A polynomial-time reduction algorithm for groups of semilinear or subfield class *

Jon F. Carlson[(1)], Max Neunhöffer[(2)] and Colva M. Roney-Dougal[(2)]

June 26, 2009

1. Department of Mathematics, University of Georgia, Athens, GA 30602, USA
2. School of Mathematics and Statistics, St Andrews, Fife KY16 9SS, Scotland, UK.
jfc@math.uga.edu, {neunhoef, colva}@mcs.st-and.ac.uk

To Derek and John, on the occasion of their accumulated 125th birthday.

Abstract

We present a Las Vegas algorithm for finding a nontrivial reduction of groups that are irreducible with m generators and either lie in the subfield class of matrix or projective groups or are semilinear or have non-absolutely irreducible derived group. Let $R_{\mathcal{A}}$ denote the cost of producing a random element from a matrix algebra $\mathcal{A}$ and $R_{\langle H^G \rangle}$ denote the cost of producing a random element in the normal closure of a group H by a group G. Then the algorithm runs in $O(d^3(m + d \log \log d \log q) + R_{\mathcal{A}} \log(\log d) + R_{\langle H^G \rangle} d \log q)$ finite field operations. We also characterise the absolutely irreducible groups G over arbitrary fields whose derived group consists only of scalars, and prove probabilistic generation results about matrix groups.

1 Introduction

The matrix group recognition project was begun some years ago by Neumann and Praeger in a groundbreaking paper [19]. Their results answered the question of how one can determine computationally whether a given set of invertible matrices with entries in a finite field $\mathbb{F}_q$ generates the group $\mathrm{SL}(d, q)$. Since then many algorithms for computing with matrix groups over finite fields have been developed. Given a collection $g_1, \ldots, g_m$ of matrices in $\mathrm{GL}(d, q)$, the basic problem is to find a composition series for the group G that they generate and to be able to express arbitrary group elements as straight line programs in the generators. An overview of the aims of the recognition project is given in [16].

The overall approach of the project relies on a fundamental theorem of Aschbacher [1] on the maximal subgroups of classical groups. The theorem says that every subgroup of $\mathrm{GL}(d, q)$ lies in at least one of nine classes $\mathcal{C}_1, \ldots, \mathcal{C}_9$. The classes $\mathcal{C}_1, \ldots, \mathcal{C}_8$ are treated by reducing to some sort of easier setting, and there are algorithms for these cases. However, the complexity of some of them has not been analysed and many do not run in polynomial time. The overall

*The first author was partially supported by a grant from NSF, the second and third by EPSRC grant EP/C523229/1 and the third by the Nuffield foundation.

project is currently in a second phase, producing provably polynomial-time algorithms for each class.

The basic approach (see [16, 20]) is first to reduce the problem by either finding a proper nontrivial homomorphism from G or finding an isomorphism to a representation with a smaller ambient group (for example to $\mathrm{GL}(d, q_0)$ for $q_0 < q$). If a homomorphism is found then the kernel and image are treated separately, eventually producing a *composition tree*, whose leaves are either simple groups or groups that can be constructively recognised by other means.

Let $\overline{G} \leq \mathrm{PGL}(d, q)$ be the corresponding projective group. In this paper we present a fully analysed, polynomial-time Las Vegas algorithm to find a reduction for the case that G or $\overline{G}$ is not in $\mathcal{C}_1$, but is in $\mathcal{C}_3$ or $\mathcal{C}_5$, or has non-absolutely-irreducible derived group. Class $\mathcal{C}_1$ (reducible groups) is completely under control using MeatAxe methods [14, 15, 21]. If G or $\overline{G}$ is in $\mathcal{C}_3$ or $\mathcal{C}_5$ we either find a proper nontrivial homomorphism from G to a permutation, matrix or projective group, or an isomorphism writing G or $\overline{G}$ over a smaller field, or in a smaller dimension. In addition, for some groups in classes $\mathcal{C}_2$, $\mathcal{C}_4$ or $\mathcal{C}_6$, we find a nontrivial reduction homomorphism: this is important as there is as yet no fully polynomial-time analysis for these classes.

Our algorithms are efficient in the sense that they use a number of field operations that is bounded by a low-degree polynomial in m (the number of generators), d and $\log q$: the input size is $O(md^2 \log q)$ (all logarithms are to base 2 unless otherwise stated). There are also some terms concerning random element construction, which will be discussed further in Section 7. We analyse the complexity of all algorithms during the course of the paper, and have implemented our work in GAP [9]. We avoid the use of a discrete logarithm oracle. We use 3 for the exponent of matrix multiplication, as although the theoretical exponent is lower than this, for practical implementations this is more realistic.

In addition to developing reduction algorithms, we characterise the groups which have a faithful absolutely irreducible module on which the derived group acts by scalar matrices. We also develop efficient Monte Carlo methods for generating subgroups of matrix groups that behave like normal subgroups. The use of Clifford's Theorem upgrades these algorithms to Las Vegas.

One of the motivations for this work is a recent article by Glasby, Leedham-Green and O'Brien [11], who develop an algorithm to recognise groups G in class $\mathcal{C}_5$, generalising [10]. The algorithm in [11] is polynomial time provided that the commutator subgroup acts absolutely irreducibly. Here we address the case where G' is not absolutely irreducible by providing fully analysed algorithms for all actions of G', including the case where G' consists only of scalars. In the paper [11] the authors appear to misstate the complexity of generating G'. To the best of our knowledge, the best published complexity for this is $O(d^7 \log^2 q)$. In this paper, we develop Las Vegas methods to generate a normal subgroup of G that is contained in G'.

Our approach in Sections 6.4 to 6.6 is heavily influenced by SMASH [12] and we have reused many subroutines. There are two main differences between these sections and the original treatment in SMASH. Firstly, we have analysed the probability of having generators for a subgroup that has the same submodule lattice as a normal subgroup of G. Secondly, we have improved algorithms and complexity estimates for finding an irreducible submodule of a normal subgroup. Hence we are able to derive tighter upper bounds on the complexity of our algorithm.

The layout of this paper is as follows. In Section 2 we present basic definitions and our main

result. In Section 3 we characterise absolutely irreducible matrix groups over arbitrary fields whose derived group is contained in the scalar matrices. In Section 4 we prove probabilistic results about the number of random elements required to generate a matrix group. In Section 5 we present an algorithm for writing irreducible matrix groups over a smaller field. In Section 6 we present the main body of our algorithm, followed by Section 7 which summarises the complexity results and Section 8 which reports on our implementation of these algorithms.

2 Definitions and main result

Throughout the paper (except for Sections 3 and 5), we assume that $g_1, \ldots, g_m \in \mathrm{GL}(d, q)$ and let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ be the corresponding matrix group. By considering each of $g_1, \ldots, g_m$ to be defined only up to scalar multiplication, we also define a group $\overline{G} = \langle \overline{g_1}, \ldots, \overline{g_m} \rangle \leq \mathrm{PGL}(d, q)$, which is the projective group generated by the given matrices. Two matrices represent the same elements of $\overline{G}$ if one is a scalar multiple of the other, so replacing any of the g_i by scalar multiples will alter the matrix group but not the projective group. We assume throughout that G acts irreducibly on the natural module $V = \mathbb{F}_q^d$.

