Convergence in shape of Steiner symmetrized line segments

by

Arthur Korneychuk

A thesis submitted in conformity with the requirements
for the degree of Master of Science
Graduate Department of Mathematics
University of Toronto

© Copyright 2014 by Arthur Korneychuk

Abstract

Convergence in shape of Steiner symmetrized line segments

Arthur Korneychuk
Master of Science
Graduate Department of Mathematics
University of Toronto
2014

Building on the recent work of Burchard, Bianchi, Gronchi and Volcic, it is shown that for every finite union of line segments, every sequence of Steiner symmetrals will converge in shape. It is also shown via counterexample that the rotation scheme used in the aforementioned paper to establish the square summable case will not work generally.

Contents

Introduction

Given a compact set K and a sequence of symmetrizations S_{u_n}, we can consider the convergence properties of the sequence of symmetrals

$$K_n = S_{u_n}...S_{u_1}K$$

Such sequences enjoy a host of monotonicity properties useful in optimization problems, but there are open questions regarding their convergence. In a recent paper [1], it is shown that for any sequence of Steiner symmetrals whose successive angles are square summables, we have convergence up to isometry. It leaves as an open question what happens when the assumption of square summability is dropped. This thesis gives a partial answer to this question. In particular, two results are established. First, the construction used in [1] does not work when the assumption of square summability is dropped. This is the subject of chapter two, where we sketch out the construction of a counterexample.

The second result is that any finite union of line segments converge in shape under any sequence of symmetrizations. This is the subject of the third chapter. Thus for this restricted class of sets, the open question is resolved affirmatively. The proof relies on the fact that symmetrization acts very similarly to projection when acting on line segments.

Chapter 1

Preliminaries

1.1 Steiner Symmetrization

Following [2], we define Steiner symmetrization as follows: let K be a compact set in $\mathbb{R}^d$, and $u \in S^{n-1}$ a direction. For each point $z \perp u$, we compute the one-dimensional measure of K with the line through z in the direction of u, and replace it with the closed interval of the same length centered on $u^\perp$. If the intersection is empty, so is the interval; if it's nonempty and of measure zero, then it is a single point. The resulting set is S_uK. Formally:

Given a compact set $K \subset R^n$ and a direction $u \in S^{n-1}$, we define the **Steiner symmetral of K in the direction u** as

$$S_uK = \{z + tu \mid z \in \pi_uK \text{ and } |t| \leq \frac{1}{2}m_1(l_z \cap K)\}$$

π_u is projection in the direction u, i.e. projection onto $u^\perp$. m_1 is one-dimensional Lebesgue measure, For a more comprehensive discussion of Steiner symmetrization, see [3]. A few remarks about this definition are in order. Since our concern is mainly line segments, note that the symmetral of a line segment is almost always equivalent to projection in the same direction, i.e. $S_ul = \pi_ul$. The two exceptions to this are when u and l are colinear, i.e. when $u = \pm\frac{l}{||l||}$. This breaks the linearity of the operation in a non-trivial way, as a sequence of directions may overlap with a line segment infinitely many times without resulting in any major degeneracy.

Another remark is that the symmetral of a finite union of line segments is almost always the union of the symmetrals of the individual line segments, i.e.

$$S_u(l_1 \cup ... \cup l_k) = S_ul_1 \cup ... \cup S_ul_k$$

The exception is when we have two colinear, non-overlapping line segments and we symmetrize in their direction. This will not be a serious obstacle since it can only happen finitely many times in a given sequence, given that the two segments will become one. Given that we're concerned only about convergence, and hence the tail end of a sequence, we can consider a new sequence starting from the point in the old sequence where we no longer have line segments merging. The two sequences will have the same limits, if they converge. We will use this idea several times in the paper to assume without loss of generality that certain special cases do not occur anywhere in a sequence of symmetrals.

Throughout the paper, we use the Hausdorff metric. We denote the ball around the origin of radius δ as B_δ. By $K + H$ for sets K and H we mean the Minkowski sum $\{k + h \mid k \in K, h \in H\}$. Given compact sets K and G, the **Hausdorff distance between K and G** is defined as

$$d_H(K, G) = \inf\{\delta > 0 \mid K \subset G + B_\delta \text{ and } G \subset K + B_\delta\}$$

The results concern the convergence of line segments, which can be characterized by their endpoints. We denote the line segment between points a and b by $\overline{ab}$. The following proof is little more than validation that the Hausdorff metric provides a reasonable notion of convergence of line segments:

Proposition 1. *Let $l_n = \overline{a_n b_n}$ be a sequence of line segments.*
If $a_n \to a$, and $b_n \to b$, then l_n converges to $l = \overline{ab}$ in the Hausdorff metric.

