Classification and Measure for Algebraic Fields

Russell Miller

Queens College & CUNY Graduate Center

Logic Seminar
Cornell University
23 August 2017

The eternal question

Goal today: explain how to classify the elements of various classes $\mathcal{C}$ of countable structures, up to isomorphism. Usually $|\mathcal{C}| = 2^\omega$.
(Primary example: $\mathcal{C} = \{$all algebraic field extensions of $\mathbb{Q}\}$.)

The eternal question

Goal today: explain how to classify the elements of various classes $\mathcal{C}$ of countable structures, up to isomorphism. Usually $|\mathcal{C}| = 2^{\omega}$. (Primary example: $\mathcal{C} = \{$all algebraic field extensions of $\mathbb{Q}\}$.)

Here is the basic difficulty with doing classifications:

WHAT DO MATHEMATICIANS WANT?

Formally, any bijection Φ from a class $\mathcal{C}$ onto another class $\mathcal{D}$ could be called a classification of the elements of $\mathcal{C}$.

The eternal question

Goal today: explain how to classify the elements of various classes $\mathcal{C}$ of countable structures, up to isomorphism. Usually $|\mathcal{C}| = 2^{\omega}$. (Primary example: $\mathcal{C} = \{$all algebraic field extensions of $\mathbb{Q}\}$.)

Here is the basic difficulty with doing classifications:

WHAT DO MATHEMATICIANS WANT?

Formally, any bijection Φ from a class $\mathcal{C}$ onto another class $\mathcal{D}$ could be called a classification of the elements of $\mathcal{C}$.

Informally, a good classification also requires that:

- We should already know $\mathcal{D}$ pretty well.
- We should be able to compute Φ and Φ^{-1} fairly readily – which starts with choosing good representations of $\mathcal{C}$ and $\mathcal{D}$.

Classes of countable structures

A structure $\mathcal{A}$ with domain ω (in a fixed language) is identified with its atomic diagram $\Delta(\mathcal{A})$, making it an element of 2^ω. We consider classes of such structures, e.g.:

$$Alg = \{D \in 2^\omega : D \text{ is an algebraic field of characteristic } 0\}.$$

$$ACF_0 = \{D \in 2^\omega : D \text{ is an ACF of characteristic } 0\}.$$

$$\mathcal{T} = \{D \in 2^\omega : D \text{ is an infinite finite-branching tree}\}.$$

$$\text{TFAb}_n = \{D \in 2^\omega : D \text{ is a torsion-free abelian group of rank } n\}.$$

On each class, we have the equivalence relation $\cong$ of isomorphism.

Topology on *Alg* and $Alg/\cong$

Alg inherits the subspace topology from 2^ω: basic open sets are

$$\mathcal{U}_\sigma = \{D \in Alg : \sigma \subset D\},$$

determined by finite fragments σ of the atomic diagram D.

We then endow the quotient space $Alg/\cong$ of $\cong$-classes $[D]$, modulo isomorphism, with the quotient topology:

$$\mathcal{V} \subseteq Alg/\cong \text{ is open } \iff \{D \in Alg : [D] \in \mathcal{V}\} \text{ is open in } Alg.$$

Thus a basic open set in $Alg/\cong$ is determined by a finite set of polynomials in $\mathbb{Q}[X]$ which must each have a root (or several roots) in the field.

Examining this topology

The quotient topology on $Alg/\cong$ is not readily recognizable. The isomorphism class of the algebraic closure $\overline{\mathbb{Q}}$ (which is universal for the class *Alg*) lies in *every* nonempty open set $\mathcal{U}$, since if $F \in \mathcal{U}$, then some finite piece of the atomic diagram of F suffices for membership in $\mathcal{U}$, and that finite piece can be extended to a copy of $\overline{\mathbb{Q}}$.

In contrast, the prime model $[\mathbb{Q}]$ lies in no open set $\mathcal{U}$ except the entire space $Alg/\cong$. If $\mathbb{Q} \in \mathcal{U}$, then some finite piece of the atomic diagram of $\mathbb{Q}$ suffices for membership in $\mathcal{U}$, and this piece can be extended to a copy of any algebraic field.

This does not noticeably illuminate the situation.

Expanding the language for *Alg*

Classifying $Alg/\cong$ properly requires a jump, or at least a fraction of a jump. For each $d > 1$, add to the language of fields a predicate R_d:

$$\models_F R_d(a_0, \ldots, a_{d-1}) \iff X^d + a_{d-1}X^{d-1} + \cdots + a_0 \text{ has a root in } F.$$

Write Alg^* for the class of atomic diagrams of algebraic fields of characteristic 0 in this expanded language.

