Describing structures and classes of structures

Julia F. Knight

University of Notre Dame

Topics to be discussed

I. Describing specific structures—Scott complexity

II. Characterizing classes—Borel complexity

III. Classification problems—Borel cardinality

Conventions

1. Structures are countable, with universe ω or a subset.
2. Languages are countable. usually computable.
3. Classes consist of structures for a fixed language, and are closed under isomorphism.

$L_{\omega_1\omega}$

In $L_{\omega_1\omega}$, we allow countably infinite disjunctions and conjunctions, but only finite strings of quantifiers.

Sample formulas

1. The following sentence says of a real closed ordered field that it is Archimedean:

$$(\forall x) \bigvee_n x < \underbrace{1 + \cdots + 1}_{n}$$

2. The following formula says of an element in an Abelian group that it has infinite order.

$$\bigwedge_n \underbrace{x + \cdots + x}_{n} \neq 0$$

Normal form

Note. We do not have prenex normal form for $L_{\omega_1\omega}$-formulas. We cannot, in general, bring the quantifiers to the front.

New normal form. For $L_{\omega_1\omega}$ formulas, we can bring the negations inside. This gives a new normal form.

For formulas in this normal form, we measure complexity by considering alternations of $\bigvee\exists$ and $\bigwedge\forall$.

Complexity of $L_{\omega_1\omega}$-formulas

1. $\varphi(\bar{x})$ is Σ_0 and Π_0 if it is finitary quantifier-free.
2. For a countable ordinal $\alpha > 0$,
   (a) $\varphi(\bar{x})$ is Σ_α if it has form $\bigvee_i (\exists \bar{u}_i)\psi_i(\bar{x}, \bar{u}_i)$, where each ψ_i is Π_{β_i} for some $\beta_i < \alpha$,
   (b) $\varphi(\bar{x})$ is Π_α if it has form $\bigwedge_i (\forall \bar{u}_i)\psi_i(\bar{x}, \bar{u}_i)$, where each ψ_i is Σ_{β_i} for some $\beta_i < \alpha$.

Further terminology. A formula is d-Σ_α if it has form $(\varphi \;\&\; \psi)$, where φ is Σ_α and ψ is Π_α.

Complexity of sample formulas

1. The sentence $(\forall x) \bigvee_n x < \underbrace{1 + \cdots + 1}_{n}$ is Π_2.
2. The formula $\bigwedge_n \underbrace{x + \cdots + x}_{n} \neq 0$ is Π_1.

There is a natural d-Σ_2 sentence saying of a $\mathbb{Q}$-vector space that it has a specific finite dimension. We will see more examples of this kind.

Computable infinitary formulas

The *computable infinitary formulas* are $L_{\omega_1\omega}$ formulas in which the infinite disjunctions and conjunctions are over computably enumerable (c.e.) sets.

We classify these formulas as *computable* Σ_α, *computable* Π_α, for *computable* ordinals α.

1. The sentence $(\forall x) \bigvee_n x < \underbrace{1 + \cdots + 1}_{n}$ is computable Π_2.
2. The formula $\bigwedge_n \underbrace{x + \cdots + x}_{n} \neq 0$ is computable Π_1.

Remark: Computable infinitary formulas seem comprehensible.

Describing a structure

Theorem (Scott, 1965). For each structure $\mathcal{A}$, there is an $L_{\omega_1\omega}$-sentence φ whose countable models are the copies of $\mathcal{A}$. (Such a sentence is called a *Scott sentence* for $\mathcal{A}$.)

Proof sketch. Scott determined a family of formulas $\varphi_{\bar{a}}$ that define the orbits of tuples $\bar{a}$. Then $\varphi = \bigwedge_{\bar{a}} \rho_{\bar{a}}$, where

$$\rho_{\emptyset} = (\forall y) \bigvee_{b} \varphi_b(y) \;\&\; \bigwedge_{b} (\exists y)\varphi_b(y) \;,$$

and for other $\bar{a}$,

$$\rho_{\bar{a}} = (\forall \bar{u})[\varphi_{\bar{a}}(\bar{u}) \rightarrow ((\forall y) \bigvee_{b} \varphi_{\bar{a},b}(\bar{u}, y) \;\&\; \bigwedge_{b} (\exists y)\varphi_{\bar{a},b}(\bar{u}, y))].$$

Note: If the formulas $\varphi_{\bar{a}}$ are all Σ_α, then φ is $\Pi_{\alpha+1}$.

