FRAÏSSÉ LIMITS OF METRIC STRUCTURES

ITAÏ BEN YAACOV

ABSTRACT. We develop *Fraïssé theory*, namely the theory of *Fraïssé classes* and *Fraïssé limits*, in the context of metric structures. We show that a class of finitely generated structures is Fraïssé if and only if it is the age of a separable approximately homogeneous structure, and conversely, that this structure is necessarily the unique limit of the class, and is universal for it.

We do this in a somewhat new approach, in which "finite maps up to errors" are coded by *approximate isometries*.

INTRODUCTION

The notions of Fraïssé classes and Fraïssé limits were originally introduced by Roland FRAÏSSÉ [Fra54], as a method to construct countable homogeneous (discrete) structures:

(i) Every Fraïssé class $\mathcal{K}$ has a Fraïssé limit, which is unique (up to isomorphism). The limit is countable and ultra-homogeneous (or, in more model-theoretic terminology, quantifier-free-homogeneous).

(ii) Conversely, every countable ultra-homogeneous structure is the limit of a Fraïssé class, namely, its *age*.

Moreover, the limit is universal for countable $\mathcal{K}$-structures, namely for countable structures whose age is contained in $\mathcal{K}$.

Similar results hold for metric structures as well. Indeed, some general theory of this form is discussed in the PhD dissertation of Schoretsanitis [Sch07]. Independently, Kubiś and Solecki [KS] treated the special case of the class of finite dimensional Banach spaces, essentially showing that their Fraïssé limit is the Gurarij space, which is therefore unique and universal, without ever actually uttering the phrase "Fraïssé limit" (and in a fashion which is very specific to Banach spaces). This multitude of somewhat incompatible approaches, reinforced by considerable nagging from Todor Tsankov convinced the author of the potential usefulness of the present paper.

There is one main novelty in the present treatment, compared with earlier treatments of back-and-forth arguments in the metric setting, in that we replace partial maps with *approximate isometries* (which is just a fancy term for bi-Katětov maps). These allow us to code in a single, hopefully natural, object, notions such as a partial isometry between metric spaces, or even a "partial isometry only known up to some error term $\varepsilon > 0$". On a technical level, approximate isometries are easier to manipulate than, say, partial isometries, and can be freely composed without loss of information. More importantly, their use simplifies arguments and dispenses with the need for several limit constructions at several crucial points:

- In the back-and-forth argument. The reader is invited to compare the proof of Theorem 2.19, which is hardly distinguishable from the argument for discrete structures, with "traditional" arguments for metric structures, involving the construction of partial isomorphisms which only extend each other up to some error, as in the proofs of Facts 1.4 and 1.5 of [BU07].
- When checking that a structure is a Fraïssé limit, e.g., when proving that such exists, or when proving that the Gurarij space is the limit of finite-dimensional Banach spaces (Theorem 3.3). Indeed, approximate isometries allow us to define a Fraïssé limit in a manner which is formally weaker than the "traditional approach" definition (namely Corollary 2.20(iv)). The limit constructions required to pass from the weaker definition to the stronger one are then entirely subsumed in the back-and-forth argument referred to above.

Of course, some preliminary work is required in order to develop these tools. However, once this is done, many arguments in metric model theory, not only those present here, can be simplified significantly, so we consider this is worth the effort. In addition, approximate isometries are essential for a generalisation of metric Fraïssé theory, to appear in a subsequent paper, in which the limit is only unique up to arbitrarily small error (e.g., a Banach space which is almost isometrically unique).

1. Approximate isometries

Finite partial isomorphisms between structures play a crucial role in classical Fraïssé theory. For example, homogeneity and uniqueness of the Fraïssé limit are proved using a back-and-forth argument, in which finite partial maps serve as better and better approximations for a desired global bijection. In the metric setting, one may expect finite partial isometries to play a similar role, coding partial information regarding a desired global isometry. However, this analogy fails, essentially on the grounds that whereas finite maps define neighbourhoods of global bijections (in the topology of point-wise convergence), finite isometries do not define neighbourhoods of global isometries. In order to define an open set of isometries we need to restrict to a finite set *and allow for a small error*: if $g\colon X \dashrightarrow X$ is a finite partial isometry and $\varepsilon > 0$, then $\big\{h \in \mathrm{Iso}(X) : hx \in B(gx, \varepsilon) \text{ for all } x \in \mathrm{dom}\, g\big\}$ is open and such sets form a basis for the point-wise convergence topology on $\mathrm{Iso}(X)$.

Another deficiency of partial isometries arises when considering compositions. Say $f\colon X \dashrightarrow Y$ and $g\colon Y \dashrightarrow Z$ are partial isometries, such that $\mathrm{img}\, f \cap \mathrm{dom}\, g = \varnothing$, and say $x \in \mathrm{dom}\, f$ is such that fx is very close to some $y \in \mathrm{dom}\, g$. Then we should like to say that gfx is very close to gy, but the composition gf is empty and cannot code this information.

In order to remedy either problem we require a more flexible object than a partial isometry, which can say where an element goes, more or less, without having to say exactly where. These objects will serve us mostly as approximations of actual isometries, whence their name. The reader may wish to compare with the treatment of bi-Katětov functions in Uspenskij [Usp08].

Definition 1.1. Let X, Y and Z denote metric spaces.
  (i) We say that a function $\psi\colon X \to [0, \infty]$ is *Katětov* if for all $x, y \in X$ we have $\psi(x) \le d(x, y) + \psi(y)$ and $d(x, y) \le \psi(x) + \psi(y)$. Unlike Uspenskij (and Katětov) we allow the value ∞, observing that a Katětov function is either finite or constantly ∞.
  (ii) We say that $\psi\colon X \times Y \to [0, \infty]$ is an *approximate isometry* from X to Y, and write $\psi\colon X \rightsquigarrow Y$, if it is bi-Katětov, i.e., separately Katětov in each argument. The special case $\psi = \infty$ is called the *empty approximate isometry*.
  (iii) Given any $\psi\colon X \times Y \to [0, \infty]$ and $\varphi\colon Y \times Z \to [0, \infty]$ we define a *composition* $\varphi\psi\colon X \times Z \to [0, \infty]$ and a *pseudo-inverse* $\psi^*\colon Y \times X \to [0, \infty]$ by
$$\varphi\psi(x, z) = \inf_{y \in Y} \psi(x, y) + \varphi(y, z), \qquad \psi^*(y, x) = \psi(x, y).$$

An approximate isometry $\psi\colon X \rightsquigarrow Y$ is meant to provide partial information regarding some isometry, possibly between larger spaces. We shall understand ψ as saying that x must be sent within $\psi(x, y)$ of y, so an isometry f is considered to satisfy the constraints prescribed by ψ if $\psi(x, y) \ge d(fx, y)$ for all x, y, i.e., if $f = \psi_f \le \psi$ in the sense of Definition 1.4 below. Accordingly, another $\varphi\colon X \rightsquigarrow Y$ imposes stronger constraints if and only if $\psi \ge \varphi$. The rest of our terminology (coarsening, refinement, etc.) should be understand in the context of this interpretation.

Remark 1.2. Let $\psi\colon X \times Y \to [0, \infty)$ be given, let $Z = X \amalg Y$, and define d_Z extending d_X and d_Y by $d(x, y) = d(y, x) = \psi(x, y)$. Then ψ is bi-Katětov (i.e., an approximate isometry) if and only if d is a pseudo-distance on Z. The reader is advised that, while this interpretation is close to Katětov's original use for such functions, it is quite distant from our intended use, and may therefore be misleading.

Lemma 1.3.   (i) *The composition and pseudo-inverse of approximate isometries are again approximate isometries.*
  (ii) *Composition is associative, and pseudo-inversion acts as an involution:* $\psi^{**} = \psi$, $(\varphi\psi)^* = \psi^*\varphi^*$.

Proof. Let $\varphi\colon X \rightsquigarrow Y$ and $\psi\colon Y \rightsquigarrow Z$. Then for each $x \in X$ and $y \in Y$, the function $z \mapsto \varphi(x, y) + \psi(y, z)$ is Lipschitz with constant 1, and therefore so is $z \mapsto \psi\varphi(x, z) = \inf_y \varphi(x, y) + \psi(y, z)$ for any fixed x. Similarly, for any $x \in X$, $y, y' \in Y$ and $z, z' \in Z$ we have

$$\varphi(x, y) + \psi(y, z) + \varphi(x, y') + \psi(y', z') \ge d(y, y') + \psi(y, z) + \psi(y', z') \ge \psi(y', z) + \psi(y', z') \ge d(z, z'),$$

whence $\psi\varphi(x, z) + \psi\psi(x, z') \ge d(z, z')$. Therefore $\psi\varphi$ is an approximate isometry, and it is clear that φ^* is one as well. The second item is even easier and is left to the reader. ∎

The first examples we give of approximate isometries are simply partial isometries, viewed as instances of the former (modulo some obvious identifications: a partial isometry and its restriction to a dense subset of its domain carry the same information, and indeed induce the same approximate isometry).

