JOURNAL OF
SYMPLECTIC GEOMETRY
Volume 20, Number 5, 1135–1158, 2022

Epsilon-non-squeezing and C^0-rigidity of epsilon-symplectic embeddings

Stefan Müller

An embedding $\varphi\colon (M_1, \omega_1) \to (M_2, \omega_2)$ (of symplectic manifolds of the same dimension) is called ϵ-symplectic if the difference $\varphi^* \omega_2 - \omega_1$ is ϵ-small with respect to a fixed Riemannian metric on M_1. We prove that if a sequence of ϵ-symplectic embeddings converges uniformly (on compact subsets) to another embedding, then the limit is E-symplectic, where the number E depends only on ϵ, and $E(\epsilon) \to 0$ as $\epsilon \to 0$. This generalizes C^0-rigidity of symplectic embeddings, and answers a question in topological quantum computing by Michael Freedman.

As in the symplectic case, this rigidity theorem can be deduced from the existence and properties of symplectic capacities. An ϵ-symplectic embedding preserves capacity up to an ϵ-small error, and linear ϵ-symplectic maps can be characterized by the property that they preserve the symplectic spectrum of ellipsoids (centered at the origin) up to an error that is ϵ-small. We also sketch an alternative proof using the shape invariant, which gives rise to an analogous characterization and rigidity theorem for ϵ-contact embeddings.

1. Introduction and main results

In this paper, we consider smooth manifolds M equipped with a symplectic structure ω and a Riemannian metric g. We do not necessarily assume that the metric is compatible with the symplectic structure, or that the induced volume forms coincide (up to a constant multiple), though some of the estimates in this article are more explicit in those cases, in particular in dimension two.

Definition 1.1 (Epsilon-symplectic and epsilon-anti-symplectic). Let (M_1, ω_1) and (M_2, ω_2) be two symplectic manifolds of the same dimension, g be a Riemannian metric on M_1, and $\epsilon \geq 0$. An embedding $\varphi\colon M_1 \to$

M_2 is called ϵ-symplectic if $\|\varphi^*\omega_2 - \omega_1\|_2 \leq \epsilon$, and ϵ-anti-symplectic if $\|\varphi^*\omega_2 + \omega_1\|_2 \leq \epsilon$.

See section 2 for the definition of the norm $\| \cdot \|_2$ and a number of general related results. A goal of this paper is to prove the following rigidity theorem.

Theorem 1.2. *Let (M_1, ω_1) and (M_2, ω_2) be two symplectic manifolds of the same dimension, and g be a Riemannian metric on M_1. Then there are constants $\delta = \delta(\omega_1, g) > 0$ and $E = E(\omega_1, g, \epsilon) \geq 0$ with $E \to 0^+$ as $\epsilon \to 0^+$ so that, if $\epsilon < \delta$ and $\varphi_k \colon M_1 \to M_2$ is a sequence of ϵ-symplectic embeddings that converges uniformly (on compact subsets) to an embedding $\varphi \colon M_1 \to M_2$, then φ is E-symplectic.*

In case (M_1, ω_1) and (M_2, ω_2) are both subsets of $(\mathbb{R}^{2n}, \omega_0)$ with its standard symplectic structure and standard flat metric, an explicit lower bound for δ and explicit upper bound for E can be derived from the proof given in this paper.

Corollary 1.3. *Let (M_1, ω_1) and (M_2, ω_2) be two symplectic manifolds of the same dimension, and $\epsilon_k \geq 0$ be a sequence of non-negative numbers so that $\epsilon_k \to 0^+$ as $k \to \infty$. Suppose that $\varphi_k \colon M_1 \to M_2$ is a sequence of embeddings that converges uniformly (on compact subsets) to another embedding $\varphi \colon M_1 \to M_2$, and that each φ_k is ϵ_k-symplectic. Then the limit φ is a symplectic embedding.*

The choice of Riemannian metric on M_1 is not relevant for the corollary. See Remark 4.9. Analogous to the symplectic case, we show that an embedding is ϵ-symplectic or ϵ-anti-symplectic if and only if it preserves the capacity of ellipsoids up to an ϵ-small error. Most of the paper is devoted to establishing its linear version on $\mathbb{R}^{2n}$ with its standard symplectic structure and Riemannian metric.

Proposition 1.4. *Let $0 \leq \epsilon < 1/\sqrt{2}$, and $\epsilon' = \sqrt{2}\,\epsilon$. Then an ϵ-symplectic linear map $\Phi \colon \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ is ϵ'-non-squeezing and ϵ'-non-expanding.*

The constant $1/\sqrt{2}$ is not optimal; see section 2 for details. By Remark 5.5 below, there is no form of non-squeezing for ϵ-symplectic embeddings with $\epsilon \geq 1$.

Theorem 1.5. *Suppose a linear map $\Phi \colon \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ has the linear ϵ-non-squeezing and linear ϵ-non-expanding property. Then for $\epsilon \geq 0$ sufficiently* *small, Φ is either ϵ'-symplectic or ϵ'-anti-symplectic, where $\epsilon' = K(\epsilon) \to 0^+$ as $\epsilon \to 0^+$.*

See sections 4 and 5 for details. Symplectic capacities are discussed in section 6. A geometric expression of ϵ-symplectic rigidity is the following generalization of Gromov's non-squeezing theorem. Consider again $\mathbb{R}^{2n}$ with its standard symplectic structure ω_0. Denote by $B_r^{2n} \subset \mathbb{R}^{2n}$ the (closed) ball of radius $r > 0$ (centered at the origin), and by $Z_R^{2n} = B_R^2 \times \mathbb{R}^{2n-2}$ the (symplectic) cylinder of radius $R > 0$.

Proposition 1.6 (Epsilon-non-squeezing). *If there is an ϵ-symplectic embedding of B_r^{2n} into Z_R^{2n}, with $0 \leq \epsilon < 1/\sqrt{2}$, then $r \leq (1 - \sqrt{2}\,\epsilon)^{-\sqrt{2n}} R$.*

Recall that Gromov's non-squeezing theorem (the case $\epsilon = 0$) can be considered as a geometric expression of the uncertainty principle [6, page 458]. Given a point $(x_1, y_1, \ldots, x_n, y_n)$ in $\mathbb{R}^{2n} = T^*\mathbb{R}^n$, think of x_j as the j-th position coordinate and y_j as the j-th momentum coordinate of some Hamiltonian system. If the state of the system is measured to lie somewhere in a subset $U \subset \mathbb{R}^{2n}$ that is (or contains) a ball of radius r, then the range of uncertainty (to the extent of our knowledge) of the values of the conjugate pair (x_j, y_j) is the area πr^2. Proposition 1.6 then means that if the system is transformed by an ϵ-symplectic diffeomorphism, this range of uncertainty can be decreased by a factor of at most $(1 - \sqrt{2}\,\epsilon)^{2n}$.

The results of this paper are of interest in symplectic integrator methods and topological quantum computing, where computations can be performed up to any prescribed level of accuracy only. The question by Michael Freedman [4] was the starting point of this paper. Corollary 1.3 is also relevant in C^0-symplectic topology.

For most of the paper, we assume that M is compact or a relatively compact subset U of $\mathbb{R}^{2n}$. In the latter case, we also assume that there exists a Riemannian metric $\overline{g}$ defined on a neighborhood of $\overline{U}$ such that $\overline{g}_{|_U} = g$. In particular, all of the supremums considered below are in fact maximums, and in particular, are finite. See the (first paragraph of the) proof of Theorem 1.2 in section 6 for the case of non-compact manifolds. An alternate argument using the shape invariant, and ϵ-contact embeddings, are discussed in the final section 7.

2. Norms of vector fields and differential forms

The Riemannian metric g induces a norm on each tangent space T_xM given by $\|v\|_2 = \sqrt{g(v,v)}$ for $v \in T_xM$. The norm of a vector field X on M is then defined by $\|X\|_2 = \sup_{x\in M} \|X(x)\|_2$. Let $\star$ denote the Hodge star of the metric g. Then for a k-covector $v^* \in \Lambda^k(T_xM)$, let $\|v^*\|_2 = \sqrt{\star(v^* \wedge \star v^*)}$, and for a differential form β, define $\|\beta\|_2 = \sup_{x\in M} \|\beta(x)\|_2$. We can also define the comass norms

$$\|v^*\|_C = \sup\{v^*(v_1,\dots,v_k) \mid \|v_1\|_2 = \dots = \|v_k\|_2 = 1\}$$

and $\|\beta\|_C = \sup_{x\in M} \|\beta(x)\|_C$. The norms $\|\cdot\|_2$ and $\|\cdot\|_C$ are in fact equivalent. We sketch a proof to the degree necessary for our purposes. See [3, Chapter 1], for instance, for details.

