Accessible categories and model theory: hits and misses

Accessible Categories Workshop, Leeds

Tibor Beke

University of Massachusetts
tibor_beke@uml.edu

July 17, 2018

epithet

“ It would still be very welcome to be told that the whole scheme coincides with, say, something known to category theorists in another context. ”

epithet

" It would still be very welcome to be told that the whole scheme coincides with, say, something known to category theorists in another context. "

Wilfred Hodges in [Hodges 2008]

prehistory

Theorem (Schreier, 1926): Every subgroup of a free group is free.

Also true for

- abelian groups
- abelian p-groups
- over a field k:
  - Lie algebras (Witt)
  - commutative (non-associative) algebras (Shirshov)
  - magmas (Kurosh)

Remark: there is still no complete characterization of these varieties known.

conversely: Does this characterize free groups?

Question Suppose G is an infinite group such that all subgroups of G of cardinality less than G, are free. (We will call such a group G *almost free*.) Is G free?

Answer No. Eklof (1970) gave examples of non-free but almost free abelian groups (and Mekler, of non-abelian groups) of cardinality $\aleph_n$, for any $n \in \mathbb{N}^+$. On the other hand ...

- Higman (1951): for singular κ of cofinality ω_0, an almost free group of cardinality κ is necessarily free.
- Hill (1970): for singular κ of cofinality ω_0, an almost free abelian group of cardinality κ is necessarily free.
- Hill (1974): this also holds for singular κ of cofinality ω_1.

Shelah: Singular Cardinal Compactness

Shelah (1974) proved three related statements, each of which has the form:

Let κ be singular and S a structure of size κ. If all substructures of S of size less than κ have property $\mathcal{P}$, then S itself has property $\mathcal{P}$.

(1) structure = abelian group; $\mathcal{P}$ = free

(2) structure = graph; $\mathcal{P}$ = having coloring number $\leqslant \mu$

(3) structure = set of countable sets; $\mathcal{P}$ = having a transversal.

Subsequently, many more examples added: completely decomposable modules, Q-filtered modules ...

what is Singular Cardinal Compactness *really*?

Not an instance of compactness in any classical sense (e.g. for propositional or first-order logic, or compact cardinals)

- Shelah (1974) axiomatizes when his proof works
- Hodges (1982), building on Shelah (unpublished), gives another, elegant axiomatization
- Eklof (2006): SCC is about *an abstract notion of "free"*

What is "free" about a graph with coloring number less than μ?

cells: topological origin

- $D_n = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| \leqslant 1\}$ closed unit n-ball
- $\partial D_n = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| = 1\}$ unit n-1-sphere
- $\partial D_n \xrightarrow{i_n} D_n$ the inclusion.

A pushout p in the category of topological spaces

$$\begin{array}{ccc} \partial D_n & \longrightarrow & X \\ \downarrow{\scriptstyle i_n} & & \downarrow{\scriptstyle p} \\ D_n & \longrightarrow & Y \end{array}$$

is called *attaching an n-cell*. A *cellular space* arises via a (possibly transfinite) sequence of cell attachments. Cellular spaces are topological generalizations of the geometrically more restricted notions of 'simplicial complex' and 'CW-complex'.

I-cellular maps

Let $\mathcal{C}$ be a cocomplete category and I a class of morphisms. The class of *I-cellular maps* is defined as the closure under isomorphisms (in the category of morphisms of $\mathcal{C}$) of well-ordered smooth colimits of pushouts of elements of I.

An object X is I-cellular if the map $\varnothing \to X$ from the initial object to X is I-cellular.

This notion was isolated around 1970 by Quillen, Kan, Bousfield etc. It has proved very handy in homotopical algebra, in constructing weak factorization systems and homotopy model categories.

Note that (with rare exceptions) a cellular map has many cell decompositions, without any preferred / canonical / functorial one.

I-cellular maps: example 1

Let $\mathcal{C}$ be an equational variety of algebras and their homomorphisms (e.g. groups, abelian groups, R-modules ...). Let $A_\varnothing$ be the free algebra on the empty set, $A_\bullet$ the free algebra on a singleton and let $I = \{A_\varnothing \to A_\bullet\}$.

Then I-cellular objects are the same as free algebras.

I-cellular maps: example 2

Fix a ring R and let S be a set of R-modules. Let $I = \{0 \to M \mid M \in S\}$.

Then I-cellular objects are the same as I-decomposable modules.

(This example would make sense in any category with coproducts.)

I-cellular maps: cartoon of example 3

consider the two families of inclusions of graphs

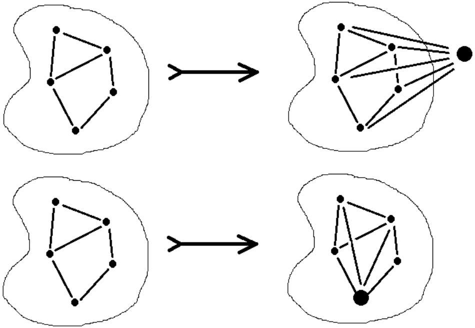

I-cellular maps: example 3 (Rosický)

Work in the category of graphs and graph homomorphisms.
Let μ be a cardinal and let I consist of all inclusions of graphs

$$\langle V, E\rangle \to \langle V \cup \bullet,\ E \cup (V \times \bullet) \cup (\bullet \times V)\rangle$$

where the set V of vertices has cardinality $< \mu$ and $\bullet$ is a singleton (we take a representative set of these morphisms).

Then I-cellular objects are precisely graphs having coloring number $\leqslant \mu$. An I-cell decomposition of a graph is the same data as a well-ordering of vertices satisfying Shelah's criterion: each vertex is connected to $< \mu$ vertices preceding it.