Definition 1.4. Let X and Y denote metric spaces.

(i) To a partial isometry $f \colon X \dashrightarrow Y$ we associate an approximate isometry $\psi_f(x,y) = \inf_{z \in \operatorname{dom} f} d(x,z) + d(fz,y)$. We shall mostly ignore the distinction between f and ψ_f, denoting the latter by f as well.

(ii) Let $i \colon X \subseteq X'$, $j \colon Y \subseteq Y'$ isometric embeddings, and let $\psi \colon X \rightsquigarrow Y$. Then $j\psi i^* \colon X' \rightsquigarrow Y'$ is called the *trivial extension* of ψ to $X' \rightsquigarrow Y'$. When there is no risk of ambiguity, we shall identify an approximate isometry with its trivial extension to any pair of larger spaces.

Lemma 1.5. (i) *If f is a partial isometry, then the corresponding ψ_f is an approximate isometry.*

(ii) *The approximate isometry $\infty = \psi_\varnothing$ is destructive for composition, and id_X, identified with $\psi_{\operatorname{id}_X} = d_X$, is neutral.*

(iii) *(Pseudo-)inversion is compatible with the identification of partial isometries with approximate ones. Similarly for composition $\psi_g \psi_f = \psi_{gf}$ when $\operatorname{dom} g \supseteq \operatorname{img} f$ or $\operatorname{dom} g \subseteq \operatorname{img} f$, and for the natural notion of trivial extension of a partial map to larger sets.*

Proof. Left to the reader. ∎

This indeed solves both problems described in the beginning of the section. If $g \colon X \dashrightarrow X$ is a finite partial isometry and $\varepsilon > 0$ then the approximate isometry $g + \varepsilon$ codes "g up to error ε", and $\{h \in \operatorname{Iso}(X) : hx \in B(gx, \varepsilon) \text{ for all } x \in \operatorname{dom} g\}$ is just $\operatorname{Iso}(X) \cap \operatorname{Apx}^{<g+\varepsilon}(X)$ in the sense of Definition 1.6 below. Similarly, in the situation of composition of partial isometries, if $x \in \operatorname{dom} f$ and $y \in \operatorname{dom} g$ then $\psi_g \psi_f$ prescribes that x be sent no more than $(\psi_g \psi_f)(x, gy) = d(fx, y)$ from gy, which is exactly the information we wanted to keep.

Definition 1.6. Let X, Y and Z denote metric spaces.

(i) The space of all approximate isometries from X to Y will be denoted $\operatorname{Apx}(X,Y)$, and equipped with the topology induced from $[0,\infty]^{X \times Y}$. When $X = Y$ we let $\operatorname{Apx}(X) = \operatorname{Apx}(X,X)$.

(ii) For $\psi, \varphi \in \operatorname{Apx}(X,Y)$ we say that $\varphi \leq \psi$ is the comparison holds point-wise, i.e., $\varphi(x,y) \leq \psi(x,y)$ for all $(x,y) \in X \times Y$. We then also say that ψ *coarsens* φ, or that φ *refines* ψ. We define $\operatorname{Apx}^{\leq \psi}(X,Y) = \{\varphi \in \operatorname{Apx}(X,Y) : \varphi \leq \psi\}$.

(iii) We define $\operatorname{Apx}^{<\psi}(X,Y)$ as the interior of $\operatorname{Apx}^{\leq \psi}(X,Y)$ in $\operatorname{Apx}(X,Y)$. If $\varphi \in \operatorname{Apx}^{<\psi}(X,Y)$ we write $\varphi < \psi$ and say that ψ *strictly coarsens* φ, or that φ *strictly refines* ψ.

(iv) For $\mathcal{A} \subseteq \operatorname{Apx}(X,Y)$ we define its *closure under coarsening* $\mathcal{A}^\uparrow = \{\psi \in \operatorname{Apx}(X,Y) : \exists \varphi \in \mathcal{A}, \psi \geq \varphi\}$. We observe that $\overline{\mathcal{A}}^\uparrow = (\overline{\mathcal{A}^\uparrow})^\uparrow$.

Lemma 1.7. (i) *The space $\operatorname{Apx}(X,Y)$ is compact, and the interpretation of actual isometries as approximate isometries yields a topological embedding $\operatorname{Iso}(X) \subseteq \operatorname{Apx}(X)$.*

(ii) *If $\varphi_\alpha \in \operatorname{Apx}(X,Y)$ is a net then $\limsup \varphi_\alpha$, calculated coordinate-wise in $[0,\infty]^{X \times Y}$, belongs to $\operatorname{Apx}(X,Y)$ as well.*

(iii) *Composition is upper semi-continuous, in the sense that the set $\{(\psi, \varphi) : \varphi\psi \in \mathcal{A}\} \subseteq \operatorname{Apx}(X,Y) \times \operatorname{Apx}(Y,Z)$ is closed whenever $\mathcal{A} = \overline{\mathcal{A}}^\uparrow \subseteq \operatorname{Apx}(X,Z)$. Equivalently, $(\limsup \psi_\alpha)(\limsup \varphi_\alpha) \geq \limsup(\psi_\alpha \varphi_\alpha)$.*

Proof. The space $\operatorname{Apx}(X,Y)$ is closed in $[0,\infty]^{X \times Y}$ and therefore compact. A sub-basic open set $U_{x,y,\varepsilon} = \{g : d(gx,y) < \varepsilon\} \subseteq \operatorname{Iso}(X)$ agrees with $\{\varphi : \varphi(x,y) < \varepsilon\} \cap \operatorname{Iso}(X)$ under $\operatorname{Iso}(X) \subseteq \operatorname{Apx}(X)$. Conversely, if $V = \{\varphi : r < \varphi(x,y) < s\} \subseteq \operatorname{Apx}(X)$ and $f \in \operatorname{Iso}(X) \cap V$ then we may assume that $r + \varepsilon < d(fx,y) < s - \varepsilon$ in which case $f \in U_{x,fx,\varepsilon} \subseteq V$. This proves the first item. That $\limsup \varphi_\alpha$ is also an approximate isometry, and that $(\limsup \psi_\alpha)(\limsup \varphi_\alpha) \geq \limsup(\psi_\alpha \varphi_\alpha)$, follow easily from the definitions. The latter, together with $\mathcal{A} = \overline{\mathcal{A}}^\uparrow$, implies that $\{(\psi, \varphi) : \varphi\psi \in \mathcal{A}\}$ is closed. ∎

Lemma 1.8. Let X, Y and Z be metric spaces.

(i) *Let $\psi \in U \subseteq \operatorname{Apx}(X,Y)$, with U a neighbourhood of ψ. Then there exists $\varphi \in U$ such that $\psi < \varphi$. In particular, if $\psi < \varphi$ in $\operatorname{Apx}(X,Y)$ and $V \ni \psi$ is open then there exists $\rho \in \operatorname{Apx}(X,Y) \cap V$ such that $\psi < \rho < \varphi$.*

(ii) *Let $\varphi, \psi \in \operatorname{Apx}(X,Y)$. Then $\varphi > \psi$ if and only if there are finite $X_0 \subseteq X$, $Y_0 \subseteq Y$ and $\varepsilon > 0$ such that $\varphi \geq \psi\restriction_{X_0 \times Y_0} + \varepsilon$. Moreover, in this case there exists $\rho \in \operatorname{Apx}(X_0, Y_0)$ which only takes rational values (on $X_0 \times Y_0$) such that $\psi < \rho < \varphi$.*

Proof. For the first item, we may assume that there are finite sets $X_0 \subseteq X$, $Y_0 \subseteq Y$ and some $\varepsilon > 0$ such that $\varphi \in U$ if and only if $|\varphi(x,y) - \psi(x,y)| < 2\varepsilon$ on $X_0 \times Y_0$. Let $\psi_0 = \psi{\restriction}_{X_0 \times Y_0} \in \mathrm{Apx}(X_0, Y_0)$, and let $\varphi = \psi_0 + \varepsilon \in \mathrm{Apx}(X_0, Y_0) \subseteq \mathrm{Apx}(X, Y)$. Let

$$V = \{\varphi \colon \varphi(x,y) < \psi(x,y) + \varepsilon \text{ on } X_0 \times Y_0\}.$$

Then $\psi \in V \subseteq \mathrm{Apx}^{\leq \varphi}(X, Y)$, so $\psi < \varphi$.

The rest is easy.   ∎

2. Metric Fraïssé limits via approximate maps

Let us start by fixing a few basic definitions.

Definition 2.1. Let $\mathcal{L}$ be denote a collection of symbols, each being either a *predicate symbol* or a *function symbol* and each having an associated natural number called its *arity*. An $\mathcal{L}$-*structure* $\mathfrak{A}$ consists of a complete metric space A, together with,

- For each n-ary predicate symbol R, a continuous interpretation $R^{\mathfrak{A}} \colon A^n \to \mathbf{R}$. It will be convenient to consider the distance as a (distinguished) binary predicate symbol.
- For each n-ary function symbol f, a continuous interpretation $f^{\mathfrak{A}} \colon A^n \to A$. A zero-ary function is also called a *constant*.

