Large cardinal axioms in category theory

Andrew Brooke-Taylor

University of Leeds

Winter School 2020

Overview

Large cardinals come up a lot in the study of *accessible categories*

- a kind of category-theoretic model theory
- (like classical model theory) has applications to other “more mainstream” areas of mathematics

Overview

Large cardinals come up a lot in the study of *accessible categories*

- a kind of category-theoretic model theory
- (like classical model theory) has applications to other “more mainstream” areas of mathematics

This lecture

A fundamental result of Isbell, characterising measurable cardinals in category theoretic terms, introducing many of the basic ideas along the way.

The next two lectures

On with accessible categories in general.

Part I: Isbell's Theorem

Category theory preliminaries

Recall

A category $\mathcal{C}$ consists of

- a class of *objects*, and
- for every pair of objects A and B of $\mathcal{C}$, a set $\mathrm{Hom}_{\mathcal{C}}(A, B)$ of *morphisms* from A to B,

with identity morphisms and a composition (partial) function of morphisms, satisfying suitable axioms.

E.g.s

- **Set** is the category with sets as objects and functions as morphisms.
- **Gp** is the category with groups as objects and group homomorphisms as morphisms.

Limits and colimits

We think of a diagram as being a set of objects and morphisms between them.

- The limit of a diagram $\mathcal{D}$ is an object L along with a *cone* δ of projection maps to the objects of $\mathcal{D}$ (such that the triangles formed with the morphisms of $\mathcal{D}$ commute) such that any other such cone from an object of $\mathcal{C}$ factors uniquely through δ.
- The colimit of a diagram is the same in reverse.

Limits and colimits

We think of a diagram as being a set of objects and morphisms between them.

- The limit of a diagram $\mathcal{D}$ is an object L along with a *cone* δ of projection maps to the objects of $\mathcal{D}$ (such that the triangles formed with the morphisms of $\mathcal{D}$ commute) such that any other such cone from an object of $\mathcal{C}$ factors uniquely through δ.
- The colimit of a diagram is the same in reverse.

The uniqueness means that any two limits are isomorphic (and likewise for colimits). So we will talk about *the* limit of a diagram (if one exists), doing everything "up to isomorphism."

E.g.

In **Set**, every diagram $\mathcal{D}$ has a limit and a colimit:

- The limit is the subset of the product of the sets in $\mathcal{D}$ consisting of all element whose coordinates "cohere" under the functions of the diagram.

E.g.

In **Set**, every diagram $\mathcal{D}$ has a limit and a colimit:

- The limit is the subset of the product of the sets in $\mathcal{D}$ consisting of all element whose coordinates "cohere" under the functions of the diagram. I.e.

$$\lim \mathcal{D} = \left\{ (d_i, d_j, \dots) \in \prod_{D_i \in \mathrm{Obj}(\mathcal{D})} D_i \;\middle|\; \forall f\colon D_i \to D_j \in \mathrm{Mor}(\mathcal{D}) \Big(f(d_i) = d_j \Big) \right\}.$$

E.g.

In **Set**, every diagram $\mathcal{D}$ has a limit and a colimit:

- The limit is the subset of the product of the sets in $\mathcal{D}$ consisting of all element whose coordinates "cohere" under the functions of the diagram. I.e.

$$\lim \mathcal{D} = \left\{ (d_i, d_j, \ldots) \in \prod_{D_i \in \mathrm{Obj}(\mathcal{D})} D_i \;\middle|\; \forall f \colon D_i \to D_j \in \mathrm{Mor}(\mathcal{D}) \Big(f(d_i) = d_j \Big) \right\}.$$

- The colimit is the disjoint union of the sets in $\mathcal{D}$, modulo identifying elements with their images under the functions in $\mathcal{D}$.

E.g.

