. These deformations move S transversely to its group orbit, allowing variation in $W(S)$ despite its invariance under G.

Our goal is to analyse the geometric structure of such deformations and establish mechanisms for descending non-smooth or discontinuous objectives under group symmetries—without reliance on model internals or differentiable approximations.

2. Results

In this work we have developed and analyzed a novel *symmetry-breaking* framework for minimizing external costs on spaces of functions.

2.1. Energy Functional

We introduced an *external energy* functional $E_S[\phi]$, in analogy with Yang–Mills energies [3], whose minimiser generates gauge flows that align the signal S with its symmetry orbit.

2.2. Variational Analysis

We established *existence of minimisers* of E_S in the Sobolev space $H^s(M)$, $s > \frac{dim(M)}{2} + 1$, via coercivity, weak lower-semicontinuity, and the direct method. We ruled out *constant* gauge fields as global minimisers (Lemma 1) by a localized "bump" perturbation argument in the spirit of the Abrikosov–Nielsen–Olesen vortex energy in [4] and Ginzburg–Landau model on compact surfaces in [5].

2.3. Alignment of Minimisers

We showed that, up to first order in the regularisation weights α, β, the resulting displacement $w(\tilde{\phi})$ remains *anti-parallel* to the gradient of the orbit-distance $d_G(S)$ (Lemma 2), ensuring that gauge flows approximate the steepest-descent direction.

2.4. Symmetry-Breaking Descent

We proved a *randomized-ray descent theorem* demonstrating that—even for a discontinuous, piecewise-constant cost W—a symmetry-breaking update $h \in N_S$ will, with positive probability (and under mild geometric conditions on the nearby downhill boundary), exit into a strictly lower-cost region. In the globally convex, C^1 case this probability equals 1.

2.5. Numerical Case Studies

We validate our framework on diverse tasks, including image segmentation, speech recognition, and black-box classification. In each setting, applying the symmetry-breaking deformation yields consistent, measurable gains in non-differentiable performance metrics, confirming the practical effectiveness of our theoretical descent guarantees.

3. Related Work

3.1. Symmetry in Modelling and Learning

Symmetries are central to both applied mathematics and machine learning. In learning, *equivariant architectures* enforce invariance through design, including group convolution networks [6], gauge-equivariant Convolutional Neural Networks (CNNs) [7], and geometric deep learning frameworks [8]. Earlier techniques include tangent propagation [9] and learned transformations via Spatial Transformer Networks [10]. These methods rely on access to gradients and full model structure during training.

3.2. Variational Symmetry Breaking in Machine Learning

Within machine learning, symmetry-breaking has been explored at training time. Bamler and Mandt [11] propose Lie-algebraic optimisation for models with continuous symmetries. Tanaka *et al.* [12] link Noether symmetries to learning dynamics, while Elesedy [13] studies symmetry in relation to generalisation and loss landscapes. In contrast, our focus is on *test-time* symmetry breaking for black-box objectives.

3.3. Gauge Fields and Symmetry Breaking in Applied Mathematics

Gauge-theoretic methods are well established in mathematical physics and geometry. Yang [14] applies gauge-invariant variational principles to mechanical systems, while Sochen *et al.* [15] develop a gauge-invariant framework for image diffusion on manifolds. Diffeomorphic and gauge-based models underpin shape analysis and image registration [16, 17], with connections to stochastic flows and optimal transport [18].

3.4. Test-Time Adaptation and Black-Box Optimisation

Test-time adaptation methods address distribution shifts without retraining, using strategies like entropy minimisation and batch norm recalibration [19, 20], or randomized defences against adversarial inputs [21–23]. Most approaches assume differentiability and access to internal model gradients—assumptions we relax in our black-box, non-smooth setting.

4. Setup

4.1. Geometric Properties

We consider a smooth, compact, boundaryless manifold M, representing the spatial domain of signals, with a symmetry group G acting on M via diffeomorphisms [24, 25]. Signals are modeled as scalar and vector fields on M, and the associated cost functional W is assumed to be invariant under the action of G.

To ensure compatibility with variational analysis, signals are represented in Sobolev spaces of sufficiently high regularity [26, 27], providing the necessary continuity and differentiability for composing with diffeomorphisms and evaluating energy functionals.