I-cellular maps: example 4 (Rosický)

Work in the category of directed bipartite graphs, i.e. triples $\langle U, V, E \rangle$ where U, V are sets and E is a relation from U to V. Let I consist of the three morphisms

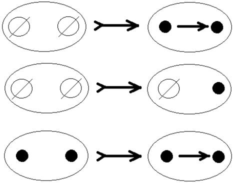

where $\varnothing$ is the empty set, $\bullet$ a singleton. Pushouts by these "create a disjoint edge", "create a vertex in the second partition", "create an edge between existing vertices" respectively. I-cellular objects are precisely directed bipartite graphs $\langle U, V, E \rangle$ that posses at least one transversal $U \rightarrowtail V$: a cell decomposition of a graph provides data for a transversal, and vice versa.

cellular SCC: cleanest version

Let $\mathcal{B}$ be a locally finitely presentable category and I a set of morphisms. Let $X \in \mathcal{B}$ be an object whose size $\|X\|$ is singular. If all subobjects of X of size less than $\|X\|$ are I-cellular, then X itself is I-cellular.

cellular SCC: clarifying remarks

- The *size* of an object X of a locally presentable category is the cardinal predecessor of the least cardinal κ such that X is κ-presentable. (The least such κ will indeed be a successor cardinal.) This notion of size is intrinsic to the category, and coincides with the naive notion of 'cardinality of the underlying set' in all cases of interest.
- It suffices to have 'enough' subobjects of X to be cellular for the conclusion to hold.
- Could have a proper class of generating maps as long as the size of their domains is bounded from above.
- The size of X should be big enough (above the Löwenheim–Skolem number).

Singular Cardinal Compactness: cellular version

Theorem [TB-Rosický 2014]

Let $\mathcal{B}$ be a locally finitely presentable category, μ a regular uncountable cardinal and I a class of morphisms with μ-presentable domains. Let $X \in \mathcal{B}$ be an object with $\max\{\mu, \mathrm{card}(\mathrm{fp}\,\mathcal{B})\} < \|X\|$. Assume

(i) $\|X\|$ is a singular cardinal

(ii) there exists $\phi < \|X\|$ such that for all successor cardinals κ^+ with $\phi < \kappa^+ < \|X\|$, there exists a dense κ^+-filter of $\mathrm{Sub}(X)$ consisting of I-cellular objects.

Then X is I-cellular.

the one that got away

[Hodges 1981: Theorem 5, Example 3]

Let k be a field and K/k a field extension with $|K|$ singular, uncountable, $|K| > |k|$. Suppose that whenever L is an intermediate extension between k and K with $|L| < |K|$ then L is a purely transcendental extension of k. Then K itself is a purely transcendental extension of k.

- This is a corollary of Hodges's version of Shelah's proof
- It is *not* a case of the cellular version of singular compactness!
- It is (I believe) vacuously true: as a consequence of the negative solution to Lüroth's problem, for any k, K as above, there exists L with $|L| < |K|$ that is not a purely transcendental extension of k.

Singular Cardinal Compactness: functorial version

Theorem [TB-Rosický 2014]

Let $\mathcal{A}$ be an accessible category with filtered colimits, $\mathcal{B}$ a finitely accessible category and $F : \mathcal{A} \to \mathcal{B}$ a functor preserving filtered colimits. Let $X \in \mathcal{B}$ be an object with $\max\{\mu_F, \mathrm{card}(\mathrm{fp}\,\mathcal{B})\} < \|X\|$. Assume

(i) $\|X\|$ is a singular cardinal

(ii) there exists $\phi < \|X\|$ such that for all successor cardinals κ^+ with $\phi < \kappa^+ < \|X\|$, the image of F contains a dense κ^+-filter of $\mathrm{Sub}(X)$

(iii) F-structures extend along morphisms.

Then X is in the image of F.

functorial SCC: clarifying remarks

- Let $F : \mathcal{A} \to \mathcal{B}$ be a functor. We say that *F-structures extend along morphisms* if, given any morphism $g : X \to Y$ and object U of $\mathcal{A}$, together with an isomorphism $i : F(U) \to F(X)$ in $\mathcal{B}$, there exists a morphism $f : U \to V$ and isomorphism $j : F(V) \to F(Y)$ such that

$$\begin{array}{ccc} F(U) & \xrightarrow{F(f)} & F(V) \\ \downarrow i & & \downarrow j \\ F(X) & \xrightarrow{F(g)} & F(Y) \end{array}$$

  commutes.
- An object of $\mathcal{B}$ is 'in the image' of the functor F if it is isomorphic to $F(X)$ for some X in $\mathcal{A}$.

about the proof

- Assume first the target category is Pre($\mathcal{C}$), the category of presheaves on a small category $\mathcal{C}$. Adapt the set-based proof of [Hodges 1981]. Fairly easy since the poset of subobjects of any object of Pre($\mathcal{C}$) is a complete distributive lattice.
- Characterize the categories $\mathcal{B}$ that posses a nice enough embedding $\mathcal{B} \to$ Pre($\mathcal{C}$) into a presheaf category so that one can deduce the conclusion for $F : \mathcal{A} \to \mathcal{B}$, given that it holds for the composite $F : \mathcal{A} \to \mathcal{B} \to$ Pre($\mathcal{C}$).
- The functorial version implies both the cellular version and Hodges's version.

question

SCC is a machine for churning out conditionals *à la*

Let κ be singular and S a structure of size κ. If all substructures of S of size less than κ have property $\mathcal{P}$, then S itself has property $\mathcal{P}$.

