Linear representations of finitely presented groups: some decidability problems

November 14, 2017

For the sake of brevity, we will use *representation* to mean 'finite-dimensional linear representation over a field'. p will denote either a prime or zero. One can ask the following questions about a group G.

1. Does G have a non-trivial representation? $\mathbf{1}(p)$: ditto in characteristic p?
2. Does G have a representation with infinite image? $\mathbf{2}(p)$: ditto in characteristic p?
3. Does G have a representation with image that is not virtually soluble? $\mathbf{3}(p)$: ditto in characteristic p?
4. Does G have a faithful representation? $\mathbf{4}(p)$: ditto in characteristic p?

Let us say that Question n is *decidable* if there exists a uniform algorithm that answers the question for every finitely presented group.

It follows from the main result of [BW], which establishes the undecidability of the question 'does G have a non-trivial finite image?', that both 1 and 2, and their p-versions, are *undecidable*. Indeed, since finitely generated linear groups are residually finite (Mal'cev's theorem, see [W] Theorem 4.2), G has a non-trivial representation if and only if G has a non-trivial finite image, and this holds if and only if $G * G$ has a representation with infinite image (if $Q \neq 1$ is a finite group then $Q * Q$ is infinite and virtually free, hence linear). A similar argument shows that Questions 3 and $3(p)$ are undecidable (for $Q*Q$ is virtually soluble only when $|Q| \leq 2$, and it is easy to determine whether or not C_2 is a quotient of G).

Questions 4 and $4(p)$ are also undecidable: this follows from the Adian-Rabin Theorem, since the properties of being linear, or linear in characteristic p, are Markoff properties (see [LS] Chapter IV, Theorem 4.1).

What if we ask instead about representations of a given degree? Let us call a representaton of degree d a *d-representation*. Replacing *representation* by *d-representation* in Questions 1 - 4 gives the new Questions 1^d, $1^d(p), \ldots,$ 4^d, $4^d(p)$. We see as before that for each $d \geq 1$, Questions 4^d and $4^d(p)$ are undecidable.

Theorem 1 *For each natural number d, each of Questions 1^d - 3^d and $1^d(p)$ - $3^d(p)$ is decidable.*

The algorithms for Questions 1^d and $1^d(p)$ are constructive. The others involve an open-ended search.

In some cases the algorithm will also decide the analogous question for representations over a given field, as long as the field has decidable elementary theory; we leave this to the reader.

Throughout we fix a natural number d. For a field K we denote its algebraic closure by $\overline{K}$ and its prime field by K_*. Set

$$K_p = \overline{\mathbb{F}_p(T)} \text{ for a prime } p,$$
$$K_0 = \overline{\mathbb{Q}(T)}$$

where $k(T)$ denotes the field of rational functions in countably many variables. It is well known (and is easily deduced from the Noether Normalization Lemma, for example) that any representation of a finitely generated group in characteristic p can be specialized to such a representation over K_p. So when convenient, we may replace 'representation in characteristc p' by 'representation over K_p'.

The fields K_p are what Fried and Jarden [FJ] call 'presented rings', that is, their elements can be specified and polynomials can be effectively evaluated. Theorem 8.4 of [FJ] now gives

Theorem 2 *Let $P(\mathbf{x})$ be a formula in the first-order language of rings. The questions*

(a) Given $\mathbf{v} \in K_p^m$, is $P(\mathbf{v})$ true?

(b) Does there exist $\mathbf{v} \in K_p^m$, for some $p \geq 0$, such that $P(\mathbf{v})$ true?

are decidable.

1 Question 1

Let $G = \langle g_1, \ldots, g_m; R\rangle$ be a finite presentation of a group G. Let's say that G is *p-trivial* if G has no non-trivial d-representation over K_p.

For convenience, we assume that R is a finite set of *positive* words in $2m$ variables, evaluated on $(g_1, g_1^{-1}, \ldots, g_m, g_m^{-1})$. The *representation variety* V_G is defined as follows:

for any field k, the set $V_G(k)$ is the subset of $\mathrm{M}_d(k)^{2m}$ consisting of points $\mathbf{v}$ such that

$$v_{2i-1} \cdot v_{2i} = \mathbf{1}_d \quad (i = 1, \ldots, m)$$
$$w(\mathbf{v}) = \mathbf{1}_d \quad \forall w \in R.$$

V_G is an algebraic variety defined over $\mathbb{Z}$. There is a $1-1$ correspondence between the set of d-representations of G over k and the points of $V_G(k)$, which to each homomorphism $f : G \to \mathrm{GL}_d(k)$ associates the point

$$\mathbf{v}_f = (f(g_1), f(g_1^{-1}), \ldots, f(g_m), f(g_m^{-1})) \in V_G(k),$$

and to each $\mathbf{v} \in V_G(k)$ associates the homomorphism $f_{\mathbf{v}} : G \to \mathrm{GL}_d(k)$ such that $\mathbf{v}_{f_{\mathbf{v}}} = \mathbf{v}$.