Definition 2.1
The group G lies in $\mathcal{C}_3$ (the class of semilinear groups) if there is a divisor e of d with $1 < e < d$ and an $\mathbb{F}_q$-vector space identification between $\mathbb{F}_q^d$ and $\mathbb{F}_{q^e}^{d/e}$ such that for $1 \leq i \leq m$ there exist automorphisms $\alpha_i \in \mathrm{Gal}(\mathbb{F}_{q^e}/\mathbb{F}_q)$ with

$$(v + \lambda w)g_i = vg_i + \lambda^{\alpha_i} wg_i$$

for all $v, w \in \mathbb{F}_{q^e}^{d/e}$ and all $\lambda \in \mathbb{F}_{q^e}$. The group $\overline{G}$ lies in $\mathcal{C}_3$ if and only if G lies in $\mathcal{C}_3$: note that multiplying $g_1, \ldots, g_m$ by scalars from $\mathbb{F}_q$ does not affect the semilinearity of G.

If the α_i generate a proper subgroup of $\mathrm{Gal}(\mathbb{F}_{q^e}/\mathbb{F}_q)$ then multiplication by elements of the corresponding invariant subfield produces $\mathbb{F}_q G$-endomorphisms that are not $\mathbb{F}_q$-scalar, so V is not absolutely irreducible. Conversely, if V is not absolutely irreducible, then there is a divisor e' of d such that we can view V as a d/e'-dimensional vector space over $\mathbb{F}_{q^{e'}}$ on which the action of G is $\mathbb{F}_{q^{e'}}$-linear. That is, G lies in class $\mathcal{C}_3$ with trivial automorphisms.

Definition 2.2
The group G lies in $\mathcal{C}_5$ if there exists a subfield $\mathbb{F}_{q_0} \subsetneq \mathbb{F}_q$, a $t \in \mathrm{GL}(d, q)$, and $\beta_1, \ldots, \beta_m \in \mathbb{F}_q^\times$ such that $t^{-1} g_i t = \beta_i h_i$ with $h_i \in \mathrm{GL}(d, q_0)$. The group $\overline{G}$ lies in $\mathcal{C}_5$ if and only if G lies in $\mathcal{C}_5$: note that multiplying $g_1, \ldots, g_m$ by scalars from $\mathbb{F}_q$ does not change the membership of G in $\mathcal{C}_5$.

If $\beta_i = 1$ for all i then G can be written over $\mathbb{F}_{q_0}$. In general, G lies in $\mathcal{C}_5$ if G can be written over $\mathbb{F}_{q_0}$ modulo scalars. Note that $\overline{G}$ being in $\mathcal{C}_5$ implies that $\overline{G} \cong \langle \overline{h_1}, \ldots, \overline{h_m} \rangle$ embeds naturally in $\mathrm{PGL}(d, q_0)$.

We assume that the input to our algorithm is an irreducible group G: see Lemma 7.1 for the complexity of proving this. The MeatAxe run which shows G to be irreducible also computes

the endomorphism ring $E = \mathrm{End}_{\mathbb{F}_q G}(V)$. If G is irreducible but not absolutely irreducible, the ring E is an extension field of $\mathbb{F}_q$. This provides an explicit E-vector space structure on V and an E-linear action of the group generators.

We now summarise our algorithm, see the relevant sections for more details.

1. Let G be irreducible with endomorphism ring E of degree $e \geq 1$ over $\mathbb{F}_q$. If $e > 1$ find an explicit base change to express the generators over $\mathbb{F}_{q^e}$.

2. Check whether G can be written over a subfield $\mathbb{F}_{q_0}$ with $\beta_i = 1$ for $1 \leq i \leq m$, using the standard basis technique described in Section 5, and find the degree f of $\mathbb{F}_q$ over $\mathbb{F}_{q_0}$.

3. If $e > f^2$ then return a homomorphism into $\mathrm{GL}(d/e, q^e)$. Otherwise, if $f > 1$ then return a monomorphism into $\mathrm{GL}(d, q_0)$.

4. Choose any nonscalar generator g_i and check whether $[g_i, g_j]$ is scalar for all j. If so, jump to step 10.

5. Compute a normal subgroup N of the derived subgroup G' of G as in Section 6.1.

6. If N is absolutely irreducible, check whether N can be written over a smaller field, as in Section 6.3. If G is not contained in $\mathcal{C}_5$, return `false` as G is not in $\mathcal{C}_3$ or $\mathcal{C}_5$.

7. If N is irreducible but not absolutely irreducible find a semilinear decomposition of G, as in Section 6.4.

8. If N is reducible with more than one homogeneous component, find an imprimitive decomposition of G as in Section 6.5.

9. If N is reducible with a single homogeneous component with irreducible N-submodules of dimension greater than 1, find a tensor decomposition of G as in Section 6.6.

10. If $[g_i, g_j]$ is scalar for some nonscalar g_i and all j, find a nontrivial homomorphism from G to $\mathbb{F}_q^\times$ as in Sections 3 and 6.7.

We will show that all of our methods can be applied to both matrix and projective groups, because the success or failure of each step is unaffected by multiplying generating matrices by scalars.

All groups that we encounter in the algorithm will have at most $O(m + d \log q + \log \delta^{-1})$ generators, and be subgroups of $\mathrm{GL}(d, q)$. We let $R_{\langle H^L \rangle}$ denote the number of finite field operations required to produce an independent, uniformly-distributed, random element of the normal closure of a group H in a group L. Furthermore we let $R_{\mathcal{A}}$, where $\mathcal{A}$ is an algebra, denote the number of finite field operations required to produce an independent, uniformly-distributed, random element of $\mathcal{A}$.

The following theorem summarises the main algorithmic results of this article.

Theorem 2.3 (Main Theorem)
Let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ or $\overline{G} = \langle \bar{g}_1, \ldots, \bar{g}_m \rangle \leq \mathrm{PGL}(d, q)$ be an absolutely irreducible group that lies in $\mathcal{C}_3$ or $\mathcal{C}_5$ or whose derived group is not absolutely irreducible. There exists an $O(d^3(m + d\log(\log d)\log q) + R_{F[K]}\log(\log d) + R_{\langle H^G \rangle} d \log q)$ Las Vegas algorithm to find a nontrivial reduction of G or $\overline{G}$, respectively. Here the group H has $O(d \log q)$ generators, and $R_{F[K]} = \max\{R_{\mathbb{F}_p[G]}, R_{\mathbb{F}_q[H]}\}$.

The complete procedure is Las Vegas in that we can prescribe an upper bound δ for the failure probability. The algorithm can succeed by returning a homomorphism or reporting `false`; or it can report `fail` with a prescribed probability bound δ. If success is reported, the result is guaranteed to be correct. If the algorithm reports `false` then some additional information may be deduced: for example, that if G lies in $\mathcal{C}_2$ then G' is transitive on all possible sets of blocks.

3 Characterisation of groups with scalar derived group

In this section we investigate groups G which satisfy the following hypothesis.

Hypothesis 3.1
The group $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, k)$ is finite and absolutely irreducible, with k an arbitrary field and $d > 1$. Furthermore, the derived group G' contains only scalar matrices.