Lemma 2.1. *Let $m = \dim M$. Then $\|\cdot\|_C \le \|\cdot\|_2 \le \sqrt{\binom{m}{k}}\|\cdot\|_C$ for any k-covector and any k-form, and $\|v^*\|_C = \|v^*\|_2$ if and only if the k-covector v^* is simple.*

Sketch of proof. Note that it suffices to prove the lemma for covectors. The natural isomorphism $\gamma\colon T_xM \to T_x^*M$ given by $\gamma(v) = g(v,\cdot)$ extends to an isomorphism $\gamma\colon \Lambda_k(T_xM) \to \Lambda^k(T_xM)$ for each k, and thus the metric g extends to a metric on the space of k-vectors given by $(v,w) \mapsto \gamma(v)(w)$. The induced norm $\|\cdot\|_2$ on k-vectors is dual to the norm $\|\cdot\|_2$ for k-covectors. In particular,

$$\|v^*\|_2 = \sup\{v^*(v) \mid v \in \Lambda_k(T_xM) \text{ with } \|v\|_2 = 1\},$$

whereas

$$\|v^*\|_C = \sup\{v^*(v_1 \wedge \ldots \wedge v_k) \mid \|v_1 \wedge \ldots \wedge v_k\|_2 = 1\},$$

i.e., the latter supremum is over all simple unit k-vectors only, where a k-(co-)vector is called simple if it is the (alternating) product of 1-(co-)vectors. That proves the first inequality, and the claim that the two norms coincide on simple k-covectors.

To prove the second inequality, choose an orthonormal basis $e_1,\dots,e_m$ of T_xM, with dual orthonormal basis $\alpha_1,\dots,\alpha_m$. Let v^* be a k-covector, and write

$$v^* = \sum_\sigma f_\sigma\, \alpha_{\sigma(1)} \wedge \ldots \wedge \alpha_{\sigma(k)}, \tag{1}$$

where the sum is over all strictly increasing functions $\sigma\colon \{1, \ldots, k\} \to \{1, \ldots, m\}$. Let $d = \binom{m}{k}$, and choose some order on the set (with d elements) of such functions. Denote by f_{v^*} the vector $(f_{\sigma_1}, \ldots, f_{\sigma_d})$ in $\mathbb{R}^d$, equipped with the standard metric $g_0 = \langle \cdot, \cdot \rangle$. Then it follows immediately from the definitions that $\|v^*\|_2 = \|f_{v^*}\|_2$.

Let $v_j = \sum_{i=1}^m \lambda_{ij} e_i$, $1 \le j \le k$, be unit vectors in T_xM, and consider the $(m \times k)$-matrix $\Lambda = [\lambda_{ij}]_{1\le i\le m, 1\le j\le k}$. Then

$$v^*(v_1, \ldots, v_k) = \langle (f_{\sigma_1}, \ldots, f_{\sigma_d}), (\det(M_{\sigma_1}), \ldots, \det(M_{\sigma_d})) \rangle\,,$$

where M_σ is the $(k \times k)$-minor obtained from Λ by deleting all but the rows in the image of the function σ. (Geometrically, the minor M_σ represents the linear transformation $\Lambda \circ \Pi_\sigma\colon \mathbb{R}^k \to \mathbb{R}^k$, where $\Pi_\sigma\colon \mathbb{R}^m \to \mathbb{R}^k$ denotes the projection to the components that belong to the image of σ. In particular, the absolute value of its determinant can be interpreted as the hyper-volume of the image of a k-dimensional face of the unit cube.) Let $N \le d$ be the number of non-zero terms in (1). If we choose $\lambda_{ij} = \delta_{i\sigma(j)}$ for some σ, then $v^*(v_1, \ldots, v_k) = f_\sigma$. Therefore $\|v^*\|_C \ge \max_\sigma |f_\sigma|$, and in particular, $\|v^*\|_2 \le \sqrt{N}\, \|v^*\|_C \le \sqrt{d}\, \|v^*\|_C$. □

We point out the following immediate consequence of the preceding proof.

Lemma 2.2. *For every k-covector v^* and every orthonormal basis B of T_xM, there exist vectors $v_1, \ldots, v_k \in B$ such that*

$$\|v^*\|_C \ge |v^*(v_1, v_2, \ldots, v_k)| \ge \binom{m}{k}^{-1/2} \|v^*\|_2.$$

Remark 2.3. The inequalities in Lemma 2.1 are not necessarily sharp for all pairs of positive integers m and k. Below we find optimal constants in the two cases $k = 1$ or 2 of interest in this paper.

Lemma 2.4. $\|\cdot\|_C = \|\cdot\|_2$ *for any k-covector and any k-form if $k = 1$ or $m - 1$.*

Proof. The proof is an immediate consequence of Lemma 2.1, since any 1-covector and $(m-1)$-covector is simple. We also give a direct argument. It again suffices to prove the lemma for covectors.

If $k=1$, write $v^* = \gamma(v)$ for a (unique) vector $v \in T_xM$. Then by definition, $\|v^*\|_C \geq \gamma(v)(v/\|v\|_2) = \|v\|_2 = \|v^*\|_2$. Conversely, for $u \in T_xM$ a unit vector, $|\gamma(v)(u)| = |g(v,u)| \leq \|v\|_2$ by the Cauchy-Schwarz inequality, so $\|v^*\|_C \leq \|v\|_2$.

For $k = m-1$ the proof is similar, once we observe that $(\det(M_{\sigma_1}), \dots, \det(M_{\sigma_m}))$ is the cross product of $v_1, \dots, v_{m-1}$, and $v^*(v_1, \dots, v_{m-1})$ is the determinant of the matrix with columns the vectors $f_{v^*}, v_1, \dots, v_{m-1}$. (Geometrically, the latter is, up to sign, the volume of the parallelepiped spanned by these vectors.) $\square$

A key ingredient in our argument in section 5 is the following lemma. We state and prove it in this section for its corollary.

Lemma 2.5. *Let ω be a two-form on an inner product space V. Then there exists an orthonormal basis B for V, $S = \{u_1, \dots, u_n, v_1, \dots, v_n\} \subset B$, $2n \leq \dim V$, and positive numbers $0 < \lambda_1 \leq \ldots \leq \lambda_n$, such that $\omega(u_j, v_k) = \lambda_j^2 \delta_{jk}$ and $\omega(u_j, u_k) = \omega(v_j, v_k) = 0$ for $1 \leq j, k \leq n$, and ω vanishes on $B \setminus S$. In other words, ω can be written in the form $\omega = \sum_{j=1}^n \lambda_j^2 \, \alpha_j \wedge \beta_j$ with one-forms α_j and β_j dual to the elements of S. Moreover, ω is non-degenerate if and only if $2n = \dim V$.*

Corollary 2.6. $\|\cdot\|_C \leq \|\cdot\|_2 \leq (\lfloor \frac{m}{2} \rfloor)^{1/2} \|\cdot\|_C$ *for any two-covector and two-form on an m-dimensional Riemannian manifold M, and these inequalities are sharp.*

Proof of Corollary. See the last three sentences of the proof of Lemma 2.1. To verify that the second inequality is also sharp, suppose that ω is non-degenerate, and that J is an almost complex structure that is compatible with g so that $\omega = g(J\cdot, \cdot)$. Then $\|\omega\|_C = \|g(J\cdot,\cdot)\|_C = 1$. On the other hand, $\|\omega_0\|_2 = \sqrt{n}$ for the standard symplectic structure ω_0 and standard Riemannian metric on $\mathbb{R}^{2n}$. $\square$

Proof of Lemma 2.5. The argument here is taken from [3, Section 1.7.3]. Let A be the skew-symmetric matrix so that $\omega(v,w) = g(Av, w)$. Decompose V into a direct sum of mutually orthogonal and A-invariant subspaces $W_1, \dots, W_s$ with $\dim W_j \leq 2$ (which exists since V has a basis of eigenvectors of the symmetric matrix A^2), and observe that $\omega(v,w) = 0$ whenever $v \in W_j$ and $w \in W_k$ with $j \neq k$. Choose an orthonormal basis u_j, v_j for each W_j that is two-dimensional, and extend to any orthonormal basis for V if $\operatorname{rank}(A) = 2n < \dim V$. Note that $g(Au_i, u_i) = \omega(u_i, u_i) = 0$, so we may choose v_i parallel to Au_i, and then $\omega(u_i, v_i) = \|Au_i\|$. Reorder the W_j if necessary. $\square$

Remark 2.7. Alternatively, one may argue as in [6, Lemma 2.4.5] in case ω is non-degenerate. Here one observes that the matrix $iA\colon \mathbb{C}^{2n} \to \mathbb{C}^{2n}$ is Hermitian, and obtains the same basis vectors u_j and v_j (up to rescaling) as real and imaginary parts of the eigenvectors corresponding to the eigenvalues $i\lambda_j^2$ of A. Note that the signs in the proof of [6, Lemma 2.4.5] are different because A is defined there with the opposite sign choice compared to the proof given above. The argument easily extends to degenerate ω.

For v a vector, denote by ι_v the interior multiplication (or contraction) of a co-vector v^* (of degree $k \geq 1$) by v, i.e., $\iota_v v^* = v^*(v, \cdot, \ldots, \cdot)$, and similarly, for a vector field X, write ι_X for interior multiplication of a differential form by X.