The group G, a subgroup of the diffeomorphism group of M, encodes natural symmetries such as reparametrizations and coordinate transformations. Its invariance property ensures that the cost functional W respects these underlying geometric symmetries.

4.2. Functional Space

Let M be a smooth, compact, connected n-dimensional Riemannian manifold, and fix a Sobolev index

$$s > \frac{n}{2} + 1,$$

so that the embedding $H^s(M) \hookrightarrow C^1(M)$ holds [26].

4.2.1. Signal space

A signal S consists of:

- A scalar field $\phi \in H^s(M)$.
- Vector fields $L_i[S] \in H^s(M, TM)$, for $i = 1, \ldots, d$ [28].

The signal space is

$$S^s(M) := H^s(M) \times (H^s(M, TM))^d.$$

For $s > \frac{n}{2} + 1$, $S^s(M)$ is a Banach algebra under pointwise multiplication, and composition with C^1 maps is well defined and differentiable.

4.2.2. Symmetry group

Define the Sobolev-class diffeomorphism group:

$$Diff^{\,s+1}(M) = \{\varphi: M \to M \mid \varphi, \varphi^{-1} \in H^{s+1}(M, M)\}.$$

Let $G \subset Diff^{\,s+1}(M)$ be a closed Lie subgroup that acts on signals isometrically via pullback:

$$\varphi^* S := (\phi \circ \varphi^{-1},\ (D\varphi^{-1}) \circ \varphi^{-1} \cdot L_i[S] \circ \varphi^{-1})_{i=1}^{d},$$

which defines a C^1 action on $S^s(M)$ [24, 25].

5. External Energy Minimisation

5.1. Energy Functional

We seek a deformation $exp(A_\phi)$ of a signal S via the infinitesimal generator

$$A_\phi(x) = \phi^i(x)\, L_i,$$

where $\{L_i\}$ are fixed admissible vector fields. Define

$$E_S[\phi] = \parallel S - exp(A_\phi) \cdot S \parallel^2 + \alpha \parallel \nabla\phi \parallel^2 + \beta \parallel |\phi|^2 - v^2 \parallel^2, \tag{1}$$

With $\alpha, \beta > 0$. The data term measures misalignment under the flow $exp(A_\phi)$, the $\parallel \nabla\phi \parallel^2$ term enforces smoothness, and the double-well potential $\parallel |\phi|^2 - v^2 \parallel^2$ prevents degeneracy.

Under the Sobolev setting $\phi \in H^s(M, R^d)$, $s > \frac{n}{2} + 1$, the map $\phi \mapsto exp(A_\phi)$ is a smooth chart near the identity in $Diff^{\,s+1}(M)$ [29, 30], making E_S well defined, locally Lipschitz, lower semicontinuous, and $G-$invariant, thereby amenable to variational analysis.

5.1.1. Existence of Minimisers

Assume $s > \frac{dim(M)}{2} + 1$, so $H^s(M, R^d) \hookrightarrow C^1(M)$. We use the direct method.

5.1.1.1. Coercivity

Since $\parallel \nabla\phi \parallel^2$ and $\parallel |\phi|^2 - v^2 \parallel^2$ dominate the H^s-norm,

$$E_S[\phi] \to \infty \quad as \quad \parallel \phi \parallel_{H^s} \to \infty.$$

5.1.1.2. Weak lower semicontinuity

Sobolev norms are weakly l.s.c. and $\phi \mapsto exp(A_\phi) \cdot S$ is continuous in H^s. Hence E_S is weakly l.s.c. on H^s.

5.1.1.3. Direct method

Any minimizing sequence $\{\phi_n\} \subset H^s$ is bounded by coercivity. By Rellich–Kondrachov, a subsequence converges $\phi_n \rightharpoonup \hat{\phi}$ in H^s and strongly in L^2. Weak l.s.c. gives

$$E_S[\hat{\phi}] \leq \lim\, inf_{n\to\infty} E_S[\phi_n] = inf_\phi E_S[\phi],$$

So $\hat{\phi}$ is a minimiser.

