Describing finite groups by first-order sentences of polylogarithmic length

André Nies
joint work with Katrin Tent

THE UNIVERSITY OF AUCKLAND
NEW ZEALAND

Invariant Kolmogorov complexity

in classes of finite structures

Invariant Kolmogorov complexity

- Fix a universal system of descriptions; say, a universal Turing machine M taking as input bit strings σ.
- The Kolmogorov complexity of a finite mathematical object x (e.g. a string) is the length of a shortest description, i.e. $\min\{|\sigma| : M(\sigma) = x\}$
- We can encode a finite structure G over a finite symbol set by a string x_G.
- The Kolmogorov complexity of x_G is not necessarily invariant under isomorphism of the structures.
- The invariant Kolmogorov complexity $K_{\texttt{inv}}(G)$ is the minimum of the Kolmogorov complexities of x_H for all structures $H \cong G$.

Compressibility for a class of finite structures

Recall: $K_{\mathtt{inv}}(G)$ is the minimum of the Kolmogorov complexities of all structures $H \cong G$.

- Let $\mathcal{C}$ be a class of finite structures for the same finite symbol set.
- Let $R\colon \mathbb{N} \to \mathbb{N}^+$ be an unbounded function.
- If $K_{\mathtt{inv}}(G) \leq R(|G|)$ for each $G \in \mathcal{C}$, what can we say about the growth of R?

I.o. lower bound of $\log n$

$\log$ denotes the logarithm in base 2.

Let $\mathcal{C}$ be a class of finite structures containing a structure of each size. If $K_{\texttt{inv}}(G) \leq R(|G|)$ for each $G \in \mathcal{C}$, then for some constant c

$$\exists^\infty n \left[\log n - c \leq R(n)\right] .$$

- To see this: for each k there is a number n with $\lfloor \log n \rfloor = k$ such that any binary description of n has at least k bits.
- If some structure of size n has too short a description, then n has a description of length $< k$, contradiction.

Lower bounds by counting: graphs

Let $\mathcal{C}$ be the class of finite graphs. If $K_{\texttt{inv}}(G) \leq R(|G|)$ for each $G \in \mathcal{C}$, then $n^2 - 6\log n = O(R(n))$. "No compression possible."

- The number of non-isomorphic undirected graphs with n vertices is at least
$$\frac{2^{\binom{n}{2}}}{n!} = \frac{1}{n}\prod_{i=1}^{n-1}\frac{2^i}{i},$$
which for large n exceeds $\frac{1}{n}2^{n^2/6}$.

- For each k there are fewer than 2^k binary descriptions of length less than k. So for some constant c, for large enough n there is an undirected graph G with n vertices such that $n^2 - 6\log n \leq c|\sigma|$, for any binary description σ of any $H \cong G$. Hence $n^2 - 6\log n \leq cK_{\texttt{inv}}(G)$.

Lower bounds by counting: p-groups

Let $\mathcal{C}$ be the class of finite p-groups (p a prime).
If $K_{\texttt{inv}}(G) \leq R(|G|)$ for each $G \in \mathcal{C}$, then $(\log n)^3 = O(R(n))$.

- Higman (1960) showed[1] that there are at least

$$p^{(\frac{2}{27}+\tau(m))m^3}$$

  non-isomorphic groups of order p^m, for some function τ with $\lim_m \tau(m) = 0$. (They are in fact all nil-2 of exponent p^2.)

- This implies that for some constant c, for each large enough n a power of p, there is a group G with n elements such that

$$\log^3 n \leq c|\sigma|$$

  for any binary description σ of any $H \cong G$.

[1] See 2007 book by Blackburn, Neumann and Venkataram

First-order compressibility within classes of finite structures

Main definition: compressibility in first-order logic

Let $\mathcal{C}$ be a class of finite structures for the same finite symbol set.
Let $R\colon \mathbb{N} \to \mathbb{N}^+$ be an unbounded function.
The class $\mathcal{C}$ is R-compressible if for any $G \in \mathcal{C}$, there exists a first-order sentence ψ_G of length $|\psi_G| = O(R(|G|))$ such that

- $G \models \psi_G$, and
- if $H \models \psi_G$ then $G \cong H$.

- The “atomic diagram” of the structure is its trivial description. (For a finite group, this is essentially the multiplication table.)
- This description has length $O(|G|^{k+1})$, where k is the maximum arity of a symbol in the set.

Remarks on decompression and encoding

- Our descriptions are now first-order sentences ϕ.
- Decompression: a machine takes input ϕ and outputs the first finite model of ϕ (if any)
- Descriptions use an infinite alphabet
- we can convert them into binary descriptions (essentially, index the variables in binary to get down to a finite alphabet).
- The length of binary description of ϕ is $O(|\phi| \log |\phi|)$ so this slightly worsens upper bounds on compressibility for classes we now give.
- Let $K_{FO}(G)$ denote the least length of a first-order description of G. We have $K_{\mathtt{inv}}(G) = O(K_{FO}(G) \log(K_{FO}(G)))$.