In **Set**, every diagram $\mathcal{D}$ has a limit and a colimit:

- The limit is the subset of the product of the sets in $\mathcal{D}$ consisting of all element whose coordinates "cohere" under the functions of the diagram. I.e.

$$\lim \mathcal{D} = \left\{ (d_i, d_j, \ldots) \in \prod_{D_i \in \mathrm{Obj}(\mathcal{D})} D_i \;\middle|\; \forall f \colon D_i \to D_j \in \mathrm{Mor}(\mathcal{D}) \Big(f(d_i) = d_j \Big) \right\}.$$

- The colimit is the disjoint union of the sets in $\mathcal{D}$, modulo identifying elements with their images under the functions in $\mathcal{D}$.

Gp has all limits & colimits too: limits are the same as in **Set**, and colimits are free products modulo identifications.

Canonical diagrams

Given a set $\mathcal{A}$ of objects in a category $\mathcal{C}$ and an object C of $\mathcal{C}$, the canonical diagram of C with respect to $\mathcal{A}$ is the diagram with

- for every object A in $\mathcal{A}$ and every morphism $f\colon A \to C$, a copy of A, which we shall denote by A_f,
- as morphisms, all morphisms $h\colon A_f \to B_g$ such that $g \circ h = f$.

Note that the morphisms $f : A_f \to C$ form a cocone to C. If this cocone makes C the colimit of its canonical diagram with respect to $\mathcal{A}$, we say that C is a canonical colimit of objects from $\mathcal{A}$.

Note that the morphisms $f : A_f \to C$ form a cocone to C. If this cocone makes C the colimit of its canonical diagram with respect to $\mathcal{A}$, we say that C is a canonical colimit of objects from $\mathcal{A}$. If *every* object is a canonical colimit of objects from from $\mathcal{A}$, we say that $\mathcal{A}$ is dense.

E.g.s

- ω is dense in **Set**: every set is the colimit of the diagram of all of its finite subsets, which are the images of functions from finite sets.
- Any set of representatives of all the isomorphism classes of finitely generated groups is dense in **Gp**: every group is the colimit of the diagram of all of its finitely generated subgroups.

Note that being a canonical colimit of objects from $\mathcal{A}$ is stronger in general than just being a colimit of *some* diagram of objects from $\mathcal{A}$.

E.g.

Let $\mathbf{Vect}_{\mathbb{R}}$ be the category of real vector spaces, with linear transformations as the morphisms. Consider the set $\mathcal{A} = \{\mathbb{R}\}$. Then every object of $\mathbf{Vect}_{\mathbb{R}}$ is a colimit of objects from $\mathcal{A}$, but $\mathcal{A}$ is *not* dense. Indeed, consider a function $\varphi\colon \mathbb{R}^2 \to \mathbb{R}^2$ respecting scalar multiplication but not addition. Then there is a cocone mapping each $\mathbb{R}_f$ to $\mathbb{R}^2$ by $\varphi \circ f$, but it doesn't factor through the canonical cocone by any linear map.

Note that being a canonical colimit of objects from $\mathcal{A}$ is stronger in general than just being a colimit of *some* diagram of objects from $\mathcal{A}$.

E.g.

Let $\mathbf{Vect}_{\mathbb{R}}$ be the category of real vector spaces, with linear transformations as the morphisms. Consider the set $\mathcal{A} = \{\mathbb{R}\}$. Then every object of $\mathbf{Vect}_{\mathbb{R}}$ is a colimit of objects from $\mathcal{A}$, but $\mathcal{A}$ is *not* dense. Indeed, consider a function $\varphi \colon \mathbb{R}^2 \to \mathbb{R}^2$ respecting scalar multiplication but not addition. Then there is a cocone mapping each $\mathbb{R}_f$ to $\mathbb{R}^2$ by $\varphi \circ f$, but it doesn't factor through the canonical cocone by any linear map.

On the other hand $\{\mathbb{R}^2\}$ *is* dense.

Note that being a canonical colimit of objects from $\mathcal{A}$ is stronger in general than just being a colimit of *some* diagram of objects from $\mathcal{A}$.

E.g.