Proof. Suppose $a_n \to a$ and $b_n \to b$.
Let $\epsilon > 0$ and choose N sufficiently large so that for all $n \geq N$,

$$|a_n - a| < \epsilon/2 \text{ and } |b_n - b| < \epsilon/2$$

Then the convexity of norm gives us the result. For all $n > N$, $t \in (0, 1)$,

$$\begin{aligned} |ta_n + (1-t)b_n - (tb_n + (1-t)b)| &= |t(a_n - a) + (1-t)(b_n - b)| \\ &\leq t|a_n - a| + (1-t)|b_n - b| \\ &< \epsilon \end{aligned}$$

So in particular, for each point on one segment (represented as a convex combination of endpoints) the corresponding convex combination of the other segment's endpoints is within ϵ. Thus

$$l_n \subset l + B_\epsilon \text{ and } l \subset l_n + B_\epsilon$$

□

As a corollary, we get that for centrally symmetric line segments (i.e. segments for which $b = -a$), it suffices to show convergence for a single endpoint.

So rather than dealing with entire sets of points, we will look at just individual endpoints for which the analysis is much simpler. Unfortunately, we don't have $S_u\overline{a(-a)} = \overline{(S_ua)(S_u(-a))}$ in the case where a lies in the direction u. So we define a new operator S'_u which gives us this relation. In particular, define S'_u as a function on R_n that acts as S_u (in particular, as projection) on points that don't lie in the direction u, and as identity on points that do. Thus we have

$$S_u\overline{a(-a)} = \overline{(S'_ua)(S'_u(-a))}$$

The upshot of all of this is that the convergence of the symmetrals of a centralized line segment has been reduced to the convergence of a single sequence of points, where each point is related to the previous by either a projection or by identity. We derive the convergence of arbitrary line segments as a corollary of the centralized case.

To get from the centralized case to the arbitrary case, the following is a useful characterization of line segments. We can represent a line segment as a parallel, centralized line segment of the same length, translated by the midpoint of the original segment.

Proposition 2. *Let* $l = \overline{ab}$, $c = \dfrac{b-a}{2}$, $L = (l - \dfrac{b-a}{2})$.
Then $l = c + L$ *and* $S_ul = S_uc + S_uL$

Proof. Immediate from definitions. □

So in particular, convergence of both the midpoint c and the centralized segment L give us convergence of l.

1.2 Rotations

Simple examples exist of sequences of symmetrals failing to converge even under strong conditions on the sequence of directions. This is explored in more detail in section 2. In all known cases, however, we have convergence in shape. Formally, a sequence of compact sets $\{K_n\}_{n\in N}$ **converges in shape** if there exists a sequence of isometries $\{R_n\}_{n\in N}$ such that R_nK_n converges.

Given a pair of orthonormal vectors v, w we can define a rotation that rotates the vw plane by θ while fixing all vectors orthogonal to this plane. Treating v, w as column vectors, it can be checked that this rotation is given by the matrix

$$R_{v,w} = I_d + \sin\theta(vw^T - wv^T) - (1 - \cos\theta)(vv^T + ww^T) \tag{1.1}$$

This type of rotation is of special interest to us as it rotates one vector to another in a minimal fashion, in the sense that the entire space orthogonal to the span of those two vectors is fixed. This allows us to derive the existence of a rotation needed for the main result:

Proposition 3. *Given a point $x = (x_1, ..., x_d)$ in $\mathbb{R}^d$, and $n < d$, there exists a rotation R such that $Rx = (x_1, ..., x_{n-1}, x'_n, 0, ..., 0)$ for some $x'_n \in R$, and R fixes the span of $e_1, ..., e_{n-1}$.*