Now we have computable reductions in both directions between $Alg^*/\cong$ and Cantor space 2^ω, and these reductions are inverses of each other. Hence $Alg^*/\cong$ is homeomorphic to 2^ω.

2^ω is far more recognizable than the original topological space $Alg/\cong$ (without the root predicates R_d). We consider this computable homeomorphism to be a legitimate classification of the class *Alg*, and therefore view the root predicates (or an equivalent) as essential for effective classification of *Alg*.

Computing this homeomorphism

$X^4 - 2$

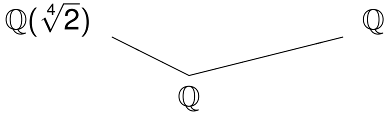

Computing this homeomorphism

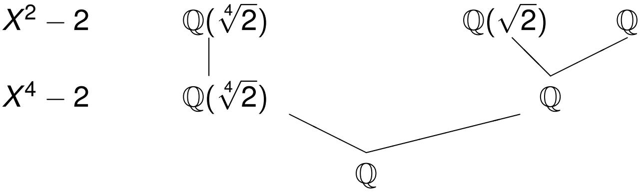

Computing this homeomorphism

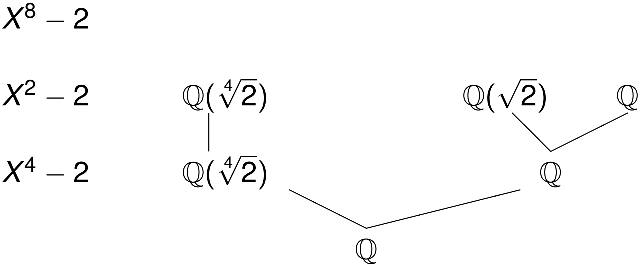

Computing this homeomorphism

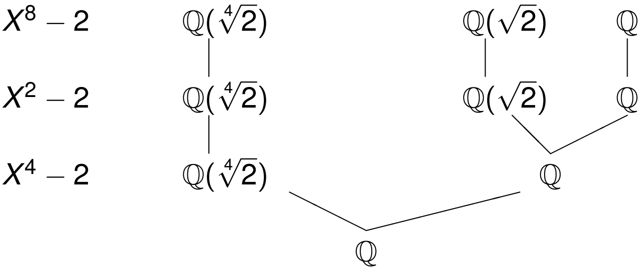

Computing this homeomorphism

$X^2 - \sqrt[4]{2}$

$X^8 - 2$

$X^2 - 2$

$X^4 - 2$

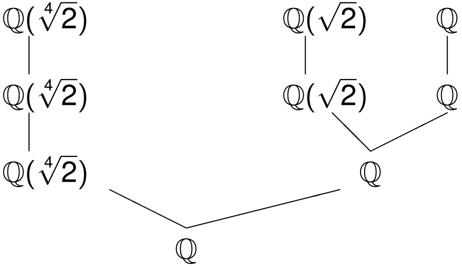

Computing this homeomorphism

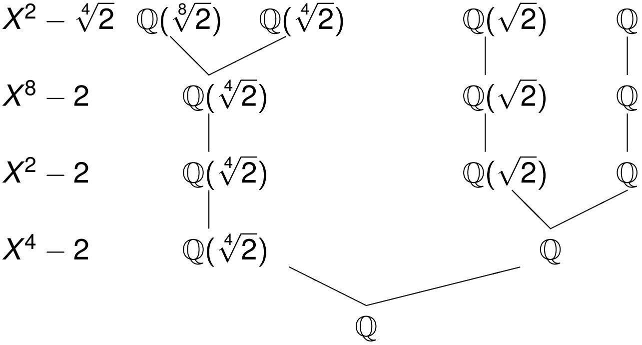

Computing this homeomorphism

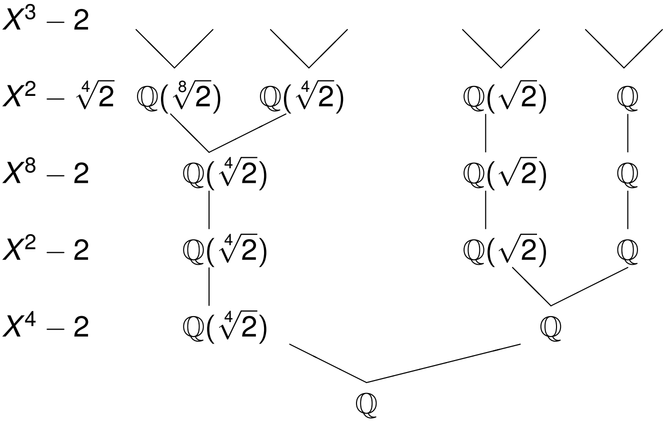

What do the R_d add?