Scott sentence for $\mathbb{Z}$

(1) For the additive group of integers, we obtain a (computable) Σ_3 Scott sentence from the conjunction of a sentence characterizing the torsion-free Abelian groups and a sentence saying that some non-zero element generates everything.

We can do better.

(2) There is a (computable) d-Σ_2 Scott sentence—the conjunction of a Π_2 sentence characterizing the torsion-free Abelian groups, a Π_2 sentence saying that for any pair x, y, there is some z that generates both, and a Σ_2 sentence saying that there is some x not divisible by $n > 1$.

The free group F_n

Theorem (Carson-Harizanov-K-Lange-McCoy-Quinn-Morozov-Safranski-Wallbaum, 2012). For all finite $n \geq 1$, the free group of rank n has a (computable) d-Σ_2 Scott sentence.

Proof sketch. We have seen the Scott sentence $n = 1$. For $n \geq 2$, we take the conjunction of:

(1) a Π_2 sentence saying that for every tuple $\bar{y}$, there is an n-tuple $\bar{x}$ that generates $\bar{y}$,

(2) a Σ_2 sentence saying that there is an n-tuple $\bar{x}$ satisfying no non-trivial relations, s.t. for all n-tuples $\bar{y}$, no "imprimitive" n-tuple of words takes $\bar{y}$ to $\bar{x}$.

Note: Nielsen (1917, 1918) described the primitive tuples of words, and showed that the set of these is computable.

Complexity of orbits

The complexity of an optimal Scott sentence for a structure is connected to the complexity of the orbits.

Theorem (Montalbán, 2015). $\mathcal{A}$ has a $\Pi_{\alpha+1}$ Scott sentence iff the orbits of all tuples in $\mathcal{A}$ are defined by Σ_α formulas.

Theorem (Alvir-K-McCoy, 2020). If $\mathcal{A}$ has a computable $\Pi_{\alpha+1}$ Scott sentence, then the orbits are defined by computable Σ_α formulas.

$Mod(L)$

For a countable language L, $Mod(L)$ is the set of L-structures with universe ω.

Identifying $Mod(L)$ **with Cantor space**. For simplicity, we suppose L is a relational language. Let C be a set of new constants representing the natural numbers. Let $(\alpha_n)_{n\in\omega}$ be a list of the atomic sentences $R\bar{a}$, where R is a relation symbol of L and $\bar{a}$ is a tuple from C.

We identify $\mathcal{A} \in Mod(L)$ with the function $f \in 2^\omega$ s.t.

$$f(n) = \begin{cases} 1 & \text{if } \mathcal{A} \models \alpha_n \\ 0 & \text{otherwise} \end{cases}$$

Borel classes

There is a natural topology on Cantor space, generated by the clopen sets $N_p = \{f \in 2^\omega : f \supseteq p\}$, for $p \in 2^{<\omega}$. The *Borel sets* are those in the σ-algebra generated by these N_p.

1. B is $\mathbf{\Sigma_0}$ and $\mathbf{\Pi_0}$ if it is a finite union of basic clopen sets.

2. For a countable ordinal $\alpha > 0$,

   (a) B is $\mathbf{\Sigma}_\alpha$ if $B = \cup_i B_i$, where each B_i is $\mathbf{\Pi}_{\beta_i}$ for some $\beta_i < \alpha$,

   (b) B is $\mathbf{\Pi}_\alpha$ if $B = \cap_i B_i$, where each B_i is $\mathbf{\Sigma}_{\beta_i}$ for some $\beta_i < \alpha$.

Axiomatizing Borel classes

Theorem (Lopez-Escobar, 1965): For $K \subset Mod(L)$, closed under automorphism, K is Borel iff it is axiomatized by a sentence of $L_{\omega_1\omega}$.

Theorem (Vaught, 1974): For $K \subseteq Mod(L)$, closed under automorphism, K is $\mathbf{\Sigma}_\alpha$ (for $\alpha \geq 1$) iff it is axiomatized by a Σ_α sentence.

Vaught's proof involved "Vaught transforms." Vanden Boom (in his senior thesis at ND), gave a different proof, and an effective version.