5.1.2. No constant minimisers

Lemma 1. (No constant minimisers). *If $S \notin Fix(G) = \{\, g \cdot S : g \in G \,\}$, then no constant field $\phi(x) \equiv c$ can minimise E_S. Hence any global minimiser is nonconstant.*

Sketch. Starting from a constant $\phi \equiv c$, one builds a compactly supported perturbation $\phi_{\delta,r}$ that:

1. lowers the data term $\parallel S - exp(A_{\phi_{\delta,r}}) S \parallel^2$ at first order,

2. tweaks the double-well potential $\| |\phi|^2 - v^2 \|^2$ linearly,
3. incurs only a higher-order increase in $\| \nabla\phi \|^2$.

The net effect is a strict energy decrease, contradicting minimality of the constant field. Full details are in Appendix.

5.2. Descent Direction

The minimiser of

$$E_S[\phi] = E_{data}[\phi] + \alpha \| \nabla\phi \|^2 + \beta \| |\phi|^2 - v^2 \|^2,$$

yields a displacement

$$w(\tilde{\phi}) = S - exp(A_{\tilde{\phi}}) \cdot S.$$

Define the orbit-distance between the original signal and the warped signal:

$$d_G(S) = inf_{g\in G} \| S - g \cdot exp(A_{\tilde{\phi}}) \cdot S \|^2.$$

Then, under weak regularity conditions, the minimiser $\tilde{\phi}$ yields the displacement $w(\tilde{\phi})$ that is nearly aligned with $\nabla d_G(S)$.

Lemma 2. (Gauge displacement alignment). *Let* $S \in H^s(M, R^k)$, $s > \frac{dim(M)}{2} + 1$. *Denote* $S_\phi = exp(tA_\phi)S$, $w(\phi) = S - S_\phi$, $V_S = span\{L_i[S]\}$. Let P_V, $P_{V\perp}$ be L^2-orthogonal projectors, $\tilde{\phi}$ a minimizer of the energy functional (1) under Neumann boundary conditions.

Then the gradient of the orbit distance $\nabla d_G(S) = 2 \cdot P_{V\perp}(w(\tilde{\phi})) + O(t \| w(\tilde{\phi}) \|)$.

Sketch. First variation of energy (1) along constant directions $\delta\phi^i = c^i$ under boundary conditions gives

$$0 = \delta E_S[\phi] = 2t \langle w, L_i[S_\phi] \rangle + 4\beta \int_M (|\phi|^2 - v^2)\phi_i + O(t^2).$$

Then $\langle tA_{\tilde{\phi}} S, L_i[S] \rangle = O(\beta\delta) + O(t)$, where δ controls the double-well potential term. *Decompose* $w = w_T + w_N$ with $w_T = P_V(w)$, $w_N = P_{V\perp}(w)$. Then the tangential component norm $\| w_T \| = O(\beta\delta t + t^2) \ll \| w_N \| = O(t)$. Full details are in Appendix.

As a corollary, if $W = f(d_G)$ with $f' < 0$, then the gauge-induced update $h = A_{\tilde{\phi}}[S]$ aligns with the steepest descent of W.

6. Cost Functional Optimisation

We now focus on discontinuous cost functionals $W(S)$, which we assume to be piecewise constant without loss of generality. In smooth regions, standard gradient descent applies; however, near discontinuities—such as decision boundaries—the gradient is undefined, and standard techniques fail to detect viable descent directions.

Our goal is to analyse the behaviour of W near such boundaries under symmetry-breaking perturbations. Specifically, we study how structured, infinitesimal changes in the signal S within the symmetry-breaking subspace can identify descent directions with respect to the true, discontinuous cost. This framework helps explain why smooth surrogates can fail and provides a principled alternative when gradient information is unavailable.

Definitions. Let $C \subset H$ be a convex region whose boundary ∂C is a C^1 co-dimension-1, locally flat sub-manifold. Suppose $W|_C < W(S)$ for any $S \notin C$. Define tangent space to the group orbit at the point S: $A_S = T_S(G \cdot S)$, normal space $N_S = (A_S)^\perp$ and metric projection of S on C: $S^* = argmin_{x\in C} \| x - S \|$, with normal vector $l := S - S^*$. At S^* write the boundary tangent $L := T_{S^*}\partial C = l^\perp$.