Lemma 2.8. *$\|\iota_v v^*\|_C \leq \|v\|_2 \|v^*\|_C$ and $\|\iota_v v^*\|_2 \leq \sqrt{k}\,\|v\|_2 \|v^*\|_2$ for a k-covector v^*, and thus $\|\iota_X \beta\|_C \leq \|X\|_2 \|\beta\|_C$ and $\|\iota_X \beta\|_2 \leq \sqrt{k}\,\|X\|_2 \|\beta\|_2$ for a k-form β.*

Proof. It is again enough to prove the lemma for covectors. For the norm $\|\cdot\|_C$, the lemma follows immediately from the definition by writing $\iota_v v^* = \|v\|_2 \cdot \iota_{(v/\|v\|_2)} v^*$. For the norm $\|\cdot\|_2$, the claim follows from the identity $\|v^*\|_2 = \|f_{v^*}\|_2$ established in the course of the proof of Lemma 2.1 and the Cauchy-Schwarz inequality. $\square$

Remark 2.9. The inequalities in the previous lemma are sharp for the comass norm, but not for the norm $\|\cdot\|_2$ when $1 < k < \dim M$.

3. Epsilon-symplectic embeddings

Let (M_1, ω_1) and (M_2, ω_2) be symplectic manifolds of the same dimension, g be a Riemannian metric on M_1, and $\epsilon \geq 0$. Recall that an embedding $\varphi\colon M_1 \to M_2$ is called ϵ-symplectic if $\|\varphi^*\omega_2 - \omega_1\|_2 \leq \epsilon$.

Remark 3.1. Let g_2 be a Riemannian metric on M_2, and $\varphi\colon M_1 \to M_2$ be an ϵ_1-symplectic diffeomorphism. Let $\beta = \varphi^*\omega_2 - \omega_1$. Then $(\varphi^{-1})^*\omega_1 - \omega_2 = -(\varphi^{-1})^*\beta$. Thus if $\epsilon_1 > 0$ and $n > 1$, the inverse diffeomorphism φ^{-1} is not necessarily ϵ_2-symplectic for some $\epsilon_2 \geq 0$ that depends only on ϵ_1 and the metric g_2. If Φ is a linear map, a similar remark holds for its transpose Φ^T. See the following example.

Example 3.2. Let $n \geq 2$, and $\omega_0 = \sum_{j=1}^n dx_j \wedge dy_j$ be the standard symplectic structure on $\mathbb{R}^{2n} = \mathbb{R}^4 \times \mathbb{R}^{2n-4}$. Let $K \neq 0$. Consider the isomorphism Ψ of $\mathbb{R}^4$ defined by $\Psi(x_1, y_1, x_2, y_2) = (x_1, y_1 + \epsilon x_2, -K^{-1}x_2, -Ky_2)$, and let $\Phi = \Psi \times \mathrm{id}$. Then $\Phi^*\omega_0 - \omega_0 = \epsilon\, dx_1 \wedge dx_2$. On the other hand, $(\Phi^{-1})^*\omega_0 - \omega_0 = (K\epsilon)\, dx_1 \wedge dx_2$ and $(\Phi^T)^*\omega_0 - \omega_0 = (-K\epsilon)\, dy_1 \wedge dy_2$.

The preceding remark and example mean that our proof of Theorem 1.5 cannot follow too closely the standard proof in the symplectic case given, for instance, in [6, Section 2.4]. We include the following result solely for the sake of completeness.

Lemma 3.3. *Let $\varphi\colon M_1 \to M_2$ be an embedding, and suppose that ψ_1 and ψ_2 are symplectic diffeomorphisms of M_1 and M_2, respectively. Then there exists a constant $C(\psi_1)$ so that $\|(\psi_2 \circ \varphi \circ \psi_1)^*\omega_2 - \omega_1\| \leq C(\psi_1)\|\varphi^*\omega_2 - \omega_1\|$ for both the norm $\|\cdot\|_2$ and the comass norm $\|\cdot\|_C$. In fact, we may choose $C(\psi_1) = \|(\psi_1)_*\|^2$, where $\|\psi_*\| = \sup_{x\in M}\|d\psi(x)\|$, and $\|\Psi\| = \max\{\|\Psi v\|_2 \mid \|v\|_2 = 1\}$.*

Proof. The lemma follows from the identity $(\psi_2 \circ \varphi \circ \psi_1)^*\omega_2 - \omega_1 = \psi_1^*(\varphi^*\omega_2 - \omega_1)$. See [3, Section 1.7.6] for the estimate $\|\psi_1^*(\varphi^*\omega_2 - \omega_1)\|_2 \leq C(\psi_1)\|\varphi^*\omega_2 - \omega_1\|_2$. (It is sufficient to prove this for two-covectors or a dual inequality for two-vectors.) The analogous inequality is obvious for the comass norm. □

We will use the following obvious remark in our argument in section 5.

Remark 3.4. If ψ is an (anti-)symplectic diffeomorphism of (M_2, ω_2), then an embedding $\varphi\colon M_1 \to M_2$ is ϵ-(anti-)symplectic if and only if the composition $\psi \circ \varphi$ is ϵ-symplectic.

4. Epsilon-symplectic embeddings into Euclidean space

In this section, we consider the case $M = \mathbb{R}^m$ with its standard Riemannian metric $g_0 = \langle\cdot,\cdot\rangle$. If $m = 2n$, we also equip $\mathbb{R}^{2n}$ with its standard symplectic structure $\omega_0 = \sum_{j=1}^n dx_j \wedge dy_j$. Recall that $\omega_0 = g_0(J_0\cdot,\cdot)$, where J_0 is the standard (almost) complex structure on $\mathbb{R}^{2n}$.

Let $U \subset \mathbb{R}^m$ be an open subset that is star-shaped with respect to the origin. The following lemma is an immediate consequence of (the proof of) the Poincaré Lemma [10, 4.18]. Let $\Omega^k = \Omega^k(U) = \Omega^k(U, \mathbb{R})$ be the space of (differential) k-forms on U, and write as usual $d\colon \Omega^{k-1} \to \Omega^k$ for the differential. The cases of greatest interest to us are the open ball $B_r^m \subset \mathbb{R}^m$ of radius $r > 0$ (centered at the origin) and ellipsoids (centered at the origin), but we also have in mind the open polydisk $P(r_1, \ldots, r_n) = B_{r_1}^2 \times \cdots \times B_{r_n}^2 \subset \mathbb{R}^{2n}$.

Lemma 4.1 (Quantitative Poincaré Lemma). *For each $k \geq 1$, there is a bounded and $\mathbb{R}$-linear (and hence continuous) transformation $h_k\colon \Omega^k \to \Omega^{k-1}$ such that $h_{k+1} \circ d + d \circ h_k = id$, and in particular, the restriction of h_k to the space of closed k-forms is a right inverse to the differential d. In fact,*

$$\|h_k(\beta)(x)\|_2 \leq \frac{\|x\|_2}{k-1} \sqrt{k \binom{m}{k-1}} \max_{0 \leq t \leq 1} \|\beta(tx)\|_2 \leq \frac{s}{k-1} \sqrt{k \binom{m}{k-1}} \|\beta\|_2$$

for $k > 1$, and $\|h_k(\beta)(x)\|_2 \leq \|x\|_2 \sqrt{m} \max_{0 \leq t \leq 1} \|\beta(tx)\|_2 \leq s\sqrt{m}\|\beta\|_2$ if $k = 1$, where $s = \sup_{x \in U} \|x\|_2$.

Proof. Let $h_k = \alpha_{k-1} \circ \iota_X = \iota_X \circ \alpha_k$, where X denotes the radial vector field $\sum_{i=1}^m x_i \cdot (\partial/\partial x_i)$, and where for each $k \geq 0$, α_k is defined by

$$\alpha_k \left(f_\sigma(x)\, dx_{\sigma(1)} \wedge \ldots \wedge dx_{\sigma(k)}\right) = \left(\int_0^1 t^{k-1} f_\sigma(tx)\, dt\right) dx_{\sigma(1)} \wedge \ldots \wedge dx_{\sigma(k)},$$

and then extended linearly to all of Ω^k. More explicitly, for each $k \geq 1$,

$$\begin{aligned} h_k \left(f_\sigma(x)\, dx_{\sigma(1)} \wedge \ldots \wedge dx_{\sigma(k)}\right) &= \int_0^1 t^{k-1} f_\sigma(tx)\, dt \\ &\quad \cdot \sum_{j=1}^k (-1)^{j+1} x_{\sigma(j)} dx_{\sigma(1)} \wedge \ldots \wedge dx_{\sigma(j-1)} \wedge dx_{\sigma(j+1)} \wedge \ldots \wedge dx_{\sigma(k)}. \end{aligned}$$

This definition is given in [10, 4.18] only for the unit ball B_1^m. However, the definition is the same for B_r^m and arbitrary star-shaped U. In fact, if $D_r\colon \mathbb{R}^m \to \mathbb{R}^m$ denotes the dilation $x \mapsto r \cdot x$, then we can define $h_k^r = (D_r^{-1})^* \circ h_k \circ D_r^*$ on $\Omega_k(B_r^m)$, and the definitions of α_k, ι_X, and h_k are invariant under conjugation by the induced isomorphism D_r^*. See [10, 4.18] for a proof that $h_{k+1} \circ d + d \circ h_k = \mathrm{id}$.