Such a statement could be 'worthless' for two reasons:

- There exists no such S.
- If all proper substructures of S have property $\mathcal{P}$ then it has property $\mathcal{P}$ (regardless of size).

Is there a way to unify the construction of 'paradoxical' i.e. almost free but non-free objects in *regular* characteristics?

Note that these constructions are fairly few and scattered. They seem to work best for $\aleph_n$ for finite n, or under $V = L$ or similar set-theoretic assumptions.

the first-order Lefschetz Principle

Let Θ be a first order sentence in the signature $+, -, \times, 0, 1$ of rings. The following are equivalent:

- $\mathbb{C} \models \Theta$
- $K \models \Theta$ for every alg. closed field K of characteristic zero
- there is an infinite set of primes P such that $\overline{\mathbb{F}}_p \models \Theta$ for all $p \in P$
- there is an integer $p > 0$ such that $\mathbb{F} \models \Theta$ whenever $\mathbb{F}$ is algebraically closed with $\mathrm{char}(\mathbb{F}) \geqslant p$.

Not so many properties of varieties (let alone schemes) can be formulated as first order(!) sentences in the language of rings.

Nonetheless, there are applications of the first order Lefschetz Principle, such as the Ax–Grothendieck theorem.

Lefschetz Principle à la Weil

" For a given value of the characteristic p, every result, involving only a finite number of points and of varieties, which has been proved for some choice of the universal domain remains valid without restriction; there is but one algebraic geometry of characteristic p for each value of p, not one algebraic geometry for each choice of the universal domain. "

A formal proof of this principle would require "*a formal metamathematical characterization of the type of proposition to which it applies; this would have to depend on the metamathematical, i.e. logical analysis of all our definitions*."

[Weil 1962]

A "universal domain" (as defined by Weil) is an algebraically closed field of infinite transcendence degree over its prime field.

Let K and k be algebraically closed, of the same characteristic, both of cardinality $> \lambda$. Then

- K and k are $\mathcal{L}_{\infty,\lambda}$-equivalent
- any field extension K/k between them is a $\mathcal{L}_{\infty,\lambda}$-elementary embedding.

In particular, any two of Weil's "universal domains" of the same characteristic are $\mathcal{L}_{\infty,\omega}$-equivalent.

This suggests that Weil's principle should mean: algebraic geometric properties that can be formulated as sentences of $\mathcal{L}_{\infty,\omega}$ are true or false across an entire universal domain; and an extension K/k of universal domains *induces an isomorphism* between "algebraic geometry over k" and "algebraic geometry over K".

But *what* are these properties? How can one recognize them in practice?

[Hodges 1973]

"word constructions" (transfinite compositions of extensions by definition, and definable quotients)

Main result: if $\mathcal{S}_1$ and $\mathcal{S}_2$ are $\mathcal{L}_{\infty,\lambda}$-equivalent, and W is a suitable word construction, then $W(\mathcal{S}_1)$ and $W(\mathcal{S}_2)$ are $\mathcal{L}_{\infty,\lambda}$-equivalent.

[Feferman 1970] [Eklof 1973]:

If F is a "λ-local functor" from str_{S_1} to str_{S_2}, then F preserves $\mathcal{L}_{\infty,\lambda}$-equivalence.

To apply this theorem, algebraic geometry has to be "encoded" as structures in a specific signature fixed in advance:

ular theorem of algebraic geometry we are interested in. For example, let us look at Weil's example ([6, pp. 307ff.]) which we rewrite in a form which makes clearer its logical structure. The superscripts on variables name the sort of variable, i.e. $v^{(n)} \in A_n$ as defined above, e.g. $v^{(5)}$ is a variable standing for an abstract variety. For each $0 \leq r \leq n$,

$$\begin{aligned}
&\forall v^{(5)} \Big[(\{v^{(5)} \text{ is complete and has no multiple points}\} \wedge \{\dim v^{(5)} = n\}) \\
&\quad \to \Big(\bigvee_m \exists v_1^{(6)} \cdots \exists v_m^{(6)} \Big[\bigwedge_{i=1}^{m} (\{v_i^{(6)} \text{ is a cycle on } v^{(5)}\} \wedge \{\dim v_i^{(6)} = r\}) \\
&\qquad \wedge \forall w_1^{(6)} \forall w_2^{(6)} \Big[\Big(\bigwedge_{i=1}^{2} \{w_i^{(6)} \text{ is a cycle on } v^{(5)}\} \\
&\qquad\qquad \wedge \{\dim w_1^{(6)} = r\} \wedge \{\dim w_2^{(6)} = n - r\} \\
&\qquad\qquad \wedge \{w_2^{(6)} \text{ intersects } w_1^{(6)} \text{ properly on } v^{(5)}\} \\
&\qquad\qquad \wedge \bigwedge_{i=1}^{m} \{w_2^{(6)} \text{ intersects } v_i^{(6)} \text{ properly on } v^{(5)}\} \Big) \\
&\qquad \to \bigvee_{(a_1 \cdots a_m) \in \omega^m} \Big\{ \deg(w_1^{(6)} \cdot w_2^{(6)}) = \sum_i a_i \deg(v_i^{(6)} \cdot w_2^{(6)}) \Big\} \Big] \Big] \Big) \Big].
\end{aligned}$$

The expressions in braces are relations which we include in the structure $\mathfrak{A} = F(U)$ where $U \in \mathscr{U}_0$. Then the above sentence is a sentence of the language $L_{\infty\omega}$ corresponding to $\mathfrak{A}$, and to be able to apply the theorem

$$\mathbf{S_1}\textbf{–structures} \xrightarrow{\textbf{nice word construction}} \mathbf{S_2}\textbf{–structures}$$

