Six theorems about injective metric spaces

by J. R. Isbell, University of Washington (Seattle)

Introduction

A metric space Y is *injective* if every mapping which increases no distance from a subspace of any metric space X to Y can be extended, increasing no distance, over X. Aronszajn and Panitchpakdi showed [1] that topologically, every injective metric space is a complete absolute retract, and asked whether the converse is true. It is obviously true in 1-dimensional spaces. But in 2-dimensional spaces there are additional necessary conditions. First, every injective metric space can be contracted to a point *freely*, i. e. by a path $\{h_t\}$ of decreasing deformation retractions. Conversely, for 2-dimensional finite polyhedra, this condition is sufficient. It is equivalent (for any triangulation) to *collapsibility* in the sense of Whitehead [5]. In infinite 2-dimensional polyhedra, collapsibility is sufficient and free contractibility necessary, and it may be that these properties are (still) equivalent.

Second topological necessary condition: a locally compact injective metric space is locally triangulable at every homotopically stable point (in the sense of Hopf and Pannwitz [4]).

Three geometric theorems. (1) Every metric space X has a smallest containing injective *envelope* εX, which is compact if X is compact. (2) A compact injective space Y has a *boundary*, the smallest closed subset B such that $\varepsilon B = Y$. (3) An n-dimensional compact injective space has at least $2n$ boundary points and has injective n-dimensional subspaces with exactly $2n$ boundary points. Those subspaces may be chosen to be isometric copies of closed cells in n-dimensional l_∞ space.

I am indebted to T. Ganea and to W. B. Woolf for some conversations concerning this material.

1. Polyhedra

By a *mapping* between metric spaces we mean a function $f: X \to Y$ such that for all x, x' in X, the distance $d(f(x), f(x')) \leq d(x, x')$. Y is an *injective* metric space if every mapping from a subspace of any space X to Y can be extended (to a mapping) over X. Aronszajn and Panitchpakdi introduced these spaces [1], calling them *hyperconvex* because of the characterizations which follow.

It suffices to require that every metric space consisting of Y and one more point can be retracted upon Y. This reduces to the condition that any set of closed solid spheres $S(y_\alpha, r_\alpha)$ in Y such that for each α and β, $r_\alpha + r_\beta \geq \geq d(y_\alpha, y_\beta)$, has a common point. Equivalently, (a) Y is convex (any 2 sufficiently large solid spheres meet) and (b) a collection of solid spheres has a common point if every 2 of its members have a common point.

Note that every solid sphere in an injective space is an injective subspace.

A *free deformation retraction* of a topological space X upon a subspace A is a homotopy $\{h_t\}$, $t \in [0,1]$, such that $h_1 : X \to X$ is the identity, $h_0 : X \to X$ is a retraction upon A, and every composition $h_s h_t$ is h_r where $r = \min(s, t)$. A free deformation retraction to a point is a *free contraction.*

1.1. Theorem. *An injective metric space is freely contractible to each of its points.*

Proof. Let p be a point of the injective space Y. We construct a free contraction $\{h_t\}$ with each h_t retracting Y upon $S(p, t)$, such that $d(h_s(x), h_t(y)) \leq \max(d(x, y), |s - t|)$, using Zorn's Lemma. With this prescribed modulus of continuity, we need only show that when $\{h_t\}$ is already defined on a subspace Z of Y, and q is a point not in Z, $\{h_t\}$ can be extended over a subspace containing q also. We may suppose Z is closed; and we may confine attention to the solid sphere $S = S(p, u)$ just large enough to include q. On S, $h_t(x) = x$ for $t \geq u$. The non-trivial homotopy $h : (Z \cap S) \times [0, u] \to Z \cap S \subset S$ can be extended to a contraction $j : S \times [0, u] \to S$, not free but having the prescribed modulus of continuity. Here $j(\{q\} \times [0, u])$ must be a shortest path J from q to p; so j yields a free contraction defined on Z and on the part of J from q to where J first meets Z.

In this paper, a *polyhedron* is a topological space which is the body of a finite-dimensional simplicial complex, with the metric topology induced by defining the distance between two points as the maximum difference in their barycentric coordinates. We remark that, since the complexes are finite-dimensional, all reasonable distance functions give the same topology; and some routine details concerning the global treatment of infinitely many simplexes will be omitted below (1.6).