Equivalently, the corresponding projective group $\overline{G}$ is abelian.
Let V be the natural G-module. The hypothesis implies the following facts.

Lemma 3.2
Suppose that G and V are as above. The following all hold.

1. *The derived group G' is contained in the centre $Z(G)$.*
2. *The group G is nilpotent of class 2 and hence is a direct product of its Sylow subgroups.*
3. *The centre and the derived group of G are cyclic.*
4. *If the characteristic of k is $p > 0$ then the order of G is not divisible by p.*
5. *Let $[g, h] = g^{-1}h^{-1}gh$ be the commutator of elements g and h in G. Then*

$$[g, h_1 h_2] = [g, h_1][g, h_2], \qquad [h_1 h_2, g] = [h_1, g][h_2, g]$$

   for all g, h_1, h_2 in G.
6. *Let $k^\times$ denote the multiplicative group of k. For any $g \in G$ there is a homomorphism $\psi_g : G \longrightarrow k^\times$ given by $\psi_g(h)I_d = [h, g]$. These homomorphisms satisfy $\psi_{g_1 g_2}(x) = \psi_{g_1}(x)\psi_{g_2}(x)$ for all $x, g_1, g_2 \in G$. Moreover, ψ_g is a constant function if and only if $g \in Z(G)$.*

PROOF: Part (1) is trivially true. Part (2) follows from (1). It is well-known that finite nilpotent groups are the direct product of their Sylow subgroups.

Part (3) is true because G is absolutely irreducible, so $Z(G)$ is a group of scalar matrices, which must be cyclic. Part (4) then follows from the fact that a Sylow p-subgroup of G would contribute a factor of p to the order of $Z(G)$.

(5) is a straightforward calculation. That is,

$$h_1h_2g[g,h_1h_2] = gh_1h_2 = h_1g[g,h_1]h_2 = h_1gh_2[g,h_1] = h_1h_2g[g,h_1][g,h_2]$$

since $[g,h_1],[g,h_2] \in Z(G)$ by (1). The second equation follows by inverting the first. Part (6) is a direct consequence of (5). □

The lemma allows us to prove the following.

Proposition 3.3
Let z be the order of $Z(G)$, and let c be the order of G'. Then:

1. *The exponent of $G/Z(G)$ is at most c.*
2. *The exponent of G is at most cz.*
3. *The order of G is a divisor of $c^l z$ where $l = |\{i \mid g_i \notin Z(G)\}|$.*

PROOF: The group G is generated by elements $g_1,\ldots,g_m$. Let $\psi_i : G \to G'$ be given by $\psi_i(h) = [h,g_i]$. Then ψ_i is a homomorphism by Lemma 3.2.6, and the kernel K_i of ψ_i has index in G at most c, since $\mathrm{Im}(\psi_i)$ has order dividing c. Let ψ_i^c be defined by $\psi_i^c(x) := \psi_i(x)^c$. Then ψ_i^c is the constant homomorphism from G to $\{I_d\}$. From this and Lemma 3.2.6 it follows that g_i^c is in the centre of G for all i. This proves (1).

Part (2) is a direct consequence of (1) since the exponent of $Z(G)$ is z. Finally, (3) is a simple count based on the fact that $\cap_i K_i \leq Z(G)$. □

We know by hypothesis that V is absolutely irreducible and by Lemma 3.2.2 that G is a direct product $G = S_1 \times S_2 \times \cdots \times S_t$ of its Sylow subgroups. From this we get the next lemma.

Lemma 3.4
The module V is a tensor product

$$V \;\cong\; V_1 \otimes \cdots \otimes V_t$$

where each V_i is an absolutely irreducible module for S_i on which S_j acts trivially for $i \neq j$.

PROOF: It is well-known that irreducible modules of direct products are tensor products (see for instance [7, 51.13]). Since V is absolutely irreducible, so is each factor. □

Proposition 3.5 (Eigenspace decomposition)
Let $g \in G$ have all eigenvalues in k. Then V is a vector space direct sum

$$V = V_{\lambda_1} \oplus V_{\lambda_2} \oplus \cdots \oplus V_{\lambda_s}$$

where $\lambda_1,\ldots,\lambda_s$ are the eigenvalues of g and V_{λ_i} is the λ_i-eigenspace. Moreover, if $h \in G$, then $V_{\lambda_i}h = V_{\lambda_j}$ where $\lambda_j = \psi_g(h)\lambda_i$.

PROOF: The space V is a direct sum of eigenspaces of g since $k[\langle g \rangle]$ is semi-simple and split by k. For the next statement, let $v \in V_{\lambda_i}$. Then as asserted

$$(vh)g = vgh[h, g] = \lambda_i v[h, g]h = \lambda_i \psi_g(h)vh.$$

□

If $g \in Z(G)$, then in Proposition 3.5, V is only a single eigenspace for the action of g.

Theorem 3.6 (Noncentral element of prime order)
Let r be a prime and let $G \leq \mathrm{GL}(d, k)$ be an r-group satisfying Hypothesis 3.1. Then either G is isomorphic to the quaternion group of order 8 or G has a noncentral element g of order r.

PROOF: Since $Z(G)$ is cyclic, it suffices to prove that either G is isomorphic to the quaternion group of order 8 or G contains more than one cyclic subgroup of order r.

It is well-known (see for instance [23, 3.15]) that the only r-groups which contain a single subgroup of order r are the cyclic groups and the generalised quaternion groups. Since G is absolutely irreducible and $d > 1$, the group G is not cyclic and so either G has a noncentral element of order r or G is isomorphic to a generalised quaternion group.

The generalised quaternion group of order 2^i for $i \geq 3$ has presentation $\langle a, b \, | a^{2^{i-1}} = b^4 = a^{2^{i-2}} b^{-2} = b^{-1}aba = 1 \rangle$. A short calculation shows that the derived group contains noncentral elements for $i > 3$, and hence if G is isomorphic to a generalised quaternion group then $|G| = 8$. □

We finish this section with our characterisation of the groups satisfying Hypothesis 3.1.

Theorem 3.7 (Characterisation Theorem)
Let $G \leq \mathrm{GL}(d, k)$ be absolutely irreducible, with $d > 1$ and k an arbitrary field, such that the derived group of G is contained in the set of scalar matrices. Then either $d = 2$ and G is isomorphic to an extension by scalars of the quaternion group of order 8 acting semilinearly, or G is imprimitive.

PROOF: By Lemma 3.4 we may consider V as a tensor product, with a distinct Sylow subgroup of G acting on each tensor factor. Let V_i be one such factor.

The first possibility is that V_i is 1-dimensional, and S_i is cyclic. Secondly, if S_i is isomorphic to Q_8 then the dimension of V_i is 2 (see for instance [7, §47]).

Otherwise, by Theorem 3.6 the group S_i has a noncentral element g of order r. The group G also contains a central element of order r, which is a scalar. This shows that the field k contains primitive rth roots of unity. Hence all of the eigenvalues of g lie in k. Since g is not central, it has more than one eigenvalue. It follows from Proposition 3.5 that the elements of G permute the eigenspaces of g. Since G is irreducible, this action on the eigenspaces is transitive and hence G is imprimitive.