Proof. We apply the rotation $R = R_{e_n,e_{n+1}}...R_{e_{d-1},e_d}$ where the angle in each rotation $R_{e_i,e_{i+1}}$ is chosen such that $(x_1, ..., x_i, x_{i+1}, 0, ..., 0)$ is sent to $(x_1, ..., x'_i, 0, ..., 0)$ for some x'_i. As we've seen, each rotation will fix every other component so we end up with the desired Rx, and since each chosen rotation fixes at least the span of $e_1, ..., e_{n-1}$, the composite rotation R will do so as well. □

1.3 Summability and convergence

Here we record two propositions relating summability and convergence. The first is a condition for convergence from basic analysis that arises frequently in this context. We omit the simple proof:

Proposition 4. *If $\sum_{i=1}^{\infty} |a_n - a_{n-1}| = M < \infty$, for $a_n \in \mathbb{R}$, then a_n converges.*

In both of the following chapters, we are interested in whether a sequence of numbers is summable or merely square summable.

Proposition 5. *Suppose $\{u_n\}_{n\in N}$ is a sequence of directions and $\{p_n\}_{n\in N}$ a sequence of points in $\mathbb{R}^d$ such that $p_n = S'_{u_n} p_{n-1}$ for $n \geq 1$. Then*

$$\sum_{i=1}^{\infty} ||p_i - p_{i-1}||^2 < \infty$$

Proof. The proof rests on showing the following claim:

$$||p_0||^2 = ||p_n||^2 + \sum_{i=1}^{n} ||p_i - p_{i-1}||^2$$

Once we show this for every n, then we have

$$\sum_{i=1}^{n} ||p_i - p_{i-1}||^2 \leq ||p_0||^2$$

This immediately gives us $\sum_{i=1}^{\infty} ||p_i - p_{i-1}||^2 \leq ||p_0||^2 < \infty$ as desired.

We prove the claim by induction. Consider the point $p_{n-1} - p_n = p_{n-1} - S'_{u_n} p_{n-1}$. If S' acts as identity, this point is the origin. Otherwise, if S' acts as projection, then $p_{n-1} - \pi_{u_n} p_{n-1}$ is orthogonal to p_{n-1}. This observation allows us to use the Pythagorean Theorem as follows:

$$\begin{aligned} ||p_0||^2 &= ||p_1 + (p_0 - p_1)||^2 \\ &= ||p_1||^2 + ||p_0 - p_1||^2 \end{aligned}$$

This proves the base case. Assuming it holds for $n - 1$, we have

$$\begin{aligned} ||p_0||^2 &= ||p_{n-1}||^2 + \sum_{i=1}^{n-1} ||p_i - p_{i-1}||^2 \\ &= ||p_n + (p_{n-1} - p_n)||^2 + \sum_{i=1}^{n-1} ||p_i - p_{i-1}||^2 \\ &= ||p_n||^2 + \sum_{i=1}^{n} ||p_i - p_{i-1}||^2 \end{aligned}$$

This proves the claim, and hence also the proposition. □

Chapter 2

Convergence of Symmetrals

In this chapter we take a closer look at convergence in shape of symmetrals, explain the main result of [1] for the special class of sequences of directions whose angles are square summable, and show how this result does not generalize beyond that class. Nothing in this chapter is used in the third chapter, and may be skipped to get to the main result.

2.1 Symmetrals do not always converge

It was mentioned in the first chapter that there exist simple counterexamples to the convergence of symmetrals. One such example involves taking a line segment in two dimensions and sequence of directions u_n such that the angle between successive rotations, defined by $\cos\theta_n = \langle u_n, u_{n-1}\rangle$, is square summable but not summable, and always rotated the same way (e.g. clockwise). Specifically, we take the line segment l from the origin to e_2, and take e_1 as the first direction u_1. We define u_n as the direction that forms an angle of $1/n$ radians with respect to u_{n-1}, in the clockwise direction. Each symmetral then shrinks the segment by a factor of $\cos(1/n)$, and rotates it by $1/n$ radians. Square summability implies that the partial products $\Pi\cos(1/n)$ converge to a positive number, so the length converges. Lack of summability implies that the partial sums $\sum 1/n$, which quantify the spinning of the segment, do not converge. Thus the line segment does not converge.

We're shrinking the norm by a monotone factor, so the length of the segments clearly converges. But we're also introducing a rotation by projecting onto a sequence of moving hyperplanes ${u_n}^{\perp}$, which is responsible for the non-convergence.