A finite simplicial complex K is called *collapsible* [5] if it can be built up from a point by successive adjunctions of single closed simplexes s such that s intersects the complex so far constructed exactly in all but one of its maximal proper faces. The ordered set of simplexes s, in the order of their adjunction, beginning with a vertex s_0, is an *expansion* of K from s_0. For infinite complexes, the definitions are the same, with the provision that the adjoined simplexes are well-ordered.

By finite combination of steps one can show, at least in dimensions <4, that every complex built up from a point by attaching simplexes s each by a contractible subcomplex of its boundary is collapsible. On the other hand, since the elementary steps never add a vertex except when a 1-simplex is adjoined, one shows by a trivial induction that a collapsible simplicial complex can be expanded from any vertex.

We define a *collapsible cubical 2-complex* as a (cell) complex built up from a point by successive adjunction of edges and 2-cubes so that (i) the intersection of any two of these cells is a common face or the empty set; (ii) no three 2-cubes $abcd$, $ade\mathfrak{f}$, $a\mathfrak{f}\mathfrak{g}b$ occur; and (iii) each edge is attached to its predecessors by one vertex, each 2-cube by two adjacent edges. Again, we could admit 2-cubes attached by one edge or one vertex, by combining steps. (It is not clear whether the concept generalizes usefully to 3-complexes composed of cubes, or of octahedra, or not at all.)

The next proposition is essentially outside the main argument, though it can be used in proving 1.3.

1.2. *A 2-dimensional polyhedron freely contractible to one of its points is freely contractible to each of its points.*

Details will be omitted; the special feature of 2-polyhedra is that every arc is tame. In particular, every arc is a free deformation retract of a neighborhood of itself. Then to change a free contraction upon p to a free contraction upon q, consider the path followed by q in the contraction upon p. It is a monotone continuous image of an arc; hence it is an arc J. Some neighborhood U can be deformed freely upon J, and some neighborhood V of q is contracted to p within U. Then using a real-valued continuous function that is 0 at q and 1 outside V, damp the contraction so that q does not move. One still has a free deformation retraction into U, and the rest is obvious.

1.3. *If a freely contractible triangulated 2-polyhedron P consists of two subcomplexes Q, R, intersecting in a point or arc, then Q and R are freely contractible.*

The proof is omitted. 2-dimensionality is not needed.

1.4. *A freely contractible 2-dimensional finite polyhedron is collapsible in any triangulation.*

Proof. Such a polyhedron P must have either a free vertex (lying on exactly one edge) or a free edge (lying on exactly one triangle); this is clear from considering a small open set in which the free contraction differs from the identity as soon as possible. Now if P is a closed simplex, it is collapsible (in that triangulation). Inductively we may suppose every freely contractible proper subcomplex of P is collapsible. We shall be done when we find a vertex or an edge-path separating P into subcomplexes Q, R, for they are freely contractible by 1.3, so collapsible by inductive hypothesis, and in particular collapsible to a vertex of $Q \cap R$. In case there is a free vertex, the vertex joned to it by an edge separates P. Otherwise there is a free edge ab; it lies on one triangle abc, and some subcomplex K of the pair of edges ac, bc separates P. If K is connected it is a point or arc. If K is disconnected, its components are points or arcs, and each component separates because P is simply connected.

I do not know whether 1.4 holds for infinite polyhedra.

1.5. *A collapsible 2-dimensional simplicial complex can be subdivided to a collapsible cubical 2-complex.*

Proof. We define the subdivision by induction relative to some expansion $\{s_\alpha\}$. The 1-simplexes s_α will not be subdivided. Each triangle s_α will be subdivided into a number of quadrilaterals, with new vertices occurring only on the edge e_α by which s_α is not attached to its predecessors. There may be finitely many new vertices v_j previously introduced on the other edges of s_α. From each v_j, and from the (old) vertex opposite e_α, draw two new edges to e_α, none of these edges meeting except at their origins v_j. This subdivides s_α into several quadrilaterals and triangles; make each triangle into a quadrilateral by introducing a new vertex on e_α. Clearly the resulting cubical complex is collapsible.