Let $\mathbf{Vect}_{\mathbb{R}}$ be the category of real vector spaces, with linear transformations as the morphisms. Consider the set $\mathcal{A} = \{\mathbb{R}\}$. Then every object of $\mathbf{Vect}_{\mathbb{R}}$ is a colimit of objects from $\mathcal{A}$, but $\mathcal{A}$ is *not* dense. Indeed, consider a function $\varphi\colon \mathbb{R}^2 \to \mathbb{R}^2$ respecting scalar multiplication but not addition. Then there is a cocone mapping each $\mathbb{R}_f$ to $\mathbb{R}^2$ by $\varphi \circ f$, but it doesn't factor through the canonical cocone by any linear map.

On the other hand $\{\mathbb{R}^2\}$ *is* dense.

Say that a set of objects $\mathcal{A}$ is colimit-dense if every object is a colimit of some diagram of objects from $\mathcal{A}$.

Opposite categories

Given a category $\mathcal{C}$, $\mathcal{C}^{op}$ is the category with the same objects as $\mathcal{C}$, and the same morphisms but in the opposite direction. Identity functions remain identity functions, and compositions of morphisms remain compositions of morphisms, just in the opposite order.

E.g.

$\mathbf{Set}^{op}$ is the category with sets as objects, and functions as morphisms, with any $f : X \to Y$ in the usual sense being considered as going from Y to X.

Opposite categories

Given a category $\mathcal{C}$, $\mathcal{C}^{op}$ is the category with the same objects as $\mathcal{C}$, and the same morphisms but in the opposite direction. Identity functions remain identity functions, and compositions of morphisms remain compositions of morphisms, just in the opposite order.

E.g.

$\mathbf{Set}^{op}$ is the category with sets as objects, and functions as morphisms, with any $f: X \to Y$ in the usual sense being considered as going from Y to X.

Questions

- Is there a colimit-dense set of objects in $\mathbf{Set}^{op}$?
- Is there a dense set of objects in $\mathbf{Set}^{op}$?

Theorem (Adámek, B-T, Campion, Positselski & Rosický, 2020)

In $\mathbf{Set}^{op}$, $\{3\}$ *is colimit-dense.*

Theorem (Adámek, B-T, Campion, Positselski & Rosický, 2020)

In $\mathbf{Set}^{op}$, $\{3\}$ *is colimit-dense.*

Proof

Let's work in **Set**. So we want to show that every set is the limit of a suitable diagram of 3-element sets.

Theorem (Adámek, B-T, Campion, Positselski & Rosický, 2020)

In $\mathbf{Set}^{op}$, $\{3\}$ *is colimit-dense.*

Proof

Let's work in **Set**. So we want to show that every set is the limit of a suitable diagram of 3-element sets.

First observe that if our diagram is just a single three-element set with one endomorphism f, the colimit is the set of fixed points of f. So that deals with sets of size at most 3.

Suppose now that X is a set of cardinality at least 3. Choose $x_0 \in X$, and take some $s \notin X$.

Idea:

X is the limit of the diagram of all of its finite subsets containing x_0, with surjections between them, with elements not in the range being mapped to x_0.

Suppose now that X is a set of cardinality at least 3. Choose $x_0 \in X$, and take some $s \notin X$.

Idea:

X is the limit of the diagram of all of its finite subsets containing x_0, with surjections between them, with elements not in the range being mapped to x_0.

Actually 2 and 3 element subsets containing x_0 suffices, and 2 elements can be simulated with 3 elements and an endomorphism as above.

So:

- for every $x \in X \setminus \{x_0\}$, let
$$K_x = \{x_0, x, s\};$$
- for every pair $x \neq y$ both in $X \setminus \{x_0\}$, let
$$Y_{\{x,y\}} = \{x_0, x, y\};$$
- for every $x \in X \setminus \{x_0\}$, let p_x be the function from K_x to itself given by
$$p_x(x_0) = x_0, \qquad p_x(x) = x, \qquad p_x(s) = x_0;$$
- for every $x \in X \setminus \{x_0\}$ and every $y \in X \setminus \{x_0, x\}$, let $f_{x,y}$ be the function from $Y_{\{x,y\}}$ to K_x given by
$$f_{x,y}(x_0) = x_0, \qquad f_{x,y}(x) = x, \qquad f_{x,y}(y) = x_0.$$

This forms our diagram $\mathcal{D}$.