If $\mathfrak{A}$ is a structure and $A_0 \subseteq A$, then the smallest substructure of $\mathfrak{A}$ containing A_0 is denoted $\langle A_0 \rangle$, the substructure *generated* by A_0. Its underlying set is just the metric closure of A_0 under the interpretations of function symbols.

An *embedding* of $\mathcal{L}$-structures $\varphi \colon \mathfrak{A} \to \mathfrak{B}$ is a map which commutes with the interpretation of the language: $R^{\mathfrak{B}}(\varphi \bar{a}) = R^{\mathfrak{A}}(\bar{a})$ and $f^{\mathfrak{B}}(\varphi \bar{a}) = \varphi f^{\mathfrak{A}}(\bar{a})$ (in particular, $d^{\mathfrak{B}}(\varphi a, \varphi b) = d^{\mathfrak{A}}(a, b)$, so an embedding is always isometric). A *partial isomorphism* $\varphi \colon \mathfrak{A} \dashrightarrow \mathfrak{B}$ is a map $\varphi \colon A_0 \to B$ where $A_0 \subseteq A$ and φ extends (necessarily uniquely) to en embedding $\langle A_0 \rangle \to \mathfrak{B}$.

Remark 2.2. The definition given here is more relaxed than definitions given in more general treatments of continuous logic, such as [BU10, BBHU08] for the bounded case and [Ben08] for the general (unbounded) case, in that we only require plain continuity (rather than uniform), and no kind of boundedness. Indeed, let us consider the following properties of a map $f \colon X \to Y$ between metric spaces, which imply one another from top to bottom:

(i) The map f is uniformly continuous.
(ii) The map f sends Cauchy sequences to Cauchy sequences (equivalently, f admits a continuous extension to the completions, $\hat{f} \colon \widehat{X} \to \widehat{Y}$). Let us call this *Cauchy continuity*.
(iii) The map f is continuous.

If X is complete then the last two properties coincide, if X is totally bounded then the first two coincide, and if X is compact then all three do. Thus Cauchy continuity is intimately connected with completeness. Similarly, uniform continuity is intimately related with compactness: on the one hand, compactness implies uniform continuity (assuming plain continuity), while on the other hand, uniform continuity of the language is a crucial ingredient in the proof of compactness for first order continuous logic (similarly, in unbounded logic, compactness below every bound corresponds to uniform continuity on bounded sets).

In light of this, and since compactness will not intervene in any way in our treatment, plain continuity on complete spaces will suffice. In situations involving incomplete spaces we shall require Cauchy continuity.

Definition 2.3. We say that a separable structure $\mathfrak{M}$ is *approximately ultra-homogeneous* if every finite partial isomorphism $\varphi \colon \mathfrak{M} \dashrightarrow \mathfrak{M}$ is arbitrarily close to the restriction of an automorphism of $\mathfrak{M}$: for every $\varepsilon > 0$ there exists $f \in \mathrm{Aut}(\mathfrak{M})$ such that $d(\varphi a, f a) < \varepsilon$ for all $a \in \mathrm{dom}\, \varphi$. Equivalently, if $\overline{\mathrm{Aut}(\mathfrak{M})^{\uparrow}} \subseteq \mathrm{Apx}(M)$ contains every (finite) partial isomorphism $\varphi \colon \mathfrak{M} \dashrightarrow \mathfrak{M}$.

Definition 2.4. The *age* of an $\mathcal{L}$-structure $\mathfrak{A}$, denoted $\mathrm{Age}(\mathfrak{A})$, is the class of finitely generated structures which embed in $\mathfrak{A}$.

Metric Fraïssé theory deals with (ages of) approximately ultra-homogeneous separable structures. One could, of course, say that a structure $\mathfrak{M}$ is (precisely, rather than approximately) ultra-homogeneous if every isomorphism of finitely generated substructures extends to an isomorphism, but this would make us lose important examples (e.g., the Gurarij space), and in any case it does not seem that a Fraïssé theory can be developed for this stronger notion. It follows that, whereas classical Fraïssé theory deals with finite partial isomorphism (and their extensions to automorphisms), metric Fraïssé theory must

deal with finite partial isomorphisms "up to some error", which is by no means a new phenomenon in metric model theory.

The standard approach so far in similar situations, say when carrying out back-and-forth arguments (see for example [BU07, Facts 1.4 and 1.5]), involves constructing a sequence of finite partial isomorphisms f_n such that each f_{n+1} only extends f_n up to some allowable error ε_n, keeping $\sum \varepsilon_n$ small. This involves a considerable amount of bookkeeping, limit constructions and other complications. Replacing "partial isometries up to error" with approximate isometries, as suggested in Section 1, we manage to avoid these complications, and the metric Fraïssé theory follows quite effortlessly, in almost perfect analogy with its discrete counterpart.

Definition 2.5. Let $\mathcal{K}$ be a class of finitely generated structures.

(i) By a $\mathcal{K}$-*structure* we mean an $\mathcal{L}$-structure $\mathfrak{A}$ such that $\mathrm{Age}(\mathfrak{A}) \subseteq \mathcal{K}$.

(ii) We say that $\mathcal{K}$ has the *HP (Hereditary Property)* if every member of $\mathcal{K}$ is a $\mathcal{K}$-structure.

(iii) Assume that $\mathcal{K}$ has HP. We say that $\mathcal{K}$ has the *NAP (Near Amalgamation Property)* if for every $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$, finite partial isomorphism $f \colon \mathfrak{A} \dashrightarrow \mathfrak{B}$ and $\varepsilon > 0$ there are $\mathfrak{C} \in \mathcal{K}$ and embeddings $g \colon \mathfrak{A} \to \mathfrak{C}$, $h \colon \mathfrak{B} \to \mathfrak{C}$ such that $d(ga, hfa) \leq \varepsilon$ for all $a \in \mathrm{dom}\, f$, or equivalently, such that (as approximate isometries) $f + \varepsilon \geq h^* g$.

Notice that an age always has HP, and if $\mathfrak{M}$ is approximately ultra-homogeneous then $\mathrm{Age}(\mathfrak{M})$ has NAP as well.

Definition 2.6. Let $\mathcal{K}$ be a class of finitely generated structures with HP, and let $\mathfrak{A}$ and $\mathfrak{B}$ be $\mathcal{K}$-structures. We define $\mathrm{Apx}_1(\mathfrak{A}, \mathfrak{B})$ to be the set of all finite partial isomorphisms $f \colon \mathfrak{A} \dashrightarrow \mathfrak{B}$, and

$$\mathrm{Apx}_{2,\mathcal{K}}(\mathfrak{A}, \mathfrak{B}) = \{ fg : g \in \mathrm{Apx}_1(\mathfrak{A}, \mathfrak{C}) \text{ and } f \in \mathrm{Apx}_1(\mathfrak{C}, \mathfrak{B}) \text{ for some } \mathfrak{C} \in \mathcal{K} \},$$

where composition is in the sense of approximate isometries. Notice that if we allowed $\mathfrak{C}$ to be an arbitrary $\mathcal{K}$-structure we would obtain the same definition, since we can always replace $\mathfrak{C}$ with $\langle \mathrm{img}\, g \cup \mathrm{dom}\, f \rangle$. Finally, following Definition 1.6, define

$$\mathrm{Apx}_{\mathcal{K}}(\mathfrak{A}, \mathfrak{B}) = \overline{\mathrm{Apx}_{2,\mathcal{K}}(\mathfrak{A}, \mathfrak{B})^\uparrow}.$$

Members of $\mathrm{Apx}_{\mathcal{K}}(\mathfrak{A}, \mathfrak{B})$ are called *($\mathcal{K}$-intrinsic) approximate isomorphisms*. When $\mathcal{K}$ is clear from the context we usually drop it.

For $\psi \in \mathrm{Apx}(\mathfrak{A}, \mathfrak{B})$, we define $\mathrm{Apx}^{<\psi}(\mathfrak{A}, \mathfrak{B}) = \mathrm{Apx}(\mathfrak{A}, \mathfrak{B}) \cap \mathrm{Apx}^{<\psi}(A, B)$. We say that ψ is a *strictly approximate isomorphism* if $\mathrm{Apx}^{<\psi}(\mathfrak{A}, \mathfrak{B}) \neq \varnothing$, and let $\mathrm{Stx}(\mathfrak{A}, \mathfrak{B})$ denote the collection of such ψ.