We do *not* have the same reductions between $Alg/\cong$ and 2^ω: these are not homeomorphic. This seems strange: all R_d are definable in the smaller language, so how can they change the isomorphism relation?

The answer is that they do not change the underlying set: we have a bijection between Alg and Alg^* which respects $\cong$. However, the relations R_d change the topology on $Alg^*/\cong$ from that on $Alg/\cong$. (These are both the quotient topologies of the subspace topologies inherited from 2^ω.)

We do have a continuous map from $Alg^*/\cong$ onto $Alg/\cong$, by taking reducts, and so $Alg/\cong$ is also compact. This map is bijective, but its inverse is not continuous.

Too much information

Now suppose that, instead of merely adding the dependence relations R_d, we add *all* computable Σ_1^c predicates to the language. That is, instead of the algebraic field F, we now have its jump F'.

Fact

$$F \cong K \iff F' \cong K'.$$

However, the class Alg' of all (atomic diagrams of) jumps of algebraic extensions of $\mathbb{Q}$, modulo $\cong$, is no longer homeomorphic to 2^ω. In particular, the Σ_1^c property

$$(\exists p \in \mathbb{Q}[X])(\exists x \in F)\ [p \text{ irreducible of degree } > 1\ \&\ p(x) = 0]$$

holds just in those fields $\not\cong \mathbb{Q}$. Therefore, the isomorphism class of $\mathbb{Q}$ forms a singleton open set in the space $Alg'/\cong$.
(Additionally, $Alg'/\cong$ is not compact.)

Related spaces

From the preceding discussion, we infer that the root predicates are exactly the information needed for a nice classification of *Alg*.

(What does "nice" mean here? To be discussed....)

For another example, consider the class $\mathcal{T}$ of all finite-branching infinite trees, under the predecessor function P. As before, we get a topological space $\mathcal{T}/\cong$, which is not readily recognizable. (There is still a prime model, with a single node at each level, but no universal model.)

The obvious predicates to add are the *branching predicates* B_n:

$$\models_T B_n(x) \iff \exists^{=n} y\ (P(y) = x).$$

Which yield...

The enhanced class $\mathcal{T}^*$, in the language with the branching predicates, again has a nice classification. Let $T_{m,0}, T_{m,1}, \ldots$ list all finite trees of height exactly m. Given $T \in \mathcal{T}^*$, we can find the unique number $f(0)$ with $T_{1,f(0)} \cong T^{<2}$, where $T^{<2}$ is just T chopped off after level 1.

Which yield...

The enhanced class $\mathcal{T}^*$, in the language with the branching predicates, again has a nice classification. Let $T_{m,0}, T_{m,1}, \ldots$ list all finite trees of height exactly m. Given $T \in \mathcal{T}^*$, we can find the unique number $f(0)$ with $T_{1,f(0)} \cong T^{<2}$, where $T^{<2}$ is just T chopped off after level 1.

Next consider those trees in $T_{2,0}, T_{2,1}, \ldots$ with $T_{2,i}^{<2} \cong T^{<2}$. Choose $f(1)$ so that $T^{<3}$ is isomorphic to the $f(1)$-th tree on this list. Continue choosing $f(2), f(3), \ldots$ recursively this way.

Which yield...

The enhanced class $\mathcal{T}^*$, in the language with the branching predicates, again has a nice classification. Let $T_{m,0}, T_{m,1}, \ldots$ list all finite trees of height exactly m. Given $T \in \mathcal{T}^*$, we can find the unique number $f(0)$ with $T_{1,f(0)} \cong T^{<2}$, where $T^{<2}$ is just T chopped off after level 1.

Next consider those trees in $T_{2,0}, T_{2,1}, \ldots$ with $T_{2,i}^{<2} \cong T^{<2}$. Choose $f(1)$ so that $T^{<3}$ is isomorphic to the $f(1)$-th tree on this list. Continue choosing $f(2), f(3), \ldots$ recursively this way.