Define the blind subspace: $BS := L \cap N_{S^*} = \{ w \in N_{S^*} : \langle w, l \rangle = 0\}$.

Theorem 1. (Descent towards a Convex Cell). Let $W: H \to R$ be piecewise-constant and invariant under a smooth isometric action of G by pullback. Fix $S \notin C$ that is close enough to the cell C.

Then:

1. If $BS = \{0\}$, one of the rays $S \pm t\,h$ meets ∂C at some $t > 0$ and W strictly decreases.
2. If $dim\ BS > 0$, sample additional perturbation δ from $\sim N(0, Cov)$ on N_{S^*} with full support: $h' = h + \varepsilon\delta, \varepsilon > 0$, Then one of the rays $S \pm t\,h'$ meets ∂C at some $t > 0$ and W strictly decreases with probability > 0; it happens almost surely over repeated independent draws of δ.

Proof Sketch.

According to Lemma 2, the gauge displacement h is aligned with the gradient of the orbit distance $\nabla d_G(S)$and lies in the normal space N_S up to linear regularization. Since S is close to C, also $h \in N_{S^*}$ up to a small error.

Generally, h will have a non-zero component towards the cell C, unless it falls into the blind subspace BS, which is composed by directions in the normal space to the group orbit that are tangential to the cell boundary ∂C.

a) If $BS = \{0\}$, for any gauge displacement h we have $h = -\kappa l + o(\| h \|)$. This situation happens if the symmetry group G fully describes the decision boundary ∂C. Gauge displacement in that case is anti-collinear to the boundary normal, so that one of the rays $S \pm t\,h$ meets ∂C at some $t > 0$ and W strictly decreases.

b) If BS is not trivial, gauge displacement h may be orthogonal to l. In that case, consider a Gaussian distribution $\sim N(0, Cov)$ on N_{S^*} with full support, draw a random perturbation δ and update the gauge displacement: $h' = h + \varepsilon\delta, \varepsilon > 0$.

In degenerate case $BS = N_{S^*}$ the following arguments hold after adding a small bias: $h' = h + b, b \notin N_{S^*}$.

If instead $BS \neq N_{S^*}$, then BS is a proper closed subspace of N_{S^*} and hence has Gaussian measure zero, so that $P[\langle h', l \rangle = 0] = 0$. By assumption, the boundary ∂C is locally flat; the probability to hit the boundary depends on the visibility cone of the locally flat patch and is positive. With repeated draws of perturbation δ one of the rays $S \pm t\,h'$ meets ∂C and W decreases almost surely.

See Appendix for the complete proof.

Via minimisation of external energy Eq. (1) we obtain symmetry-breaking updates

$$h = A_{\tilde{\phi}}[S] \in N_S = (T_S O)^{\perp}$$

That automatically align with the normals to the discontinuities of a piecewise-constant loss W. Thus, when S lies near a boundary across which W drops, these updates become natural descent directions. The probability that a random $h \in N_S$ escapes a plateau is governed by two factors:

1. *Blind subspace*: a nontrivial subspace $BS \subset N_S$ can leave W unchanged, and gauge displacement h will generically have a component tangential to the desired descent direction (normal to the cell boundary)
2. *Visibility of the flat patch*: if the gauge displacement h is not perfectly aligned with the normal to the cell boundary, corresponding ray $S \pm t\,h$ can miss the boundary, unless the starting point S is sufficiently close to the cell C.

7. Numerical Experiments

The core idea of our method is to solve a variational problem for an auxiliary energy functional, coupled with the actual data, and apply the resulting input transformation to reduce the task's cost metric.

We demonstrate this principle in settings where the target metric is discontinuous and thus cannot be directly optimised via gradient descent. Typically, such tasks rely on a differentiable proxy metric during training.

For example, in image classification and segmentation, the true task metrics—such as accuracy or Intersection over Union (IoU)—are discontinuous and piecewise constant. Therefore, training relies on proxies like the cross-entropy loss, which measures model confidence but not correctness. As a result, gradient-based methods may become stuck in "confident but wrong" local minima.