Intuitively, approximate isomorphisms are to partial isomorphisms (between members of $\mathcal{K}$) as approximate isometries are to partial isometries, so in particular every member of $\mathrm{Apx}_1(\mathfrak{A}, \mathfrak{B})$ should then be considered an approximate isomorphism. The reason for taking the two-iterate is that Apx_1 may "miss" some information: for example, it may happen that $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$ are "close", as witnessed by some embeddings $\mathfrak{A} \to \mathfrak{C}$ and $\mathfrak{B} \to \mathfrak{C}$ with close images (which is captured by Apx_2), even though they have no non-trivial common substructure (so Apx_1 sees nothing). We also require $\mathrm{Apx}(\mathfrak{A}, \mathfrak{B})$ to be compact and closed under coarsening (as is $\mathrm{Apx}(A, B)$), whence the definition. *Strictly* approximate isomorphisms are analogous to *finite* partial isomorphisms in the classical setting, in that they do not fix too much information, leaving an open set of possibilities (clearly, $\mathrm{Apx}^{<\psi}(\mathfrak{A}, \mathfrak{B})$ contains the relative interior of $\mathrm{Apx}^{\leq\psi}(\mathfrak{A}, \mathfrak{B})$ in $\mathrm{Apx}(\mathfrak{A}, \mathfrak{B})$, and one can check that in fact, the two agree).

Lemma 2.7. *Let $\mathcal{K}$ be a class of finitely generated structures with HP. Let $\mathfrak{A}$ and $\mathfrak{B}$ be $\mathcal{K}$-structures. Then*

(i) $\mathrm{Apx}(\mathfrak{A}, \mathfrak{B}) = \overline{\mathrm{Stx}(\mathfrak{A}, \mathfrak{B})}$.

(ii) *Every partial isomorphism between $\mathfrak{A}$ and $\mathfrak{B}$ belongs to $\mathrm{Apx}(\mathfrak{A}, \mathfrak{B})$ (see Remark 2.13 below for a converse of this).*

(iii) *If $\psi \in \mathrm{Stx}(\mathfrak{A}, \mathfrak{B})$ then $\mathrm{Apx}_2^{<\psi}(\mathfrak{A}, \mathfrak{B}) \neq \varnothing$.*

Proof. First, let $\varphi \in \mathrm{Apx}(\mathfrak{A}, \mathfrak{B})$. For finite $A_0 \subseteq A$ and $B_0 \subseteq B$ and for $\varepsilon > 0$ we have $\varphi\!\restriction_{A_0 \times B_0} + \varepsilon > \varphi$ by Lemma 1.8(ii), so $\varphi\!\restriction_{A_0 \times B_0} + \varepsilon \in \mathrm{Stx}(\mathfrak{A}, \mathfrak{B})$. It follows that $\varphi \in \overline{\mathrm{Stx}(\mathfrak{A}, \mathfrak{B})}$, and the converse inclusion is clear. Similarly, every finite partial isomorphism is an approximate isomorphism. Taking limits, every partial isomorphism is an approximate isomorphism. The last item follows from the definitions. ∎

Lemma 2.8. *Let $\mathcal{K}$ be a class of finitely generated structures with HP and NAP. Then*

(i) *For any $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$, $\varphi \in \mathrm{Apx}_2(\mathfrak{A}, \mathfrak{B})$ and $\varepsilon > 0$ there exist $\mathfrak{C} \in \mathcal{K}$ and embeddings $f \colon \mathfrak{A} \to \mathfrak{C}$, $g \colon \mathfrak{B} \to \mathfrak{C}$ such that $g^* f < \varphi + \varepsilon$.*

(ii) *For any $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$ and $\varphi \in \mathrm{Stx}(\mathfrak{A}, \mathfrak{B})$ there exist $\mathfrak{C} \in \mathcal{K}$ and embeddings $f \colon \mathfrak{A} \to \mathfrak{C}$, $g \colon \mathfrak{B} \to \mathfrak{C}$ such that $g^* f < \varphi$.*

(iii) *Let $\mathfrak{A} \in \mathcal{K}$, let $\mathfrak{B}$ be a $\mathcal{K}$-structure, and let $\varphi \in \mathrm{Stx}(\mathfrak{A}, \mathfrak{B})$. Then there exists an extension $\mathfrak{A} \subseteq \mathfrak{C} \in \mathcal{K}$ and a finite partial isomorphism $f \colon \mathfrak{C} \dashrightarrow \mathfrak{B}$ such that $f < \varphi$.*

(iv) *The composition of any two (strictly) approximate isomorphisms between $\mathcal{K}$-structures is one as well.*

Proof. For the first item, there exists $\mathfrak{C}_0 \in \mathcal{K}$ and finite partial isomorphisms $f_0 \colon \mathfrak{A} \dashrightarrow \mathfrak{C}_0$, $g_0 \colon \mathfrak{C}_0 \dashrightarrow \mathfrak{B}$ such that $\varphi = g_0 f_0$. By NAP there are $\mathfrak{C}_1, \mathfrak{C}_2 \in \mathcal{K}$ and embeddings as in the diagram below such that $f_2^* f_1 \leq f_0 + \varepsilon$ and $g_2^* g_1 \leq g_0 + \varepsilon$. Let $X = \mathrm{img}\, f_0 \subseteq C_0$, a finite set, and we let $h_0 \colon \mathfrak{C}_1 \dashrightarrow \mathfrak{C}_2$ be the finite partial isomorphism sending $f_2 X \mapsto g_1 X$, i.e., $h_0 = (g_1 f_2^*)\!\restriction_{f_2 X \times g_1 X} = g_1 \, \mathrm{id}_X \, f_2^*$. Applying NAP once more we complete the diagram with $h_2^* h_1 \leq h_0 + \varepsilon$.


Now,

$$\varphi + 3\varepsilon = g_0 \, \mathrm{id}_X \, f_0 + 3\varepsilon \geq g_2^* g_1 \, \mathrm{id}_X \, f_2^* f_1 + \varepsilon = g_2^* h_0 f_1 + \varepsilon \geq g_2^* h_2^* h_1 f_1.$$

Letting $f = h_1 f_1$ and $g = h_2 g_2$ we obtain $\varphi + 4\varepsilon > g^* f$ (by Lemma 1.8(ii), with $X_0 \times Y_0 = \mathrm{dom}\, f_0 \times \mathrm{img}\, g_0$), which is enough.

For the second item, by Lemma 2.7(iii) there exists $\psi \in \mathrm{Apx}_2^{<\varphi}(\mathfrak{A}, \mathfrak{B})$ and we may assume that $\psi + \varepsilon < \varphi$ for some $\varepsilon > 0$. We then apply the first item.

For the third item, by definition of Apx_2 there exists a finitely generated $\mathfrak{B} \supseteq \mathfrak{B}_0 \in \mathcal{K}$ such that $\varphi \in \mathrm{Apx}_2(\mathfrak{A}, \mathfrak{B}_0)$. Applying the second item, there exists $\mathfrak{C} \in \mathcal{K}$ and embeddings $f \colon \mathfrak{A} \to \mathfrak{C}$ and $g \colon \mathfrak{B} \to \mathfrak{C}$ such that $g^* f < \varphi$. We may assume that f is an inclusion, and by Lemma 1.8(ii) there exists a finite set $C_0 \subseteq C$ such that $g^{-1}\!\restriction_{C_0} \in \mathrm{Apx}^{<\varphi}(\mathfrak{C}, \mathfrak{B})$. Finally, by definition of the trivial extension, we may assume that $C_0 \subseteq \mathrm{img}\, g$ so $g^{-1}\!\restriction_{C_0}$ is a finite partial isomorphism.

For the last item, let $\varphi_i \in \mathrm{Stx}(\mathfrak{A}_i, \mathfrak{A}_{i+1})$ for $i = 0, 1$, where $\mathfrak{A}_i$ are $\mathcal{K}$-structures. By definition there exist $\psi_i \in \mathrm{Apx}_2^{<\varphi_i}(\mathfrak{A}_i, \mathfrak{A}_{i+1})$, say $\psi_i = g_i * f_i$ where $f_i \colon \mathfrak{A}_i \dashrightarrow \mathfrak{B}_i$ and $g_i \colon \mathfrak{A}_{i+1} \dashrightarrow \mathfrak{B}_i$ are finite partial isomorphisms, and $\mathfrak{B}_i \in \mathcal{K}$. By the first item (and the fact that only a finitely generated sub-structure of $\mathfrak{A}_1$ is actually involved) there are $\mathfrak{C} \in \mathcal{K}$ and embeddings $h_i \colon \mathfrak{B}_i \to \mathfrak{C}$ such that $h_1^* h_0 < f_1 g_0^* + \varepsilon$:

It follows that $(h_1 g_1)^* (h_0 f_0) \leq \psi_1 \psi_0 + \varepsilon$, and $\rho = (h_1 g_1)^* (h_0 f_0)$ belongs to $\mathrm{Apx}_2(\mathfrak{A}_0, \mathfrak{A}_2)$. Now, since ε was arbitrary, we could have chosen it so that $\psi_i + \varepsilon < \varphi_i$. Again by Lemma 1.8(ii) it follows that $\psi_1 \psi_0 + \varepsilon < \psi_1 \psi_0 + 2\varepsilon \leq \varphi_1 \varphi_0$, so $\rho < \varphi_1 \varphi_0$ and $\varphi_1 \varphi_0 \in \mathrm{Stx}(\mathfrak{A}_0, \mathfrak{A}_1)$.