There is a natural cone η from X to $\mathcal{D}$: for each object Z of $\mathcal{D}$ (i.e. Z is some K_x or $Y_{\{x,y\}}$) define the function $\eta_Z \colon X \to Z$ by

$$\eta_Z(w) = \begin{cases} w & \text{if } w \in Z \\ x_0 & \text{otherwise.} \end{cases}$$

Clearly this commutes with the maps p_x and $f_{x,y}$.

Suppose we have a cone ζ from a set A to $\mathcal{D}$. We want to show that it factors uniquely through η, that is, there is some $g\colon A \to X$ such that $\eta_Z \circ g = \zeta_Z$ for all objects Z of $\mathcal{D}$. Now for all $a \in A$:

- For all $x \in X$, $\zeta_{K_x}(a) \neq s$, as $p_x \circ \zeta_{K_x} = \zeta_{K_x}$.
- If there is some $x \in X \smallsetminus \{x_0\}$ such that $\zeta_{K_x}(a) = x$, it is unique (for any other y, to commute with $f_{x,y}$ we must have $\zeta_{Y_{\{x,y\}}}(a) = x$, and so to commute with $f_{y,x}$ we must have $\zeta_{K_y}(a) = x_0$). So let $g(a)$ be this x.
- If $\zeta_{K_x}(a) = x_0$ for all $x \in X \smallsetminus \{x_0\}$, we must also have $\zeta_{Y_{\{x,y\}}}(a) = x_0$ for all pairs $\{x, y\}$. So in this case let $g(a) = x_0$.

Clearly this g provides the factorisation, and is unique in this. $\square$

Towards the second question

For any cardinal κ (treated as the set of all lesser ordinals) and any set X, consider the canonical diagram $\mathcal{D}$ in $\mathbf{Set}^{op}$ of X with respect to κ.

Towards the second question

For any cardinal κ (treated as the set of all lesser ordinals) and any set X, consider the canonical diagram $\mathcal{D}$ in $\mathbf{Set}^{op}$ of X with respect to κ.

Since morphisms are reversed, this is the diagram with an object for every function from X to an ordinal less than κ, with a function h from α_f to β_g if $h \circ f = g$.

Towards the second question

For any cardinal κ (treated as the set of all lesser ordinals) and any set X, consider the canonical diagram $\mathcal{D}$ in $\mathbf{Set}^{op}$ of X with respect to κ.

Since morphisms are reversed, this is the diagram with an object for every function from X to an ordinal less than κ, with a function h from α_f to β_g if $h \circ f = g$.

We can think about such functions $f : X \to \alpha$ in terms of the partitions $\{f^{-1}\{\gamma\} \mid \gamma \in \alpha\}$ that they define. In this context, the functions in the diagram represent coarsening maps.

Towards the second question

For any cardinal κ (treated as the set of all lesser ordinals) and any set X, consider the canonical diagram $\mathcal{D}$ in $\mathbf{Set}^{op}$ of X with respect to κ.

Since morphisms are reversed, this is the diagram with an object for every function from X to an ordinal less than κ, with a function h from α_f to β_g if $h \circ f = g$.

We can think about such functions $f : X \to \alpha$ in terms of the partitions $\{f^{-1}\{\gamma\} \mid \gamma \in \alpha\}$ that they define. In this context, the functions in the diagram represent coarsening maps.

Let's again switch to talking about things in terms of $\mathbf{Set}$ rather than $\mathbf{Set}^{op}$. So we want to see whether the canonical cone from X to $\mathcal{D}$ makes X the limit of the diagram.

The limit of $\mathcal{D}$

The elements of the limit are elements $\mathbf{u} = (u_f)_{\alpha_f \in \mathcal{D}}$ of the product of the ordinals α_f in $\mathcal{D}$ — in the α_f coordinate, the element u_f of α_f is chosen — such that the choices cohere with the coarsening maps.