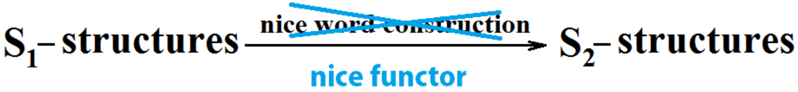
S_1– structures
~~nice word construction~~
nice functor
S_2– structures

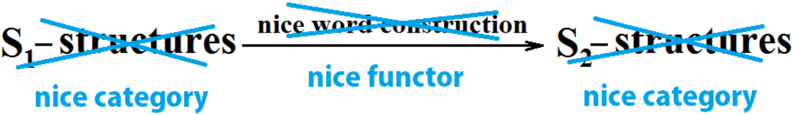
S1– structures
nice word construction
S2– structures
nice category
nice functor
nice category

goals

Create a calculus of “λ–equivalence” for objects in (suitable) categories that

- is applicable to a wide class of categories
- specializes to $\mathcal{L}_{\infty,\lambda}$-equivalence for categories of structures
- possesses all formal properties of $\mathcal{L}_{\infty,\lambda}$-equivalence, but can be formulated without reference to signature or language (using *just* a category as background)
- comes with a matching notion of “λ-embedding” that specializes to $\mathcal{L}_{\infty,\lambda}$-embedding
- extends the Feferman–Eklof theorem to a situation $\mathcal{C} \xrightarrow{F} \mathcal{D}$ where the categories $\mathcal{C}$, $\mathcal{D}$ need not be assumed to be categories of structures.

goals

Create a calculus of "λ–equivalence" for objects in (suitable) categories that

- is applicable to a wide class of categories
- specializes to $\mathcal{L}_{\infty,\lambda}$-equivalence for categories of structures
- possesses all formal properties of $\mathcal{L}_{\infty,\lambda}$-equivalence, but can be formulated without reference to signature or language (using *just* a category as background)
- comes with a matching notion of "λ-embedding" that specializes to $\mathcal{L}_{\infty,\lambda}$-embedding
- extends the Feferman–Eklof theorem to a situation $\mathcal{C} \xrightarrow{F} \mathcal{D}$ where the categories $\mathcal{C}$, $\mathcal{D}$ need not be assumed to be categories of structures.

This is 90% done [TB-Rosický 2017].

Karp's theorem

Let $\mathcal{S}_1$, $\mathcal{S}_2$ be structures for a language in $\mathcal{L}_{\infty,\lambda}$. A partial isomorphism $\mathcal{S}_1 \leftarrowtail D \rightarrowtail \mathcal{S}_2$ is a structure D equipped with embeddings into $\mathcal{S}_1$, $\mathcal{S}_2$. A set I of partial isomorphisms $\mathcal{S}_1 \leftarrowtail D_i \rightarrowtail \mathcal{S}_2$ satisfies the "$< \lambda$ back and forth property" if

- for every $i \in I$ and set $X \subseteq \mathcal{S}_1$ with $|X| < \lambda$, there exists $j \in I$ such that $\mathcal{S}_1 \leftarrowtail D_j \rightarrowtail \mathcal{S}_2$ extends $\mathcal{S}_1 \leftarrowtail D_i \rightarrowtail \mathcal{S}_2$ and the image of D_j in $\mathcal{S}_1$ contains X, and
- for every $i \in I$ and set $Y \subseteq \mathcal{S}_2$ with $|Y| < \lambda$, there exists $j \in I$ such that $\mathcal{S}_1 \leftarrowtail D_j \rightarrowtail \mathcal{S}_2$ extends $\mathcal{S}_1 \leftarrowtail D_i \rightarrowtail \mathcal{S}_2$ and the image of D_j in $\mathcal{S}_2$ contains Y.

Karp's theorem

$\mathcal{S}_1$ and $\mathcal{S}_2$ are "$< \lambda$ back and forth equivalent" if there is a set of partial isomorphisms between them satisfying this property.

Theorem (Karp) $\mathcal{S}_1$ and $\mathcal{S}_2$ are $< \lambda$ back and forth equivalent if and only if they are $\mathcal{L}_{\infty,\lambda}$-equivalent.

There are variants for $\mathcal{L}_{\infty,\lambda}$-embeddings and for fragments of $\mathcal{L}_{\infty,\lambda}$ with prescribed quantifier prefixes.

For $\mathcal{L}_{\omega,\omega}$, Ehrenfreucht-Fraïssé games are a *little bit* analogous.

mono-generated categories

Let $\mathcal{C}$ be a category and λ a regular cardinal.

- object $X \in \mathcal{C}$ is λ-*generated* if $\hom(X, -)$ preserves those λ-directed diagrams of monos that exist in $\mathcal{C}$
- $\mathcal{C}$ is λ-*mono-generated* if every object can be written as the colimit of a λ-directed diagram of monos and λ-generated objects.

Remark: no assumption that *all* λ-directed diagrams of monos have a colimit, and no assumption that there is a *set* of λ-generated objects whose λ-directed colimits generate all objects.

Logical implications

$$\begin{array}{ccccc}
\text{AEC} & \longrightarrow & \text{accessible} & \longrightarrow & \text{mono-accessible} \\
 & & \downarrow & & \downarrow \\
 & & \text{class-accessible} & \longrightarrow & \text{mono-generated}
\end{array}$$

none of which is reversible.

example

Let FreeAbMono be the subcategory of abelian groups with objects the free abelian groups and morphisms the injective homomorphisms.