We evaluate our symmetry-breaking framework on tasks where the cost function is a composition of a black-box neural network and a piecewise constant metric:

$$W = Net \circ Metric.$$

On the Pascal VOC2012 segmentation benchmark, we use a frozen DeepLab-v3–ResNet backbone and optimize mean IoU. Our method yields a statistically significant improvement of +3% in IoU, while the cross-entropy loss remains nearly unchanged—suggesting that symmetry-breaking perturbations cross decision boundaries invisible to standard gradients.

A similar effect is observed in image classification on the CUB-200-2011 dataset. Using a frozen ResNet and top-1 accuracy as the evaluation metric, we achieve a +1 pp gain in accuracy, again with negligible change in cross-entropy.

In scenarios with high model uncertainty, the benefit is even greater. Following Osipov [31], we apply local Lie-group deformations (e.g., time warps, frequency shifts, amplitude modulations) to the log–Mel spectrograms of dysarthric speech. A ResNet is trained to predict the deformation field ϕ, regularised via reconstruction and symmetry energy [32]. At test time, the predicted flow corrects pathological speech, reducing Word Error Rate (WER) by up to 17 points on TORGO [33] and UA-Speech [34]—without using gradients or labels.

Across speech, vision, and classification tasks, our framework consistently reduces non-differentiable cost metrics via symmetry-aware gauge flows, with no gradient access or task supervision.

8. Conclusion

We have developed a variational framework for test-time symmetry breaking in structured signal spaces. Given any cost functional W invariant under a global symmetry group G, we introduce a local gauge-field energy whose minimiser $\hat{\phi}$ defines a symmetry-breaking update, which provably decreases W under mild geometric and variational assumptions—even when W is discontinuous or piecewise constant. Our randomized-ray descent theorem further guarantees positive-probability (and, in certain cases, almost-sure) reduction of such non-differentiable costs.

This methodology applies broadly across variational signal models—from classical Sobolev-space formulations to modern neural network representations and metric-based losses in machine learning. We demonstrate its effectiveness on speech and vision tasks with non-differentiable, discontinuous metrics.

Appendix A: Proof of Lemma 1

Proof. Suppose for contradiction that a constant gauge field $\phi(x) \equiv c \in R^d$ minimizes E_S. Since $S \notin Fix(G)$, the residual

$$R(c, x) = \phi(x) - c^i L_i[S](x) = c^i(L_i[S](x))$$

is not identically zero, so there is some $x_0 \in M$ with $R(c, x_0) \neq 0$.

1. Bump perturbation. Let $\eta \in C_c^\infty(B_r(x_0))$ satisfy $0 \leq \eta \leq 1$, and set

$$\delta\phi(x) = \varepsilon \frac{R(c, x_0)}{\| R(c, x_0) \|} \eta(x),$$

which lies in $H^s(M)$ for $s > \frac{dim(M)}{2} + 1$. Write

$$b_i = \int_{B_r(x_0)} \eta(x)\,\langle R(c,x), L_i[S](x)\rangle\,dx, \quad b = (b_1, \dots, b_d).$$

2. First variation. Expanding $E_S[\phi + \delta\phi]$ to first order in ε gives

$$\delta E = \varepsilon[-2\langle c, b\rangle + 4\beta(|c|^2 - v^2)\ \| c \|\,] + O(\varepsilon^2).$$

Since $\phi \equiv c$ is assumed minimal, the bracket must vanish for every choice of bump η.

3. Contradiction via tangential perturbation. Writing $b = b_\parallel + b_\perp$ with $b_\parallel \parallel c$ and $b_\perp \perp c$, the bracket condition forces $\langle c, b_\perp\rangle = 0$. But $b_\parallel$ is already parallel to c, so one can choose the bump direction proportional to $b_\perp$ itself. Then

$$\delta E = -2\varepsilon\ \| b_\perp \|^2 < 0,$$

for small $\varepsilon > 0$, contradicting minimality.

Hence no constant field $\phi \equiv c$ can be a global minimiser.

Appendix B: Proof of Lemma 2

Proof. Let $\tilde{\phi} = arg\ min\ E_S$ be the minimizer of the energy functional. In what follows we assume that the double well potential term keeps the minimiser close to the vacuum value: $\delta := \|\ |\tilde{\phi}|^2 - v^2\ \|^2 \ll 1$ and that the orbit tangent is not degenerate: the Gram matrix $G_{i\,j} := \langle L_i[S],\ L_j[S]\rangle$ is well-conditioned on its range.