This yields a computable reduction of $\mathcal{T}^*/\cong$ to Baire space ω^ω, whose inverse is also a computable reduction.

So $\mathcal{T}^*/\cong$ and $Alg^*/\cong$ are *not homeomorphic*. In fact, there are computable reductions in both directions between these spaces, but none is bijective.

Back to Alg^*

Since $Alg^*/\cong$ is homeomorphic to 2^ω it seems natural to transfer the Lebesgue measure from 2^ω to $Alg/\cong$. But this requires care.

Fix a computable $\overline{\mathbb{Q}}$, and enumerate $\overline{\mathbb{Q}}[X] = \{f_0, f_1, \ldots\}$. Let $F_\lambda = \mathbb{Q}$. Given $F_\sigma \subset \overline{\mathbb{Q}}$, we find the least i, with f_i irreducible in $F_\sigma[X]$ of prime degree, for which it is not yet determined whether f_i has a root in F_σ. Adjoin such a root to $F_{\sigma\hat{}1}$, but not to $F_{\sigma\hat{}0}$. This gives a homeomorphism from 2^ω onto $Alg^*/\cong$, via $h \mapsto \bigcup_n F_{h\restriction n}$.

If we transfer standard Lebesgue measure to $Alg^*/\cong$, we get a measure in which the odds of 2 having a 1297-th root are $\frac{1}{2}$, but the odds of 2 having a 16-th root are much smaller.

Even worse, the odds of 2 having a square root depend on the ordering $f_0, f_1, f_2, \ldots$ we choose!

Haar measure on $Alg^* / \cong$

The worst problem is solved by considering only polynomials f_σ which are irreducible *of prime degree* over the existing field F_σ.

Haar measure on $Alg^*/\cong$

The worst problem is solved by considering only polynomials f_σ which are irreducible *of prime degree* over the existing field F_σ.

A further improvement is to use *Haar measure* μ on $Alg^*/\cong$. Here the probability of f_σ having a root is deemed to equal $\frac{1}{\deg(f_\sigma)}$. This idea (and the name) are justified by:

Proposition

For every algebraic field F_0 which is normal of finite degree d over $\mathbb{Q}$,

$$\mu(\{[K] \in Alg/\cong : \ F_0 \subseteq K\}) = \frac{1}{d}.$$

Notice that $\frac{1}{d}$ is precisely the measure of the pointwise stabilizer of F_0 within the group $\mathrm{Aut}(\overline{\mathbb{Q}})$, under the usual Haar measure on this compact group.

Measuring properties of algebraic fields

Using either of these measures, for (the isomorphism type of) an algebraic field, the property of being normal has measure 0. So does the property of having relatively intrinsically computable predicates R_d.

Measuring properties of algebraic fields

Using either of these measures, for (the isomorphism type of) an algebraic field, the property of being normal has measure 0. So does the property of having relatively intrinsically computable predicates R_d.

In Alg^*, the property of being relatively computably categorical has measure 1: given two roots x_1, x_2 of the same irreducible polynomial, one can wait for them to become distinct, since with probability 1 there will be an f for which $f(x_1, Y)$ has a root in the field but $f(x_2, Y)$ does not. This allows computation of isomorphisms between copies of the field. The process works uniformly except on a measure-0 set of fields.

Measuring properties of algebraic fields

Using either of these measures, for (the isomorphism type of) an algebraic field, the property of being normal has measure 0. So does the property of having relatively intrinsically computable predicates R_d.

In Alg^*, the property of being relatively computably categorical has measure 1: given two roots x_1, x_2 of the same irreducible polynomial, one can wait for them to become distinct, since with probability 1 there will be an f for which $f(x_1, Y)$ has a root in the field but $f(x_2, Y)$ does not. This allows computation of isomorphisms between copies of the field. The process works uniformly except on a measure-0 set of fields.

Surprisingly, measure-1-many fields in *Alg* remain relatively computably categorical even when the root predicates are removed from the language. However, the procedures for computing isomorphisms are not uniform. A single procedure can succeed only for measure-$(1-\epsilon)$-many fields.

Randomness and computable categoricity

Theorem (Franklin & M.)

For every Schnorr-random real $h \in 2^\omega$, the corresponding field F_h is relatively computably categorical, even in the language without the root predicates. However, there exists a Kurtz-random h for which F_h is not r.c.c. (in the language without the root predicates).

Randomness and computable categoricity

Theorem (Franklin & M.)