1.6. Remark. To draw topological conclusions from 1.5 (which we mean to do), one should define a standard metric on these cubical complexes (e.g. as in 1.7) and add some details to the proof of 1.5.

The next construction involves a standard 2-cube Q which it is convenient to present as the square in l_∞ space spanned by the four points $(\pm 1,0)$, $(0,\pm 1)$. We may mention the center $(0,0)$ and the 1-skeleton Q^1 (the boundary) and the 0-skeleton Q^0 (the four vertices). Note that every two points of Q^0 are joined by a segment in Q^1.

1.7. *Every collapsible cubical 2-complex admits an injective metric.*

Proof. Metrize the complex L so that each edge is a segment of length 1 and each 2-cube a copy of Q; define the distance between two points not in a common cell as the length of the shortest path joining them. Then L is at least a convex metric space. It will be convenient to note that the 1-skeleton L^1 is an even graph in which any two vertices are joined by a segment; thus the sum of the distances among any triple of vertices is an even integer. Further preliminaries: a subspace of L isometric with Q^0 lies in L^0 if three of its points are in L^0 (even if one of its points is in L^0, but we do not need that); and it is then the 0-skeleton of a 2-cube of L. (Easy inductions prove all this.)

We must prove that a collection of solid spheres has a common point if every two of its members have a common point. Suppose first that L is a finite complex. Then it suffices to consider spheres with rational radii. More: since regular rectangular subdivision of all cells of L with mesh $1/n$ yields again a collapsible cubical complex, homothetic to L, it suffices to consider spheres whose centers are vertices and whose radii are integers. We shall prove by induction that when such spheres meet pairwise, they have a common point which is a vertex or the center of a 2-cube.

L, being collapsible and finite, consists of a last cell q attached by half of its boundary to a collapsible subcomplex M. As the induced metric on M agrees with the metric defined by applying the present construction to M, the inductive method is applicable. The inductive step is trivial if q is 1-dimensional. Then suppose q is a 2-cube $abcz$, attached to M by ab and bc. Given an integer-valued function f on some of the vertices of L, satisfying $f(x) + f(y) \geq d(x, y)$, we want a vertex or center within $f(x)$ of every x. Clearly it exists in case $f(z) = 0$.

Consider the case $f(z) \geq 2$. The sphere $S(z, f(z))$ meets M just in the union of $S(a, f(z) - 1)$ and $S(c, f(z) - 1)$. Replacing $S(z, f(z))$ with either of these subsets of it, we get a family of spheres in M which would have a common point (vertex or center) if every two of them met. We may suppose, then, that there are vertices u and v with $d(u, c) > f(u) + f(z) - 1$ and $d(v, a) > f(v) + f(z) - 1$. Of course $d(u, a)$ and $d(v, c)$ are smaller, and therefore smaller by 2. Then there is a vertex or center m within distance $d(u, a)$ of u, $d(v, c)$ of v, and 1 of b, in view of $d(u, a) = d(u, c) - - 2 \geq f(u) + f(z) - 2 \geq f(u)$. Because of the large distances $d(u, c)$, $d(v, a)$, the distances of m from u, v, and b are exactly the numbers indicated, and $d(m, a) = d(m, c) = 2$.

If m is not a vertex, we can replace it by a vertex. For m is the center of a 2-cube $brst$ and is closer to u and v than b is. $brst$ has a vertex even closer to u, which can only be s; we get $d(u, s) = d(u, b) - 2$ and likewise $d(v, s) = d(v, b) - 2$. Then r is exactly $d(u, a)$ from u, $d(v, c)$ from v, and 1 from b.

There is a vertex or center x within $d(u, a) - 1$ of u, within 1 of a, and within 1 of m. It follows that $\{m, x, a, b\}$ is a copy of Q^0 and thus that there is a 2-cube $mxab$ in L. Similarly there is a 2-cube $mycb$ in L. With $abcz$, this violates condition (ii) of the definition of a collapsible complex.