When $\varphi_i \in \mathrm{Apx}(\mathfrak{A}_i, \mathfrak{A}_{i+1})$, we have $\varphi_i = \lim_\alpha \varphi_{i,\alpha}$ for nets $\varphi_{i,\alpha} \in \mathrm{Stx}(\mathfrak{A}_i, \mathfrak{A}_{i+1})$, by Lemma 2.7. Possibly passing to a sub-net we may assume $\psi = \lim_\alpha \varphi_{1,\alpha} \varphi_{0,\alpha}$ exists (by compactness). Then $\psi \in \mathrm{Apx}(\mathfrak{A}_0, \mathfrak{A}_2)$ and $\psi \leq \varphi_1 \varphi_0$ (by Lemma 1.7(iii)), which is enough. ∎

Convention 2.9. We equip products of metric spaces with the supremum distance, so for two n-tuples $\bar{a}$ and $\bar{b}$ we have $d(\bar{a}, \bar{b}) = \max_i d(a_i, b_i)$.

Definition 2.10. Let $\mathcal{K}$ be a class of finitely generated $\mathcal{L}$-structures. For $n \geq 0$, we let $\mathcal{K}_n$ denote the class of all pairs $(\bar{a}, \mathfrak{A})$, where $\mathfrak{A} \in \mathcal{K}$ and $\bar{a} \in A^n$ generates $\mathfrak{A}$. By an abuse of notation, we shall refer to $(\bar{a}, \mathfrak{A}) \in \mathcal{K}_n$ by $\bar{a}$ alone, and denote the generated structure $\mathfrak{A}$ by $\langle \bar{a} \rangle$.

By $\mathrm{Apx}(\bar{a}, \mathfrak{B})$ we shall mean those members $\mathrm{Apx}(\langle \bar{a} \rangle, \mathfrak{B})$ which extend trivially from an approximate isometry $\bar{a} \rightsquigarrow B$, and similarly for $\mathrm{Stx}(\bar{a}, \mathfrak{B})$, $\mathrm{Apx}(\bar{a}, \bar{b})$, and so on. Under HP and NAP, these still compose correctly as per Lemma 2.8.

Definition 2.11. Let $\mathcal{K}$ be a class of finitely generated structures with NAP. We equip $\mathcal{K}_n$ with a pseudo-distance $d^{\mathcal{K}}$ defined by

$$d^{\mathcal{K}}(\bar{a}, \bar{b}) = \inf_{\psi \in \mathrm{Stx}(\bar{a}, \bar{b})} d(\psi) = \inf_{\psi \in \mathrm{Apx}(\bar{a}, \bar{b})} d(\psi), \qquad \text{where } d(\psi) = \max_i \psi(a_i, b_i).$$

Equivalently, $d(\bar{a}, \bar{b})$ is the infimum of all possible $d(\bar{a}, \bar{b})$ under embeddings of $\langle \bar{a} \rangle$ and $\langle \bar{b} \rangle$ into some $\mathfrak{C} \in \mathcal{K}$. The triangle inequality is a consequence of Lemma 2.8.

Definition 2.12. A *Fraïssé class* (of $\mathcal{L}$-structures) is a class $\mathcal{K}$ of finitely generated $\mathcal{L}$-structures having the following properties:

- *HP.*
- *JEP (Joint Embedding Property)*: Every two members of $\mathcal{K}$ embed in a third one.
- *NAP.*
- *PP (Polish Property)*: The pseudo-metric $d^{\mathcal{K}}$ is separable and complete on $\mathcal{K}_n$ for each n.
- *CP (Continuity Property)*: Every symbol is continuous on $\mathcal{K}$. For an n-ary predicate symbol P, this means that the map $K_n \to \mathbf{R}$, $\bar{a} \mapsto P^{\langle \bar{a} \rangle}(\bar{a})$, is continuous. For an n-ary function symbol P, this means that for each m, the map $K_{n+m} \to \mathcal{K}_{n+m+1}$, $(\bar{a}, \bar{b}) \mapsto (\bar{a}, \bar{b}, f^{\langle \bar{a}, \bar{b} \rangle}(\bar{a}))$, is continuous.

We say that $\mathcal{K}$ is an *incomplete Fraïssé class* if instead of PP & CP we have:

- *WPP (Weak Polish Property)*: The pseudo-metric $d^{\mathcal{K}}$ is separable on $\mathcal{K}_n$ for each n.
- *CCP (Cauchy Continuity Property)*: Every symbol is Cauchy continuous on $\mathcal{K}$ (as per Remark 2.2).

Remark 2.13. We observe that:

(i) CP implies that the kernel of $d^{\mathcal{K}}$ on $\mathcal{K}_n$ is exactly the isomorphism relation: $d^{\mathcal{K}}(\bar{a}, \bar{b}) = 0$ if and only if exists a (necessarily unique) isomorphism $\varphi \colon \langle \bar{a} \rangle \to \langle \bar{b} \rangle$ sending $\bar{a} \mapsto \bar{b}$. It follows that a partial isometry between $\mathcal{K}$-structures is an approximate isomorphism if and if it is a partial isomorphism.

(ii) Together with PP this implies that a $\mathcal{K}$-structure generated by a set of cardinal κ has density character at most $\kappa + \aleph_0$ (even if the language contains more than κ symbols). In particular, every member of $\mathcal{K}$ is separable.

(iii) Every Fraïssé class is in particular an incomplete Fraïssé class, and conversely, every incomplete Fraïssé class $\mathcal{K}$ admits a unique *completion* $\widehat{\mathcal{K}}$, consisting of all limits of Cauchy sequences in $\mathcal{K}$ (that is, in $\mathcal{K}_n$, as n varies), which is a Fraïssé class.

(iv) JEP is equivalent to saying that the empty approximate isometry is always a (strictly) approximate isomorphism. Modulo NAP, JEP is further equivalent to there being a unique $\varnothing$-generated (empty, if there are no constant symbols) structure in $\mathcal{K}$.

Definition 2.14. We say that an approximate isometry $\psi \colon X \rightsquigarrow Y$ is *r-total* for some $r > 0$ if $\psi^* \psi \leq \mathrm{id}_X + 2r$, or equivalently, if for all $x \in X$ and $s > r$ there is $y \in Y$ such that $\psi(x, y) < s$. If $\psi \psi^* \leq \mathrm{id}_Y + 2r$ then we say that ψ is *r-surjective* and if it is both then it is *r-bijective*.

Definition 2.15. Let $\mathcal{K}$ be a Fraïssé class. By a *limit* of $\mathcal{K}$ we mean a separable $\mathcal{K}$-structure $\mathfrak{M}$, satisfying that for every $\mathcal{K}$-structure $\mathfrak{A}$, finite $A_0 \subseteq A$, $\psi \in \mathrm{Stx}(\mathfrak{A}, \mathfrak{M})$ and $\varepsilon > 0$ there exists $\varphi \in \mathrm{Stx}^{<\psi}(\mathfrak{A}, \mathfrak{M})$ which is ε-total on A_0.

Lemma 2.16. *Let $\mathcal{K}$ be a Fraïssé class, $\mathfrak{M}$ a separable $\mathcal{K}$-structure. For each n let $\mathcal{K}_{n,0} \subseteq \mathcal{K}_n$ be $d^{\mathcal{K}}$-dense, and let $M_0 = \{a_i\}_{i \in \mathbf{N}} \subseteq M$ be dense. We shall use the notation $a_{<m}$ for the tuple $(a_i)_{i<m}$.*

Then in order for $\mathfrak{M}$ to be a limit of $\mathcal{K}$, is enough that for every $n, m \in \mathbf{N}$, $\varepsilon > 0$, $\bar{b} \in \mathcal{K}_{n,0}$ and $\psi \colon \bar{b} \times a_{<m} \to \mathbf{Q}$, if $\psi \in \mathrm{Stx}(\bar{b}, \mathfrak{M})$ (so in particular, $\psi \colon \bar{b} \rightsquigarrow a_{<m}$ is an approximate isometry) then there exist $\varphi \in \mathrm{Apx}^{\leq \psi}(\bar{b}, \mathfrak{M})$ which is ε-total on $\bar{b}$.