If at least one of the induced actions is imprimitive then G is imprimitive. Not all of the Sylow subgroups can be cyclic since $d > 1$. □

All representations of extraspecial r-groups over finite fields $\mathbb{F}_q$ with $d = r^n$ and r dividing $(q - 1)$ lie in this class, but so do other r-groups. For example, the subgroup G of $\mathrm{GL}(3, 19)$ of

order 3^4 generated by

$$\begin{bmatrix} 0 & 0 & 1 \\ 17 & 0 & 0 \\ 0 & 11 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 7 & 0 \\ 0 & 0 & 11 \end{bmatrix}$$

satisfies $G' \leq Z(\mathrm{GL}(3,19))$ but does not contain an extraspecial group of order 3^{1+2}.

4 Generation of matrix groups by random elements

In this section we analyse the generation of a subgroup $H = \langle g_1, \ldots, g_n \rangle$ of a normal subgroup N of a matrix group G, and in particular provide bounds on n for H to have the same submodule structure and endomorphism ring as N, with probability at least $1 - \delta$. Perhaps surprisingly, we do this via results for permutation groups.

Lemma 4.1
Let $X_1, X_2, \ldots$ be a sequence of 0-1 valued random variables such that $\mathrm{Prob}(X_i = 1) \geq p$ for any values of the previous X_j (but the distribution of X_i may depend on the outcome of the X_j for $j < i$).

Then, for all integers t and all $0 < \epsilon < 1$,

$$\mathrm{Prob}\left(\sum_{i=1}^{t} X_i \leq (1-\epsilon)pt\right) \leq e^{-\epsilon^2 pt/2}.$$

PROOF: See [2, Corollary 2.2] or [22, Lemma 2.3.3]. □

The following proposition is based on [22, 2.3.7], where it is proved for the case G transitive. Note that the hypotheses here are slightly more general than in [22, 2.3.7], where G is given as a group of permutations. Our Proposition 4.2 can for example be applied to any finite group equipped with a permutation action.

Proposition 4.2 (Correct orbits of subgroup)
Suppose that a finite group G acts on a finite set Ω, with α orbits. Let $1 > \delta > 0$ be arbitrary. Then with probability at least $1-\delta$, a sequence of $\max\{24 \log_e \delta^{-1}, 45 \log_e |\Omega|\} = O(\log \delta^{-1} + \log |\Omega|)$ uniformly distributed random elements of G generates a subgroup of G that has the same orbits on Ω as G.

PROOF: Let $t = c \log_e |\Omega|$ where $c \geq \max\{24 \log_e \delta^{-1} / \log_e |\Omega|, 45\}$. Let $g_1, \ldots, g_t$ be uniformly distributed random elements of G. For $1 \leq i \leq t$ let $G_i = \langle g_1, \ldots, g_i \rangle$, let N_i be the number of G_i-orbits on Ω, and let M_i be the number of G_i-orbits that coincide with G-orbits. Let $k_i = N_i - M_i$. Note that $N_i \geq \alpha$ for $1 \leq i \leq t$, that $M_i \leq \alpha$ for $1 \leq i \leq t$, and that $N_i = \alpha$ if and only if $M_i = \alpha$. Hence k_i is either 0 or at least 2.

We claim that if $N_{i-1} > \alpha$, then

$$\mathrm{Prob}\left(k_i \leq \frac{7}{8} k_{i-1}\right) \geq \frac{1}{3}.$$

To see this, let $k = k_{i-1}$ and let $\Delta_1, \ldots, \Delta_k$ be the G_{i-1} orbits on Ω that are *not* G-orbits. For $1 \leq j \leq k$, let $X_j = 1$ if $\Delta_j^{g_i} \neq \Delta_j$ and let $X_j = 0$ otherwise. Now, Δ_j lies in an orbit of length at least two in the action of G on subsets of Ω. Therefore at most half of the elements of G fix Δ_j, and so $E(X_j) \geq 1/2$ for $1 \leq j \leq k$. Let $X = \sum_{j=1}^{k} X_j$, then $E(X) \geq k/2$.

Let p be the probability that $X \leq k/4$. Then with probability p the variable X takes value at most $k/4$ whilst X takes value greater than $k/4$ and less than or equal to k with probability $1 - p$. Therefore

$$\frac{pk}{4} + (1-p)k \geq E(X) \geq k/2,$$

so $p \leq 2/3$. Hence, with probability at least $1/3$, at least $k/4$ of the G_{i-1}-orbits that are not G-orbits are proper subsets of orbits of G_i. Thus, with probability at least $1/3$, the number of orbits of G_i which are not orbits of G is at most $7k/8$, and the claim follows.

Define $Y_1, \ldots, Y_t$ by

$$\begin{array}{ll} Y_i = 0 & \text{if } k_i > 0 \text{ and } k_i > \frac{7}{8}k_{i-1} \\ Y_i = 1 & \text{if } k_i = 0 \text{ or } k_i \leq \frac{7}{8}k_{i-1}. \end{array}$$

By the previous claim, $\mathrm{Prob}(Y_i = 1) \geq 1/3$.

Now, $N_0 = |\Omega|$, so $k_0 \leq |\Omega|$. Clearly, $k_t \leq k_0 \left(\frac{7}{8}\right)^{\sum_{i=1}^{t} Y_i}$. The group G_t has the same orbits as G if and only if $k_t \leq 1$, which will follow if $|\Omega| \left(\frac{7}{8}\right)^{\sum_{i=1}^{t} Y_i} \leq 1$. In turn this simplifies to $|\Omega| \leq \left(\frac{8}{7}\right)^{\sum_{i=1}^{t} Y_i}$, which gives

$$\sum_{i=1}^{t} Y_i \geq \frac{\log_e |\Omega|}{\log_e \frac{8}{7}}.$$

Then by Lemma 4.1, with $p = 1/3$, $t = c \log_e |\Omega|$ and $\epsilon = 1 - 3/(c \log_e(8/7))$ we get

$$\mathrm{Prob}(\sum_{i=1}^{t} Y_i \leq \frac{\log_e |\Omega|}{\log_e \frac{8}{7}}) \leq e^{-\frac{1}{6}(1-\frac{3}{c \log_e(8/7)})^2 c \log_e |\Omega|} \leq e^{-c \log_e |\Omega|/24}$$

for $c \geq 45$. In turn this is less than or equal to δ. □

We now apply the previous proposition to matrix groups, by considering their action on vectors. Let $s_{\delta,d,q} := \max\{24(1+\log_e \delta^{-1}), 45d \log_e q\} + \max\{22 \log_e d, 16(1+\log_e \delta^{-1})/3\}$.

Theorem 4.3 (Correct action of subgroup)
Let $G \leq \mathrm{GL}(d,q)$, and let $1 > \delta > 0$ be arbitrary. With probability at least $1 - \delta$, a sequence of $s_{\delta,d,q} = O(\log \delta^{-1} + d \log q)$ uniformly distributed random elements of G generate a subgroup H of G with the same submodule lattice as G on $V = \mathbb{F}_q^d$. Furthermore, if G is irreducible then $\mathrm{End}_{\mathbb{F}_q G}(V) = \mathrm{End}_{\mathbb{F}_q H}(V)$ with probability at least $1 - \delta$.