Now G is p-trivial if and only if

$$V_G(K_p) = \{(\mathbf{1}_d, \ldots, \mathbf{1}_d)\}. \tag{1}$$

By Theorem 2 we can decide (a) for a given p whether (1) is true, and (b) whether (1) is true for all $p \geq 0$; the latter holds if and only if G has no non-trivial d-representation. Thus Questions $1^d(p)$ and 1^d are (constructively) decidable.

2 Uniform bounds

Let G be a finitely generated group. Let us say G is *small* if *every d-representation of G has finite image*, and G is *p-small* if this holds for d-representations in characteristic p.

Proposition 3 *If G is small (*resp. *p-small) then there exists $q \in \mathbb{N}$ such that $|f(G)| \leq q$ for every representation $f : G \to \mathrm{GL}_d(K)$, K a field (*resp. *K a field with* $\mathrm{char} K = p$*).*

Proof. Suppose G is a counterexample. For each $n \in \mathbb{N}$ there exist a field $K(n)$ (resp. a field $K(n)$ with $\mathrm{char}(K(n)) = p$) and $f_n : G \to \mathrm{GL}_d(K(n))$ with $|f_n(G)| > n$. Let

$$E = \left(\prod_{n \in \mathbb{N}} K(n) \right) / \mathcal{U}$$

be the ultraproduct over a non-principal ultrafilter $\mathcal{U}$. Then E is a field (resp. a field of characteristic p), and the maps f_n induce a representation $\Phi : G \to \mathrm{GL}_d(E)$. Now $|\Phi(G)|$ is finite, so $D = \ker \Phi$ has finite index in G and D is generated by a finite set Y. For each $y \in Y$ the set $\{n \in \mathbb{N} \mid f_n(y) = 1\}$ is in $\mathcal{U}$. As Y is finite it follows that $X := \{n \in \mathbb{N} \mid f_n(D) = 1\}$ is in $\mathcal{U}$. But

$$n \in X \Longrightarrow n < |f_n(G)| \leq |G : D|,$$

so X is finite. Contradiction. ■

Proposition 4 *If every d-representation of G (in characteristic p) has virtually soluble image, then there exists $q \in \mathbb{N}$ such that the image of every d-representation of G (in characteristic p) has a soluble normal subgroup of index at most q.*

Proof. For a linear group H let H_s denote the soluble radical of H. Say $G = \langle g_1, \ldots, g_m \rangle$. Suppose that for each $n \in \mathbb{N}$ there exist a field $K(n)$ (resp. a field $K(n)$ with $\mathrm{char}(K(n)) = p$) and $f_n : G \to \mathrm{GL}_d(K(n))$ such that $|f_n(G) : f_n(G)_s| > n$. We may assume that each $K(n)$ is algebraically closed.

Define E and $\Phi : G \to \mathrm{GL}_d(E)$ as in the preceding proof. Then E is algebraically closed, and $\mathrm{char}(E) = p$ if $\mathrm{char}(K(n)) = p$ for all n.

Suppose that $\Phi(G)$ is virtually soluble. Then $\Phi(G)$ has a triangularizable normal subgroup L of finite index k, say (the Lie-Kolchin Theorem, [W] Theorem 5.8). Put $H = \Phi^{-1}(L)$. Then $|G : H| \leq k$ and H is generated by some finite subset $\{y_1, \ldots, y_l\}$ of G.

Now let $T(z_1, \ldots, z_s; E)$ be the formula

$$T(z_1, \ldots, z_s; E): \quad \exists u, v \in \mathrm{M}_d(E).\ u.v = 1_d \ \& \ (uz_sv)_{ij} = 0 \ \ (1 \leq s \leq l,\ 1 \leq j < i \leq d). \tag{2}$$

For $z_1, \ldots, z_s \in \mathrm{GL}_d(E)$, this says that the group $\langle z_1, \ldots, z_s \rangle$ can be conjugated to triangular form over E.

As E is algebraically closed, L is triangularizable over E, so $T(\Phi(y_1), \ldots, \Phi(y_s); E)$ is true. It follows that $T(f_n(y_1), \ldots, f_n(y_s); K(n))$ is true for infinitely many values of n (Loś's Theorem, see e.g. [FJ] Proposition 6.11). For each such n the group $f_n(H)$ is triangularizable, hence soluble, whence

$$n < |f_n(G) : f_n(G)_s| \leq |f_n(G) : f_n(H)| \leq k,$$

contradiction!