In case $f(z) = 1$, consider the three possible points in $S(z, 1)$, namely a, c, and the center h of $abcz$. Every $S(x, f(x))$ includes at least one of them. The subset of $\{a, h, c\}$ lying in $S(x, f(x))$ is order-convex in this order, i.e. it is not $\{a, c\}$. If it were, we would get an impossible 2-cube $abcz'$ from the requirements that z' is within 1 of a and c and within $f(x) - 1$ of x. (Condition (i) of the definition is violated in that case.)

If $S(u, f(u))$ meets $S(z, 1)$ in a and $S(v, f(v))$ meets $S(z, 1)$ in c, the argument of the case $f(z) \geq 2$ can be repeated, for we still have $d(u, a) = d(u, c) - 2 = f(u)$ and the corresponding conditions on v. The remaining cases to consider are $\{a\}$ and $\{h, c\}$ and the similar case $\{a, h\}$ and $\{c\}$; it will suffice to treat the first of them. Then we have u closest to a, v closest to c, $f(u) = d(u, a)$, $f(v) = d(v, c) + 1$. There is a center or vertex m within $f(u)$ of u, $f(v)$ of v, 1 of a and 1 of b; and these distances are exact because of $d(u, b)$ and $d(v, a)$. Since m is equidistant from a and b, it is not a vertex but the center of a 2-cube $abnx$. Since v is equidistant from m and b, and further from a, it is closer to n; $d(v, n) = f(v) - 1$. Then there is y within $f(v) - 2$ of v, 1 of n, and 1 of c; and these distances are exact. This makes $cbny$ a 2-cube of L, violating (ii) of the definition, and proving 1.7 for finite complexes.

For the general case, there is a finiteness lemma.

In the 1-skeleton of a collapsible cubical 2-complex metrized as above, any two vertices are joined by only finitely many shortest paths.

The maximum number is the maximum number of maximal chains between two plane lattice points. To prove merely the italicized assertion, it suffices to show that if a and b are vertices at distance $n + 1$ there cannot be three vertices c, d, e at distance 1 from a and n from b. If there were, there would be a vertex or center of a 2-cube $\mathfrak{f}$ at distance 1 from each of c, d, e and $n - 1$ from b. Then $\{a, c, d, \mathfrak{f}\}$, $\{a, c, e, \mathfrak{f}\}$, $\{a, d, e, \mathfrak{f}\}$ would all be 0-skeletons of 2-cubes, any two of which have too many common faces.

To apply this, we want two more lemmas. First, it suffices to establish the intersection property for spheres with integral radii centered at vertices. That will imply, by subdivision as before, that any set of conditions $d(p, x) \leq f(x)$ (where $f(x) + f(y) \geq d(x, y)$) can be satisfied to within an arbitrarily small error $\varepsilon > 0$. To reduce the error to 0, find p_1 with error ε_1; adjoin $d(p, p_1) \leq \varepsilon_1$ to the conditions; find p_2 satisfying all these conditions to within half as large an error, and so on to the limit.

Second, if a family of spheres meets pairwise then every finite subfamily meets. Perhaps the simplest way to prove this is to use the sublemma: if $\mathfrak{s} = \{q_0, q_1, \ldots\}$ is an expansion of L and $\mathfrak{t} = \{q_0, q_{\alpha_1}, \ldots\}$ is a subsequence of $\mathfrak{s}$ and an expansion of a subcomplex M, then $\mathfrak{t}$ followed by the remainder of $\mathfrak{s}$ in the given order is an expansion of L. Then it is clear that M is an injective subspace. As every finite set of points of L lies in a unique smallest such M, which is finite, the intersection property follows.

Now given a function f on the vertices of any collapsible cubical 2-complex L such that $f(x) + f(y) \geq d(x, y)$, replace f by a minimal function satisfying these inequalities; we shall still call it f. For each x, there is y such that $f(x) + f(y) = d(x, y)$, since f is minimal and integral. From the finiteness lemma, the set $H = S(x, f(x)) \cap S(y, f(y))$, for any such x and y, is compact. The traces of the other spheres on H are a family of closed sets having the finite intersection property; so the total intersection is not empty.

The principal conclusion:

1.8. Theorem. *A 2-dimensional finite polyhedron is injectively metrizable if and only if it is freely contractible, and this is if and only if it is collapsible* (*in any triangulation*).