Proof. Let $\mathfrak{B}$ be a $\mathcal{K}$-structure, $B_0 \subseteq B$ finite, $\psi \in \mathrm{Stx}(\mathfrak{B}, \mathfrak{M})$ and $\varepsilon > 0$. There exist a finite tuple $\bar{b} \in B^n$ and $\psi_0 \in \mathrm{Stx}(\bar{b}, \mathfrak{M})$ such that $\psi_0 < \psi$, and we may assume that $\bar{b}$ contains B_0. Let $0 < \delta \leq \varepsilon/3$ be small enough that $\psi_0 + 3\delta < \psi$. Let $\bar{c} \in \mathcal{K}_{n,0}$ with $d^{\mathcal{K}}(\bar{c}, \bar{b}) < \delta$, and let $\rho \in \mathrm{Stx}(\bar{c}, \bar{b})$ witness this, namely

satisfy $d(\rho) < \delta$ as per Definition 2.11. Then $\psi_0\rho \in \mathrm{Stx}(\bar{c}, \mathfrak{M})$, so there exists $\psi_1 \in \mathrm{Apx}^{<\psi_0\rho}(\langle\bar{c}\rangle, \mathfrak{M})$. Replacing ψ_1 with its restriction to $\bar{c} \times M_0$ we still have $\psi_1 < \psi_0\rho$. By Lemma 1.8(ii) there exist some m and $\psi': \bar{c} \times a_{<m} \to \mathbf{Q}$ such that $\psi_1 < \psi' < \psi_0\rho$, i.e., $\psi' \in \mathrm{Stx}^{<\psi_0\rho}(\bar{c}, \mathfrak{M})$. By assumption there exists $\varphi' \in \mathrm{Apx}^{\leq\psi'}(\bar{c}, \mathfrak{M})$ which is δ-total on $\bar{c}$, and we may further assume that $\varphi' \in \mathrm{Apx}^{\leq\psi_0\rho}(\bar{c}, a_{<k})$ for some k. Thus $\varphi'\rho^* < \psi'\rho^* + \delta \leq \psi_0\rho\rho^* + \delta \leq \psi_0 + 3\delta < \psi$, so $\psi'\rho^* + \delta \in \mathrm{Stx}^{<\psi}(\mathfrak{B}, \mathfrak{M})$ and $\psi'\rho^* + \delta$ is moreover ε-total on $\bar{b}$, as desired. $\blacksquare$

Lemma 2.17. *Every Fraïssé class $\mathcal{K}$ admits a limit.*

Proof. We construct an increasing chain of $\mathfrak{A}_n \in \mathcal{K}$, starting with $\mathfrak{A}_0$ being the unique $\varnothing$-generated structure in $\mathcal{K}$, letting $i_{n,m}: \mathfrak{A}_n \to \mathfrak{A}_m$ denote the inclusion maps. For each n we fix a countable $d^{\mathcal{K}}$-dense subset of $\mathcal{K}_n$, call it $\mathcal{K}_{n,0}$, and a countable dense subset $A_{n,0} \subseteq A_n$, such that $A_{n,0} \subseteq A_{n+1,0}$.

By Lemma 2.8(ii) we can construct the chain $\mathfrak{A}_n$ so that for each $\bar{b} \in \mathcal{K}_{n,0}$, finite subset $B \subseteq A_{m,0}$ and $\psi: \bar{b} \times B \to \mathbf{Q}$, if $\psi \in \mathrm{Stx}(\bar{b}, \mathfrak{A}_m)$ then there exists $k > m$ and an embedding $h: \langle\bar{b}\rangle \to \mathfrak{A}_{k+1}$ such that $i^*_{k,k+1}h < \psi$, and in particular $h < \psi$. By PP and CP, the chain $\mathfrak{A}_0 \subseteq \mathfrak{A}_1 \subseteq \ldots$ admits a unique limit in the category of $\mathcal{K}$-structures, which we denote by $\mathfrak{M} = \bigcup \mathfrak{A}_n$, in which $M_0 = \bigcup A_{n,0} \subseteq M$ is dense. By Lemma 2.16, $\mathfrak{M}$ is a limit. $\blacksquare$

In fact, we can do better. For $\bar{a} \in \mathcal{K}_n$ let $[\bar{a}]$ denote the equivalence class $\bar{a}/\ker d^{\mathcal{K}}$, and let $\overline{\mathcal{K}}_n = \mathcal{K}_n/\ker d^{\mathcal{K}}$ denote the quotient space, equipped with the quotient metric (which is separable and complete, by PP). For each n we have a natural map $\overline{\mathcal{K}}_{n+1} \to \overline{\mathcal{K}}_n$, sending $[a_0, \ldots, a_n] \mapsto [a_0, \ldots, a_{n-1}]$, giving rise to an inverse system with a limit $\overline{\mathcal{K}}_\omega = \varprojlim \overline{\mathcal{K}}_n$, equipped with the topology induced from $\prod_n \overline{\mathcal{K}}_n$. A member of $\overline{\mathcal{K}}_\omega$ will be denoted by ξ, represented by a compatible sequence $(\xi_n)_{n \in \mathbf{N}}$. Considering limits of increasing chains as in the proof of Lemma 2.17, we see that for every $\xi \in \overline{\mathcal{K}}_\omega$ there exists a $\mathcal{K}$-structure $\mathfrak{M}^\xi$ along with a generating sequence $\bar{a}^\xi = (a_i^\xi)_{i \in \mathbf{N}} \subseteq M^\xi$, such that $\xi_n = [a_{<n}^\xi]$ for all n, and this pair $(\mathfrak{M}^\xi, \bar{a}^\xi)$ is determined by ξ up to a unique isomorphism. Conversely, any pair of a separable $\mathcal{K}$-structure $\mathfrak{M}$ and a generating $\mathbf{N}$-sequence is of this form.

Theorem 2.18. *Let $\mathcal{K}$ be a Fraïssé class, and let $\overline{\mathcal{K}}_\omega$ be as above. Let Ξ be the set of $\xi \in \overline{\mathcal{K}}_\omega$ for which $\mathfrak{M}^\xi$ is a limit of $\mathcal{K}$ and every tail of the sequence (a_i^ξ) is dense in $\mathfrak{M}^\xi$. Then $\overline{\mathcal{K}}_\omega$ is a Polish space and $\Xi \subseteq \overline{\mathcal{K}}_\omega$ is a dense G_δ.*

Proof. That $\overline{\mathcal{K}}_\omega$ is a Polish space is clear.

Let $\mathcal{K}_{n,0} \subseteq \mathcal{K}_n$ be countable dense as earlier, and let $\bar{b} \in \mathcal{K}_{n,0}$, $\varepsilon > 0$ (say rational) and $\psi: \bar{b} \times m \to \mathbf{Q}^{>0}$. Define $X_{\bar{b},\varepsilon,\psi} \subseteq \overline{\mathcal{K}}_\omega$ to consist of all ξ such that one of the following holds:

- either there is no $\varphi \in \mathrm{Stx}(\bar{b}, \mathfrak{M}^\xi)$ such that $\varphi(b_i, a_j^\xi) < \psi(b_i, j)$ for all $i < n$, $j < m$ (let us call such a φ *good*),
- or there exists a good φ such that, moreover, for each $i < n$ there is $k \geq m$ with $\varphi(b_i, a_k^\xi) < \varepsilon$.

It is easy to check using Lemma 2.16 that Ξ is the intersection of all such $X_{\bar{b},\varepsilon,\psi}$, of which there are countably many, so all we need to show is that each $X_{\bar{b},\varepsilon,\psi}$ is a dense G_δ set.

The first possibility defines a closed set and the second an open one, so $X_{\bar{b},\varepsilon,\psi}$ is indeed a G_δ set. For density, let $U \subseteq \overline{\mathcal{K}}_\omega$ be open and $\xi \in U$. If there is no good $\varphi \in \mathrm{Stx}(\bar{b}, \mathfrak{M}^\xi)$ then $\xi \in X_{\bar{b},\varepsilon,\psi} \cap U$ and we are done. Otherwise, let us fix a good φ, and let $\varphi_0 \in \mathrm{Stx}(\bar{b}, a_{<m}^\xi)$ be the restriction of φ to $\bar{b} \times a_{<m}^\xi$. We may assume that U is the inverse image in $\overline{\mathcal{K}}_\omega$ of an open set $V \subseteq \overline{\mathcal{K}}_\ell$, with $\ell \geq m$ and $\xi_\ell \in V$. By Lemma 2.8(ii) there exists an extension $\langle a_{<\ell}^\xi \rangle \subseteq \mathfrak{C} \in \mathcal{K}$ and an embedding $\varphi_0 > h: \langle\bar{b}\rangle \to \mathfrak{C}$, and we may assume that $\mathfrak{C} = \langle\bar{c}\rangle$ where $\bar{c} = a_{<\ell}^\xi, h\bar{b}$, so $\bar{c} \in \mathcal{K}_{\ell+n}$. Let $\zeta \in \overline{\mathcal{K}}_\omega$ be any such that $\zeta_{\ell+n} = [\bar{c}]$. Then $\zeta \in U \cap X_{\bar{b},\varepsilon,\psi}$, as desired. $\blacksquare$

Theorem 2.19. *Let $\mathcal{K}$ be a Fraïssé class, $\mathfrak{M}$ and $\mathfrak{N}$ separable $\mathcal{K}$-structures, and let $\psi \in \mathrm{Stx}(\mathfrak{M}, \mathfrak{N})$.*

(i) *If $\mathfrak{N}$ is a limit of $\mathcal{K}$ then ψ strictly coarsens an embedding $\theta: \mathfrak{M} \to \mathfrak{N}$.*

(ii) *If both $\mathfrak{M}$ and $\mathfrak{N}$ are limits of $\mathcal{K}$ then ψ strictly coarsens an isomorphism $\theta: \mathfrak{M} \cong \mathfrak{N}$.*

In particular (with $\psi = \infty$), the limit of $\mathcal{K}$ is unique up to isomorphism.