For every Schnorr-random real $h \in 2^\omega$, the corresponding field F_h is relatively computably categorical, even in the language without the root predicates. However, there exists a Kurtz-random h for which F_h is not r.c.c. (in the language without the root predicates).

Lemma

Let $\alpha, \beta \in \overline{\mathbb{Q}}$ be algebraic numbers conjugate over $\mathbb{Q}$. Then, for every finite algebraic field extension $E \supseteq \mathbb{Q}(\alpha, \beta)$, there is a set $D = \{q_0 < q_1 < \cdots\} \subseteq \mathbb{Q}$, decidable uniformly in E, such that for every k, both of the following hold:

$$\sqrt{\alpha + q_k} \notin E(\sqrt{\alpha + q_l}, \sqrt{\beta + q_l} \ : \ l \neq k)(\sqrt{\beta + q_k});$$

$$\sqrt{\beta + q_k} \notin E(\sqrt{\alpha + q_l}, \sqrt{\beta + q_l} \ : \ l \neq k)(\sqrt{\alpha + q_k}).$$

Proving the theorem

Given an $\epsilon > 0$, and a polynomial $f \in \mathbb{Q}[X]$ with two roots α, β, fix the set D from the lemma and choose N so large that the odds are $> 1 - \epsilon$ that, in an arbitrary field $\supseteq \mathbb{Q}(\alpha, \beta)$, all of the following hold:

- For at least $0.4N$ of the numbers $q_0, \ldots, q_{N-1}$ in D, $\alpha + q_i$ has a square root in the field.
- For at most $0.35N$ of these numbers, $\alpha + q_i$ and $\beta + q_i$ both have square roots in the field.

The procedure for mapping $\alpha, \beta \in F$ to the right images in a copy $\widetilde{F}$ waits until at least $0.4N$ elements $\sqrt{\alpha + q_i}$ with $i < N$ have appeared in F. Then it maps α to the first $\tilde{\alpha} \in \widetilde{F}$ it finds for which corresponding elements $\sqrt{\tilde{\alpha} + q_i}$ all appear in $\widetilde{F}$.

Proving the theorem

Given an $\epsilon > 0$, and a polynomial $f \in \mathbb{Q}[X]$ with two roots α, β, fix the set D from the lemma and choose N so large that the odds are $> 1 - \epsilon$ that, in an arbitrary field $\supseteq \mathbb{Q}(\alpha, \beta)$, all of the following hold:

- For at least $0.4N$ of the numbers $q_0, \ldots, q_{N-1}$ in D, $\alpha + q_i$ has a square root in the field.
- For at most $0.35N$ of these numbers, $\alpha + q_i$ and $\beta + q_i$ both have square roots in the field.

The procedure for mapping $\alpha, \beta \in F$ to the right images in a copy $\widetilde{F}$ waits until at least $0.4N$ elements $\sqrt{\alpha + q_i}$ with $i < N$ have appeared in F. Then it maps α to the first $\tilde{\alpha} \in \widetilde{F}$ it finds for which corresponding elements $\sqrt{\tilde{\alpha} + q_i}$ all appear in $\widetilde{F}$.

For polynomials of larger degree, use a similar procedure considering each possible pair of roots of the polynomial.

What about trees?

For the class $\mathcal{T}$ of finite-branching trees, one must first decide on a probability measure for ω^ω. The canonical choice is that, for $\sigma = (n_0, \ldots, n_k)$, we set $\mu(\mathcal{U}_\sigma) = 2^{-(1+k+n_0+\cdots+n_k)}$.

What about trees?

For the class $\mathcal{T}$ of finite-branching trees, one must first decide on a probability measure for ω^ω. The canonical choice is that, for $\sigma = (n_0, \ldots, n_k)$, we set $\mu(\mathcal{U}_\sigma) = 2^{-(1+k+n_0+\cdots+n_k)}$.

With this or most other reasonable measures, measure-1-many trees in $\mathcal{T}^*$ are r.c.c. However, in the language without branching predicates, measure-1-many trees in $\mathcal{T}$ fail to be relatively computably categorical.

The problem in $\mathcal{T}$ is that two siblings, $\alpha\hat{}0$ and $\alpha\hat{}1$, could both be terminal, with probability $\frac{1}{4}$. So we cannot fix any sort of N by which they will have (almost certainly) distinguished themselves from each other – but without knowing the branching, we cannot be too certain that they are automorphic either.

What constitutes a nice classification?

With both *Alg* and $\mathcal{T}$, we found very satisfactory classifications, by adding just the right predicates to the language. But it is not always so simple.