It may be that this generalizes to infinite polyhedra. Straightforward combinatorics prove that a simplicial complex is collapsible if and only if every finite subcomplex is in a collapsible finite subcomplex (by means of the lemma: any expansion of a subcomplex of a collapsible complex K is an initial segment of an expansion of K). Perhaps straight forward, delicate simplicial approximation will prove a corresponding reduction for free contractibility, and that would complete the generalization.

It may be that $1 \cdot 8$ generalizes to arbitrary polyhedra, but the present results scarcely suffice to suggest such a conjecture.

2. Envelope and boundary

We call a mapping of metric spaces $e: X \to E$ an *injective envelope* of X if E is injective, e is an isometric embedding, and no injective proper subspace of E contains $e(X)$. Two injective envelopes $e: X \to E$, $f: X \to F$ are *equivalent* if they are related by an isometry $i: E \to F$.

2.1. Theorem. *Every metric space has an injective envelope and all of its injective envelopes are equivalent.*

Proof. We define an *extremal* function on X as a real-valued function f which is pointwise minimal subject to

$$f(x) + f(y) \geq d(x, y) \tag{2.2}$$

for all x, y in X. Then f also satisfies

$$f(x) + d(x, y) \geq f(y) \tag{2.3}$$

for all x and y. If this were false, one would define g to coincide with f except at y, where $g(y) = f(x) + d(x, y)$. By the triangle inequality, g satisfies (2.2); as $g \leq f$, we must conclude $g = f$.

Therefore the difference between any two extremal functions f, g is bounded; any number $f(x) + g(x)$ is a bound. Thus th› set εX of all extremal functions on X is a metric space with $d(f, g) = \sup | f(x) - g(x) |$. An isometric embedding $e : X \to \varepsilon X$ is defined by $e(x)(y) = d(x, y)$.

By (2.3), every extremal function is continuous; in fact, all extremal functions are equicontinuous. As also every limit of extremal functions is extremal, hence εX is compact if X is compact.

(2.2) and (2.3) together are equivalent to

$$f(x) = d(f, e(x)) \tag{2.4}$$

for all x.

2.5. Every function satisfying (2.2) is greater than or equal to some extremal function.

2.6. If X is compact, then for any f in εX and x in X there is, by minimality, some y in X such that $f(x) + f(y) = d(x, y)$. In general we have only $f(x) + f(y) < d(x, y) + \delta$, where δ is any positive number and y depends on δ.

2.7. *If s is an extremal function on the metric space εX, then se is extremal on X.*

Proof. Suppose the contrary. We get $h \in \varepsilon X, h \leq se, h(x) < se(x)$. Define t on εX by $t(f) = s(f)$ except at $e(x)$; $te(x) = h(x)$. To show t satisfies (2.2), it suffices to show

$$te(x) + t(f) \geq d(f, e(x)) \tag{2.8}$$

for all f in εX (as t agrees with s elsewhere in εX). For any $\delta > 0$, pick a y such that $f(x) + f(y) < d(x, y) + \delta$. If $y = x$ or $f = e(x)$, then (within an error of δ) 2.8) holds. Otherwise $te(x) + te(y) = h(x) + se(y) \geq$ $\geq h(x) + h(y) \geq d(x, y) > f(x) + f(y) - \delta = d(f, e(x)) + f(y) - \delta$. Moreover, since t coincides with s at f and at $e(y)$, (2.4) and (2.3) imply $t(f) + f(y) \geq te(y)$. Adding, $te(x) + te(y) + t(f) + f(y) > d(f, e(x)) + f(y) -$ $- \delta + te(y)$. Since δ is arbitrary, the proof is complete.

2.9. *εX is injective.*

Proof. We use the criterion from [1]; any closed spheres $S(f_\alpha, r_\alpha)$ such that always $r_\alpha + r_\beta \geq d(f_\alpha, f_\beta)$ must have a point in common. We may suppose r is a function defined on all of εX, satisfying (2.2). Let s be an extremal function $\leq r$. Then se belongs to every $r(f)$-sphere about f. In fact $se(x) - f(x) = se(x) - d(f, e(x)) \leq s(f)$ for each x, by (2.3); and $f(x) - - se(x) = d(f, e(x)) - se(x) \leq s(f)$, by (2.2).