PROOF: First consider G as a permutation group on $|\Omega| = q^d$ points. By Proposition 4.2, any group H generated by $\max\{24(1 + \log_e \delta^{-1}), 45d \log_e q\}$ uniformly distributed random elements of G has the same orbits as G with probability at least $1 - \delta/2$.

A submodule for G is a union of orbits of G in its action on vectors that is closed under addition and scalar multiplication, so the first claim follows.

For the second claim, let G be irreducible and $\mathrm{End}_{\mathbb{F}_q G}(V) = \mathbb{F}_{q^f}$. Let H_0 be generated by $\max\{24(1+\log_e \delta^{-1}), 45d\log_e q\}$ random elements of G, so that H_0 is irreducible with probability $1-\delta/2$. Let $t = c\log_e d$ for $c \geq \max\{22, 16\log_e(2\delta^{-1})/(3\log_e d)\}$ and let $h_1, \ldots, h_t$ be further random elements of G. For $1 \leq i \leq t$ let $H_i = \langle H_0, h_1, \ldots, h_i\rangle$. Notice that $H = H_t$ is generated by $O(\log \delta^{-1} + d\log q)$ elements of G.

Since H_0 is irreducible, $\mathrm{End}_{\mathbb{F}_q H_0}(V) = \mathbb{F}_{q^s}$ for some s that is a multiple of f and divides d. For $1 \leq i \leq t$ let N_i be the degree of $\mathrm{End}_{\mathbb{F}_q H_i}(V)$ over $\mathbb{F}_{q^f}$.

We claim first that if $N_{i-1} > 1$ then $\mathrm{Prob}(N_i \leq N_{i-1}/2) \geq 1/2$. To see this let x generate $\mathrm{End}_{\mathbb{F}_q H_{i-1}}(V)$. Then since x is not centralised by G, at most half of the elements of G commute with x. If $[h_i, x] \neq 1$ then N_i is a proper divisor of N_{i-1} so the claim follows.

Now for $1 \leq i \leq t$ define $Y_i = 0$ if $N_{i-1} > 1$ and $N_i = N_{i-1}$ and $Y_i = 1$ otherwise. If $\sum_{i=1}^t Y_i \geq \log_2 d \geq \log_2(d/f)$ then $\mathrm{End}_{\mathbb{F}_q H}(V) = \mathrm{End}_{\mathbb{F}_q G}(V)$. Now, $\mathrm{Prob}(Y_i = 1) \geq 1/2$ by the claim, so by Lemma 4.1 with $p = 1/2$, $t = c\log_e d$ and $\epsilon = 1 - 2/(c\log_e 2)$

$$\mathrm{Prob}(\sum_{i=1}^{t} Y_i \leq \log_2 d) \leq e^{-\frac{1}{4}(1-\frac{2}{c\log_e 2})^2 c\log_e d} \leq e^{-\frac{3}{16}c\log_e d}$$

since $c \geq 22$. This is at most $\delta/2$ so the result follows. □

Recall the definition of $s_{\delta,d,q}$ given before the previous theorem.

Corollary 4.4 (Correct action of normal subgroup)
For $G \leq \mathrm{GL}(d,q)$, let $N \trianglelefteq G$ and let $1 > \delta > 0$ be arbitrary. With probability $1-\delta$ any group H generated by $s_{\delta,d,q} = O(\log \delta^{-1} + d\log q)$ uniformly distributed elements of N has the same submodule lattice as N, the same homogeneous components as N and, if N is irreducible, then $\mathrm{End}_{\mathbb{F}_q H}(V) = \mathrm{End}_{\mathbb{F}_q N}(V)$.

5 Writing G over a smaller field

In this section unless indicated otherwise we let K be a finite field, let $G = \langle g_1, \ldots, g_m\rangle \leq \mathrm{GL}(d,K)$, and let $V = K^{1\times d}$ be the natural right KG-module. We assume that V is irreducible but not necessarily absolutely irreducible. We want to determine whether there exists a $t \in \mathrm{GL}(d,K)$ such that for $1 \leq i \leq m$ the matrices $t^{-1}g_i t$ have entries over some proper subfield of K. If such a t exists, we want to construct it for the smallest possible subfield F of K. We first analyse when such a t exists.

The K-algebra $K \otimes_F FG$ is isomorphic as a K-algebra to the group algebra KG by the F-linear map given by $x \otimes g \mapsto xg$ for $x \in K$ and $g \in G$. This isomorphism makes the tensor product $K \otimes_F \tilde{V}$ into a KG-module, for any FG-module $\tilde{V}$.

Lemma 5.1
There exists a $t \in \mathrm{GL}(d,K)$ such that $t^{-1}Gt \in \mathrm{GL}(d,F)$ if and only if there exists an irreducible FG-module $\tilde{V}$ such that $V \cong K \otimes_F \tilde{V}$ as KG-modules.

PROOF: If there exists a $\tilde{V}$ such that $V \cong K \otimes_F \tilde{V}$ as KG-modules then there is an irreducible representation of G over F which is K-equivalent to the natural representation of G on V, hence there is a t as required.

On the other hand, such a t gives rise to a representation of G over F and thus to an FG-module $\tilde{V}$. The extension of scalars $K \otimes_F \tilde{V}$ of $\tilde{V}$ to K is isomorphic to V. If $\tilde{V}$ had a nontrivial FG-invariant subspace then V would have a nontrivial KG-invariant subspace, thus $\tilde{V}$ is irreducible. □

For a subfield F of K we denote by $F[G]$ the set of F-linear combinations of the elements of G as an F-subalgebra of $K^{d\times d}$. This is also called the *F-enveloping algebra of G*. We denote the prime field of K by K_0.

Proposition 5.2 (Prime field enveloping algebra I)
Let $G \leq \mathrm{GL}(d, F)$ for F finite, and let $V := F^{1\times d}$ be irreducible. Let $E := \mathrm{End}_{F[G]}(V) = \mathrm{End}_{FG}(V)$ with $e = [E : F]$ and $d' = d/e$. Identify V with $E^{1\times d'}$ so that $F[G] = E^{d'\times d'}$. Set $L := F_0[G] \cap (E \cdot \mathbf{1})$. Then $F_0[G] \cong L^{d'\times d'}$ as an F_0-algebra.

PROOF: Clearly $e = 1$ and $E = F$ if and only if V is absolutely irreducible.

Choosing an F-basis $(c_1, \ldots, c_e)$ of E we can express each element of E as an $(e \times e)$-matrix over F. The set V is an E-vector space and if $(b_1, \ldots, b_{d'})$ is an E-basis of V, then $(c_i b_j)_{i,j}$ is an F-basis of V. This choice of basis fixes an embedding $E^{d'\times d'} \subseteq F^{d\times d}$. Since the action of G on V is E-linear, we may assume that $G \leq \mathrm{GL}(d', E) \leq \mathrm{GL}(d, F)$. By the Density Theorem (see [6, (3.27)]), since V is an irreducible $F[G]$-module, $F[G] = E^{d'\times d'}$.