Vanden Boom

Effective Borel hierarchy. The *effective Borel sets* are obtained from the basic clopen neighborhoods using c.e. unions and intersections.

Theorem (Vanden Boom, 2007). A set $B \subseteq Mod(L)$, closed under isomorphism, is effective Σ_α (for $\alpha \geq 1$) iff it is axiomatized by a computable Σ_α formula.

Vanden Boom's proof used forcing. The formulas that define forcing substitute for Vaught transforms. Relativizing Vanden Boom's Theorem, we get Vaught's.

Borel embeddings and Borel cardinality

Definition (H. Friedman & Stanley, 1989). Let $K \subseteq Mod(L)$, $K' \subseteq Mod(L')$, both closed under isomorphism. A *Borel embedding* of K in K' is a Borel function $\Phi : K \to K'$ s.t. for $\mathcal{A}, \mathcal{B} \in K$, $\mathcal{A} \cong \mathcal{B}$ iff $\Phi(\mathcal{A}) \cong \Phi(\mathcal{B})$.

Notation: We write $K \leq_B K'$ if there is such an embedding. We write $K <_B K'$ if $K \leq_B K'$ and $K' \not\leq_B K$, and we write $K \equiv_B K'$ if $K \leq_B K'$ and $K' \leq_B K$.

Definition. The *Borel cardinality* of K is its $\equiv_B$-class.

On top

Theorem (Lavrov, Maltsev, Mekler, Friedman-Stanley, Marker). The following classes lie on top under $\leq_B$:

1. undirected graphs
2. fields
3. 2-step nilpotent groups
4. linear orderings
5. real closed ordered fields

New result (Paolini and Shelah). Torsion-free Abelian groups also lie on top.

Embedding $Mod(L)$ in undirected graphs

Theorem (Lavrov, 1963). $Mod(L) \leq_B$ undirected graphs.

There are slightly different embeddings due to Marker, Nies. We follow Marker.

We start with the case where L has just one binary relation symbol–$Mod(L)$ is the class of directed graphs. The embedding Φ takes a directed graph $(A, \rightarrow)$ to an undirected graph $(B, -)$ with a special point b_a representing each $a \in A$ and a special point $p_{(a,a')}$ representing each ordered pair (a, a'). The following picture shows how the embedding works.

Picture

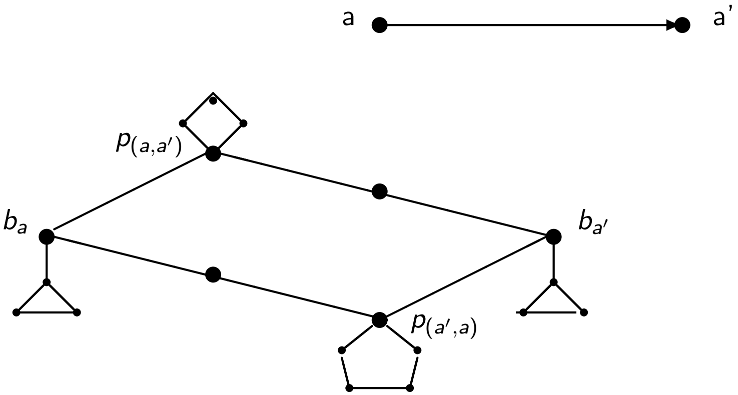

More

To see that Φ is $1 - 1$ on isomorphism types, we note that there is a copy of $\mathcal{A}$ defined in $\Phi(\mathcal{A})$—the definition uses finitary existential formulas.

To embed $Mod(L)$ in undirected graphs, we use more special points and more n-gons.

fields $\leq_B$ 2-step nilpotent groups

Maltsev, 1960. Let Φ take each field F to its Heisenberg group $H(F)$, consisting of matrices of form

$$\begin{bmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{bmatrix},$$

where $a, b, c \in F$.

Maltsev gave finitary existential formulas that define a copy of F in $H(F)$ with an arbitrary non-commuting pair as parameters.

Theorem (Alvir-Calvert-Goodman-Harizanov-K-Miller-Morozov-Soskova-Weisshaar). There are finitary existential formulas that, for all fields F, effectively interpret F in $H(F)$.

graphs $\leq_B$ linear orderings

Friedman and Stanley defined an embedding Φ of graphs in linear orderings. The proof that Φ is $1-1$ does not involve a definition or interpretation.