2.2 Rotation scheme for convergence in shape

One idea of how to fix the unwanted rotational component is to rotate the hyperplane to a fixed place at each iteration. Specifically, we apply the following rotation scheme:

Let R'_1 send u_1 to e_1, fixing $(\text{span}\{e_1, u_1\})^{\perp}$.
Let R'_n send $R'_{n-1}...R'_1 u_n$ to e_1, fixing $(\text{span}\{e_1, R'_{n-1}...R'_1 u_n\})^{\perp}$.
Let $R_n = R'_n...R'_1$.

This gives us a sequence of rotated symmetrals $K_n = R_n S_{u_n}...S_{u_1} K$ for arbitrary compact sets K. Using the unproven commutation relation $RS_uH = S_{Ru}RH$, we can rewrite the expression for K_n as

$$K_n = S_{e_1} R'_n S_{e_1} R'_{n-1}...S_{e_1} R'_1 K$$

In particular, this gives us the relation

$$K_n = S_{e_1} R'_n K_{n-1} \tag{2.1}$$

So effectively we're implementing the simplest rotations that force the symmetrizations to be in the same direction, so that the hyerplane being projected against is always ${e_1}^{\perp}$. It's easy to see that this gives us convergence in shape of the two-dimensional line segment, as each iteration yields a line segment from the origin to a shrinking point along e_2. What's remarkable is that if the sequence of directions has square summable angles, this exact rotation scheme works in every dimension for every compact set K. This is the main result of [1]. This constructively proves that with the square summable assumption, every sequence of symmetrals converges in shape.

An open question is to what extent we can relax the assumption of square summability. If we permit orthogonal directions, this same rotation scheme clearly won't work, as can be seen by taking the line segment from $-e_1$ to e_1 and the sequence of directions $u_{2n-1} = e_1, e_{2n} = e_2$. Every symmetrization fixes the set, but the rotations rotate the line segment by $\pi/2$ each time. It's clear where we go wrong: symmetrizing in the direction of a line segment produces a dramatically different effect from symmetrizing in an arbitrarily close direction, and the rotation scheme does not reflect this. But even if we restrict the angles to not being orthogonal, a counterexample can still be produced (albeit far less trivially).

2.3 Limitation of rotation scheme

Now we sketch the construction of a counterexample in three dimensions. We start with a line segment in the yz plane (i.e. ${e_1}^\perp$). By (2.1) the symmetral at each iteration can be computed by performing a rotation on such a segment, and then killing the x-component, since symmetrization reduces to projection. So in particular, we want to see what happens to a line segment (or more precisely, its endpoint) in the yz plane, $l = (0, b, c)$, when we apply a rotation of the form specified, i.e. a rotation from $u = (u_x, u_y, u_z)$ to e_1. Without loss of generality, we suppose that $\langle u, e_1 \rangle$ is between 0 and $\pi/2$. This is because u and $-u$ define the same symmetrals, so we can replace a direction with its negative without affecting connvergence. In order to use (1.1), we need a pair of orthonormal vectors. So in the notation of (1.1) we can take $w = e_1$ and

$$v = \frac{u - \langle u, e_1 \rangle e_1}{||u - \langle u, e_1 \rangle e_1||} = \frac{1}{\sin\theta}(0, u_y, u_z)$$

where θ is the angle between u and e_1.

Computing the corresponding rotation matrix R and applying it to the segment l yields

$$Rl = (-\langle u, l \rangle, b - \frac{u_y}{1+\cos\theta}\langle u, l \rangle, c - \frac{u_z}{1+\cos\theta}\langle u, l \rangle)$$

Thus we have

$$RS_u l = S_{e_1} Rl = (0, b - \frac{u_y}{1+\cos\theta}\langle u, l \rangle, c - \frac{u_z}{1+\cos\theta}\langle u, l \rangle)$$

In order to determine convergence, we are particularly interested in the angle α between successive symmetrals. Of special interest is therefore the following inner product:

$$||RS_u l||\ ||l|| \cos\alpha = \langle RS_u l, l \rangle = ||l||^2 - \frac{\langle u, l \rangle^2}{1+\cos\theta}$$

The left equality is the definition of α, the right equality is the computation of the dot product. Rearranging, we get

$$\cos\alpha = \frac{||l||^2}{||l||\ ||RS_u l||} - \frac{\langle u, l \rangle^2}{(1+\cos\theta)||l||\ ||RS_u l||}$$