We conclude that Φ is a d-representation of G (in characteristic p) such that $\Phi(G)$ is not virtually soluble. ■

Combined with Mal'cev's theorem [W], Theorem 3.6, this gives

Corollary 5 *If every d-representation of G (in characteristic p) has virtually soluble image, then there exists $q' \in \mathbb{N}$ such that the image of every d-representation of G over an algebraically closed field (of characteristic p) has a triangularizable normal subgroup of index at most q'.*

Remark When $p = 0$, Proposition 4 is immediate from Platonov's extension of Jordan's theorem, [W] **10.11**. In this case the bounds q and q' depend only on d, for example $q = (d+1)!$ will do if $d \geq 71$ [C].

3 Question 2

Let $G = \langle g_1, \ldots, g_m; R \rangle$ be as in Section 1.

Let $F_m = \langle x_1, \ldots, x_m \rangle$ be free of rank m. For each $n \in \mathbb{N}$ set

$$M_n = \bigcap \{N \lhd F_m \mid |F_m/N| \leq n\}, \tag{3}$$

put $n^* = |F_m/M_n|$, and let W_n be a finite set of positive words on $x_1^{\pm 1}, \ldots, x_m^{\pm 1}$ such that $M_n = \langle W_n \rangle$. For any given n we can effectively find a set W_n.

For any homomorphism $f : G \to \mathrm{GL}_d(K)$ we have

$$\begin{aligned} |f(G)| \leq n &\Longrightarrow w(f(g_1), f(g_1)^{-1}, \ldots, f(g_m), f(g_m)^{-1}) = 1 \text{ for each } w \in W_n \\ &\Longrightarrow |f(G)| \leq n^*. \end{aligned}$$

The middle statement '$w(f(g_1), \ldots, f(g_m)^{-1}) = 1$ for each $w \in W_n$' is equivalent to

$$S_n(\mathbf{v}) : \bigwedge_{w \in W_n} w(v_1, \ldots, v_{2m}) = 1$$

where $\mathbf{v}_f = \mathbf{v} = (v_1, \ldots, v_{2m}) \in V_G(K)$. Thus if $|f(G)| \leq n$ for every $f : G \to \mathrm{GL}_d(K)$ then

$$S_n(K) : (\mathbf{v} \in V_G(K)) \to S_n(\mathbf{v})$$

is true. Conversely, if $S_n(K)$ is true then $|f(G)| \leq n^*$ for every $f : G \to \mathrm{GL}_d(K)$.

The statements $S_n(K_p)$, for each $p \geq 0$, and the statement '$S_n(*) :\ S_n(K_p)$ holds for every $p \geq 0$' are decidable by Theorem 2.

Now we describe two procedures which together answer the question 'is G small?' (case i) or the question 'is G p-small?' (case ii)

Procedure 3.1.

For $n = 1, 2, \ldots$ decide: is $S_n(*)$ (in case i), $S_n(K_p)$ (in case ii) true or false?
If **false**, go to the next n.
If **true**, stop.

Suppose that G is small (resp. p-small). By Proposition 3, there exists n such that $S_n(*)$ (respectively $S_n(K_p)$) is true, so Procedure 3.1 terminates.

Conversely, suppose that Procedure 3.1 terminates at step n. Then $S_n(*)$ (respectively $S_n(K_p)$) is true, so $|f(G)| \leq n^*$ for every $f : G \to \mathrm{GL}_d(K_p)$ and all p (resp. for every $f : G \to \mathrm{GL}_d(K_p)$), and so G is small (resp. p-small).

Procedure 3.2.

Enumerate the points of $\bigcup_{p \geq 0} V_G(K_p)$ (in case i), the points of $V_G(K_p)$ (in case ii).
For $\mathbf{v} \in V_G(K_p)$ decide: is $|f_{\mathbf{v}}(G)|$ finite? (see Algorithm 3.3 below).
If **yes**, go to the next point.
If **no**, stop.

Clearly Procedure 3.2 in case i terminates if and only if G is not small, and it terminates in case ii if and only if G is not p-small.

Conclusion In each case (i or ii), either Procedure 3.1 terminates or Procedure 3.2 terminates, and this determines whether or not G is small (case i), respectively p-small (case ii).

Thus Questions 2^d and $2^d(p)$ are decidable.

Algorithm 3.3. Let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}_d(K_p)$.

1) When $p \neq 0$.