Proof. We only prove the second assertion, the first being similar and easier. Let $\{a_n\}$ and $\{b_n\}$ enumerate dense subsets of $\mathfrak{M}$ and $\mathfrak{N}$, respectively. We construct a decreasing sequence of $\theta_n \in \mathrm{Stx}(\mathfrak{M}, \mathfrak{N})$, starting with $\theta_0 = \psi$. For even n we choose $\theta_{n+1} \in \mathrm{Stx}^{<\theta_n}(\mathfrak{M}, \mathfrak{N})$ which is 2^{-n}-total on $a_{<n}$. For odd n we similarly choose $\theta_{n+1} \in \mathrm{Stx}^{<\theta_n}(\mathfrak{M}, \mathfrak{N})$, which is 2^{-n}-surjective on $b_{<n}$ (i.e., $\theta_{n+1}^* \in \mathrm{Stx}^{<\theta_n^*}(\mathfrak{N}, \mathfrak{M})$ which is 2^{-n}-total on $b_{<n}$). Then $\theta = \lim \theta_n$ is the desired isomorphism. $\blacksquare$

The unique limit of $\mathcal{K}$ will be denoted by $\lim \mathcal{K}$. It can also be characterised in terms of actual maps.

Corollary 2.20. *Let $\mathcal{K}$ be a Fraïssé class and $\mathfrak{M}$ a separable $\mathcal{K}$-structure. Then the following are equivalent:*

    (i) *The structure $\mathfrak{M}$ is a limit of $\mathcal{K}$.*

    (ii) *Theorem 2.19*(i) *holds: for every separable $\mathcal{K}$-structure $\mathfrak{B}$ and $\psi \in \mathrm{Stx}(\mathfrak{B}, \mathfrak{M})$, there is an embedding $f : \mathfrak{B} \to \mathfrak{M}$, $f < \psi$.*

    (iii) *For a separable $\mathcal{K}$-structure $\mathfrak{B}$, finite tuple $\bar{a} \in B$, embedding $h : \langle \bar{a} \rangle \to \mathfrak{M}$ and $\varepsilon > 0$, there is an embedding $f : \mathfrak{B} \to \mathfrak{M}$ such that $d(f\bar{a}, h\bar{a}) < \varepsilon$. (Equivalently, we can take h to be a finite partial isomorphism and $\bar{a}$ to enumerate $\mathrm{dom}\, h$.)*

    (iv) *Same, where $\mathfrak{B}$ is finitely generated (i.e., $\mathfrak{B} \in \mathcal{K}$).*

Proof. (i) $\implies$ (ii). By Theorem 2.19(i).

    (ii) $\implies$ (iii) $\implies$ (iv). Clear.

    (iv) $\implies$ (i). Let $\mathfrak{B} \in \mathcal{K}$ and $\psi \in \mathrm{Stx}(\mathfrak{B}, \mathfrak{M})$. By Lemma 2.8(iii), possibly increasing $\mathfrak{B}$ we may assume there is a finite partial isomorphism $h : \mathfrak{B} \dashrightarrow \mathfrak{M}$ such that $h < \psi$, and so $h + \varepsilon < \psi$ for some $\varepsilon > 0$. By hypothesis we obtain $f : \mathfrak{B} \to \mathfrak{M}$ such that $d(fa, ha) < \varepsilon$ for $a \in \mathrm{dom}\, h$, or $f < h + \varepsilon$. In particular, f is total and $f < \psi$. Thus $\mathfrak{M}$ is a limit. $\blacksquare$

Theorem 2.21. *Let $\mathcal{K}$ be a class of finitely generated structures. Then the following are equivalent:*

    (i) *The class $\mathcal{K}$ is a Fraïssé class.*

    (ii) *The class $\mathcal{K}$ is the age of a separable approximately ultra-homogeneous structure $\mathfrak{M}$.*

Moreover, such a structure $\mathfrak{M}$ is necessarily a limit of $\mathcal{K}$, and thus unique up to isomorphism and universal for separable $\mathcal{K}$-structures.

Proof. The second item clearly implies the first, as well as the moreover part. Conversely, if $\mathcal{K}$ is a Fraïssé class then by Lemma 2.17 it has a limit $\mathfrak{M}$. By Theorem 2.19(i) we have $\mathrm{Age}(\mathfrak{M}) = \mathcal{K}$, and homogeneity follows from Theorem 2.19(ii). $\blacksquare$

Remark 2.22. Let $\mathcal{K}$ be a Fraïssé class, and let $\theta : [0, \infty] \to [0, 1]$ be any increasing sub-additive map which is continuous and injective near zero. For example, plain truncation $x \mapsto x \wedge 1$ will do, or if one wants a homeomorphism, one may take $x \mapsto 1 - e^{-x}$ or $x \mapsto \frac{x}{x+1}$. The important point is that for any distance function d, θd is a bounded distance function, uniformly equivalent to d.

We define a new language $\mathcal{L}_\mathcal{K}$, consisting of one n-ary predicate symbol $P_{[\bar{a}]}$ for each equivalence class $[\bar{a}]$ in $\overline{\mathcal{K}}_n$ (or in a dense subset thereof). Then every $\mathcal{K}$-structure $\mathfrak{A}$ gives rise to an $\mathcal{L}_\mathcal{K}$-structure $\mathfrak{A}'$, with the same underlying set, where

$$d^{\mathfrak{A}'} = \theta d^{\mathfrak{A}}, \qquad P_{[\bar{a}]}^{\mathfrak{A}}(\bar{b}) = \theta d^{\mathcal{K}}(\bar{a}, \bar{b}).$$

Let $\mathcal{K}' = \bigcup_{\mathfrak{A} \in \mathcal{K}} \mathrm{Age}(\mathfrak{A}')$. Since $\mathcal{L}'$ is purely relational, all members of $\mathcal{K}'$ are necessarily finite, while members of $\mathcal{K}$ are merely finitely generated, and in general $\mathcal{K}' \neq \{\mathfrak{A}' : \mathfrak{A} \in \mathcal{K}\}$. However, for each n we do have canonical identification between $\mathcal{K}_n$ and $\mathcal{K}'_n$, with $d^{\mathcal{K}'} = \theta d^{\mathcal{K}}$. Then one checks that $\mathcal{K}'$ is a Fraïssé class, and that a $\mathcal{K}$-structure $\mathfrak{M}$ is a limit of $\mathcal{K}$ if and only if $\mathfrak{M}'$ is a limit of $\mathcal{K}'$.

We conclude that up to a change of language, any Fraïssé class or approximately ultra-homogeneous structure can be assume to be in a 1-Lipschitz, $[0, 1]$-valued relational continuous language, and that our more relaxed definitions (see Remark 2.2), while convenient for some concrete examples, do not in truth add any more generality.

Another curious property of this construction is that $(\lim \mathcal{K})' = \lim \mathcal{K}'$ is always an atomic model of its continuous first order theory (since all distances to types are definable), and therefore a prime model.

Notice that in Remark 2.22 all isolated types are isolated by quantifier-free formulae, but non-isolated types need not be determined by they quantifier-free restriction, so the theory need not eliminate quantifiers.

3. Examples of metric Fraïssé classes

3.1. **Standard examples.** Let $\mathcal{K}_M$ be the class of finite metric spaces; $\mathcal{K}_{M,1}$ the class of finite metric spaces of diameter at most one; $\mathcal{K}_H$ the class of finite dimensional Hilbert spaces; and $\mathcal{K}_P$ the class of finite probability algebra, each in the appropriate language. We leave it to the reader to check that these are all Fraïssé classes. We claim that the Urysohn space, the Urysohn sphere, ℓ^2, and the (probability algebra of the) Lebesgue space $([0, 1], \lambda)$, are, respectively, limits of these classes. In fact, in each of these cases, the limits satisfy a strong version of Corollary 2.20(iv):

        For each extension $\mathfrak{A} \subseteq \mathfrak{B}$ of members of $\mathcal{K}$, every embedding $\mathfrak{A} \to \mathfrak{M}$ extends to an embedding $\mathfrak{B} \to \mathfrak{M}$.

3.2. **An incomplete example.** Fix $1 \le p < \infty$, and let $\mathcal{K}$ be the class of (real) atomic L^p lattices with finitely many (see [Mey91] for a formal definition and [BBH11] for a model-theoretic treatment).

Then $\mathcal{K}$ is *not* a Fraïssé class, since it is incomplete (this is in contrast with the class of finite probability algebras, which are all atomic, and do form a complete class). Indeed, working inside $E = L^p[0,1]$, let $f(x) = 1$ and $g(x) = x$. Then on the one hand, $E = \langle f, g \rangle$ is non atomic, while on the other hand, approximating g by step functions, the pair (f, g) can be arbitrarily well approximated by pairs which do generate an atomic lattice.

The class $\mathcal{K}$ is an incomplete Fraïssé class, though, and its completion is the class of *all* separable L^p lattices, whose limit is the unique separable atomless L^p lattice. This is somewhat uninteresting, since the limit already belongs to $\mathcal{K}$.