The limit of $\mathcal{D}$

The elements of the limit are elements $\mathbf{u} = (u_f)_{\alpha_f \in \mathcal{D}}$ of the product of the ordinals α_f in $\mathcal{D}$ — in the α_f coordinate, the element u_f of α_f is chosen — such that the choices cohere with the coarsening maps.

This corresponds to the choice of a piece from each of the partitions ($f^{-1}\{u_f\}$ in the partition corresponding to $f\colon X \to \alpha$), in a way that the coarsening maps respect — we can think of this as choosing a "big" piece from each partition.

The limit of $\mathcal{D}$

The elements of the limit are elements $\mathbf{u} = (u_f)_{\alpha_f \in \mathcal{D}}$ of the product of the ordinals α_f in $\mathcal{D}$ — in the α_f coordinate, the element u_f of α_f is chosen — such that the choices cohere with the coarsening maps.

This corresponds to the choice of a piece from each of the partitions ($f^{-1}\{u_f\}$ in the partition corresponding to $f \colon X \to \alpha$), in a way that the coarsening maps respect — we can think of this as choosing a "big" piece from each partition.

Claim

These choices form a κ-complete ultrafilter on X!

A κ-complete ultrafilter from an element of the limit

Proof of the claim

First, by coarsening, if Y is chosen in any partition, it is chosen in the partition $\{Y, X \smallsetminus Y\}$, from which it can be seen that Y is chosen in every partition containing it.

A κ-complete ultrafilter from an element of the limit

Proof of the claim

First, by coarsening, if Y is chosen in any partition, it is chosen in the partition $\{Y, X \setminus Y\}$, from which it can be seen that Y is chosen in every partition containing it. So let $\mathcal{U}$ be the set of $Y \subseteq X$ such that Y is chosen in some (any) partitition in which it appears as a piece (i.e., if $Y = f^{-1}(u_f)$).

Let $\chi_Y : X \to 2$ be the characteristic function of Y, $\chi_Y(x) = 1 \leftrightarrow x \in Y$. Then

$$\begin{aligned} \mathcal{U} &= \{Y \subseteq X \mid \exists \alpha < \kappa \exists f : X \to \alpha (Y = f^{-1}\{u_f\})\} \\ &= \{Y \subseteq X \mid u_{\chi_Y} = 1\}. \end{aligned}$$

We shall show that this is a κ-complete ultrafilter on X.

$\mathcal{U}$ is a κ-complete ultrafilter

$$\begin{aligned}\mathcal{U} &= \{Y \subseteq X \mid \exists\alpha < \kappa \exists f : X \to \alpha (Y = f^{-1}\{u_f\})\} \\ &= \{Y \subseteq X \mid u_{\chi_Y} = 1\}.\end{aligned}$$

$\mathcal{U}$ is a κ-complete ultrafilter

$$\mathcal{U} = \{Y \subseteq X \mid \exists \alpha < \kappa \exists f : X \to \alpha (Y = f^{-1}\{u_f\})\}$$
$$= \{Y \subseteq X \mid u_{\chi_Y} = 1\}.$$

- If $Y \in \mathcal{U}$ and $Y \subseteq Z \subseteq X$, then $\{Y, Z \setminus Y, X \setminus Z\}$ coarsens to $\{Z, X \setminus Z\}$, so $Z \in \mathcal{U}$.

$\mathcal{U}$ is a κ-complete ultrafilter

$$\begin{aligned}\mathcal{U} &= \{Y \subseteq X \mid \exists\alpha < \kappa \exists f \colon X \to \alpha(Y = f^{-1}\{u_f\})\} \\ &= \{Y \subseteq X \mid u_{\chi_Y} = 1\}.\end{aligned}$$

- If $Y \in \mathcal{U}$ and $Y \subseteq Z \subseteq X$, then $\{Y, Z \setminus Y, X \setminus Z\}$ coarsens to $\{Z, X \setminus Z\}$, so $Z \in \mathcal{U}$.
- If $Y \in \mathcal{U}$ and $\{Z_\gamma \mid \gamma < \alpha\}$ is a partition of Y into fewer than κ many pieces, then since $\{X \setminus Y\} \cup \{Z_\gamma \mid \gamma < \alpha\}$ coarsens to $\{Y, X \setminus Y\}$, one of the Z_γ is in $\mathcal{U}$, so $\mathcal{U}$ is κ-complete.