Now consider $B := F_0[G]$, which is an F_0-subalgebra of $F[G]$ such that $FB = E^{d'\times d'} = F[G]$. We first show that B is a simple algebra. If J is a nilpotent two-sided ideal of B, then FJ is a nilpotent two-sided ideal in $FB = E^{d'\times d'}$, contradicting its simplicity. So B has no nilpotent two-sided ideals and hence is semi-simple. It follows that B is a direct sum of simple algebras. The identity elements in these simple summands form an orthogonal set of central idempotents in B. A central idempotent in B is also central in $FB = E^{d'\times d'}$, and hence is the identity. Consequently, B is a simple algebra.

By the usual Wedderburn Theorems there exists an isomorphism $\psi : L^{s\times s} \to B$ for some s, and some extension L of F_0. The field L need not contain F. However, the elements of B corresponding to scalar matrices in $L^{s\times s}$ are central in B and hence also central in $FB = E^{d'\times d'}$. Therefore we can identify L with the centre of B and thus with some subfield of E. That is, $L = F_0[G] \cap (E \cdot \mathbf{1}) \subseteq E^{d'\times d'}$.

This produces a homomorphism of rings

$$\varphi : E^{s\times s} \cong E \otimes_L L^{s\times s} \cong E \otimes_L F_0[G] \to E^{d'\times d'} = F[G]$$

given by $\varphi(x \otimes b) = xb$. Since φ is surjective and $E^{s\times s}$ is simple, φ is an isomorphism and thus $s = d'$. □

Proposition 5.3 (Prime field enveloping algebra II)
Let G be as in Proposition 5.2, and suppose additionally that there is no proper subfield D of F such that there exists $t \in \mathrm{GL}(d, F)$ with $G^t \leq \mathrm{GL}(d, D)$. Then $F_0[G] = F[G]$.

PROOF: Let $\psi : L^{d'\times d'} \to F_0[G]$ be the isomorphism given by Proposition 5.2. Let $e_{i,j} \in F_0[G]$ for $1 \leq i, j \leq d'$ be the image under ψ of a set of matrix units in $L^{d'\times d'}$. Then $e_{i,j}e_{k,l} = \delta_{j,k}e_{i,l}$ and the $e_{i,j}$ are an L-basis of $F_0[G]$.

We claim that the $e_{i,j}$ are an E-basis of $E^{d'\times d'}$. To see linear independence, let

$$\sum_{i,j=1}^{d'} \lambda_{i,j}e_{i,j} = 0,$$

then multiplying on the left by $e_{k,k}$ and on the right by $e_{l,l}$ shows that $\lambda_{k,l}e_{k,l} = 0$ and thus $\lambda_{k,l} = 0$ for all k, l. On the other hand, since $F_0[G]$ is the F_0-span of the elements of G, the $e_{i,j}$ span $E[G] = E^{d'\times d'}$ as an E-vector space. Thus they are an E-basis of $E^{d'\times d'}$.

Since $\mathbf{1}_{\mathbf{d'}\times\mathbf{d'}} = \sum_{i=1}^{d'} e_{i,i}$ gives rise to a decomposition of $E^{1\times d'}$ as an E-vector space in which the direct summands are the row spaces of the $e_{i,i}$, it follows that these row-spaces are all one-dimensional.

Let $b'_1 \in E^{1\times d'}$ such that $\langle b'_1\rangle_E$ is the row space of $e_{1,1}$ and set $b'_i := b'_1e_{1,i}$. Then $b'_ie_{j,k} = \delta_{i,j}b'_k$. If $t^{-1} \in E^{d'\times d'}$ has rows $b'_1, b'_2, \ldots, b'_{d'}$, then $t^{-1}e_{i,j}t$ has 1_E in position (i, j) and zeroes elsewhere. Thus $F_0[G^t] = L^{d'\times d'}$.

If V is absolutely irreducible then $E = F$ and so $L \leq F$. By assumption F is the smallest possible field over which G can be written, so $L = F = E$ as required.

Now consider the case $F < E$ and let $l := [E : L]$. Since $F[G^t] = F[G] = E^{d'\times d'}$ it follows that $F \cdot L^{d'\times d'} = E^{d'\times d'}$. Therefore, E is the smallest field containing both F and L implying that l and $e = [E : F]$ are coprime. Let $D := F\cap L$. Then D is a field with $[F : D] = l$ and $[L : D] = e$.

We claim that G can be written over D. Let $(c'_1, \ldots, c'_e)$ be a D-basis of L. Then it is also an F-basis of E, since every element of E is an F-linear combination of elements of L. Now change basis in $V = F^{1\times d}$ from the F-basis $(c_ib'_j)_{i,j}$ to $(c'_ib'_j)_{i,j}$ using a base change $s \in \mathrm{GL}(d, F)$, to get $G^{ts} \leq D^{d\times d}$.

However, our assumption that F is the smallest subfield of F over which G can be written implies that $D = F \cap L = F$ and thus $F \leq L$. From $F \cdot L^{d'\times d'} = E^{d'\times d'}$ it now follows immediately that $L = E$.

Since we have proved $L = E$ in both cases and we already know that $F_0[G]$ is isomorphic to $L^{d'\times d'}$ as F_0-algebras and contained in $F[G] = E^{d'\times d'}$, the proposition follows. □

Theorem 5.4 (Characterisation of smallest possible field)
Let $G \leq \mathrm{GL}(d, K)$ act irreducibly on its natural module. Then there is a unique smallest subfield F of K such that there exists $t \in \mathrm{GL}(d, K)$ with $G^t \leq \mathrm{GL}(d, F)$. This F is uniquely determined by $K_0[G] \cap (K \cdot \mathbf{1}_{d\times d}) = F \cdot \mathbf{1}_{d\times d}$. Furthermore, $K_0[G] \cong E^{(d/e)\times(d/e)}$ where $E = \mathrm{End}_{F[G^t]}(F^{1\times d})$ is an extension field of F of degree e.

PROOF: Since K is finite there is a smallest subfield F of K such that there exists a $t \in \mathrm{GL}(d, K)$ with $G^t \leq \mathrm{GL}(d, F)$.

By Lemma 5.1 the natural $F[G^t]$-module $V := F^{1\times d}$ is irreducible. Then Proposition 5.3, applied to G^t, shows that $F_0[G^t] = F[G^t] = E^{(d/e)\times(d/e)}$. Since the F-scalar matrices are central in $F^{d\times d}$, the theorem follows immediately. □

From now on we assume that the equivalent statements in Lemma 5.1 hold for some subfield F of K. We now develop some theory which leads to an algorithm that finds a t and the smallest possible subfield F, or proves that none exists.

Let $B = \{b_1, \ldots, b_f\}$ be an F-basis for K, such that $b_1 = 1$. We start by noting that if the natural KG-module V is isomorphic to $K \otimes_F \tilde{V}$ as KG-modules, then $V \cong \bigoplus_{i=1}^{f} b_i \otimes_F \tilde{V}$ as FG-modules. We therefore identify V with $K \otimes_F \tilde{V}$ and $\tilde{V}$ with $1 \otimes_F \tilde{V}$ respectively via this second isomorphism, and thus write $b_i v$ instead of $b_i \otimes_F v$ and v for $1 \otimes_F v \in \tilde{V}$.

Lemma 5.5
Let F be a subfield of K such that the equivalent statements in Lemma 5.1 hold and let $\tilde{V}$ be as above. Then the F-dimension of $\mathrm{End}_{FG}(\tilde{V})$ is equal to $e = \dim_F(\mathrm{End}_{KG}(V))$.