Now consider a sequence $\{u_n\}_{n\in N}$ in S^{n-1}, with $\theta_n \in (0, \pi/2)$, for $\cos\theta_n = \langle u_n, u_{n-1} \rangle$ and a line segment identified by its endpoint $l = (0, b, c)$, with l_n being the nth rotated symmetral. If the norm converges to a positive number A, then

$$\cos\alpha_n \approx 1 - \frac{\langle u_n, l_{n-1}\rangle^2}{(1+\cos\theta_n)A^2}$$

This tells us that

$$|\alpha_n| \approx |A_n\langle u_n, l_{n-1}\rangle|$$

Where the set of $A_n = \dfrac{1}{\sqrt{(1+\cos\theta_n)}A}$ is bounded by $1/A$ and $1/\sqrt{2}A$. Thus the upshot of this computation is that the summability of α_n, which tells us whether the net rotation converges or not, is equivalent to the summability of the inner products $\langle u_n, l_{n-1}\rangle$. From this computation, we see that a counterexample would need to satisfy two conditions: the norm must converge to a positive number and the sequence of inner products $\langle u_n, l_{n-1}\rangle$ must not be summable.

Satisfying the above conditions is not difficult. Start with some line segment $(0, 1, 0)$, then pick a sequence of directions $u_n = (e^{-n}, y_n, z_n)$ such that $\langle u_n, l_{n-1}\rangle = \dfrac{C}{n}$ for a constant C chosen to ensure the norm converges to a positive number. There are only two such directions satisfying this equality, and we can pick the one that makes the line segment spin in the same direction each time to avoid cancellation of the net angle.

Chapter 3

Main Result

This chapter is structured as follows. Lemma 1 shows that inner products converge under the operation S'. This gives us convergence of angles between line segments that don't collapse to the origin. Lemma 2 shows that a maximal linearly independent set is preserved under S'. Theorem 4 establishes convergence in shape for finite unions of centralized line segments. The proof of the theorem looks at a set of points consisting of one endpoint from each line segment. From this it selects a maximal linearly independent set, which acts as a 'frame' for the shape of the set. In particular, by fixing the orientation of this frame, we fix the rest of the linearly dependent points. By Lemma 2 this frame is preserved at every iteration. By Lemma 1 we have convergence of the points in these frames, and hence of the linearly dependent points as well. This gives us the convergence of the union.

Finally, we use Theorem 4 to derive the general case, and also to show that the result holds if we use projections instead of symmetrizations.

Lemma 1. *Let* $a, b \in R^n$.
Let $\{u_n\}_{n \in N}$ *be a sequence of directions.*
Let $a_n = S'_{u_n} \ldots S'_{u_1} a$, $b_n = S'_{u_n} \ldots S'_{u_1} b$.
Then $P_n = \langle a_n, b_n \rangle$ *converges.*

Proof. If a_n, b_n ever fall on the same line through the origin, the result becomes trivial as the inner product reduces to the product of norms, each of which converges by monotonicity.

So suppose this never happens. We first prove the claim under the assumption that the S' operator acts as projection at all but finitely many iterations. Without loss of generality, we assume it acts as projection for every iteration. We show that the quantity $D_n = P_n - P_{n-1} = \langle a_n, b_n \rangle - \langle a_{n-1}, b_{n-1} \rangle$ is absolutely summable, which proves the lemma by Proposition 4. First, consider the following two quantities:

$$\langle a_n, b_n - b_{n-1}\rangle \text{ and } \langle b_n, a_n - a_{n-1}\rangle$$

a_n is being projected onto the subspace ${u_n}^\perp$, while $b_n - b_{n-1}$ lies entirely on u_n. Thus the first inner product is 0, and by the same reasoning, so is the second. Thus we have

$$\begin{aligned}\langle a_n, b_n\rangle &= \langle a_{n-1} + (a_n - a_{n-1}), b_{n-1} + (b_n - b_{n-1})\rangle \\ &= \langle a_{n-1}, b_{n-1}\rangle + \langle a_{n-1}, b_n - b_{n-1}\rangle + \langle a_n - a_{n-1}, b_{n-1}\rangle + \langle a_n - a_{n-1}, b_n - b_{n-1}\rangle \\ &= \langle a_{n-1}, b_{n-1}\rangle + \langle a_n - a_{n-1}, b_n - b_{n-1}\rangle\end{aligned}$$

Rearranging, we end up with

$$D_n = \langle a_n - a_{n-1}, b_n - b_{n-1}\rangle$$

Taking absolute values we have

$$|D_n| \leq ||a_n - a_{n-1}||\ ||b_n - b_{n-1}|| \leq ||a_n - a_{n-1}||^2 + ||b_n - b_{n-1}||^2$$

Thus by Proposition 5, $|D_n|$ is summable.