To verify the claimed estimates, first note that by Lemma 2.8,

$$\|(\iota_X\beta)(x)\|_2 \le \sqrt{k}\,\|x\|_2\,\|\beta(x)\|_2 \le \sqrt{k}\,s\,\|\beta\|_2.$$

Moreover, with the notation from section 2,

$$\begin{aligned}
\|\alpha_k(\beta)(x)\|_2 &= \sqrt{\sum_\sigma \left(\int_0^1 t^{k-1} f_\sigma(tx)dt\right)^2} \\
&\le \frac{1}{k}\sqrt{\sum_\sigma \max_{0\le t\le 1} f_\sigma^2(tx)} \qquad (2)\\
&\le \frac{1}{k}\sqrt{\binom{m}{k}}\max_\sigma \max_{0\le t\le 1} |f_\sigma(tx)| \qquad (3)\\
&\le \frac{1}{k}\sqrt{\binom{m}{k}}\max_{0\le t\le 1} \|\beta(tx)\|_2, \qquad (4)
\end{aligned}$$

which yields the inequality for $h_k = \alpha_{k-1} \circ \iota_X$ if $k > 1$ and for $h_1 = \iota_X \circ \alpha_1$. □

Corollary 4.2. *If all of the coefficients f_σ of a differential k-form β, $k \ge 1$, are constant along rays through the origin, then*

$$\|h_k(\beta)(x)\|_2 \le \frac{\|x\|_2}{\sqrt{k}}\|\beta(x)\|_2 \le \frac{s}{\sqrt{k}}\|\beta\|_2.$$

Proof. Apply the argument in the previous proof to $h_k = \iota_X \circ \alpha_k$ and observe that $\|\alpha_k(\beta)(x)\|_2 = \frac{1}{k}\|\beta(x)\|_2$. □

Remark 4.3. Note that the following estimates in the course of the proof of Lemma 4.1 are sharp: (2), for example, when all f_σ are constant and equal, (3) when all f_σ are equal and nonzero, and (4) when at most one f_σ is nonzero. The combination of the inequalities however may not be sharp (as in Corollary 4.2).

Remark 4.4. On a general closed and oriented Riemannian manifold, one has the Hodge decomposition $d \circ \delta \circ G + \delta \circ G \circ d + H = \mathrm{id}$, where $\delta = \pm \star d \star$, G is Green's operator, and H is the projection to harmonic forms [10, Chapter 6]. An explicit estimate as in Lemma 4.1 for the norm of $(\delta \circ G)(\beta) = (G \circ \delta)(\beta)$ in terms of the norm of an exact form β is much more challenging than on Euclidean space.

Remark 4.5. In general, a family of linear transformations $\sigma_k\colon d(\Omega^{k-1}) \to \Omega^{k-1}$ such that $d \circ \sigma_k = \mathrm{id}$ for all $k \geq 1$ is called a splitting (of the de Rham complex). In that language, the maps h_k in Lemma 4.1 and $\delta \circ G$ in Remark 4.4 are splittings. These splittings are smooth in the sense that the constructions depend smoothly on the differential form. A splitting also exists in the case of non-compact manifolds (that are countable at infinity), see [1, Section 1.5] for a summary.

Remark 4.6. Even more generally, given two maps φ and ψ between smooth manifolds, a family of linear transformations $h_k\colon \Omega^k \to \Omega^{k-1}$ such that $\varphi^* - \psi^* = h_{k+1} \circ d + d \circ h_k$ is called a homotopy operator between φ and ψ. In that language, the linear transformations in Lemma 4.1 define a homotopy operator between the identity map and a constant map [10, Remark 4.19].

A. Banyaga [1, Section 3.1] constructed a homotopy operator $I_{\{\varphi_t\}}$ between a (compactly supported) diffeomorphism that is isotopic to the identity and the identity: let $\{\varphi_t\}_{0 \leq t \leq 1}$ be an isotopy with $\varphi_0 = \mathrm{id}$ and $\varphi_1 = \varphi$, and $X = \{X_t\}_{0 \leq t \leq 1}$ the unique vector field that generates this isotopy. Then $I_{\{\varphi_t\}}(\beta) = \int_0^1 \varphi_t^*(\iota_{X_t}\beta)dt$. In particular, if β is a closed k-form, then $\varphi^*\beta - \beta = dI_{\{\varphi_t\}}(\beta)$. However, the $(k-1)$-form $I_{\{\varphi_t\}}(\beta)$ is not necessarily small when β is small, unless the isotopy $\{\varphi_t\}$ is already known to be C^1-small (or $k = 1$ and the C^0-norm $\|X\|$ is small).

Lemma 4.7. *Suppose that $\varphi\colon B_r^{2n} \to \mathbb{R}^{2n}$ is an ϵ-symplectic embedding such that $0 \leq \epsilon < 1/\sqrt{2}$, and define $\rho > 0$ by*

$$\rho = (1 - \sqrt{2}\,\epsilon)^{\sqrt{2n}} \leq 1.$$

Then there exists an embedding $\psi\colon B_{\rho r}^{2n} \to B_r^{2n}$ such that $B_{\rho s}^{2n} \subset \psi(B_s^{2n}) \subset B_{\rho^{-1}s}^{2n}$ for all $s \leq \rho r$, $\|\psi(x) - x\|_2 \leq (\rho^{-1} - 1)\|x\|_2$ for all $x \in B_{\rho r}^{2n}$, and $(\varphi \circ \psi)^*\omega_0 = \omega_0$. Moreover, ψ is isotopic to the canonical inclusion. If φ is (the restriction of) a linear map $\mathbb{R}^{2n} \to \mathbb{R}^{2n}$, then we may choose

$$\rho = \sqrt{1 - \sqrt{2}\epsilon} \leq 1.$$

Proof. We construct the embedding ψ as the time-one map of an isotopy ψ_t defined on the smaller ball $B_{\rho r}^{2n}$ and with $\psi_0 = \mathrm{id}$ using Moser's argument, and so that $(\varphi \circ \psi)^*\omega_0 = \psi^*(\varphi^*\omega_0) = \omega_0$.

Let $\omega_t = \omega_0 + t(\varphi^*\omega_0 - \omega_0)$. Since $\epsilon < 1$, each two-form ω_t is symplectic, and $\frac{d}{dt}\omega_t = \varphi^*\omega_0 - \omega_0$ is a closed (and hence exact) two-form. Let h_2 be the linear transformation in Lemma 4.1, and consider the one-form $\sigma = h_2(\varphi^*\omega_0 - \omega_0)$. By Lemma 4.1 and by hypothesis,

$$\|\sigma(x)\|_2 \leq 2\sqrt{n}\,\|x\|_2\|\varphi^*\omega_0 - \omega_0\|_2 \leq 2\sqrt{n}\,\epsilon\|x\|_2.$$

Following Moser's idea, define a family of vector fields X_t by $\iota_{X_t}\omega_t = -\sigma$. Then

$$\begin{aligned}
\|(\iota_{X_t}\omega_t)(x)\|_2 &\geq \|(\iota_{X_t}\omega_0)(x)\|_2 - t\,\|(\iota_{X_t}(\varphi^*\omega_0 - \omega_0))(x)\|_2 \\
&\geq \|(\iota_{X_t}\omega_0)(x)\|_2 - \sqrt{2}\,t\,\|X_t(x)\|_2\|\varphi^*\omega_0 - \omega_0\|_2 \\
&= \|X_t(x)\|_2 - \sqrt{2}\,t\,\|X_t(x)\|_2\|\varphi^*\omega_0 - \omega_0\|_2 \\
&\geq (1 - \sqrt{2}\,\epsilon\,t)\|X_t(x)\|_2.
\end{aligned}$$

We used Lemma 2.8 for the second inequality. Thus $\|X_t(x)\|_2 \leq C(t)\|x\|_2$, where

$$C(t) = \frac{\sqrt{2n}\cdot\sqrt{2}\,\epsilon}{1 - \sqrt{2}\,\epsilon\,t}.$$

The trajectories $x\colon [0,1] \to \mathbb{R}^{2n}$ of X_t therefore satisfy (away from the fixed point at the origin) the differential inequality

$$\left|\frac{d}{dt}\|x(t)\|_2\right| = \frac{|2\langle x(t), x'(t)\rangle|}{2\|x(t)\|_2} \leq \|x'(t)\|_2 = \|X_t(x(t))\|_2 \leq C(t)\|x(t)\|_2,$$

and thus $\|x(0)\|_2\,(1 - \sqrt{2}\,\epsilon\,t)^{\sqrt{2n}} \leq \|x(t)\|_2 \leq \|x(0)\|_2\,(1 - \sqrt{2}\,\epsilon\,t)^{-\sqrt{2n}}$. Finally,

$$\begin{aligned}
\|x(1) - x(0)\| \leq \int_0^1 \|x'(t)\|dt &\leq \int_0^1 \frac{C(t)\|x(0)\|}{(1 - \sqrt{2}\epsilon t)^{\sqrt{2n}}}dt \\
&= \|x(0)\|((1 - \sqrt{2}\epsilon)^{-\sqrt{2n}} - 1),
\end{aligned}$$

where we dropped the subscript from the norm $\|\cdot\|_2$ for better readability. Thus the map ψ has all of the stated properties. For the statement about linear maps, substitute the estimate in Corollary 4.2 into the above argument. $\square$

Example 4.8. Suppose $\varphi^*\omega_0 = \sum_{j=1}^n c_j^2\, dx_j \wedge dy_j,\ c_j > 0$. Then the construction in the previous lemma yields

$$\psi(r_1, \theta_1, \dots, r_n, \theta_n) = (c_1 r_1, \theta_1, \dots, c_n r_n, \theta_n),$$

where $x_j = r_j \cos\theta_j$ and $y_j = r_j \sin\theta_j$ are polar coordinates on each $\mathbb{R}^2$-factor.