Cycle graphs are log-compressible

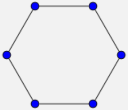

Let C_n be the undirected cycle graph with n vertices.

The class $\mathcal{C} = \{C_n : \; n \in \mathbb{N}\}$ is log-compressible.

Let $n = 2k$ or $n = 2k + 1$. Let ϕ_n be the sentence

- the graph is undirected, every vertex has degree 2,
- $\forall x, y \; P_k(x, y)$,
- $\exists x, y \; \neg P_{k-1}(x, y)$.

$P_k(x, y)$ is a formula of length $O(\log k)$ saying that there is a path with $\leq k$ edges from x to y. (It uses $O(\log k)$ quantifiers.)
To distinguish whether n is even or odd, note that
$\forall u \forall v \forall z [Euv \to (P_{k-1}uz \vee P_{k-1}vz)]$ holds iff $n = 2k$.

Cyclic groups are log-compressible

Recall:

The class $\mathcal{C}$ is R-compressible if for any $G \in \mathcal{C}$, there exists a first-order sentence ψ_G of length $|\psi_G| = O(R(|G|))$ such that

- $G \models \psi_G$, and
- if $H \models \psi_G$ then $G \cong H$.

To say that a group G is cyclic of order n, express that

$$\exists g\,[\text{ the undirected Cayley graph given by } g \text{ is } C_n].$$

Finite (difference) fields are log-compressible

A difference field is a field with a distinguished automorphism.

- (i) For any finite field $\mathbb{F}_q$, there is a Σ_3-sentence φ_q of length $O(\log q)$ in the language $L(+, \times, 0, 1)$ describing $\mathbb{F}_q$.
- (ii) For any finite difference field $(\mathbb{F}_q, \sigma)$ there is a Σ_3-sentence $\psi_{q,\sigma}$ of length $O(\log q)$ in $L(+, \times, 0, 1, \sigma)$ describing $\langle \mathbb{F}_q, \sigma \rangle$.

Proof (i): Let $q = p^n$ for a prime p.

- The sentence ϕ_q says that the structure is a field of characteristic p such that $\forall x\, [x^{p^n} = x]$ and $\exists y\, [y^{p^{n-1}} \neq y]$.
- These formulas can be replaced by short formulas using a method from the theory of algorithms known as "exponentiation via repeated squaring".

First-order compressibility of finite simple groups:

$$\log |G|$$

Examples of first-order sentences for groups

Let $[x, y]$ denote the commutator $x^{-1}y^{-1}xy$.

- The first-order sentence $\forall x \forall y \; [x, y] = 1$ expresses that the group is abelian.
- The following first-order sentence expresses that every commutator is a product of three squares:

$$\forall u \forall v \exists r \exists s \exists t \; [u, v] = rrsstt.$$

Short formulas for defining a generated subgroup

For each $n, k > 0$ we can find a first-order formula $\alpha_{\text{gen}}(g; x_1, \ldots, x_k)$ of length $O(k + \log n)$ such that $G \models \alpha_{\text{gen}}(g; x_1, \ldots, x_k)$ if and only if $g \in \langle x_1, \ldots, x_k \rangle$ for $|G| = n$.

To build the formulas α_{gen} we use a technique that originated in computational complexity, e.g. to show that the set of true quantified boolean formulas is complete for polynomial space.

Drawback: this leads to an unbounded number of quantifier alternations. We can avoid them at the cost of somewhat longer formulas:

For each each $n, k > 0$ we can find existential f.o. formula $\beta_{\text{gen}}(g; x_1, \ldots, x_k)$ of length $O(k \log^2 n)$ such that $G \models \beta_{\text{gen}}(g; x_1, \ldots, x_k)$ if and only if $g \in \langle x_1, \ldots, x_k \rangle$ for $|G| = n$.

Ree groups[2]

- Let $G_2(q)$ be the automorphism group of the octonion algebra over a q-element field $\mathbb{F}_q$, where q has the form 3^{2k+1}, $k > 0$.
- τ is the automorphism of $G_2(q)$ arising from the symmetry of the underlying (undirected) Dynkin diagram $\cdot \overset{6}{—} \cdot$
- σ is the automorphism of $G_2(q)$ given by the square root of the Frobenius automorphism $x \to x^3$ on F_q (so $\sigma(x) = x^{3^{k+1}}$).
- The Ree group ${}^2G_2(q)$ is a subgroup of $G_2(q)$: it consists of the elements g such that $\tau(g) = \sigma(g)$.