Next we prove the lemma under the assumption that S' acts as identity infinitely many times. We show the inner product must converge to 0. Consider an index n where S'_{u_n} acts as identity on one of the points (and only one, by assumption). One point is fixed on the direction u_n, while the other gets projected into ${u_n}^\perp$, so they're orthogonal. In particular, $P_n = 0$ for every index where S'_{u_n} acts as identity.

Consider an index for which S' does not act as identity on either point. Denote this index by $k+n$, where k is the biggest previous index such that S' last operated as identity. In particular, this means S' acts as projection for $k+1, ..., k+n$. We know that $P_k = 0$. Thus it follows that

$$P_{k+n} = P_k + \sum_{i=k+1}^{k+n} D_i = \sum_{i=k+1}^{k+n} D_i$$

Taking absolute values and using the computation for D_i from above, we get

$$|P_{k+n}| \leq \sum_{i=k+1}^{k+n} |D_i| \leq \sum_{i=k+1}^{\infty} (||a_i - a_{i-1}||^2 + ||b_i - b_{i-1}||^2)$$

This is the tail end of a convergent series, so it goes to zero as k grows large, which happens as $k+n$ is taken large. □

Remark: it may amuse the reader to note that Lemma 1 can be proved immediately for projections using the polarization identity and the fact that projections are linear operators with operator norm not greater than one. Nonetheless the proof used above was chosen since we need the computation for $|D_n|$.

Lemma 2. *If $A = \{b_1, ..., b_k\}$ is a linearly dependent set in $\mathbb{R}^d$, and u is a direction, then $A' = \{S'_u b_1, ..., S'_u b_k\}$ is linearly dependent.*

Proof. It is clear that projection preserves linear dependence in this way, as projection is linear. S'_u reduces to projection except when it acts on a point in the direction u, so if A contains no points in this direction, then A' is dependent. Suppose without loss of generality that b_0 is in the direction u. If another point b_i is in this direction, then b_0 and b_i are dependent (being equal up to a scalar multiple), and since S' fixes both, A' remains dependent. So suppose b_0 is the only point in this direction. Furthermore, if $b_1, ..., b_k$ are dependent, then since none of them lie in the direction u, S' acts as projection on them and they remain dependent. This leaves the case where b_0 is a linear combination of $b_1, ..., b_k$. Thus

$$b_0 = c_1 b_1 + ... + c_k b_k$$

where c_i are not all zero. Now if we project along u and use linearity we get

$$0 = \pi_u b_0 = c_1 \pi_u b_1 + ... + c_k \pi_u b_k$$

This is a non-trivial linear combination of points in A' equal to 0, showing that A' is dependent. □

Corollary 3. *Given any finite set of points $a_1, ..., a_k \in R^d$ and any sequence of directions u_n, there exists an $m \in N$ such that*

$$\textit{For all } n > m, \{a_{k_1}^{(n)}, ..., a_{k_p}^{(n)}\} \textit{ is linearly independent iff } \{a_{k_1}^{(n+1)}, ..., a_{k_p}^{(n+1)}\} \textit{ is.}$$

Proof. By Lemma 2, S' cannot create a new linearly independent set. Thus the number of linearly independent sets, which is finite, is monotonically decreasing and hence is eventually fixed after some index. That index is the m for the claim. □

Theorem 4. *Let $l_1, ..., l_k$ be centralized line segments, $\{u_n\}_{n \in N}$ a sequence of directions. Then $L_n = S_{u_n} ... S_{u_1}(l_1 \cup ... \cup l_k)$ converges in shape.*

Proof. Let a_i be an endpoint of the l_{k_i} which don't collapse to the origin. To show that the union converges under a suitable sequence of isometries, we show that each line segment converges, using the fact that the symmetral of a union of centralized line segments is the union of the symmetrals. To show that each line segment converges, it suffices to show that a single endpoint from each segment converges, by the corollary to Proposition 1.