Remark 4.9. Suppose that $g = \langle \cdot, A(x) \cdot \rangle$ is a Riemannian metric on $\mathbb{R}^m$, where $A(x)$ is a non-singular symmetric matrix that depends smoothly on $x \in \mathbb{R}^m$. Then

$$\|A^{-1}(x)\|^{-k/2}\|v^*\|_{g_0} \leq \|v^*\|_g \leq \|A(x)\|^{k/2}\|v^*\|_{g_0}$$

for any k-covector $v^* \in \Lambda^k(T_x\mathbb{R}^m)$, where $\|A(x)\| = \sup\{\|A(x)v\|_2 \mid \|v\|_2 = 1\}$, and

$$\|A^{-1}\|^{-k/2}\|\beta\|_{g_0} \leq \|\beta\|_g \leq \|A\|^{k/2}\|\beta\|_{g_0}$$

for any k-form β, where $\|A\| = \sup_{x \in U} \|A(x)\|$. In particular, all estimates with respect to the standard metric g_0 hold for an arbitrary Riemannian metric g up to a constant factor that depends on the metric g only.

5. Linear epsilon-non-squeezing and non-expanding

We will show in this section that linear ϵ-symplectic maps are characterized by the property that they preserve the linear symplectic width of ellipsoids up to an error that depends continuously on ϵ and converges to zero as $\epsilon \to 0^+$. The key observation is that the failure to be symplectic can be expressed quantitatively in terms of the symplectic spectrum of ellipsoids (centered at the origin).

We identify $\mathbb{R}^{2n}$ with $\mathbb{C}^n$ in the usual way with $z = (x, y)$ corresponding to $x + iy$ for $x, y \in \mathbb{R}^n$. Recall that with this identification, $\mathrm{Sp}(2n) \cap \mathrm{O}(2n) = \mathrm{U}(n)$, where $\mathrm{Sp}(2n)$, $\mathrm{O}(2n)$, and $\mathrm{U}(n)$ denote the groups of symplectic, orthogonal, and unitary matrices, respectively. We do not distinguish between a matrix and the linear map $\mathbb{R}^{2n} \to \mathbb{R}^{2n}$ or $\mathbb{C}^n \to \mathbb{C}^n$ it represents.

Remark 5.1. For a singular matrix Φ we have $\|\Phi^*\omega_0 - \omega_0\|_2 \geq 1$ (cf. the proof of Lemma 4.7). Thus an ϵ-symplectic matrix with $\epsilon < 1$ is always non-singular.

For a non-singular matrix A, denote the ellipsoid AB_1^{2n} (the image of the closed unit ball) by $E(A) = \{z \in \mathbb{R}^{2n} \mid \langle z, ((A^{-1})^T A^{-1})\, z\rangle \leq 1\}$. For $0 <$

$r_1 \leq \cdots \leq r_n$, consider the diagonal matrix $\Delta(r_1, \ldots, r_n) \colon \mathbb{C}^n \to \mathbb{C}^n$ whose diagonal entries are $r_1, \ldots, r_n$ (in that order), and abbreviate $E(r_1, \ldots, r_n) = E(\Delta(r_1, \ldots, r_n))$, i.e.,

$$E(r_1, \ldots, r_n) = \left\{ z \in \mathbb{C}^n \mid \sum_{j=1}^{n} \left| \frac{z_j}{r_j} \right|^2 \leq 1 \right\}.$$

Recall that for each ellipsoid $E(A)$, there exists a symplectic matrix Ψ such that $\Psi E(A) = E(r_1, \ldots, r_n)$ for some n-tuple $(r_1, \ldots, r_n)$ with $0 < r_1 \leq \cdots \leq r_n$, and that is uniquely determined by A. It is called the symplectic spectrum of $E(A)$, and the number r_1 is its linear symplectic width [6, Section 2.4]. In fact, $r_j^2 = \alpha_j$, where $\pm i\alpha_j$ are the (purely imaginary) eigenvalues (counted with multiplicities) of the matrix $A^T J_0 A$ [6, Lemma 2.4.6]. Recall in this context that A is symplectic if and only if $A^T J_0 A = J_0$, and the latter has eigenvalues $\pm i$ (with multiplicity n). We will generalize the following lemma to a quantitative result for ϵ-symplectic matrices.

Theorem 5.2 ([6]). *A linear map $\Phi \colon \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ is symplectic or anti-symplectic if and only if it preserves the symplectic spectrum of ellipsoids (centered at the origin).*

Proof. By [6, Theorem 2.4.4], a linear map is symplectic or anti-symplectic if and only if it preserves the linear symplectic width of ellipsoids (centered at the origin). Moreover, a linear symplectic map in fact preserves the entire symplectic spectrum of ellipsoids (centered at the origin) [6, Lemma 2.4.6], and the same is true for a linear anti-symplectic map, since for each ellipsoid there exists an anti-symplectic linear map to itself (compose with complex conjugation on $E(r_1, \ldots, r_n)$). □

We will later make use of the following lemma.

Lemma 5.3. *If A is a non-singular matrix, $E(A)$ has linear symplectic width r_1, and $a > 0$, then the ellipsoid $E(aA)$ has linear symplectic width ar_1.*

Proof. The linear map $x \mapsto ax$ is conformally symplectic. □

The following definitions of ϵ-non-squeezing and ϵ-non-expanding are motivated by Lemma 4.7 and Example 4.8. Their meaning is that such a linear map has a non-squeezing property for ellipsoids up to a (small) error that depends on ϵ and the ellipsoid. This error is made precise in the next definition, and is tailored to the proofs of Propositions 1.4 and 6.2 below. In light of Remark 3.1, ϵ-non-expanding is a substitute for a non-squeezing property of the inverse.

Definition 5.4 (Epsilon-non-squeezing and non-expanding). Let $0 \leq \epsilon < 1$, and $0 < \rho = \sqrt{1-\epsilon} \leq 1$. For a non-singular matrix A, let

$$s_A = (1 + \|A^{-1}\|\|A\|(\rho^{-1} - 1))^{-1} \leq \rho \leq 1.$$

Note that $s_A = \rho$ if (and only if) $E(A)$ is a ball (or $\epsilon = 0$). If $\|A^{-1}\|\|A\|(\rho^{-1} - 1) < 1$, define $e_A = (1 - \|A^{-1}\|\|A\|(\rho^{-1} - 1))^{-1} \geq 1$.

(a) A linear map Φ has the linear ϵ-non-squeezing property if for each ellipsoid $E(A)$ with linear symplectic width r_1 such that the image ellipsoid $\Phi E(A)$ has linear symplectic width R_1, the inequality $s_A\, r_1 \leq R_1$ holds.

(b) The linear map Φ has the linear ϵ-non-expanding property if (for $r > 0$) the linear symplectic width of the ellipsoid ΦB^{2n}_r is at most $\rho^{-1} r$, and moreover, for each ellipsoid $E(A)$ with linear symplectic width r_1 and $\|A^{-1}\|\|A\|(\rho^{-1} - 1) < 1$, the linear symplectic width R_1 of the image ellipsoid $\Phi E(A)$ satisfies the inequality $R_1 \leq e_A\, r_1$.

Proof of Proposition 1.4. Let $\rho = \sqrt{1-\epsilon'}$, and $\psi \colon \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ be the embedding from Lemma 4.7, so that the composition $\Phi \circ \psi$ is symplectic. Let $E(A)$ be as in Definition 5.4(a) with ϵ replaced by ϵ' everywhere. Then $x \in E(s_A A)$ implies

$$\begin{aligned}\|A^{-1}(\psi(x))\|_2 &\leq \|A^{-1}x\|_2 + \|A^{-1}\|\|\psi(x) - x\|_2 \\ &\leq s_A + \|A^{-1}\|\|A\|(\rho^{-1} - 1)s_A = 1,\end{aligned}$$

i.e., $\psi(E(s_A A)) \subset E(A)$, and in particular, $(\Phi \circ \psi)E(s_A A) \subset \Phi E(A)$. Note that ψ need not be linear in general. However, the restriction of any (relative) symplectic capacity to ellipsoids (centered at the origin) equals (up to a factor π) the square of the linear symplectic width [6, Example 12.1.7] (see also section 6), and hence the above inclusion implies $s_A\, r_1 \leq R_1$. Thus Φ is ϵ'-non-squeezing.

In the two situations of Definition 5.4(b), $\psi(B^{2n}_{\rho^{-1}r}) \supset B^{2n}_r$, and $x \in \partial E(e_A A)$ implies $\|A^{-1}(\psi(x))\|_2 \geq 1$, respectively. The argument for ϵ'-non-expanding is then analogous to the above argument for ϵ'-non-squeezing. $\square$

Remark 5.5. There exists no non-squeezing result in any form for $\epsilon \geq 1$. Indeed, for any $\delta > 0$, the linear map

$$(x_1, y_1, \ldots, x_n, y_n) \mapsto (\delta x_1, \delta y_1, x_2, y_2, \ldots, x_n, y_n)$$

is 1-symplectic and maps the unit ball to a (symplectic) cylinder of radius δ. More generally, let ω be a symplectic form and δ be the non-degeneracy radius around ω, i.e., the supremum over all numbers d so that $\|\omega' - \omega\|_2 \leq d$ implies that ω' is non-degenerate. Again use Moser's argument for any $d < \delta$. Thus ϵ-symplectic does not guarantee any form of non-squeezing beyond (and at) the threshold $\epsilon = \delta$.