FreeAbMono is finitely mono-generated. Indeed, any directed colimit of monos in FreeAbMono — if it exists! — is "standard" (computed on underlying sets). Any finitely generated free abelian group is finitely generated as an object of FreeAbMono. Any free abelian group is directed colimit of its finitely generated (free) subgroups.

Whether FreeAbMono is accessible or not depends on set theory.

spans, dense sets of spans, λ-equivalence

A *span* between $X, Y \in \mathcal{C}$ is $X \leftarrowtail U \rightarrowtail Y$ where the arrows are mono. A set of spans $X \leftarrowtail U_i \rightarrowtail Y$ is λ-*dense* if (0) it is non-empty and (1) for all $i \in I$ and monomorphism $X \leftarrowtail G$ with λ-generated G, there exist $j \in I$ and morphisms $G \to U_j$ and $U_i \to U_j$ such that

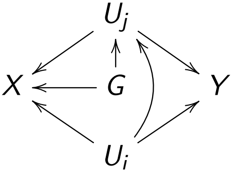

commutes and (2) the symmetric condition, with respect to "test objects" $G \rightarrowtail Y$ with λ-generated G, holds too.

$X \sim_\lambda Y$, i.e. X and Y are λ-*equivalent*, if there exists a λ-dense set of spans between them.

λ-embeddings

A morphism $f : X \to Y$ is called a λ-*embedding* if there is a λ-dense set $\mathcal{S}$ of spans between X and Y such that for any monomorphism $g : G \rightarrowtail X$ with λ-generated G there exist $X \xleftarrow{u} U \xrightarrow{v} Y \in \mathcal{S}$ and $t : G \to U$ such that in the diagram

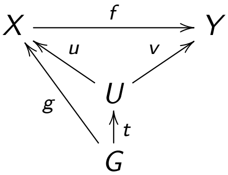

$ut = g$ and $vt = fg$.

(A bit unlikely sounding, perhaps, but motivated by a variant of Karp's theorem for $\mathcal{L}_{\infty,\lambda}$-embeddings. It works.)

as expected

Work in a λ-mono-generated category.

- $\sim_\lambda$ is an equivalence relation
- if $X \sim_\lambda Y$ and X is λ-generated, then X and Y are isomorphic
- any λ-embedding is a monomorphism
- if $f : X \to Y$ is a λ-embedding with X, Y λ-generated, then f is an isomorphism
- λ-embeddings are closed under composition
- if g, gf are λ-embeddings, so is f.

The proofs play with spans.

elementary chains

The Tarski–Vaught theorem states that if $\mathcal{S}$ is the union of a well-ordered, continuous chain of structures $\mathcal{S}_\alpha$, where each $\mathcal{S}_\alpha \subseteq \mathcal{S}_{\alpha+1}$ is an elementary embedding, then $\mathcal{S}_\alpha \subseteq \mathcal{S}$ is an elementary embedding for each α.

A better way of saying this is the following: let str_S be the category of S-structures and homomorphisms (in some fixed first-order signature S), and let elem_S be the category of S-structures and elementary embeddings. Then the canonical inclusion

$$\mathit{elem}_S \rightarrowtail \mathit{str}_S$$

creates filtered colimits.

Tarski–Vaught type of results

Work in a λ-mono-generated category $\mathcal{A}$ and assume that colimits referred to exist. Let λ-*emb*$(\mathcal{A})$ denote the subcategory of $\mathcal{A}$ whose morphisms are λ-embeddings.

- λ-directed colimit of natural transformations consisting of λ-embeddings between diagrams of λ-embeddings, is itself a λ-embedding
- the colimit cocone of a λ-directed diagram of λ-embeddings consists of λ-embeddings
- the inclusion λ-*emb*$(\mathcal{A}) \rightarrowtail \mathcal{A}$ creates λ-directed colimits
- in a finitely mono-generated category, a transfinite composite of finitary embeddings is a finitary embedding.

The proofs play with spans a little longer.

consistency with $\mathcal{L}_{\infty,\lambda}$-equivalence

Let S be a language for $\mathcal{L}_{\infty,\lambda}$ and let str_S denote the category of S-structures and homomorphisms, while emb_S is the category of S-structures and embeddings. For S-structures X and Y, the following are equivalent:

(1) X and Y are $\mathcal{L}_{\infty\lambda}$-elementary equivalent

(2) there is a set of partial isomorphisms between X and Y satisfying the $< \lambda$ back-and-forth property

(3) $X \sim_\lambda Y$ in $emb(\Sigma)$

(4) $X \sim_\lambda Y$ in $str(\Sigma)$.

The equivalence of (3) and (4) is not quite trivial. Need to play with factorizations of spans.

consistency with $\mathcal{L}_{\infty,\lambda}$-embedding

For an embedding $f : X \to Y$ of S-structures the following are equivalent:

(1) f is an $\mathcal{L}_{\infty\lambda}$-elementary embedding

(2) there is a set I of partial isomorphisms between X and Y satisfying the $< \lambda$ back-and-forth property, and such that for every subset Z of X of cardinality less than λ there is $h \in I$ such that $f(z) = h(z)$ for every $z \in Z$

(3) f is a λ-embedding in emb_S

(4) f is a λ-embedding in str_S.

Let $\mathcal{A}$ be a λ-mono-generated category and $F : \mathcal{A} \to \mathcal{B}$ a functor preserving monomorphisms and λ-directed colimits of monomorphisms.

- If $X \sim_\lambda Y$ then $F(X) \sim_\lambda F(Y)$.
- If $f \in \mathcal{A}$ is a λ-embedding, so is $F(f)$.