1. Energy variation

The first variation of the energy functional with respect to ϕ_i under Neumann boundary conditions is

$$\delta_i E_S[\phi] = 2t\,\langle w, L_i[S_\phi]\rangle + 4\beta \int_M (|\phi|^2 - v^2)\phi_i + O(t^2).$$

Write the displacement as $w = S - S_\phi = -t\ a + O(t^2)$, where $a := \tilde{\phi}^i\ L_i[S]$.

Since $L_i[S_\phi] = L_i[S] + O(t)$, one has $\langle a,\ L_i[S]\rangle = O(\beta\delta) + O(t)$ for $i = 1, \dots, d$. Thus the tangential component $a_T := P_V(a)$ satisfies $\|\ a_T\ \|^2 = D^T\ G^+ D\ = O((\beta\delta + t))^2$, where $D_i := \langle a,\ L_i[S]\rangle$.

2. The gradient of the orbit distance

Define the orbit distance between the original signal and the warped signal as follows:

$$d_G(S) = inf_{g\in G}\ \| S - g \cdot exp(tA_{\tilde{\phi}}) \cdot S \|^2.$$

We will assume that the minimising group element $\tilde{g}$ exists, is unique in a neighborhood of the identity, and $\tilde{g}\ = O(t)$. Denote the residual $R := S - \tilde{g} \cdot exp(tA_{\tilde{\phi}}) \cdot S$. Since the G action is isometric,

$$\nabla d_G(S) = 2(I - \tilde{g}\ exp(tA_{\tilde{\phi}}))\ R.$$

Using the fact that $\tilde{g}\ = O(t)$ and $exp(tA_{\tilde{\phi}}) = I + O(t)$, we obtain $\nabla d_G(S) = 2R\ +\ O(t\ \| R \|)$.

Since $R = P_{V\perp}\ (w)\ + O(t\ \| w \|)$, we get $\nabla d_G(S) = 2\ P_{V\perp}\ (w) + O(t\ \| w \|)$.

3. Inner-Product estimate

Decompose $w = w_T\ + w_N$ with $w_T = P_V\ (w)$, $w_N = P_{V\perp}\ (w)$.

Then $\| w_T \| = t\ \| a_T \|\ + O(t^2)\ =\ O(\beta\delta t\ + t^2)$, whereas $\| w \| = t\ \| a \|\ + O(t^2)$.

Hence, $\frac{\|w_T\|}{\|w\|} \le C_1\delta + C_2 t$ (displacement w is nearly aligned with ∇d_G), which proves Lemma 2.

Appendix C: Proof of Theorem 1

Proof. **Gauge displacement.** According to Lemma 2, the gauge displacement h can be written as $h = -\kappa\xi + r$, with $\xi \in N_S$ and $r = O(\delta + t)$. Since the group G action is smooth and isometric, the map $X \mapsto A_X$ is continuous, and so is $X \mapsto N_X$ (gap topology), hence the distance between N_S and N_{S^*} (gap topology)will be arbitrarily small as S approaches S^*.

Local flatness. Assume there exists a ball B with $r_b > 0$ such that $\partial C \cap B(S*, r_b) = (S* + l\perp) \cap B(S*, r_b)$ (local flat patch). Let u be unit with decomposition $u = u_N\, n + u_T$. The ray $S + tu$ meets the flat patch of ∂C whenever $u_N < 0$ and $\| u_T \| \leq r_b \| u_N \| / \| l \|$.

(a) Deterministic case. Recall the definition of the blind subspace: $BS := L \cap N_{S^*} = \{ w \in N_{S^*} : \langle w, l \rangle = 0\}$.

Suppose $BS = \{0\}$ and let $u = h/\| h \|$. In that case, $\langle h, l \rangle < 0$ and $\| u_T \| \leq r_b \| u_N \| / \| l \|$ holds automatically once S is close enough to S^*, and by double-sign test, one ray enters the cell C.

(b) Probabilistic case. Assume $dim\ BS > 0, BS \subsetneq N_{S*}$. Let $\delta \sim N(0, Cov)$ be a centred Gaussian on N_{S*} with trace-class, strictly positive covariance and full support. For $\varepsilon > 0$ set $h' := h + \varepsilon\delta, u := h'/\| h' \|$.