2.10. *$e: X \to \varepsilon X$ is an injective envelope of X and is equivalent to every injective envelope of X.*

Proof. A mapping of εX into itself leaving X pointwise fixed must take each f to some g such that $d(g, e(x)) = g(x) \leq f(x)$ for all x; thus it is the identity. Then εX cannot be retracted upon any proper subset S containing X; S is not injective. Finally, for any injective envelope $f: X \to F$, f can be extended over εX and e can be extended over F. The composed mapping $\varepsilon X \to F \to \varepsilon X$ is the identity. Hence $\varepsilon X \to F$ is an isometry upon its image. Hence the image is injective; so it is all of F. This completes the proof of 2.10 and of 2.1.

The proof has shown also

2.11. *The injective envelope of a compact space is compact; the injective envelope of a finite space is a polyhedron.*

We define an *end point* of a compact metric space X containing more than one point as a point x such that for some point y the equation $d(w, x) + + d(x, y) = d(w, y)$ implies $w = x$. (This implies $y \neq x$).

2.12. *If x is an end point of X, then x is an end point of the injective envelope εX, and every closed subset of εX not containing x lies in an injective subspace of εX not containing x.*

Proof. The definition gives us a certain point y of X. If x were between y and w in εX, we would apply 2.6 to get u in X such that w is between u and y. Then $d(u, x) + d(x, y) = d(u, y)$, a contradiction. If H is a closed subset of εX not containing x, so is $K = H \cup \{y\}$. The embedding of K in εX can be extended to an embedding of εK in εX; but x cannot be in εK, for the function $d(x, k)$ on K is not extremal.

This proposition is vacuously true in a space of 0 points, false in a space of 1 point. It is true for non-compact spaces if we define an end point x by $(\alpha)(\exists\, y)(\exists\, \beta)[\alpha > 0 \Longrightarrow \beta > 0] \wedge [d(w, x) + d(x, y) < d(w, y) + \beta \Longrightarrow \Longrightarrow d(w, x) < \alpha]$. However, the notion of end point is less interesting in non-compact spaces, because the following theorem fails.

2.13. Theorem. *In a compact injective space Y containing more than one point, the closure B of the set of end points is the smallest closed subset which is not contained in any injective proper subspace of Y.*

Proof. A closed set lying in no injective proper subspace must contain B, by 2.12. Supposing εB to be a proper subspace of Y, let y be a point not in εB and consider $f(x) = d(x, y)$ on εB. By the triangle inequality, f satisfies (2.2). Since y is not an end point, for each x in εB there is z in Y such that y is between x and z; choosing z at maximum distance past y, z also is in εB because it is an end point. Hence f is extremal; $y \in \varepsilon B$.

This set B will be called the *boundary*.

3. Some other results

The following remark will presumably be of central importance in any general theory of injective metric spaces.

3.1. Remark. Let Y be an injective metric space, and S a subspace such that every point of Y is within δ of some point of S; then Y contains εS (by a non-unique embedding) and there is a retraction $r: Y \to \varepsilon S$ which moves no point more than δ. If Y is compact, S can be taken to be finite.

A converse:

3.2. *A complete metric space Y is injective if for every $\varepsilon > 0$ there is an injective subspace S of Y such that every point of Y is within ε of some point of S.*

Proof. Let f be an extremal function on Y. Select injective subspaces S_n coming within ε_n of every point of Y, where $\Sigma \varepsilon_n < \infty$. There is p_n in S_n within $f(s)$ of each point s of S_n; hence p_n is within $f(y) + 2\varepsilon_n$ of every point y of Y. Then since f is extremal, $f(p_n) \leq 2\varepsilon_n$, and the points p_n form a CAUCHY sequence converging to the required point p.

Recall next that a point x is *homotopically labile* [4] if for every $\varepsilon > 0$ there is a deformation of the identity mapping to a non-onto mapping, with the ε-neighborhood of x deformed in itself and the rest of the space remaining pointwise fixed. The weaker requirement that no point moves more than ε (rather, $\varepsilon/2$) has the same effect, by an obvious damping argument. There is a stronger requirement in which "non-onto" becomes "omitting the value x"; if this is satisfied, we call x *freely labile.* A non-labile point is *homotopically stable;* a point that is not freely labile is *weakly stable*, or "stable in the sense of BORSUK and JAWOROWSKI" [2].