---

[2]Rimhak Ree, 1961

Short presentations for finite simple groups

A finite presentation of a group has the form

$$\langle x_1, \ldots, x_k \mid r_1, \ldots, r_m \rangle.$$

E.g. for dihedral groups we have $D_{2n} = \langle x, y \mid x^2, y^n, x^{-1}yxy \rangle$.
This presentation has length $n + constant$.

Theorem (Guralnick et al., JAMS, 2008)

For some constant C_0: the nonabelian finite simple groups, possibly except the Ree groups of type ${}^2G_2$, have a presentation with

- *at most C_0 generators and relators*
- *length at most $C_0(\log q + \log n)$.*

q is the size of the underlying field, n the Lie rank of the group.

$\log n + \log q \leq \log |G|$, so the presentations are $O(\log |G|)$ long.

Conversion to a short first-order description

Suppose that a finite simple group G has a presentation $\langle x_1, \ldots, x_k \mid r_1, \ldots, r_m \rangle$. Let g_i be the image of x_i in G. There is a formula $\psi(x_1, \ldots, x_k)$ of length $O(\log |G| + k + \sum_i |r_i|)$ describing the structure $(G, g_1, \ldots, g_k)$.

- The formula is $x_1 \neq 1 \wedge \bigwedge_{1 \leq i \leq m} r_i = 1 \wedge \forall y \, \alpha_{\mathrm{gen}}(y; x_1, \ldots, x_k)$.
- α_{gen} is the formula of length $O(k + \log |G|)$ from a previous lemma, expressing that y is generated by the x_i within G.
- The models of ψ are the nontrivial quotients of $(G, g_1, \ldots, g_k)$.
- Then, since G is simple, ψ describes $(G, g_1, \ldots, g_k)$.

Compression for finite groups: first result

Theorem

Suppose a finite group G has a presentation of length N.
Then G has a first-order description of length $N + O(\log^2 |G|)$.

The proof follows the argument in the case of simple groups. The group is determined by:

- generated by a sequence $g_1, \ldots, g_k$ satisfying the relators
- the length of a composition series. Using the formulas α_{gen} for generation, this extra information takes length $O(\log^2 |G|)$.

Corollary (using Guralnick et al, JAMS 2008)

The class of finite groups not containing a Ree group ${}^2G_2$ as a composition factor is $\log^3$*-compressible.*

Short first-order descriptions for Ree groups

- ${}^2G_2(q)$ is bi-interpretable with the difference field $(\mathbb{F}_q, \sigma)$. The formulas don't depend on q. (See Ryten's 2007 PhD thesis at the University of Leeds for a proof.)
- The class of finite difference fields is log-compressible.
- log-compressibility is preserved under bi-interpretability.
- So we can find short first-order descriptions of length $O(\log n)$ for the Ree groups.

There is a slight complication because the interpretation of the appropriate difference fields in the Ree groups needs parameters. To deal with this, we actually want that the class of difference fields is strongly log-compressible, i.e. we can add a list of constants of fixed length and still get a description of length $O(\log)$.

First-order compressibility of all finite groups:

$$(\log |G|)^3$$

Preliminary: straight line programs

Let G be a finite group, $S \subseteq G$ and $g \in G$.

- A straight line program (SLP) $\mathcal{L}$ over S is a sequence of group elements such that each element of $\mathcal{L}$ is in S, an inverse of an earlier element, or a product of two earlier elements.
- The reduced length is the number of entries not in S.
- $\mathcal{L}$ computes B from S if $\mathcal{L}$ is an SLP over S containing B.

Reachability Lemma of Babai and Szemerédi (1984)

For each set $S \subseteq G$, there is a straight line program $\mathcal{L}$ over S of reduced length at most $(\log|G| + 1)^2$ that computes a preprocessing set $A = \{z_1, \ldots, z_n\}$ as follows:
each g in $\langle S \rangle$ is of the form $q^{-1}p$ where p, q are products of some of the z_i in ascending order; so its red'd length over A is $\leq 2\log|G|$.

Proof that finite groups G are $(\log|G|)^3$-compressible

We fix a composition series

$$1 = G_0 \lhd G_1 \lhd \ldots \lhd G_r = G$$

with simple factors $H_i := G_i/G_{i-1}$, $i = 1, \ldots r$.
The length r is bounded by $\log|G|$. Pick an "appropriate" sequence

$$\emptyset = T_0 \subset T_1 \subset \ldots \subset T_r = T \text{ with } \langle T_i \rangle = G_i.$$

Define the G_i from the sets T_i using the formulas α_{gen}.