Alternatively, one could add structure to atomic L^p lattices making embeddings preserve atoms. With this added structure, the class of L^p lattices over finitely many atoms is a Fraïssé class, with limit the unique atomic L^p with $\aleph_0$ atoms. The automorphism group of the latter is S_∞, the permutation group of $\mathbf{N}$, so in a sense this fails to produce something truly new.

3.3. **The Gurarij space.** We recall that

Definition 3.1. A *Gurarij space* is a separable Banach space $\mathbf{G}$ having the property that for any $\varepsilon > 0$, finite dimensional Banach space $E \subseteq F$, and isometric embedding $\psi \colon E \to \mathbf{G}$, there is a linear embedding $\varphi \colon F \to \mathbf{G}$ extending ψ such that in addition, for all $x \in F$, $(1 - \varepsilon)\|x\| < \|\varphi x\| < (1 + \varepsilon)\|x\|$.

Gurarij [Gur66] proved the existence and almost isometric uniqueness of such spaces, while actual (i.e., isometric) uniqueness of $\mathbf{G}$ was shown by Lusky [Lus76]. This uniqueness was more recently re-proved by Kubiś and Solecki [KS], in what essentially amounts to showing that it was the Fraïssé limit of the class of all finite dimensional Banach spaces, an observation we now have the tools to state and prove formally. From here on, $\mathcal{K} = \mathcal{K}_B$ is the class of finite dimensional Banach space. Then this is a Fraïssé class. In particular, it is separable since a separable universal Banach space exists.

Let us also recall the following fact, hitherto unpublished, due to Henson:

Fact 3.2 (See also [BH]). *Let* $\bar{a}, \bar{b} \in \mathcal{K}_n$. *Then*

$$(1) \qquad d^{\mathcal{K}}(\bar{a}, \bar{b}) = \sup_{\sum |s_i| = 1} \left| \left\| \sum s_i a_i \right\| - \left\| \sum s_i b_i \right\| \right|.$$

Proof. The inequality $\ge$ is clear. For $\le$, let r denote the right hand side of (1). Let $E = \langle \bar{a} \rangle \oplus \langle \bar{b} \rangle$ in the category of vector spaces over $\mathbf{R}$, and for $x \in \langle \bar{a} \rangle$, $y \in \langle \bar{b} \rangle$ define:

$$\|x - y\|' = \inf_{\bar{s}} \left\| x - \sum s_i a_i \right\|^{\langle \bar{a} \rangle} + \left\| y - \sum s_i b_i \right\|^{\langle \bar{b} \rangle} + r \sum |s_i|.$$

This is clearly a semi-norm on E, and $\|a_i - b_i\|' \le r$. For $x \in \langle \bar{a} \rangle$ we have $\|x\|' \le \|x\|^{\langle \bar{a} \rangle}$, while on the other hand, for any $\bar{s}$ we have by choice of r:

$$\|x\|^{\langle \bar{a} \rangle} \le \left\| x - \sum s_i a_i \right\|^{\langle \bar{a} \rangle} + \left\| \sum s_i a_i \right\|^{\langle \bar{a} \rangle}$$
$$\le \left\| x - \sum s_i a_i \right\|^{\langle \bar{a} \rangle} + \left\| \sum s_i b_i \right\|^{\langle \bar{b} \rangle} + r \sum |s_i|.$$

It follows that $\|x\|' = \|x\|^{\langle \bar{a} \rangle}$, and similarly for $y \in \langle \bar{b} \rangle$, whence the desired amalgam. ∎

Theorem 3.3. *A Banach space G is a Gurarij space if and only if it is the Fraïssé limit of the class of all finite dimensional Banach space. In particular, the Gurarij space exists, is unique, and is universal for separable Banach spaces.*

Proof. Assume first that $G = \lim \mathcal{K}$. Let $E \subseteq F$ be two finite dimensional Banach spaces, with bases $\bar{a} \subseteq \bar{b}$, respectively, and let $\psi \colon E \to G$ be an isometric embedding. By Corollary 2.20 there exists an isometric $\varphi' \colon F \to G$ with $d(\bar{a}, \varphi \bar{a}) = \delta$ arbitrarily small. Define $\varphi \colon F \to G$ as ψ on $\bar{a}$ and φ' on $\bar{b} \smallsetminus \bar{a}$. Taking δ sufficiently small, φ is injective, and both $\|\varphi\|$ and $\|\varphi^{-1}\|$ (with φ restricted to its image) arbitrarily close to one, so G is Gurarij.

Conversely, assume that G is Gurarij, and let $F = \langle \bar{b} \rangle \in \mathcal{K}$, $\psi \in \mathrm{Stx}(\bar{b}, G)$ and $\varepsilon > 0$ be given. By Lemma 2.8(iii), possibly extending F and decreasing ε we may assume that there are a finite tuple $\bar{c} \in F^m$ and an isometric embedding $\psi' \colon \langle \bar{c} \rangle \to G$ such that $\psi \ge \psi' {\restriction_{\bar{c}}} + \varepsilon$. By assumption there exists a linear $\varphi \colon F \to G$ extending ψ', with $\|\varphi\|, \|\varphi^{-1}\|$ arbitrarily close to one. By Fact 3.2 we can then have $d^{\mathcal{K}}(\bar{b}\bar{c}, \varphi(\bar{b}\bar{c})) < \varepsilon$. Then there exists $\varphi' \in \mathrm{Apx}(\bar{b}\bar{c}, \varphi(\bar{b}\bar{c})) \subseteq \mathrm{Apx}(F, G)$ with $\varphi'(b_i, \varphi b_i) < \varepsilon$, $\varphi'(c_j, \psi'c_j) < \varepsilon$. This φ' is ε-total on $\bar{b}$ and $\psi \ge \psi' {\restriction_{\bar{c}}} + \varepsilon > \varphi' {\restriction_{\bar{c}}} \ge \varphi'$, so G is a limit. ∎

References

[BBH11] Itaï Ben Yaacov, Alexander Berenstein, and C. Ward Henson, *Model-theoretic independence in the Banach lattices $L_p(\mu)$*, Israel Journal of Mathematics **183** (2011), 285–320, doi:10.1007/s11856-011-0050-4, arXiv:0907.5273. 3.2

[BBHU08] Itaï Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov, *Model theory for metric structures*, Model theory with Applications to Algebra and Analysis, volume 2 (Zoé Chatzidakis, Dugald Macpherson, Anand Pillay, and Alex Wilkie, eds.), London Math Society Lecture Note Series, vol. 350, Cambridge University Press, 2008, pp. 315–427. 2.2

[Ben08] Itaï Ben Yaacov, *Continuous first order logic for unbounded metric structures*, Journal of Mathematical Logic **8** (2008), no. 2, 197–223, doi:10.1142/S0219061308000737, arXiv:0903.4957. 2.2

[BH] Itaï Ben Yaacov and C. Ward Henson, *Generic orbits and type isolation in the Gurarij space*, research notes – under revision. 3.2

[BU07] Itaï Ben Yaacov and Alexander Usvyatsov, *On d-finiteness in continuous structures*, Fundamenta Mathematicae **194** (2007), 67–88, doi:10.4064/fm194-1-4. (document), 2

[BU10] _____, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society **362** (2010), no. 10, 5213–5259, doi:10.1090/S0002-9947-10-04837-3, arXiv:0801.4303. 2.2

[Fra54] Roland Fraïssé, *Sur l'extension aux relations de quelques propriétés des ordres*, Annales Scientifiques de l'École Normale Supérieure. Troisième Série **71** (1954), 363–388. (document)

[Gur66] Vladimir I. Gurarij, *Spaces of universal placement, isotropic spaces and a problem of Mazur on rotations of Banach spaces*, Sibirsk. Mat. Ž. **7** (1966), 1002–1013. 3.3

[KS] Wiesław Kubiś and Sławomir Solecki, *A proof of the uniqueness of the Gurariĭ space*, Israel Journal of Mathematics, to appear, arXiv:1110.0903. (document), 3.3

[Lus76] Wolfgang Lusky, *The Gurarij spaces are unique*, Archiv der Mathematik **27** (1976), no. 6, 627–635. 3.3

[Mey91] Peter Meyer-Nieberg, *Banach lattices*, Universitext, Springer-Verlag, Berlin, 1991. 3.2

[Sch07] Konstantinos Schoretsanitis, *Fraïssé theory for metric structures*, Ph.D. thesis, University of Illinois at Urbana-Champaign, 2007. (document)

[Usp08] Vladimir V. Uspenskij, *On subgroups of minimal topological groups*, Topology and its Applications **155** (2008), no. 14, 1580–1606, doi:10.1016/j.topol.2008.03.001, arXiv:math/0004119. 1

Itaï Ben Yaacov, Université Claude Bernard – Lyon 1, Institut Camille Jordan, CNRS UMR 5208, 43 boulevard du 11 novembre 1918, 69622 Villeurbanne Cedex, France

URL: http://math.univ-lyon1.fr/~begnac/