We will use the following obvious remark in the proof of Theorem 1.2 below.

Remark 5.6. Let Φ be a linear map, and Ψ be a symplectic or anti-symplectic linear map. Then Φ has the linear ϵ-non-squeezing (ϵ-non-expanding) property if and only if $\Psi \circ \Phi$ does. Compare to Remark 3.4.

Lemma 5.7. *A bounded subset $U \subset \mathbb{R}^{2n}$ that is contained in a hyperplane H is contained in an ellipsoid of arbitrarily small linear symplectic width. In particular, a singular matrix Φ does not have the linear ϵ-non-squeezing property for any ϵ.*

Proof. Let u be a vector that is orthogonal to H, and v be a vector that belongs to H so that $\omega_0(u, v) > 0$. After rescaling v if necessary, we may assume U is contained in the ball $\{w \in \mathbb{R}^{2n} \mid \|w\| \leq \|v\|\}$. Let $R > 0$. After rescaling u if necessary, we may assume that $\omega_0(u, v) = R^2$. Choose a symplectic basis $\{u_1, \ldots, u_n, v_1, \ldots, v_n\}$ of $\mathbb{R}^{2n}$ so that $u_1 = R^{-1}u$ and $v_1 = R^{-1}v$, and a symplectic matrix Ψ that maps this basis to the standard basis of $\mathbb{R}^{2n}$. Then $\Psi U \subset Z_R^{2n} = B_R^2 \times \mathbb{R}^{2n-2}$. That proves the first claim. The second claim follows by considering $U = \Phi B_1^{2n}$ (so in Definition 5.4 (a), A is the identity matrix) and any positive number $R < \rho$. $\square$

Proof of Theorem 1.5. The proof is given in six steps.

Step 1. Apply Lemma 2.5 to the two-form $\omega = \Phi^* \omega_0$ to find an orthonormal basis $B = \{u_1, \ldots, u_n, v_1, \ldots, v_n\}$ of $\mathbb{R}^{2n}$ and numbers $0 \leq \lambda_1 \leq \cdots \leq \lambda_n$ so that $(\Phi^* \omega_0)(u_j, v_k) = \delta_{jk} \lambda_j^2$ and $(\Phi^* \omega_0)(u_j, u_k) = (\Phi^* \omega_0)(v_j, v_k) = 0$ for $1 \leq j, k \leq n$. (Note that the matrix that is called A in the proof of Lemma 2.5 is $\Phi^T J_0 \Phi$ here.) By Lemma 5.7, we may assume that $\lambda_1 > 0$, and therefore ΦB is (up to rescaling) a symplectic basis of $\mathbb{R}^{2n}$. By composing Φ (on the left) with a symplectic matrix, we may assume that $\Phi u_j = \lambda_j e_j$ and $\Phi v_j = \lambda_j f_j$, where $\{e_1, \ldots, e_n, f_1, \ldots, f_n\}$ denotes the standard symplectic basis of $\mathbb{R}^{2n}$. In particular, Φ maps the unit ball $B_1^{2n} = E(I_{2n})$, where $I_{2n} = I$ denotes the identity matrix, to the standard ellipsoid $E(\lambda_1, \ldots, \lambda_n)$. The hypotheses imply that $\rho = s_I \leq \lambda_1 \leq \rho^{-1}$.

Step 2. For $j = 1, \ldots, n$, write $\mu_j = \sqrt{|\omega_0(u_j, v_j)|} \leq 1$. Fix an index j with $1 \leq j \leq n$, and abbreviate $u = u_j$, $v = v_j$, $\lambda = \lambda_j$, and $\mu = \mu_j$. Let $0 < a \leq 1$. For the remainder of this proof, let A denote the linear map defined by $Au = au$, $Av = av$, and A is the identity on $S = \operatorname{span}(B \backslash \{u, v\})$. Write r_1 for the linear symplectic width of the ellipsoid $E(A)$. The volume of $E(A)$ yields the constraints $a \leq r_1 \leq \sqrt[n]{a}$. We will improve these estimates to symplectic estimates and in terms of the number $0 \leq \mu \leq 1$ as follows.

There exists a unit vector $w \in S$ and $0 \leq s, t \leq 1$ with $s^2 + t^2 = 1$ so that the vectors u and $J_0 u = sv + tw$ span a unitary disk $D \subset B_1^{2n}$ of radius 1, and in particular, $|\omega_0(u, sv + tw)| = 1$. By the Cauchy-Schwarz inequality, the latter implies that $s = |\omega_0(u, v)| = \mu$ and $t = |\omega_0(u, w)| = \sqrt{1 - \mu^2}$. Therefore

$$a^2 \leq r_1^2 = |\omega_0(au, a\mu v + \sqrt{1 - \mu^2}\, w)| = a^2 \mu^2 + a(1 - \mu^2) \leq a. \tag{5}$$

On the other hand, the linear symplectic width R_1 of the ellipsoid $\Phi E(A)$ is the smaller of the two numbers λ_1 and $a\lambda$.

Step 3. By the ϵ-non-expanding hypothesis, $\min(\lambda_1, a\lambda) \leq e_A r_1$ for all numbers a with $\rho^{-1} - 1 < a \leq 1$ (so that the number e_A is well-defined). We will show that for appropriate choices of a, $\rho > e_A \sqrt{a}$, and thus by step 1 and by (5), $\lambda_1 \geq \rho > e_A \sqrt{a} \geq e_A r_1$. Then $a\lambda \leq e_A r_1 \leq e_A a$. Therefore $\lambda_j \leq e_A < \rho\, a^{-1/2}$ for all $j = 1, \ldots, n$ and for all numbers a as above.

Consider the function $f(a) = a^{3/2} - \rho a + 1 - \rho$, $\rho^{-1} - 1 < a \leq 1$. Then the condition $\rho > e_A \sqrt{a}$ translates to the inequality $f(a) < 0$. The (absolute) minimum of the function $f(a)$ is achieved at the point $a = \frac{4}{9}\rho^2$. Therefore the inequality $f(a) < 0$ has a solution if and only if $f(\frac{4}{9}\rho^2) = -\frac{4}{27}\rho^3 + 1 - \rho < 0$. The cubic equation $z^3 + \frac{27}{4} z = \frac{27}{4}$ has a single real root $z_0 = \frac{3}{2}((1 + \sqrt{2})^{1/3} + (1 - \sqrt{2}))^{1/3})$. Thus for $\epsilon < 1 - z_0^2$, the inequality $\rho > e_A \sqrt{a}$ can be solved for some $a \geq \frac{4}{9}\rho^2$. Note that $\frac{4}{9}\rho^2 > \rho^{-1} - 1$ is equivalent to $\frac{4}{9}\rho^3 + \rho - 1 > 0$, and the latter is greater or equal to $-f(\frac{4}{9}\rho^2)$ and thus indeed is positive.

The equation $f(a) = 0$ can be solved in closed form by making the substitution $a = (c\rho)^2$, which leads to the cubic equation $c^3 - c^2 + \rho^{-3}(1 - \rho) = 0$. Since $f(\frac{4}{9}\rho^2) < 0$, $f(\rho^2) = 1 - \rho \geq 0$, and $f'(a) > 0$ for $a > \frac{4}{9}\rho^2$, there exists a single real root c_ρ with $2/3 < c_\rho \leq 1$. By the last sentence of the first paragraph of this step, $\lambda_j \le c_\rho^{-1}$ for all $j = 1, \ldots, n$. For a more explicit estimate in terms of ρ, observe that $f(\rho^6) = (\rho^7(\rho+1) - 1)(\rho - 1)$, and $\rho^7(\rho+1) - 1 \ge 2\rho^8 - 1 > 0$, provided that $\rho > (\frac{1}{2})^{1/8}$. Therefore if $\epsilon < 1 - (\frac{1}{2})^{1/4}$, then $\lambda_j \le \rho^{-2}$ for all $j = 1, \ldots, n$. In particular, $c_\rho \to 1^-$ as $\rho \to 1^-$.

Step 4. As in step 2, fix an index j, and drop the subscripts from the notation. We will prove in this step that $\mu \to 1^-$ as $\rho \to 1^-$.

Define the function $g(a) = (\mu^2 - \lambda^2)a^2 + (1 - \mu^2 - 2\lambda^2(\rho^{-1} - 1))a - (\lambda(\rho^{-1} - 1))^2$. Again by (5), the ϵ-non-squeezing hypothesis $s_A r_1 \le \min(\lambda_1, a\lambda) \le a\lambda$ for all $a > 0$ guarantees that $g(a) \le 0$ for all $0 < a \le 1$. If $\mu \ge \lambda$, then $\mu \ge \rho \to 1^-$ as $\rho \to 1^-$, and there is nothing more to prove. Thus we assume henceforth that $\mu < \lambda$.