(a) introduce ascending preprocessing sets A_i for the T_i

(b) describe each H_i, together with the image of $T_i \setminus T_{i-1}$

(c) describe each G_i as a group extension

$$1 \to G_{i-1} \to \boxed{G_i} \to H_i \to 1 \text{ (exact sequence).}$$

Introduce preprocessing sets A_i for the T_i

Recall: A is preprocessing set for S if $\langle A \rangle = \langle S \rangle$ and for each g in $\langle S \rangle$ the reduced length of g over A is $\leq 2 \log |G|$.

- Our sentence describing G starts with a block of existential quantifiers for the T, and another block referring to a preprocessing set A for the generating set T of G.
- It states how A has been obtained from T via a SLP of reduced length $(\log |G| + 1)^2$. It uses at most that many further existential quantifiers.

We build A in levels $A_0 \subseteq \ldots \subseteq A_s = A$, where A_i is a preprocessing set for T_i. To do so we successively extend SLPs computing A_i from T_i. The A_i will allow rapid access to particular elements of G_i (at a cost of $2 \log |G_i|$).

Describe H_i and the image of $T_i \setminus T_{i-1}$

Recall that $H_i = G_i/G_{i-1}$. By the case of simple groups (and the right choice of the T_i) we have an $O(\log|G|)$ sentence ϕ_i describing (H_i, T_i). We now need to express within G that $(H_i, T_i) \models \phi_i$.

- Via the formulas α_{gen} we express that G_{i-1} is a normal subgroup of G_i, using a length of $O(\log|G_i|)$.
- We restrict the quantifiers in ϕ_i to G_i using α_{gen} and replace each occurrence of "$u = v$" in ϕ_i by "$uv^{-1} \in G_{i-1}$".
- Since we replace the equality symbols in ϕ_i by strings of length $O(\log|G_{i-1}|)$, the resulting formula χ_i has length $O(\log|H_i| \log|G_{i-1}|)$. Then $\bigwedge_i \chi_i$ has length $O(\log^2|G|)$.

Describe G_i as a group extension of G_{i-1} by H_i

$$1 \to G_{i-1} \to \boxed{G_i} \to H_i \to 1 \text{ (exact sequence).}$$

Conjugation action of G_i on G_{i-1}:

- Since $\langle T_{i-1} \rangle = G_{i-1}$, it suffices to determine $g^{-1}wg$, for each pair $g \in T_i \setminus T_{i-1}$ and $w \in T_{i-1}$, as an element $h_{g,w} \in G_{i-1}$.
- $h_{g,w}$ has length at most $2 \log |G_{i-1}|$ over A_{i-1}.
- There are at most $C_0 \cdot \log |G_{i-1}|$ such pairs g, w. (C_0 is a bound on the number of generators of H_i. We picked the T_i so that $|T_i \setminus T_{i-1}| \leq C_0$.)
- So this can be done by a formula of length $O(\log^2 |G_{i-1}|)$.

Describe G_i as a group extension of G_{i-1} by H_i

Use result of Lubotzky and Segal that there is a “profinite” presentation for H_i of length $O(\log |H_i|)$. Also use:

Lemma: there is $d \le r \cdot \log |Z(G_{i-1})|$, and there are words $w_1, \dots w_d$ in $\overline{a}_i = A_i \setminus A_{i-1}$ of length at most $3 \log |H_i|$ such that the values $w_m(\overline{a}_i) \in G_{i-1}$ determine G_i.

This is proved via some cohomology describing possible group extensions, and the following fact suggested originally by Alex Lubotzky at Hebrew U.

Let A be a finite abelian group (in our case it is the centre of G_{i-1}). Let X be a set. Let $V \le A^X$ be a subgroup generated by d elements. There is a set $Y \subseteq X$ of size at most $d \log |A|$ such that for each $g \in V$, $g \restriction_Y = 0 \Rightarrow g = 0$.

Further directions and open questions

- Is the compression we obtain optimal for subclasses of the finite simple groups, such as the alternating groups?
- Compress classes of groups close to simple, such as the almost simple groups ($S \leq G \leq \texttt{Aut}(S)$ for some finite simple group S), or the central extensions of a simple group.
- Find short f.o. descriptions of the simple Lie algebras over $\mathbb{C}$. Descriptions must work within the class of Lie algebras over $\mathbb{C}$.
- Fix a constant c. Develop a model theory for classes of finite structures where the language consists of the first-order formulas of size $O(\log^c)$.

References

- A. Nies and K. Tent, Describing finite groups by short first-order sentences. Israel J. of Mathematics, to appear. http://arxiv.org/abs/1409.8390
- Report by Yuki Maehara under Nies' supervision, http://arxiv.org/abs/1305.0080

Also see the Wikipedia page on straight line programs.