Let TFAb_1 be the class of torsion-free abelian groups G of rank exactly 1. We usually view these as being classified by tuples $(\alpha_0, \alpha_1, \ldots)$ from $(\omega + 1)^\omega$, saying that an arbitrary nonzero $x \in G$ is divisible by p_n exactly $f(n)$ times. To account for the arbitrariness of x, we must identify tuples $\vec{\alpha}$ and $\vec{\beta}$ with only finite differences:

$$\exists k[(\forall j > k\ \alpha_j = \beta_j)\ \&\ (\forall j\ |\alpha_j - \beta_j| < k)].$$

What constitutes a nice classification?

With both *Alg* and $\mathcal{T}$, we found very satisfactory classifications, by adding just the right predicates to the language. But it is not always so simple.

Let TFAb_1 be the class of torsion-free abelian groups G of rank exactly 1. We usually view these as being classified by tuples $(\alpha_0, \alpha_1, \ldots)$ from $(\omega + 1)^\omega$, saying that an arbitrary nonzero $x \in G$ is divisible by p_n exactly $f(n)$ times. To account for the arbitrariness of x, we must identify tuples $\vec{\alpha}$ and $\vec{\beta}$ with only finite differences:

$$\exists k[(\forall j > k\ \alpha_j = \beta_j)\ \&\ (\forall j\ |\alpha_j - \beta_j| < k)].$$

The space $\text{TFAb}_1/\cong$ has the indiscrete topology: no finite piece of an atomic diagram rules out any isomorphism type. More info needed!

If, for all primes p, we add $D_p(x)$ and $D_{p^\infty}(x)$, saying that x is divisible by p and infinitely divisible by p, then we get the classification above. However, it is not homeomorphic to Baire space itself.

Classification using equivalence relations

With D_p and D_{p^∞} added to the language of groups, we now have $\text{TFAb}_1/\cong$ computably homeomorphic to ω^ω/E_0^* (or to $(\omega+1)^\omega/E_0^*$, with the right topology) where E_0^* denotes differing on only finitely many columns and by only finitely much:

$$A\ E_0^*\ B \iff \exists k[(\forall n > k)A(n) = B(n)\ \&\ (\forall n)|A(n) - B(n)| < k].$$

In turn, ω^ω is computably homeomorphic to Baire space under the usual E_0 relation, denoting finite symmetric difference. So we have a classification using a standard equivalence relation.

But what sort of measure could one put on $(\omega+1)^\omega$?

An alternative

If we add just the D_p relations to the language of groups, then $\text{TFAb}_1/\cong$ is homeomorphic to $2^\omega/E_0$. The initial segment $\sigma = 0111001$, for example, denotes that some nonzero $x \in G$ is:

- not divisible by 2;
- divisible by 3;
- divisible by 5;
- divisible by 3^2;
- not divisible by 7;
- not divisible by 5^2;
- divisible by 3^3;
- etc.

An alternative

If we add just the D_p relations to the language of groups, then $\text{TFAb}_1/\cong$ is homeomorphic to $2^\omega/E_0$. The initial segment $\sigma = 0111001$, for example, denotes that some nonzero $x \in G$ is:

- not divisible by 2;
- divisible by 3;
- divisible by 5;
- divisible by 3^2;
- not divisible by 7;
- not divisible by 5^2;
- divisible by 3^3;
- etc.

Here infinite divisibility by p is a measure-0 property. Thus almost all structures here are r.c.c. in this language, and relatively Δ^0_2-categorical even without the D_p predicates.

One more example

An *equivalence structure* simply consists of an equivalence relation on the domain. Isomorphism is Π^0_4-complete for computable equivalence structures. The natural classification maps a structure E to $(\alpha_0, \alpha_1, \alpha_2, \ldots) \in (\omega + 1)^\omega$, where E has exactly α_n classes of size n, along with α_0 infinite classes.

One more example

An *equivalence structure* simply consists of an equivalence relation on the domain. Isomorphism is Π^0_4-complete for computable equivalence structures. The natural classification maps a structure E to $(\alpha_0, \alpha_1, \alpha_2, \ldots) \in (\omega + 1)^\omega$, where E has exactly α_n classes of size n, along with α_0 infinite classes.

Making this classification effective requires adding some less-than-natural predicates to the language. Even with a class of such simple structures, it is difficult to decide on the "best" classification. We are brought back to the original question:

WHAT DO MATHEMATICIANS WANT?