PROOF: This result is a consequence of the fact that

$$\mathrm{End}_K(K \otimes_F \tilde{V}) \cong K \otimes_F \mathrm{End}_F(\tilde{V}),$$

see for example [6, (2.38)]. To assist the reader and set up some notation, we first prove that $\mathrm{End}_K(V) \cong K \otimes_F \mathrm{End}_F(\tilde{V})$. If $(m_1, \ldots, m_d)$ is an F-basis of $\tilde{V}$, then $(b_i m_j)_{1 \le i \le f, 1 \le j \le d}$ is an F-basis of V and $(m_j)_{1 \le j \le d} = (1 \otimes_F m_j)_{1 \le j \le d}$ is a K-basis of V. Hence every $\varphi \in \mathrm{End}_K(V)$ can be written in a unique way as

$$\varphi = \sum_{i=1}^{f} b_i \varphi_i$$

with $\varphi_i \in \mathrm{End}_F(\tilde{V})$. Therefore $\mathrm{End}_K(V) \cong K \otimes_F \mathrm{End}_F(\tilde{V})$.

Next we show that $\mathrm{End}_{KG}(V) \cong K \otimes_F \mathrm{End}_{FG}(\tilde{V})$. If $\psi \in \mathrm{End}_K(V)$ then $\psi \in \mathrm{End}_{KG}(V)$ if and only if $\psi(vg) - \psi(v)g = 0$ for all $g \in G$ and $v \in V$. By K-linearity, it suffices to check this for $v \in (m_j)_{1 \le j \le d}$. Writing $\psi = \sum_{i=1}^{f} b_i \psi_i$ with $\psi_i \in \mathrm{End}_F(\tilde{V})$ shows that

$$\psi \in \mathrm{End}_{KG}(V) \Leftrightarrow \sum_{i=1}^{f} b_i(\psi_i(m_j g) - \psi_i(m_j)g) = 0$$

for all $1 \le j \le d$ and all $g \in G$ and by the uniqueness above this in turn is equivalent to $\psi_i(m_j g) - \psi_i(m_j)g = 0$ for all i, j, and g. This proves that $\mathrm{End}_{KG}(V) \cong K \otimes_F \mathrm{End}_{FG}(\tilde{V})$. Now the F-dimension of $K \otimes_F \mathrm{End}_{FG}(\tilde{V})$ is equal to $f \dim_F(\mathrm{End}_{FG}(\tilde{V}))$ and the F-dimension of $\mathrm{End}_{KG}(V)$ is ef, so $e = \dim_F(\mathrm{End}_{FG}(\tilde{V}))$, as required. □

Lemma 5.6
Let F be a subfield of K such that the equivalent statements in Lemma 5.1 hold, and let $\tilde{V}$ and $B = \{b_1, \ldots, b_f\}$ be as above. Let $w \in \tilde{V}$ and $x_1, \ldots, x_k \in FG$. The set of vectors $\{\sum_{i=1}^{f} b_i w x_j \mid 1 \le j \le k\}$ is linearly independent over K if and only if it is linearly independent over F.

PROOF: For $1 \leq j \leq k$ let $c_j = \sum_{i=1}^{f} b_i w x_j$. Let $a = \sum_{i=1}^{f} b_i$ and for $1 \leq j \leq k$ let $d_j = a^{-1} c_j = w x_j \in \tilde{V}$. The d_j are linearly independent over K if and only if the c_j are linearly independent over K. Since each c_j has been multiplied by the *same* element $a^{-1} \in K$, the same statement is true over F.

The set $\{d_j : 1 \leq j \leq k\}$ is linearly independent over F if and only if $\{1 \otimes d_j : 1 \leq j \leq k\}$ is linearly independent over K. The result follows from the identification of the two sets. □

We are now in a position to attack the original problem of this section. The method for finding the matrix t as in Lemma 5.1 is an instance of the "standard basis method" which is usually used for finding homomorphisms from irreducible modules into arbitrary modules. In fact, we use the FG-module isomorphism $V \cong \bigoplus_{i=1}^{f} b_i \tilde{V}$ and then find an FG-homomorphism ι from $\tilde{V}$ to V. We will show that the K-span of the image $\iota(\tilde{V})$ is V such that every F-basis of $\tilde{V}$ is mapped to a K-basis of V by ι. The representing matrices with respect to such a basis are the same as those on $\tilde{V}$ and thus are over F.

To describe this, we first define the term "standard basis", which is most easily done by means of an algorithm. Note that this concept was described by Parker in [21, Section 6].

Definition 5.7 (Standard basis)
Let $G = \langle g_1, \ldots, g_m \rangle$ be a group, let V be a right FG-module, and let $0 \neq v \in V$. The ***standard basis*** *starting at v with respect to $(g_1, \ldots, g_m)$ is a list of vectors.*

Starting with (v), successively apply each of $g_1, \ldots, g_m$ in this order to each vector in the list, finding all m images of one vector before progressing to the next. Whenever the result is not contained in the F-linear span of the previous vectors, add it to the end of the list. This produces a basis $SB(V, v, (g_1, \ldots, g_m))$ for a nontrivial G-invariant subspace, which is V itself if V is irreducible.

We next present a theorem which is useful for isomorphism testing with an irreducible module. Although the ideas are described in [21, Section 6], we include the exact formulation and a proof, since these arguments are used in an intricate way later in the determination of the matrix t from Lemma 5.1.

Theorem 5.8 (Isomorphism test)
Let $G = \langle g_1, \ldots, g_m \rangle$ be a group, V a finite-dimensional, irreducible FG-module, $E := \mathrm{End}_{FG}(V)$ its endomorphism ring, W a finite-dimensional FG-module, and $c \in FG$ an element such that $\dim_F(\ker_V(c)) = \dim_F(E)$. Let $N := \ker_W(c) = \{w \in W : wc = 0\}$. There are two possibilities:

- *If $N = \{0\}$, then $V \ncong W$ as FG-modules.*
- *If $N \neq \{0\}$, let $0 \neq w \in N$. Then $SB(W, w, (g_1, \ldots, g_m))$ spans W if and only if $V \cong W$. If $V \cong W$ as FG-modules, then the F-linear map $\varphi : V \to W$ mapping $SB(V, v, (g_1, \ldots, g_m))$ to $SB(W, w, (g_1, \ldots, g_m))$ is an FG-isomorphism for every nonzero $v \in \ker_V(c)$.*

  Hence, if $V \cong W$, then for any nonzero $w_1, w_2 \in N$ there is an FG-automorphism of W mapping w_1 to w_2.

REMARK: This provides an efficient algorithm to test whether $V \cong W$ as FG-modules and if so to construct an explicit isomorphism, provided V is known to be irreducible and $\dim_F(E)$ is known. The algorithm finds c, computes $SB(V, v, (g_1, \ldots, g_m))$, and then computes N, looking for $0 \neq w \in N$. If an appropriate c is found and $N \neq 0$ then the algorithm computes $SB(W, w, (g_1, \ldots, g_m))$. This computation verifies whether φ is an FG-isomorphism. Thus the algorithm either computes an isomorphism φ or proves that none exists.