By Corollary 3, we have that linear dependence is eventually preserved so we assume without loss of generality that it's preserved across the entire sequence for notational simplicity.

Now we construct the isometries. Let $b_1, ..., b_m \in \{a_1, ..., a_k\}$ denote a maximal linearly independent set, chosen arbitrarily. By the previous paragraph, it remains a maximal linearly independent set after every finite iteration. For a given n, we construct a sequence of basic isometries $R_1^n, .., R_k^n$ as follows. Let R_1^n be the isometry sending $b_1^{(n)}$ to a positive multiple of e_1. In general, define R_i^n as the isometry embedding $b_i^{(n)}$ into $\text{span}\{e_1, ..., e_i\}$ such that $\langle b_i^{(n)}, e_i\rangle > 0$. We do this by using Proposition 3 and then reflecting across e_i if necessary. Define $R^n = R_k^n ... R_1^n$. This will be the sequence of isometries used to determine convergence in shape. From this point on, when we refer to a point, we mean its image under the corresponding R^n. Since rotations are isometries, this will not affect any inner products or linear dependencies.

To see that this process determines a unique position for each point $b_i^{(n)}$, note that we certainly have a fixed location for $b_1^{(n)}$ given that it has a particular norm, and is on the positive e_1 axis. For $b_i^{(n)}$, it is contained in the i-dimensional space $\text{span}\{e_1, ..., e_i\}$ and has a fixed inner product with each $b_j^{(n)}$ for $j < i$. Since $b_j^{(n)}$ form a spanning set for $\{e_1, ..., e_{i-1}\}$, the first $i-1$ coordinates of $b_i^{(n)}$ are fixed. Fixing the norm determines two solutions for the ith coordinate equal up to sign change, and $\langle b_i^{(n)}, e_i\rangle > 0$ fixes the sign.

A similar inductive argument yields convergence: it is easy but tedious to show that each coordinate of each point is a continuous function of converging inner products and norms, with only norms appearing in the denominator (inner products may go to zero, but by assumption norms do not). All that's left to show is convergence of the linearly dependent a_i not among the b's. But note that since $a_i 6(n)$ are always linearly dependent, that means S' never acts as identity on them, as that would make them independent. In particular, S' acts as projection, a continuous function. Thus the convergence of the $b_i^{(n)}$ immediately imply convergence of the a_i. Thus every one of the endpoints converges. Thus each rotated line converges, giving us the result. □

Corollary 5. *The result holds if we replace the symmetrization operations with projections in the same directions.*

Proof. Symmetrization and projection on line segments (points) differ only when the operation occurs in the direction of the segment (point) in question. Projecting in this direction simply kills the segment (point) by collapsing it to the origin. Since there are only finitely many segments (points) in question, this can only happen finitely many times, after which we can replace the projections with symmetrizations and apply Theorem 4 to get the result. □

Corollary 6. *Theorem 4 holds without the assumption of centrality.*

Proof. As we've seen by Proposition 2, we can view a line segment as a centralized segment translated by the midpoint of the original segment. Let $L_1, ..., L_k$ denote the not necessarily centralized segments. Let $a_1, ..., a_k$ denote their midpoints, and $l_1, ..., l_k$ the centralized segments, so that $L_i = a_i + l_i$. If we symmetrize in the direction of an a_i, it collapses to the origin and the L_i becomes centralized. This can happen only finitely many times. So without loss of generality, we assume it never happens. But then S_{u_n} applied to a_i is equivalent to S'_{u_n}, so we can treat a_i like the endpoint of a centralized line segment a'_i, and then apply Theorem 4 to $a'_1, ..., a'_k, l_1, ..., l_k$ to get convergence. □

Bibliography

[1] Gabriele Bianchi, Almut Burchard, Paolo Gronchi, and Aljosa Volcic. Convergence in shape of steiner symmetrizations. *Indiana Univ. Math. J.*, 61:1695–1710, 2012.

[2] Almut Burchard and Gregory Chambers. Perimeter under multiple steiner symmetrizations. *Journal of Geometric Analysis*, pages 1–12, 2013.

[3] Peter M. Gruber. *Convex and Discrete Geometry.* Springer, 2007.