The map $v \mapsto \langle v, l \rangle$ is a nonzero continuous linear functional on N_{S*} because $BS \subsetneq N_{S*}$ (and hence $l \notin A_{S^*}$).

Since δ has a nondegenerate Gaussian law on N_{S*}, the scalar $\langle \delta, l \rangle$ is a (one–dimensional) nondegenerate Gaussian; therefore $P[\langle \delta, \ell \rangle = 0] = 0$. As h is deterministic, $\langle h', l \rangle = \langle h, l \rangle + \varepsilon \langle \delta, l \rangle$ also has a continuous density; hence

$$P[\langle h', \ell \rangle = 0] = 0.$$

Thus, with probability 1, exactly one of the two signs $\pm h'$ has a negative normal component.

Choose an orthonormal basis of N_{S*},

$$e_0 := n, \quad e_1, e_2, \ldots \in L \cap N_{S*}.$$

Write

$$X := \langle h', e_0 \rangle, \quad Y_j := \langle h', e_j \rangle\ (j \geq 1), \quad {Z_K}^2 := \sum_{j > K} {Y_j}^2$$

Diagonalise covariance in this basis; then the coordinates $(X, Y_1, \ldots, Y_K)$ are jointly nondegenerate Gaussian in R^K, independent of the tail $Y_j \ , j > K$. The random variable Z_K has a continuous density on $[0, \infty)$.
Fix $\alpha > 0, \eta > 0$, $K \in N$ so that

$$P[\, X \leq -\alpha, |\, Y_j \,| \leq \eta \,, Z_K \leq \eta \,] > 0 \qquad (1 \leq j \leq K)$$

Then, with a proper choice of constants, a single trial succeeds with probability at least p>0. Independent repetitions succeed almost surely because the failure probability after N trials is $(1-p)^N \to 0$.

Remark. In degenerate case $l \in A_{S^*}$ one has $BS = N_{S*}$, so that $\langle v, l \rangle = 0$ for all $v \in N_{S*}$. No perturbation confined to N_{S*} can cross the boundary. A remedy is to add a tiny bias $b \notin N_{S*}$ (or to enlarge G) and repeat the argument with $h + b$.

Conflict of Interest

The author declares no conflict of interest.

References

[1] Fonseca, I., & Leoni, G. (2007). *Modern Methods in the Calculus of Variations: Lp Spaces*. Springer.

[2] Evans, L. C. (2010). *Partial Differential Equations* (2nd ed.). American Mathematical Society.

[3] Yang, C. N., & Mills, R. L. (1954). Conservation of isotopic spin and isotopic gauge invariance. *Physical Review, 96*, 191–195. https://doi.org/10.1103/PhysRev.96.191

[4] Jaffe, A., & Taubes, C. H. (1980). *Vortices and Monopoles: Structure of Static Gauge Theories*. Boston: Birkhäuser.

[5] Bethuel, F., Brézis, H., & Hélein, F. (1994). *Ginzburg–Landau Vortices*. Boston: Birkhäuser.

[6] Cohen, T., & Welling, M. (2016). Group equivariant convolutional networks. *Proceedings of the International Conference on Machine Learning (ICML)*.

[7] Cohen, T., Geiger, M., Köhler, J., & Welling, M. (2019). Gauge equivariant convolutional networks and the icosahedral CNN. *Proceedings of the International Conference on Learning Representations (ICLR)*.

[8] Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., & Vandergheynst, P. (2021). Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *Nature*, 590, 197–205. https://doi.org/10.1038/s41586-021-03242-1

[9] Simard, P. Y., LeCun, Y., Denker, J. S., & Victorri, G. B. (1992). Tangent propagation: A formalism for learning invariant maps. *Neural Computation, 4(4)*, 650–669. https://doi.org/10.1162/neco.1992.4.4.650

[10] Jaderberg, M., Simonyan, K., Zisserman, A., & Kavukcuoglu, K. (2015). Spatial transformer networks. In *Advances in Neural Information Processing Systems (NeurIPS)* (Vol. 28). Curran Associates.