The (absolute) maximum of $g(a)$ is achieved at the point

$$a_0 = \frac{1 - \mu^2 - 2\lambda^2(\rho^{-1} - 1)}{2(\lambda^2 - \mu^2)}.$$

We derive at a contradiction if the maximum $g(a_0)$ is positive and $0 < a_0 \le 1$. We distinguish three cases:

Case (i). $a_0 \le 0$. Since $\mu < \lambda$, this is equivalent to $1 - \mu^2 - 2\lambda^2(\rho^{-1} - 1) \le 0$. Then $\mu \ge \sqrt{1 - 2\lambda^2(\rho^{-1} - 1)} \to 1^-$ as $\rho \to 1^-$ since λ is bounded by step 3.

Case (ii). $a_0 > 1$, or equivalently, $\mu > \sqrt{2\lambda^2\rho^{-1} - 1}$. Since $\lambda \ge \lambda_1 \ge \rho$, the latter is bounded from below by $\sqrt{2\rho - 1} \to 1^-$ as $\rho \to 1^-$.

Case (iii). $0 < a_0 \le 1$ and $g(a_0) \le 0$. The latter is equivalent to the inequality

$$\mu^2 - 2\lambda(\rho^{-1} - 1)\sqrt{\lambda^2 - \mu^2} \ge 1 - 2\lambda^2(\rho^{-1} - 1),$$

and in particular, $\mu \ge \sqrt{1 - 2\lambda^2(\rho^{-1} - 1)} \to 1^-$ as $\rho \to 1^-$ as in case (i).

In particular, from cases (i) to (iii) we deduce that $\mu > 0$; since $c_\rho > 2/3$, it suffices that the condition $\epsilon < 1 - z_0^2$ from step 3 implies that $\rho \ge 9/11$.

Step 5. In this step, we use a standard non-squeezing argument to show that the numbers $\omega_0(u_j, v_j) = \pm\mu_j^2$ all have the same sign for $1 \le j \le n$.

By composing Φ (on the left) with a diagonal matrix Ψ with entries equal to ± 1, we may assume that for each j the pairs of numbers $\omega_0(u_j, v_j)$ and $\omega_0(\Phi u_j, \Phi v_j)$ have the same sign, and by rearranging the basis B from step 1 if necessary, that these numbers are all positive. We will see in the next step that then $\Psi \circ \Phi$ is ϵ'-symplectic for some $\epsilon' \ge 0$ with $\epsilon' \to 0^+$ as $\epsilon \to 0^+$. In particular, for $\epsilon' < 1/\sqrt{2}$, $\Psi \circ \Phi$ is ϵ''-non-squeezing for some $\epsilon'' \ge 0$ by Proposition 1.4. But by a standard squeezing argument (see the proof of [6, Theorem 2.4.2]), Ψ squeezes the unit ball B_1^{2n} into a symplectic cylinder of arbitrarily small radius unless its diagonal entries all have the same sign. That shows that Ψ must be symplectic or anti-symplectic.

Step 6. By post-composing with the anti-symplectic matrix Ψ from the previous step if necessary, we may assume that $\omega_0(u_j, v_j) = \mu_j^2$ for all $1 \leq j \leq n$. Recall that $\|\omega_0\|_2 = \sqrt{n}$. Thus

$$\|\Phi^*\omega_0 - \omega_0\|_2^2 = \sum_{j=1}^{n}(\lambda_j^2 - \mu_j^2)^2 + (n - \sum_{j=1}^{n}\mu_j^4) \to 0^+$$

as $\rho \to 1^-$, or equivalently, as $\epsilon \to 0^+$. That proves the theorem. $\square$

Remark 5.8. Let the matrix A and basis B be as in the preceding proof, and also write B for the matrix with columns the vectors in B. Since B is orthogonal, it is symplectic if and only if it is also complex, i.e., commutes with J_0. The deviation of A from being conformally symplectic is measured by the symplectic spectrum of $E(A)$, see (5). A measure of the failure of Φ to be symplectic is therefore the collection of numbers λ_j (conformality) and $\pm\mu_j$ (failure to commute with J_0).

6. Symplectic capacities and rigidity

In this section we prove Theorem 1.2. The proof follows closely the argument in the symplectic case ($\epsilon = 0$) given in [6, Section 12.2].

Recall (from [6, Section 12.1]) that a (normalized symplectic) capacity on $\mathbb{R}^{2n}$ is a functor c that assign to an (arbitrary) subset $U \subset \mathbb{R}^{2n}$ a non-negative (possibly infinite) number $c(U)$ such that the following axioms hold:

- (monotonicity) if there exists a symplectic embedding $\psi\colon U \to \mathbb{R}^{2n}$ such that $\psi(U) \subset V$, then $c(U) \leq c(V)$,
- (conformality) $c(aU) = a^2c(U)$, and
- (normalization) $c(B_1^{2n}) = \pi = c(Z_1^{2n})$.

Moreover, the restriction of any capacity c to ellipsoids (centered at the origin) equals $c(E) = \pi r_1^2$, where r_1 denotes as before the linear symplectic width of E [6, Example 12.1.7]. By the monotonicity and conformality axioms, if $a \geq 1$ and $U \subset \mathbb{R}^{2n}$ are such that $a^{-1}E \subset U \subset aE$, then $a^{-2}c(E) \leq c(U) \leq a^2c(E)$. More generally, the restriction of a capacity to compact convex sets is continuous with respect to the Hausdorff metric [6, Exercise 12.1.8]. Therefore, an ϵ-symplectic embedding preserves the capacity of ellipsoids up to an error that converges to zero as $\epsilon \to 0^+$, see Proposition 6.2. Note that translations in $\mathbb{R}^{2n}$ are symplectic and thus preserve capacity, so we can consider ellipsoids with arbitrary center.

On the other hand, different capacities can have different values on non-ellipsoids, and thus it is necessary for the remainder of the argument to fix a capacity. From now on, denote by c the Gromov width.

Definition 6.1. Let $U \subset \mathbb{R}^{2n}$ and $\varphi\colon U \to \mathbb{R}^{2n}$ be an embedding. Let $\epsilon \geq 0$, and $s_A \leq 1$ and $e_A \geq 1$ be as in Definition 5.4. Then φ is said to preserve the capacity of ellipsoids up to ϵ if $s_A^2 c(E) \leq c(\varphi(E)) \leq e_A^2 c(E)$ for every ellipsoid $E = E(A) \subset U$ (where the second inequality holds whenever the number e_A is defined).

Proposition 6.2. *Let $0 \leq \epsilon \leq 1/\sqrt{2}$, and $\epsilon' = 1 - (1 - \sqrt{2}\epsilon)^{2\sqrt{2n}}$. Then an ϵ-symplectic embedding preserves the capacity of ellipsoids up to ϵ'.*

Proof. The proof is verbatim the same as the proof of Proposition 1.4, but with the constant ρ from Lemma 4.7 for embeddings instead of for linear maps. □

Proposition 6.3. *Let $\varphi_k\colon B_r^{2n} \to \mathbb{R}^{2n}$ be embeddings that preserve the capacity of ellipsoids up to ϵ, and converge uniformly (on compact subsets) to an embedding $\varphi\colon B_r^{2n} \to \mathbb{R}^{2n}$. Then the limit φ again preserves the capacity of ellipsoids up to ϵ.*

Proof. This follows from the definition and the continuity properties of symplectic capacities. Indeed, let $E = E(A)$ be an ellipsoid and $a > 1$, and choose k sufficiently large so that $\varphi(E) \subset \varphi_k(aE)$. Then

$$c(\varphi(E)) \leq c(\varphi_k(aE)) \leq e_A^2 c(aE) = a^2 e_A^2 c(E).$$

Since this holds for all $a > 1$, we have $c(\varphi(E)) \leq e_A^2 c(E)$ as claimed. The proof of the other inequality is similar. □

Remark 6.4. It is actually not necessary to assume that the maps φ_k in the preceding proposition are embeddings. See [6, Lemma 12.2.3].