These are easy to prove, and specialize to the main result of Feferman–Eklof.

Not so useful in practice, since functors preserving monos are rare.

basic universal theories

Let $\lambda \geqslant \kappa$ be regular cardinals. Let $\forall_{\lambda\kappa}$ denote the class of theories in $\mathcal{L}_{\infty\kappa}$ that can be axiomatized by a set of sentences of the form

$$\forall \mathbf{x}(\phi \implies \psi)$$

where ϕ and ψ are built from atomic formulas using conjunction and disjunction only, where $\bigwedge_{i\in I}$ is permitted for $|I| < \lambda$ only, while $\bigvee_{j\in J}$ is permitted for J of any cardinality.

If T is a $\forall_{\lambda\kappa}$ theory, let $\mathrm{Mod}(T)$ denote the category of models and homomorphisms of T.

basic universal theories

- more or less, models of Horn theories in $\mathcal{L}_{\lambda,\kappa}$
- contains all infinitary equational varieties, and quasi-varieties
- if T is $\forall_{\lambda\kappa}$ then Mod(T) is accessible with well-behaved image factorizations of morphisms
- possible to characterize intrinsically categories equivalent to Mod(T) for some basic universal T, but does not seem worth the trouble to spell it out.

theorem

Let $\mathcal{A}$ be a λ-mono-generated category, $T \in \forall_{\lambda\kappa}$, and $F : \mathcal{A} \to \mathrm{Mod}(T)$ a functor that takes λ-directed colimits of monos to colimits.

- $X \sim_\lambda Y$ implies $F(X) \sim_\lambda F(Y)$
- If $f \in \mathcal{A}$ is a λ-embedding, so is $F(f)$.

– extends a beautiful result of Eklof

– useful since λ-directed colimit preserving functors abound

– proof uses accessibility mixed with image factorizations.

problem

Prove or disprove: if $\mathcal{A}$ and $\mathcal{B}$ are λ-accessible categories and $F : \mathcal{A} \to \mathcal{B}$ is a functor preserving λ-directed colimits, then $X \sim_\lambda Y$ implies $F(X) \sim_\lambda F(Y)$.

This is the "missing 10%".

Note that the statement is true under the (somewhat) restrictive hypotheses that F preserves monos, or that $\mathcal{B}$ is the category of models of a basic λ-universal theory (even if $\mathcal{A}$, $\mathcal{B}$ are only λ-mono-generated).

sample application

Let S_1, S_2 be languages for $\mathcal{L}_{\infty,\lambda}$ and let $E : \mathrm{Mod}(S_1) \to \mathrm{Mod}(S_2)$ be an equivalence of categories.

- If X and Y are $\mathcal{L}_{\infty,\lambda}$-equivalent S_1-structures, then $E(X)$ and $E(Y)$ are $\mathcal{L}_{\infty,\lambda}$-equivalent S_2-structures.
- If f is an $\mathcal{L}_{\infty,\lambda}$-elementary embedding of S_1-structures, then $E(f)$ is an $\mathcal{L}_{\infty,\lambda}$-elementary embedding of S_2-structures.

This seems hard to prove syntactically or with Karp's theorem. The assumption (that the mere *categories* $\mathcal{C}_1 = \mathrm{Mod}(S_1)$ and $\mathcal{C}_2 = \mathrm{Mod}(S_2)$ are equivalent) does not allow one to relate the S_1-structure of an object of $\mathcal{C}_1$ to its S_2-structure when thought of as an object of $\mathcal{C}_2$.

sample application

If two groups are $\mathcal{L}_{\infty,\lambda}$-equivalent then so are their abelianizations.

Not easy to prove syntactically, since abelianization is a quotient (without definable representatives).

Similarly, any morphism $T_1 \to T_2$ between (essentially) algebraic theories induces an adjunction between the locally presentable categories $\mathrm{Mod}(T_1)$ and $\mathrm{Mod}(T_2)$. Both the left and the right adjoints preserve (filtered enough) colimits, a fortiori, λ-equivalence for λ big enough.

to do

- Can one characterize functors preserving λ-equivalence?
- Analogue of the calculus of λ-equivalence where spans $\bullet \leftarrowtail \bullet \rightarrowtail \bullet$ are not required to consist of monos?
- Analogue of the calculus of λ-equivalence for 2-categories, where pseudo-limits ("homotopy limits and colimits") are much better behaved than ordinary limits?
- Characterization of λ-equivalence in terms of zig-zag of λ-embeddings?
- There are categories that, for each regular cardinal λ, possess only a *set* of λ-equivalence classes of objects. How is this related to Shelah's classification theory (of models in terms of cardinal invariants)?

and even more to do

- Understand "back-and-forth equivalence" as "local isomorphism" in the sense of a suitable Grothendieck topology.
- Are categories of structured objects in algebraic geometry (e.g. varieties, schemes, ringed spaces, algebraic spaces, possibly over a base) λ-mono-generated?
- What accounts for the (observed) invariance of algebraic geometry with respect to base extensions between algebraically closed fields (not just between universal domains)?

problem: elementary equivalence as category theory?

Ultraproducts of structures need not be invariant even under *isomorphisms* of categories. The following is an unpublished example due to Börger.

Consider the theory T in the one-sorted signature $+, *, 0, 1, \mathcal{P}$ where $\mathcal{P}$ is a propositional constant. The axioms of T are the axioms of fields. Consider the functor $F : \mathrm{Mod}(T) \to \mathrm{Mod}(T)$ defined on objects by

$$F\langle \mathbb{F}, \mathcal{P}\rangle = \begin{cases} \langle \mathbb{F}, \mathcal{P}\rangle & \text{if } \mathrm{char}(\mathbb{F}) > 0 \\ \langle \mathbb{F}, \neg\mathcal{P}\rangle & \text{if } \mathrm{char}(\mathbb{F}) = 0 \end{cases}$$

On field homomorphisms, F is the identity. This makes sense since they preserve characteristics.

problem: elementary equivalence as category theory?