Let R be the subring of K_p generated by the entries of $g_1^{\pm 1}, \ldots, g_m^{\pm 1}$.

Find an epimorphism $\pi : R \to \mathbb{F}_q$ for some $q = p^e$. Put $n = q^{d^2}$.

Let $W_n = \{w_1, \ldots, w_s\} \subseteq F_m$ be the set of words specified at the beginning of this section, so $\langle W_n \rangle = M_n \lhd F_m$ (see (3)). The verbal subgroup

$$D = \gamma_d(F_s) \cdot F_s^{p^d}$$

has finite index in $F_s = \langle y_1, \ldots, y_s \rangle$. Find a finite set U of words on $\mathbf{y}$ such that $\langle U \rangle = D$.

Decide whether the following holds:

$$u(w_1(\mathbf{g}), \ldots, w_s(\mathbf{g})) = 1 \ \forall u \in U. \tag{4}$$

Claim: if (4) is **true** then G is finite; if (4) is **false** then G is infinite.

Proof. Let H be the kernel of the homomorphism $G \to \mathrm{GL}_d(\mathbb{F}_q)$ induced by π. Then H is residually a finite p-group (see the proof of [W], Theorem 4.7). So if G is finite then H is a finite p-group.

The choice of W_n ensures that $w(\mathbf{g}) \in H$ for each $w \in W_n$, so the group $H_1 = \langle w_1(\mathbf{g}), \ldots, w_s(\mathbf{g}) \rangle$ is a finite p-group, hence unipotent. It follows that $\gamma_d(H_1) \cdot H_1^{p^d} = 1$. This now implies (4).

Suppose conversely that (4) is true. Then H_1 is an image of the finite group F_s/D, so H_1 is finite. But H_1 is the image of M_n under the epimorphism $F_m \to G$ $(x_i \mapsto g_i)$, so $|G : H_1| \leq |F_m/M_n| < \infty$. It follows that G is finite. ■

2) When $p = 0$.

Let R be the subring of K generated by the entries of $g_1^{\pm 1}, \ldots, g_m^{\pm 1}$.

Find epimorphisms $\pi_j : R \to \mathbb{F}_{q_j}$ $(j = 1, 2)$ such that $q_j = p_j^{e_j}$ and p_1, p_2 are distinct primes. Put $n = q_1^{d^2} q_2^{d^2}$.

Find W_n as above.

Decide whether the following holds:

$$w_1(\mathbf{g}) = \ldots = w_s(\mathbf{g})) = 1. \tag{5}$$

Claim: if (5) is **true** then G is finite; if (5) is **false** then G is infinite.

Proof. As above, the group $H_1 = \langle w_1(\mathbf{g}), \ldots, w_s(\mathbf{g}) \rangle$ is now residually a p_i-group for $i = 1, 2$, hence torsion-free. Thus if G is finite then $H_1 = 1$. Conversely, if $H_1 = 1$ then, as above,

$$G = |G : H_1| \leq |F_m/M_n| < \infty.$$

■

4 Question 3

We repeat the method of the preceding section. Recall that the formula $T(z_1, \ldots, z_s; E)$ given in (2) expresses the statement: the group $\langle z_1, \ldots, z_s \rangle \leq \mathrm{GL}_d(E)$ can be

conjugated to triangular form over E. Let $W_n = \{w_1, \ldots, w_s\}$ be the set of words introduced above, and set

$$R_n(\mathbf{v}; E) := T(w_1(\mathbf{v}), \ldots, w_s(\mathbf{v}); E).$$

Thus for $v_1, \ldots, v_m \in \mathrm{GL}_d(K)$ and $\mathbf{v} = (v_1, v_1^{-1}, \ldots, v_m, v_{m^{-1}})$, $R_n(\mathbf{v}; K)$ is true if and only if the subgroup

$$M_n(\langle \mathbf{v} \rangle) := \langle w_1(\mathbf{v}), \ldots, w_s(\mathbf{v}) \rangle$$

is triangularizable over K. Recall that $M_n(\langle \mathbf{v} \rangle) \lhd \langle \mathbf{v} \rangle$, that $M_n(\langle \mathbf{v} \rangle) \leq N$ for every normal subgroup N of index at most n in $\langle \mathbf{v} \rangle$, and that $|\langle \mathbf{v} \rangle : M_n(\langle \mathbf{v} \rangle)| \leq n^* < \infty$.