Proposition 6.5. *Suppose that an embedding $\varphi\colon B_r^{2n} \to \mathbb{R}^{2n}$ preserves the capacity of ellipsoids up to ϵ. Then for $\epsilon \geq 0$ sufficiently small, φ is ϵ'-symplectic or ϵ'-anti-symplectic, where $\epsilon' = K(\epsilon)$ is as in Theorem 1.5 (and converges to zero as $\epsilon \to 0^+$).*

Proof. Let $x \in B_r^{2n}$. By composing with translations (which are symplectic), we may assume that $\varphi(x) = x = 0$ is the origin in $\mathbb{R}^{2n}$. Recall that $D_h \colon \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ denotes rescaling by the factor $h \neq 0$. Note that for $|h| < 1$, the numbers s_A and e_A in Definition 5.4 are rescaling invariant, i.e., $s_{hA} = s_A$ and $e_{hA} = e_A$. Thus $(D_h^{-1} \circ \varphi \circ D_h)$ preserves the capacity of ellipsoids up to ϵ. By Proposition 6.2, the derivative $d\varphi(0) = \lim_{h \to 0}(D_h^{-1} \circ \varphi \circ D_h)$ also preserves the capacity of ellipsoids up to ϵ. Then by Theorem 1.5, $d\varphi(0)$ is either ϵ'-symplectic or ϵ'-anti-symplectic. Since $\epsilon' < 1$, continuity of $d\varphi(x)$ implies that the latter is either ϵ'-symplectic for all x or ϵ'-anti-symplectic for all x. □

Proof of Theorem 1.2. Suppose that $\varphi_k \colon (M_1, \omega_1) \to (M_2, \omega_2)$ is a sequence of ϵ-symplectic embeddings that converges uniformly (on compact subsets) to the smooth embedding $\varphi \colon (M_1, \omega_1) \to (M_2, \omega_2)$. Let $x \in M_1$. Since ϵ-symplectic is a pointwise condition, we may assume that $(M_1, \omega_1) = B_r^{2n}$ is a (symplectic) ball in $(\mathbb{R}^{2n}, \omega_0)$ centered at the origin (corresponding to x), for some small $r > 0$, and $(M_2, \omega_2) = (\mathbb{R}^{2n}, \omega_0)$. We will prove the theorem for the restriction of φ to B_r^{2n} and two constants $\delta(x)$ and $E(x)$ that depend continuously on x (by Remark 4.9). If M_1 is compact, let δ and E be the minimum and maximum, respectively, over all $x \in M_1$. The theorem continues to hold for non-compact manifolds if one replaces the constant ϵ in the definition of ϵ-symplectic by a continuous positive function $\epsilon(x)$ on M_1, and likewise for the constants δ and E of the theorem.

By Proposition 6.2, the limit φ preserves the capacity of ellipsoids up to ϵ'. Then by Proposition 6.5, φ is E-symplectic or E-anti-symplectic, where $E = K(\epsilon') \to 0^+$ as $\epsilon \to 0^+$. It only remains to show that the former alternative holds.

We observe that for ϵ sufficiently small, φ_k is orientation preserving, and thus so is the limit φ. If n is odd, and E is sufficiently small, then φ cannot be E-anti-symplectic. If n is even, the same argument applies to the map $\varphi \times \mathrm{id} \colon M_1 \times \mathbb{R}^2 \to M_2 \times \mathbb{R}^2$. □

Proof of Corollary 1.3. Since $E \to 0$ as $\epsilon \to 0$, taking a subsequence of the sequence φ_k and applying Theorem 1.2 shows that φ is E-symplectic for any $E > 0$. □

7. Shape invariant and epsilon-contact embeddings

This section outlines an alternative proof of Theorem 1.2 and Corollary 1.3 using the shape invariant, and an adaptation of the proof to ϵ-contact embeddings. See [2, 7–9] for details on the shape invariant.

Let $(M, \omega = d\lambda)$ be an exact symplectic manifold of dimension $2n$, and T^n be an n-dimensional torus. An embedding $\iota\colon T^n \hookrightarrow M$ is called Lagrangian if $\iota^*\omega = 0$; the cohomology class $[\iota^*\lambda] \in H^1(T^n, \mathbb{R}) = \mathbb{R}^n$ is called its λ-period.

Definition 7.1 (Shape invariant [2]). Let $\tau\colon H^1(M, \mathbb{R}) \to H^1(L, \mathbb{R})$ be a homomorphism. Then the shape $I(M, \omega, \tau)$ is the subset of $H^1(T^n, \mathbb{R})$ that consists of all points $z \in H^1(T^n, \mathbb{R})$ such that there exists a Lagrangian embedding $\iota\colon T^n \hookrightarrow M$ with $\iota^* = \tau$ and $z = [\iota^*\lambda]$, defined up to translation.

Theorem 7.2 ([2, 9]). *For $A \subset \mathbb{R}^n$ open and connected, $I(T^n \times A, \lambda_{can}, \iota_0^*) = A$.*

The role of ellipsoids $E(r_1, \ldots, r_n)$ in this paper can therefore be replaced by products of annuli $A(a_1, b_1, \ldots, a_n, b_n) = (S^1 \times [a_1, b_1]) \times \cdots \times (S^1 \times [a_n, b_n])$, $0 < a_i < b_i$, and the spectrum $(r_1, \ldots, r_n)$ of the ellipsoid $E(r_1, \ldots, r_n)$ by the shape $[a_1, b_1] \times \cdots \times [a_n, b_n]$ of the annulus $A(a_1, b_1, \ldots, a_n, b_n)$. Similar to the proof given above, ϵ-symplectic embeddings preserve the shape invariant up to an error that converges to zero as $\epsilon \to 0^+$, and this property is preserved by uniform limits (on compact subsets). See [7] for the $\epsilon = 0$ case. Details are forthcoming.

We indicate how to prove Corollary 1.3 based on properties of the shape invariant. An embedding $\iota\colon T^n \hookrightarrow M$ is non-Lagrangian if $\iota^*\omega \neq 0$ (at at least one point), or equivalently, its image is a non-Lagrangian submanifold.

Theorem 7.3 ([5, 7]). *Let $\iota\colon T^n \hookrightarrow (\mathbb{R}^{2n}, \omega_0)$ be a non-Lagrangian embedding. Then there exists a tubular neighborhood N of $\iota(T^n)$ that admits no Lagrangian embedding $\jmath\colon T^n \hookrightarrow N$ so that the homomorphism $\jmath_*\colon H_1(T^n, \mathbb{R}) \to H_1(N, \mathbb{R})$ is injective. In particular, the shape $I(N, T^n, \iota^*)$ is empty.*

Sketch of proof of Corollary 1.3. Let $\iota\colon T^n \hookrightarrow M_1$ be a Lagrangian torus, and N be an arbitrary tubular neighborhood of $(\varphi \circ \iota)(T^n)$. Let ψ_k be as in Lemma 4.7 (defined on a polydisk $P(r_1, \ldots, r_n)$) so that $\varphi_k \circ \psi_k$ is symplectic. Then for k sufficiently large, the torus $(\varphi_k \circ \psi_k \circ \iota)(T^n)$ is Lagrangian and contained in N, and $\varphi_k \circ \psi_k$ is homotopic to φ. By Theorem 7.3, $(\varphi \circ \iota)(T^n)$ must be Lagrangian. Thus φ maps Lagrangian tori to Lagrangian tori, and hence must be conformally symplectic [7], i.e., $\varphi^*\omega_2 = c\,\omega_1$. That $c = 1$ can be proved using [7, Proposition 2.29]. $\square$

Let (M_1, ξ_1) and (M_2, ξ_2) be cooriented contact manifolds of the same dimension, and g_1 be a Riemannian metric on M_1. An embedding $\varphi\colon M_1 \to M_2$

is called ϵ-contact if there are contact forms α_1 on M_1 and α_2 on M_2 so that $\|\varphi^*\alpha_2 - \alpha_1\| \leq \epsilon$ and $\|\varphi^* d\alpha_2 - d\alpha_1\| \leq \epsilon$. This definition allows the Moser argument in the proof of Lemma 4.7 to go through in the contact setting [6, page 135f], and the proof of C^0-rigidity of ϵ-symplectic embeddings in this section can be adapted to ϵ-contact embeddings. Note that the proof using capacities does not generalize, since the capacity of the symplectization of a contact manifold is infinite. Compare to [7].

Acknowledgments

We would like to thank Michael Freedman for formulating the question about ϵ-symplectic rigidity answered in this paper and for subsequent email communication, and Yasha Eliashberg for relaying the question and for stimulating discussions.

References

[1] Augustin Banyaga. *The structure of classical diffeomorphism groups*, volume 400 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1997.

[2] Yakov Eliashberg. New invariants of open symplectic and contact manifolds. *J. Amer. Math. Soc.*, 4(3):513–520, 1991.

[3] Herbert Federer. *Geometric measure theory*. Die Grundlehren der mathematischen Wissenschaften, Band 153. Springer-Verlag New York Inc., New York, 1969.

[4] Michael Freedman. Private email communication with Yasha Eliashberg. 2017.

[5] F. Laudenbach and J.-C. Sikorav. Hamiltonian disjunction and limits of Lagrangian submanifolds. *Internat. Math. Res. Notices*, (4):161 ff., approx. 8 pp. (electronic), 1994.

[6] Dusa McDuff and Dietmar Salamon. *Introduction to symplectic topology*. Oxford Graduate Texts in Mathematics. Oxford University Press, Oxford, third edition, 2017.

[7] Stefan Müller. C^0-characterization of symplectic and contact embeddings and Lagrangian rigidity. *Internat. J. of Math.*, 30(9):1950035, 2019.

[8] J.-C. Sikorav. Quelques propriétés des plongements lagrangiens. *Mém. Soc. Math. France (N.S.)*, (46):151–167, 1991. Analyse globale et physique mathématique (Lyon, 1989).

[9] Jean-Claude Sikorav. Rigidité symplectique dans le cotangent de T^n. *Duke Math. J.*, 59(3):759–763, 1989.

[10] Frank W. Warner. *Foundations of differentiable manifolds and Lie groups*, volume 94 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1983. Corrected reprint of the 1971 edition.

Department of Mathematics, Stanford University
Stanford, CA 94305, USA
Department of Mathematical Sciences, Georgia Southern University
Statesboro, GA 30460, USA
E-mail address: `stefan.x.mueller@gmail.com`

Received June 22, 2018
Accepted April 16, 2022