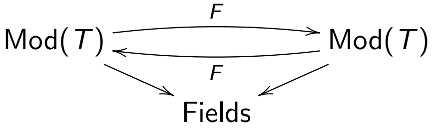

commutes, where the slanted arrows forget $\mathcal{P}$.

F is an isomorphism(!) of categories.

problem: elementary equivalence as category theory?

Let $\mathcal{U}$ be a non-principal ultrafilter on the set P of primes, and $\mathbb{F}_p$ any field of characteristic p.

$$\prod_{\mathcal{U}/\mathrm{P}} \langle \mathbb{F}_p, \top \rangle = \langle \mathbb{F}_0, \top \rangle$$

where $\mathbb{F}_0$ is a certain field of characteristic 0.

$$\prod_{\mathcal{U}/\mathrm{P}} F\langle \mathbb{F}_p, \top \rangle = \prod_{\mathcal{U}/\mathrm{P}} \langle \mathbb{F}_p, \top \rangle = \langle \mathbb{F}_0, \top \rangle$$

$$F(\prod_{\mathcal{U}/\mathrm{P}} \langle \mathbb{F}_p, \top \rangle) = F\langle \mathbb{F}_0, \top \rangle = \langle \mathbb{F}_0, \bot \rangle$$

so F is an automorphism of $\mathrm{Mod}(T)$ that does not respect ultraproducts.

problem: elementary equivalence as category theory?

- Note that F is not an $\mathcal{L}_{\omega,\omega}$-definable functor (though it is $\mathcal{L}_{\omega_1,\omega}$-definable).
- For any category with products and filtered colimits, a good analogue of ultraproducts *can* be constructed.
- Börger's counterexample does not exclude that ultrapowers (or ultralimits) can be defined in the language of category theory.
- Imitating the Keisler-Shelah theorem, Frayne's theorem, or the ultralimit theorem, maybe a good analogue of elementary equivalence can be introduced for a class of categories including Mod(T) for coherent theories T (much as $\sim_\lambda$ is a good analogue of $\mathcal{L}_{\infty,\lambda}$-equivalence for a class of categories including Mod(T) for basic $\mathcal{L}_{\infty,\lambda}$ theories T).

problem: first order recognition principle

Characterize (solely in the language of category theory) categories that are equivalent to Mod(T), the category of models and homomorphisms of some first order theory T.

- Necessary conditions known, but no necessary and sufficient conditions.
- Recognition principles exist for certain fragments of first order logic, such as varieties, quasi-varieties, essentially algebraic theories etc.
- For fragments of $\mathcal{L}_{\omega,\omega}$ where a recognition theorem is known, it is well understood to what extent the theory T can be recovered from Mod(T).

problem: semantics of language change

There exist syntactic (algorithmic) notions of

- *Skolemization*
- *Herbrandization*
- *Morley-fication*

of theories.

What are their analogues for categorical logic?

Note that a change of underlying signature changes the notion of homomorphism of models. E.g. “the” Skolemization of $\mathcal{C}$ is likely to be a certain category $\mathcal{D}$ (axiomatizable in a smaller fragment of logic) and a functor $\mathcal{D} \to \mathcal{C}$ that is bijective on isomorphism classes, but neither full nor faithful. Is there a functorial version of Skolemization etc., with a corresponding universal property?

problem: rank spectrum

Any object X of an accessible category $\mathcal{A}$ has a *rank*:

$$\mathrm{rk}(X) = \inf \{\kappa \mid \mathrm{hom}_{\mathcal{A}}(X, -) \text{ commutes with } \kappa\text{-filtered colimits in } \mathcal{A}\}$$

Agrees for successor of the cardinality of 'underlying set' for finitely accessible categories.

Behaves in a 'smoother', less set theory dependent way for κ-accessible categories.

problem (AEC)

Any object X of an accessible category $\mathcal{A}$ has a *rank*:

$$\mathrm{rk}(X) = \inf \left\{ \kappa \mid \mathrm{hom}_{\mathcal{A}}(X, -) \text{ commutes with } \kappa\text{-filtered colimits in } \mathcal{A} \right\}$$

Say that $\mathcal{A}$ is *not Łoś-like* if its spectrum fails Łoś's conjecture, i.e.

There exist a proper class C_0 and a proper class C_1 of cardinals such that for any $\kappa \in C_0$, $\mathcal{A}$ has a unique isomorphism type of κ-presentable objects but for any $\kappa \in C_1$, $\mathcal{A}$ has non-isomorphic κ-presentable objects.

Is there any accessible category that is not Łoś-like?

Their existence would not violate Shelah's conjecture if they do not arise from an AEC.

problem (AEC)

The analogue of *Hanf numbers* for accessible categories would be: given cardinals κ, λ, let

$$H_{\kappa,\lambda} = \sup \{ \text{rk}(X) \mid X \in \mathcal{A} \}$$

ranging over categories $\mathcal{A}$ that are

- κ-accessible
- (essentialy) small
- satisfy card $\text{pres}_\kappa(\mathcal{A}) \leqslant \lambda$,

where $\text{pres}_\kappa(\mathcal{A})$ is the (skeleton of) the full subcategory of $\mathcal{A}$ consisting of its κ-presentable objects, and card is the cardinality of its morphisms.