Now let $G = \langle g_1, \ldots, g_m; R \rangle$ be as above and let $f : G \to \mathrm{GL}_d(K)$ be a homomorphism. The discussion at the beginning of the preceding section shows the following:

1. If $f(G)$ has a triangularizable normal subgroup of index at most n then $R_n(\mathbf{v}_f; K)$ is true;
2. if $R_n(\mathbf{v}_f; K)$ is true then $f(G)$ has a triangularizable normal subgroup of index at most n^*.

Let us call the group G *smallish* (*p-smallish*) if every d-representation of G (in characteristic p) has virtually soluble image.

Corollary 5 shows that G is smallish if and only there exists $n = n(G) \in \mathbb{N}$ such that $R_n(\mathbf{v}; K_p)$ is true for every $\mathbf{v} \in V_G(K_p)$ and every $p \geq 0$, and that G is p-smallish if and only there exists $n \in \mathbb{N}$ such that $R_n(\mathbf{v}; K_p)$ is true for every $\mathbf{v} \in V_G(K_p)$.

The next two procedures decide whether or not G is smallish (case i), or p-smallish' (case ii).

Procedure 4.1

For $n = 1, 2, \ldots$ decide whether the following is true or false:
Case *i:*

$$\text{for every } p \geq 0\text{: } (\mathbf{v} \in V_G(K_p)) \to R_n(\mathbf{v}; K_p).$$

Case *ii:*

$$(\mathbf{v} \in V_G(K_p)) \to R_n(\mathbf{v}; K_p).$$

These are decidable by Theorem 2.
If **false**, go to the next n.
If **true**, stop.

The preceding discussion shows that Procedure 4.1 terminates in case i if and only if G is smallish, and terminates in case ii if and only if G is p-smallish

Procedure 4.2.

Enumerate the points of $\bigcup_{p\geq 0} V_G(K_p)$ (in case *i*), the points of $V_G(K_p)$ (in case *ii*).

For $\mathbf{v} \in V_G(K_p)$ decide: is $f_{\mathbf{v}}(G)$ virtually soluble? (see Algorithm 4.3 below).

If **yes**, go to the next point.

If **no**, stop.

Clearly Procedure 4.2 in case *i* terminates if and only if G is not smallish, and it terminates in case *ii* if and only if G is not p-smallish.

Conclusion In each case (*i* or *ii*), either Procedure 4.1 terminates or Procedure 4.2 terminates, and this determines whether or not G is smallish (case *i)*, respectively p-smallish (case *ii*).

Thus Questions 3^d and $3^d(p)$ are decidable.

Mal'cev's Theorem [W], Theorem 3.6, says that a soluble linear group of degree d over an algebraically closed field has a triangularizable normal subgroup of index at most $\mu(d)$, a finite number depending only on d.

Algorithm 4.3. Let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}_d(K_p)$.

Let R be the subring of K_p generated by the entries of $g_1^{\pm 1}, \ldots, g_m^{\pm 1}$.

Find an epimorphism $\pi : R \to \mathbb{F}_q$ for some $q = p^e$. Put $n = (\mu(d)q^{d^2})!$.

Decide: is $R_n(\mathbf{g}; K_p)$ true or false?

Claim: G is virtually soluble if and only if $R_n(\mathbf{g}; K_p)$ is true.

Proof. If $R_n(\mathbf{g}; K_p)$ is true then $M_n(G)$ is triangularizable, so G is virtually soluble.

Suppose conversely that G is virtually soluble. Let H be the kernel of the homomorphism $G \to \mathrm{GL}_d(\mathbb{F}_q)$ induced by π. Theorems 1 and 2 of [W2] show that H is soluble, so H has a triangularizable normal subgroup H_1 of index at most $\mu(d)$. Then H_1 contains a normal subgroup N of G with

$$|G/N| \leq |G : H_1|! \leq n.$$

It follows that N contains $M_n(G)$, so $M_n(G)$ is triangularizable, and so $R_n(\mathbf{g}; K_p)$ is true. ■

References

[BW] M. Bridson and H. Wilton, The triviality problem for profinite completions, *Invent. Math.* **202** (2015), 839-874.

[C] M. J. Collins, On Jordan's theorem for complex linear groups, *J. Group Theory* **10** (2007), 411–423.

[FJ] M. D. Fried and M. Jarden, *Field Arithmetic,* Springer-Verlag, Berlin/Heidelberg, 1986.

[W] B. A. F. Wehrfritz, *Infinite Linear Groups*, Springer-Verlag, Berlin, 1973.

[LS] R. C. Lyndon and P. E. Schupp, *Combinatorial Group Theory,* Springer-Verlag, Berlin/Heidelberg, 1977.

[W2] B. A. F. Wehrfritz, Conditions for linear groups to have unipotent derived subgroups, *J. Algebra* **323** (2010), 3147–3154