$\mathcal{U}$ is a κ-complete ultrafilter

$$\begin{aligned}\mathcal{U} &= \{Y \subseteq X \mid \exists\alpha < \kappa \exists f : X \to \alpha (Y = f^{-1}\{u_f\})\} \\ &= \{Y \subseteq X \mid u_{\chi_Y} = 1\}.\end{aligned}$$

- If $Y \in \mathcal{U}$ and $Y \subseteq Z \subseteq X$, then $\{Y, Z \smallsetminus Y, X \smallsetminus Z\}$ coarsens to $\{Z, X \smallsetminus Z\}$, so $Z \in \mathcal{U}$.
- If $Y \in \mathcal{U}$ and $\{Z_\gamma \mid \gamma < \alpha\}$ is a partition of Y into fewer than κ many pieces, then since $\{X \smallsetminus Y\} \cup \{Z_\gamma \mid \gamma < \alpha\}$ coarsens to $\{Y, X \smallsetminus Y\}$, one of the Z_γ is in $\mathcal{U}$, so $\mathcal{U}$ is κ-complete.
- For any $Y \subseteq X$, $Y \in \mathcal{U}$ if $u_{\chi_Y} = 1$ and $X \smallsetminus Y \in \mathcal{U}$ if $u_{\chi_Y} = 0$, so $\mathcal{U}$ is ultra.

$\square$

Note that, conversely, if $\mathcal{V}$ is a κ-complete ultrafilter on X, there is an element $\mathbf{u}_{\mathcal{V}}$ of the limit which for each partition funtion $f\colon X \to \alpha$ chooses the piece of the partition that lies in $\mathcal{V}$ (this clearly coheres with the coarsening maps). Moreover the ultrafilter $\mathcal{U}$ corresponding as above to $\mathbf{u}_{\mathcal{V}}$ is just $\mathcal{V}$.

So we can identify $\lim \mathcal{D}$ with the set of κ-complete ultrafilters on X.

The canonical cone from X to $\mathcal{D}$ factors through the limit cone by the map $x \mapsto \mathbf{u}_x$, where the α_f component of $\mathbf{u}_x$ is $f(x)$.

The canonical cone from X to $\mathcal{D}$ factors through the limit cone by the map $x \mapsto \mathbf{u}_x$, where the α_f component of $\mathbf{u}_x$ is $f(x)$.

Note that the $\chi_{\{x\}}$ component of $\mathbf{u}_x$ is 1, so $\{x\}$ is in the corresponding ultrafilter — it is the principal ultrafilter defined by x.

The canonical cone from X to $\mathcal{D}$ factors through the limit cone by the map $x \mapsto \mathbf{u}_x$, where the α_f component of $\mathbf{u}_x$ is $f(x)$.

Note that the $\chi_{\{x\}}$ component of $\mathbf{u}_x$ is 1, so $\{x\}$ is in the corresponding ultrafilter — it is the principal ultrafilter defined by x.

So there is a non-principal κ-complete ultrafilter on X if and only if this map $X \to \lim \mathcal{D}$ is not a bijection

The canonical cone from X to $\mathcal{D}$ factors through the limit cone by the map $x \mapsto \mathbf{u}_x$, where the α_f component of $\mathbf{u}_x$ is $f(x)$.

Note that the $\chi_{\{x\}}$ component of $\mathbf{u}_x$ is 1, so $\{x\}$ is in the corresponding ultrafilter — it is the principal ultrafilter defined by x.