Theorem (Harrison-Trainor-Montalbán, 2020; K-Soskova-Vatev, 2020). There do not exist $L_{\omega_1\omega}$ formulas that, for all graphs G, define an interpretation of G in $\Phi(G)$.

Below the top

The following classes lie strictly below the top, for different reasons:

1. $\mathbb{Q}$-vector spaces—only $\aleph_0$ isomorphism types,
2. subfields of the algebraic numbers—isomorphism relation is Borel,
3. Abelian p-groups—subtler reason.

Turing computable embeddings

Kechris suggested that my students and I consider effective embeddings.

Definition (Calvert-Cummins-K-Quinn, 2004). For classes K, K', closed under isomorphism, a *Turing computable embedding* of K in K' is a Turing operator $\Phi : K \to K'$ s.t. for $\mathcal{A}, \mathcal{B} \in K$, $\mathcal{A} \cong \mathcal{B}$ iff $\Phi(\mathcal{A}) \cong \Phi(\mathcal{B})$. We write $K \leq_{tc} K'$.

The Borel embeddings of Friedman-Stanley, Lavrov, Mekler, Maltsev, Marker are Turing computable.

Pullback Theorem

Theorem (K-Quinn-Vanden Boom, 2007). Suppose $K \leq_{tc} K'$ via Φ. For any computable infinitary sentence φ in the language of K', we can find a computable infinitary sentence φ^* in the language of K s.t. for all $\mathcal{A} \in K$, $\mathcal{A} \models \varphi^*$ iff $\Phi(\mathcal{A}) \models \varphi$. Moreover, for $0 < \alpha < \omega_1^{CK}$, if φ is computable Σ_α, then so is φ^*.

Example. For each $n \geq 1$, there is a Σ_2-sentence φ_n saying of a $\mathbb{Q}$-vector space that the dimension is at least n. If $K \leq_{tc}$ $\mathbb{Q}$-vector spaces, then the pullbacks of the sentences φ_n describe invariants for K.

Relativizing the Pullback Theorem

Corollary. Let Φ be a continuous embedding of K in K'. For $0 < \alpha < \omega_1$, for any Σ_α sentence φ in the language of K', there is a Σ_α sentence φ^* in the language of K s.t. for $\mathcal{A} \in K$, $\mathcal{A} \models \varphi^*$ iff $\Phi(\mathcal{A}) \models \varphi$.

Proof: There is a set X s.t. Φ is X-computable, the ordinal α is X-computable, and φ is X-computable Σ_α. Then we have a pullback φ^* that is X-computable Σ_α.

Invariants

Suppose $K \leq_B K'$ via Φ. The embedding Φ reduces the classification problem for K to that for K'. If we have useful invariants for K', then we do for K as well.

Having the same Borel cardinality means essentially having the same invariants. Exactly what counts as useful invariants is vague.

1. $\mathbb{Q}$-**vector spaces**: dimension—universally accepted as useful.
2. **Abelian** p-**groups**: Ulm sequence plus dimension of the divisible part—complicated, but accepted as useful by some people.

Torsion-free Abelian groups

Let TFA_n be the class of torsion-free Abelian groups of rank n. These are the groups isomorphic to subgroups of $\mathbb{Q}^n$ with n $\mathbb{Z}$-linearly independent elements.

We can describe a group in TFA_1 by taking a non-zero element a and saying which prime powers divide a. For a different non-zero element a', the sets we obtain differ only finitely. Baer gave invariants based on this fact. According to Hjorth and Thomas, the invariants for TFA_1 are generally accepted as useful

For $n \geq 2$, Maltsev and Kurosh gave invariants for groups in TFA_n. Hjorth and Thomas report that Fuchs dismissed these as no better than the group itself—they are not generally accepted as useful.

Borel cardinality increases with n

Theorem (Hjorth, 1999). $TFA_1 <_B TFA_2$

Theorem (Thomas, 2001). For $n \geq 2$, $TFA_n <_B TFA_{n+1}$.

Current project, joint with Ho, Miller: Use tools from computability to show, in simpler way, that

$$TFA_1 <_{tc} TFA_2 <_{tc} TFA_3 <_{tc} \cdots \ .$$

Also, show that the Paolini-Shelah embedding of graphs in torsion-free Abelian groups does not come from a uniform interpretation.