[11] Bamler, R., & Mandt, S. (2018). Improving optimization for models with continuous symmetries via Lie algebras. *Proceedings of the International Conference on Machine Learning (ICML)* (pp. 411–419).

[12] Tanaka, H., Kunin, D., & Ganguli, S. (2021). Noether's learning dynamics: Role of symmetries in generalization and optimization. *Proceedings of the International Conference on Learning Representations (ICLR)*.

[13] Elesedy, H. (2025). Symmetry and generalisation in machine learning. arXiv preprint, arXiv:2501.03858, 2025.

[14] Yang, Y. (2010). Gauge theory and variational principles in continuum mechanics. *Mathematics and Mechanics of Solids, 15(3),* 307–321. https://doi.org/10.1177/1081286509354348

[15] Sochen, N., Kimmel, R., & Malladi, R. (1998). A general framework for low level vision. *IEEE Transactions on Image Processing, 7(3)*, 310–318. https://doi.org/10.1109/83.661190

[16] Trouvé, A. (1998). Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision, 28(3)*, 213–221. https://doi.org/10.1023/A:1008005210584

[17] Younes, L. (2010). *Shapes and Diffeomorphisms*. Springer.

[18] Arnaudon, M., & Cruzeiro, A. B. (2016). Geometry and symmetry of optimal transport. *Journal of the European Mathematical Society, 18(8)*, 1855–1881. https://doi.org/10.4171/JEMS/640

[19] Sun, Q., Zhang, D., Liu, C., Shi, H., & Chua, T.-S. (2020). Test-time training for out-of-distribution generalization. *Proceedings of the International Conference on Machine Learning (ICML)* (pp. 8324–8333).

[20] Yang, J., Li, J., Zhang, P., & Wang, L. (2021). Test-time adaptation for generalization. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 1–10). IEEE.

[21] Xie, C., Wang, Y., Zhang, Z., Ren, W., & Yuille, A. (2018). Mitigating adversarial effects through randomization. *Proceedings of the International Conference on Learning Representations (ICLR)*.

[22] Goodfellow, I., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. *Proceedings of the International Conference on Learning Representations (ICLR)*.

[23] Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2019). Deep learning models resistant to adversarial attacks. *Proceedings of the International Conference on Learning Representations (ICLR)*.

[24] Ebin, D. G., & Marsden, J. (1970). Groups of diffeomorphisms and the motion of an incompressible fluid. *Annals of Mathematics, 92(1)*, 102–163. https://doi.org/10.2307/1970699

[25] Kriegl, A., & Michor, P. W. (1997). *The Convenient Setting of Global Analysis*. American Mathematical Society.

[26] Adams, R. A., & Fournier, J. J. F. (2003). *Sobolev Spaces* (2nd ed.). Academic Press.

[27] Hebey, E. (1999). *Nonlinear Analysis on Manifolds: Sobolev Spaces and Inequalities*. American Mathematical Society.

[28] Triebel, H. (1983). *Theory of Function Spaces*. Birkhäuser.

[29] Bruveris, M., & Vialard, F.-X. (2016). On completeness of groups of diffeomorphisms. *Journal of Mathematical Analysis and Applications, 433(1)*, 1–21. https://doi.org/10.1016/j.jmaa.2015.07.064

[30] Bauer, M., Bruveris, M., & Michor, P. W. (2014). Overview of the geometries of shape spaces and diffeomorphism groups. *Journal of Mathematical Imaging and Vision, 50(1)*, 60–97. https://doi.org/10.1007/s10851-014-0507-9

[31] Osipov, M. (2025). Dysarthria normalization via local Lie group transformations for robust ASR. arXiv preprint, arXiv:2504.12279, 2025.

[32] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 770–778).

[33] Rudzicz, F., Namasivayam, A., & Wolff, T. (2012). The TORGO database of acoustic and articulatory speech from speakers with dysarthria. *Lang Resources & Evaluation, 46*, 523–541.

[34] Kim, H.-Y., Hasegawa-Johnson, M., Perlman, M., Gunderson, J., Huang, S.-E., Watkin, K., & Wealthy, E. (2008). Dysarthric speech databases for universal access research. *Interspeech, 2008*, 1741–1744.