So there is a non-principal κ-complete ultrafilter on X if and only if this map

$X \to \lim \mathcal{D}$ is not a bijection
i.e. not an isomorphism in **Set**

The canonical cone from X to $\mathcal{D}$ factors through the limit cone by the map $x \mapsto \mathbf{u}_x$, where the α_f component of $\mathbf{u}_x$ is $f(x)$.

Note that the $\chi_{\{x\}}$ component of $\mathbf{u}_x$ is 1, so $\{x\}$ is in the corresponding ultrafilter — it is the principal ultrafilter defined by x.

So there is a non-principal κ-complete ultrafilter on X if and only if this map

$X \to \lim \mathcal{D}$ is not a bijection
i.e. not an isomorphism in **Set**
i.e. X is not the limit of $\mathcal{D}$.

(N.B. X may well have the same cardinality as the limit of $\mathcal{D}$, in which case there is *some* isomorphism between them, but it won't be one that makes the canonical cone into the colimit cone.)

Note that there is a dense set in $\mathbf{Set}^{op}$ if and only if for some κ, every X is the limit of its canonical diagram with respect to κ,

Note that there is a dense set in $\mathbf{Set}^{op}$ if and only if for some κ, every X is the limit of its canonical diagram with respect to κ, if and only if there are no non-principal κ-complete ultrafilters on any set.

Note that there is a dense set in $\mathbf{Set}^{op}$ if and only if for some κ, every X is the limit of its canonical diagram with respect to κ, if and only if there are no non-principal κ-complete ultrafilters on any set.

Recall that a cardinal κ is *measurable* if it admits a non-principal κ-complete ultrafilter. A κ as above clearly can't be measurable, or have any measurables above it (since for $\lambda > \kappa$, λ-complete implies κ-complete).

Note that there is a dense set in $\mathbf{Set}^{op}$ if and only if for some κ, every X is the limit of its canonical diagram with respect to κ, if and only if there are no non-principal κ-complete ultrafilters on any set.

Recall that a cardinal κ is *measurable* if it admits a non-principal κ-complete ultrafilter. A κ as above clearly can't be measurable, or have any measurables above it (since for $\lambda > \kappa$, λ-complete implies κ-complete).

For the converse:

Lemma

For any cardinal μ, the least cardinal κ admitting a non-principal μ-complete ultrafilter is measurable (i.e. it admits a κ-complete ultrafilter).

Lemma

For any cardinal μ, the least cardinal κ admitting a non-principal μ-complete ultrafilter is measurable (i.e. it admits a κ-complete ultrafilter).

Proof.

Let $\mathcal{U}$ be a μ-complete ultrafilter on κ that is not κ-complete, and suppose $f : \kappa \to \alpha$ defines a partition of κ into $\alpha < \kappa$ many pieces $f^{-1}(\gamma)$ none of which is in $\mathcal{U}$. Then

$$\mathcal{V} = \{X \subseteq \alpha \mid f^{-1}X \in \mathcal{U}\}$$

is a non-principal μ-complete ultrafilter on α, violating the minimality of κ. $\square$

Lemma

For any cardinal μ, the least cardinal κ admitting a non-principal μ-complete ultrafilter is measurable (i.e. it admits a κ-complete ultrafilter).

Proof.

Let $\mathcal{U}$ be a μ-complete ultrafilter on κ that is not κ-complete, and suppose $f : \kappa \to \alpha$ defines a partition of κ into $\alpha < \kappa$ many pieces $f^{-1}(\gamma)$ none of which is in $\mathcal{U}$. Then

$$\mathcal{V} = \{X \subseteq \alpha \mid f^{-1}X \in \mathcal{U}\}$$

is a non-principal μ-complete ultrafilter on α, violating the minimality of κ. $\square$

So if there are only boundedly many measurable cardinals, then there is some κ such that there are no non-principal κ-complete ultrafilters on any set.

So we have shown

Theorem (Isbell, 1960)

There is a dense set in $\mathbf{Set}^{op}$ *if and only if there are only boundedly many measurable cardinals.*

Happy Australia Day!