One sees at once

3.3. *Every end point of an injective metric space is freely labile. Thus in a compact injective metric space, every boundary point is homotopically labile.*

On the other hand, every finite-dimensional separable metric space has weakly stable points [3]. This shows

3.4. *A finite-dimensional compact injective space containing more than one point is not topologically homogeneous.*

Even in the infinite-dimensional case, a compact injective space cannot be geometrically homogeneous. If the diameter is (for convenience) $2 = d(x, y)$, there is a point z distant no more than 1 from any point, and no autoisometry can take x to z. However, it is not clear how one can find structure in the inhomogeneity. Easy examples show that every point may be a boundary point. I do not know whether every point can be an end point.

Applying the fundamental remark 3.1, we get

3.5. Theorem. *A locally compact injective metric space is locally triangulable at each homotopically stable point.*

Proof. Let x be a point of the injective space Y having a compact neighborhood N. We may suppose N is injective but not locally triangulable at x. Then for every $\varepsilon > 0$, there is a finite subset of N coming within ε of every point of N, and there is an injective polyhedron $P \subset N$ coming within ε of every point of N. There is a deformation retraction of N upon P which moves no point more than ε; as x cannot be interior to P, points arbitrarily near x are uncovered. We can modify the deformation retraction to a deformation of the identity to a non-onto mapping, affecting only the 2ε-neighborhood of x. Thus x is labile in N and in Y.

The theorem leaves something to be desired, particularly since there need not be any stable points. The proof establishes a trifle more than was stated. However, it is easy to see that the stronger statement that the space must be locally triangulable except at boundary points is not a theorem.

An n-dimensional compact injective space admits ε-deformations upon its subpolyhedra for all $\varepsilon > 0$, and therefore contains n-dimensional polyhedra – in particular, it contains n-cells. Evidently an n-dimensional locally compact injective space has n-dimensional compact injective subspaces, so that these spaces also contain n-cells. We can say more about some of these n-cells. By construction, they occur in injective envelopes of finite sets of points $x_1, \ldots, x_m$. If the set of extremal functions on $\{x_1, \ldots, x_m\}$ is n- limensional, there must be an extremal function f for which there are n linearly independent functions g such that $f + \lambda g$ is extremal whenever $|\lambda| < 1$. There are not less than $m/2$ constraints $f(x_j) + f(x_k) = d(x_j, x_k)$ (since each x_j occurs in one, by 2.6); and they imply $g(x_j) = -g(x_k)$. Moreover, these are the only constraints on g near 0. If n of the variables $g(x_{j_i})$ are independent, there are n other variables $g(x_{k_i}) = -g(x_{j_i})$. But this means that the restriction of f to a subset of $2n$ points is extremal and there are n degrees of freedom for extremal functions near it. A neighborhood of f is isometric with an open set in the l_∞ space of all functions on the set $\{x_{j_i}\}$.

3.6. Theorem. *An n-dimensional locally compact injective metric space contains n-cells, some of which are injective envelopes of sets of* $2n$ *points and are isometrically embeddable in n-dimensional* l_∞ *space.*

An n-dimensional compact injective metric space has at least $2n$ *boundary points.*

REFERENCES

[1] N. Aronszajn and P. Panitchpakdi: *Extension of uniformly continuous transformations and hyperconvex metric spaces.* Pacific J. Math. 6 (1956), 405–439.

[2] K. Borsuk and J. Jaworowski: *On labil and stabil points.* Fund. Math. 39 (1952), 159–175.

[3] T. Ganea: *Homotopically stable points.* Com. Acad. R. P. Romîne 9 (1959), 1113–1116.

[4] H. Hopf and E. Pannwitz: *Über stetige Deformationen von Komplexen in sich.* Math. Ann. 108 (1933), 433–465.

[5] J. H. C. Whitehead: *Simplicial spaces, nuclei, and m-groups.* Proc. London Math. Soc. (2) 45 (1939), 243–327.