$h_{\kappa,\lambda}$ exists, since the sup is over a *set*.

problem (AEC)

Note that for some κ, there may exist a proper class of inequivalent small κ-accessible categories. For example, any group (considered as a category) is finitely accessible. Its (unique) object is finitely presentable.

- In the above definition, can the condition 'card $\mathrm{pres}_\kappa(\mathcal{A}) \leqslant \lambda$' be omitted?
- Any relation to Hanf numbers of $\mathcal{L}_{\kappa,\lambda}$?

problem (AEC)

Recall the constructions of [Shelah-Hart 2008] and [Shelah-Villaveces 2008]: AEC with unique isomorphism classes in many cardinalities.

[Friedman-Korwien 2010]: AEC whose being categorical (in certain cardinalities) depends on set theory.

Do these constructions give examples of (finitely) accessible categories with 'interesting' rank spectrum?

- T is your favorite theory, of your preferred logical strength, formulated in the signature $\mathcal{S}$ of the logic $\mathcal{L}$ of your favorite flavor
- you will be able to form $\mathrm{Mod}(T)$, the category whose objects are models of T (in your favorite semantics)
- and whose morphisms are
    - $\mathcal{S}$-homomorphisms, or maybe
    - $\mathcal{L}$-elementary morphisms, or maybe
    - $\mathcal{S}$-homomorphism preserving formulas in some fragment $\mathcal{F}_1$, and reflecting formulas in some fragment $\mathcal{F}_2$ of $\mathcal{L}$.

If you're lucky, the model-theoretic problem you're concerned with can be expressed in terms of the category (*qua* category) $\mathcal{C} = \mathrm{Mod}(T)$, in the language of category theory.

To be answered ...

- So what? *Why* could it be advantageous to go over to the category side?
- Just *what* is the language of category theory?
- And what properties involving models can be so expressed?

If you're lucky, $\mathcal{C} = \mathrm{Mod}(T)$ (as an abstract category) can be analyzed in ways that have no analogues in model theory; that are not easy to locate within the Tarski-style 'toolbox'.

Every accessible category is equivalent to . . .

- the free κ-filtered cocompletion of a small category
- the category of functors $\mathcal{D} \to \mathrm{Set}$ that preserve a prescribed set of limits and colimits
- the category of models of an $\mathcal{L}_{\kappa,\lambda}$ theory with axioms

Every locally presentable category is equivalent to ...

- a full reflexive subcategory of a functor category $\mathrm{Set}^{\mathcal{C}}$
- the category of limit-preserving functors $\mathcal{D} \to \mathrm{Set}$
- the category of models and homomorphisms of an essentially algebraic theory.

language of category theory (finitary case)

- two-sorted signature (objects O, morphisms M)
- maps $s, t : M \to O$, $i : O \to M$
- tertiary relation $C(-,-,-) \subseteq M^3$ (composability)

$\mathcal{L}_{\mathrm{fin}}(\mathrm{Cat})$ is the first order language of this structure. (Permit O, M to be proper classes in some class theory like Gödel-Bernays.)

Let $\mathcal{C}$ be a category. The following can be expressed in $\mathcal{L}_{\mathsf{fin}}(\mathsf{Cat})$

- $\mathcal{C}$ has pullbacks; finite limits
- has a projective generator; an injective cogenerator
- has a subobject classifier; is an elementary topos [Lawvere 1970]

The following cannot *(prove this!)*

- $\mathcal{C}$ is small; locally small
- has countable coproducts; is complete
- has a small dense subcategory
- every object of $\mathcal{C}$ has a rank.

• $\mathcal{L}_{\mathsf{fin}}(\mathsf{Cat})$ is the language of categories as an essentially algebraic theory; the algebra homomorphisms are exactly the functors.

• $\mathcal{L}_{\mathsf{fin}}(\mathsf{Cat})$ permits the statement of properties that category theorists regard as "un-category-theory-like": properties that are not invariant under equivalences of categories. (For example, "there exists exactly two initial objects".)

• [Freyd 1972] characterizes those sentences of $\mathcal{L}_{\mathsf{fin}}(\mathsf{Cat})$ that are invariant under equivalences of categories. See [Makkai 1996] for a much more sophisticated treatment.

language of category theory (infinitary case)

Not clear! Could experiment with

- $\mathcal{L}_{\kappa,\lambda}$(Cat)
- $\mathcal{L}_{\infty,\lambda}$(Cat)
- $\mathcal{L}_{\infty,\infty}$(Cat)
- $\mathcal{L}_{\infty^+,\infty}$(Cat) (permitting class-size $\bigwedge$ and $\bigvee$)
- monadic second order $\mathcal{L}_{\text{fin}}$(Cat)
- $\mathcal{L}_{\text{fin}}$(Cat) with cardinality quantifiers
- categories with 'exactness conditions'
- generalized sketches [Makkai 1994]

[Makkai 2000] claims to characterize the fragment of $\mathcal{L}_{\infty,\omega}$(Cat) that is invariant under equivalences of categories, but it is unclear (to me) if details were ever published.

(re)construction

You're given Mod(T) as an abstract category. Can you reconstruct

- the logical fragment T was axiomatized in?
- T itself (to some extent)?
- the size of a model, as an object of Mod(T)?
- types?
- elementary equivalence of two models?

references

Beke, T.: Sheafifiable homotopy model categories.
Math. Proc. Camb. Phil. Soc., 2000

Hodges, W,: Remarks on Compositionality. In: Dependence Logic, Birkhäuser 2008

Joyal, A.: Letter to Alexander Grothendieck, dated 11/4/84

Weil, *Foundations of Algebraic Geometry*, 1962