PROOF: $\ker_V(c)$ is E-invariant and thus a vector space over E. By assumption its E-dimension is 1. Every FG-module isomorphism between V and W maps $\ker_V(c)$ into N. If $W \cong V$, then $\dim_F(N) = \dim_F(\ker_V(c))$ and so N is a 1-dimensional vector space over $E' := \mathrm{End}_{FG}(W)$. Therefore, for all $(w, w') \in N \times N$ with $w \neq 0 \neq w'$ there is an automorphism $e' \in E'$ with $e'(w) = w'$. Thus, if we pick any $0 \neq v \in \ker_V(c)$ and any $0 \neq w \in N$, then there is an isomorphism $\varphi : V \to W$ that maps v to w. This isomorphism necessarily maps $SB(V, v, (g_1, \ldots, g_m))$ to $SB(W, w, (g_1, \ldots, g_m))$ proving our claims. □

Recall that $R_{K[G]}$, where K is a field and G is a group, denotes the number of finite field operations required to produce a single uniformly distributed random element from the K-enveloping algebra $K[G]$.

Theorem 5.9 (Writing G over a smaller field)
Let the global assumptions for this section on G apply. As before, let $e := \dim_K(\mathrm{End}_{KG}(V))$ be the degree of the splitting field. We assume that e is already computed.

There exists a $c \in K_0G$ such that $\dim_K \ker_V(c) = e$. Let $w \in \ker_V(c)$ with $w \neq 0$, let $B := SB(V, w, (g_1, \ldots, g_m))$ and let $t^{-1} \in \mathrm{GL}(d, K)$ have the vectors in B as rows.

Let F be the smallest subfield of K for which there is an $r \in \mathrm{GL}(d, K)$ such that $r^{-1}Gr \leq \mathrm{GL}(d, F)$ (see Theorem 5.4). Then $t^{-1}Gt \leq \mathrm{GL}(d, F)$ as well. That is, $t^{-1}Gt$ writes G over the smallest possible field.

Let $1 > \delta > 0$ be arbitrary. There is a Las Vegas algorithm with failure probability bounded by δ that finds c and constructs t in $O((d^3 + R_{K_0[G]}) \log \delta^{-1} + md^3)$ field operations. If the algorithm is allowed to run indefinitely, then it finishes with probability 1 and the expected number of attempts to find c is bounded above by a constant which does not depend on d or $|K|$. Each attempt needs $O(R_{K_0[G]} + d^3)$ field operations.

PROOF: Let $B = \{b_1, \ldots, b_f\}$ and $\tilde{V}$ be as in the paragraph before Lemma 5.5 and let $E := \mathrm{End}_{F[G^r]}(\mathbb{F}^{1\times d})$. Then by Lemma 5.5, the index e is equal to $[E : F]$. We know e in advance, since we know $\mathrm{End}_{KG}(V)$ by a MeatAxe run.

We apply the standard basic technique to $V \cong \bigoplus_{i=1}^{f} b_i\tilde{V}$, where $\{b_1, \ldots, b_f\}$ is an F-basis for K, as before. Note that we assume that the isomorphism exists, but do not yet have it explicitly! We attempt to compute an FG-homomorphism $\varphi : \tilde{V} \to \bigoplus_{i=1}^{f} b_i\tilde{V}$, and will either succeed or show that $f = 1$.

1. First we look for $c \in FG$ such that $\dim_F(\ker_{\tilde{V}}(c)) = \dim_F(\mathrm{End}_{FG}(\tilde{V}))$. We do not know F nor have $\tilde{V}$, but by Lemma 5.5, $e = \dim_F(\mathrm{End}_{FG}(\tilde{V}))$ and

$$f \cdot \dim_F(\ker_{\tilde{V}}(c)) = \dim_F(\ker_V(c)) = f \cdot \dim_K(\ker_V(c)).$$

Given a possible c, we can compute $\dim_K(\ker_V(c))$, and stop if this is equal to e.

To find c we repeatedly produce random elements of $K_0[G]$ in $R_{K_0[G]}$ elementary field operations, and stop if $\dim_K(\ker_V(c)) = e$. (In practice we make c by producing random K_0-linear (and hence F-linear) combinations of elements of G.) The results in [14] and [15] show that there is an upper bound b not depending on $|K|$, $|F|$ and d for the probability that a random element $c \in K_0[G] = F_0[G] \cong F[G^t] \cong E^{(d/e)\times(d/e)}$ (compare Theorem 5.4) has $\dim_K(\ker_V(c)) \neq e$. Thus $\log_b \delta^{-1}$ tries will succeed with probability at least $1 - \delta$. Note that these arguments prove that such a $c \in K_0[G]$ actually exists!

2. Assume that we have found such a $c \in FG$, given in its action on V. We compute a nonzero $w \in \ker_V(c)$. This has the form $w = \sum_{i=1}^{f} b_i w_i$ for some $w_i \in \tilde{V}$, and since $\ker_V(c) = \bigoplus_{i=1}^{f} b_i \ker_{\tilde{V}}(c)$ all of the w_i lie in $\ker_{\tilde{V}}(c)$.

We can now use the standard basis algorithm from Definition 5.7 with w in place of v, testing for K-linear independence of the resulting vectors. In fact, this will test for F-linear independence as is required — note that this works without knowing F explicitly, provided we only take linear combinations over K_0 to find $c \in FG$! We need to prove these claims.

By Theorem 5.8 and the fact that all summands $b_i\tilde{V}$ are isomorphic to $\tilde{V}$ as FG-modules, we conclude that for all w and all nonzero $v \in \ker_{\tilde{V}}(c)$, there is a unique FG-monomorphism from $\tilde{V}$ into V, mapping $SB(\tilde{V}, v, (g_1, \ldots, g_m))$ to $SB(V, w, (g_1, \ldots, g_m))$. Note that the latter is a basis for the image of this FG-homomorphism, which is an F-subspace, and that these standard bases are defined using F-linear independence.

We do yet know F, so we cannot yet test for F-linear independence. We now make some observations which allow us to apply Lemma 5.6. By the last statement of Theorem 5.8, for $1 \leq i \leq f$ there is an automorphism $\alpha_i \in \mathrm{End}_{FG}(\tilde{V})$ mapping w_i to w_1. Thus, there is an automorphism α of the FG-module $V \cong \oplus_{i=1}^{f} b_i\tilde{V}$ with $\alpha(b_i w_i) = b_i w_1$ for all i and thus $\alpha(w) = \sum_{i=1}^{f} b_i w_1$.

By Lemma 5.6, with $x_1, x_2, \ldots, x_k \in FG$, the tuple $t := (\sum_{i=1}^{f} b_i w_1 x_j)_{1\leq j\leq k}$ is F-linearly independent if and only if it is K-linearly independent. This in turn holds if and only if the tuple $(wx_j)_{1\leq j\leq k}$ is K-linearly independent, since it is mapped to t by α.

This proves our claim that we in fact compute $SB(V, w, (g_1, \ldots, g_m))$ with testing for F-linear independence.

3. By the above arguments, the result $SB(V, w, (g_1, \ldots, g_m))$ is an F-basis of an FG-submodule of V that is isomorphic to $\tilde{V}$. In particular, the representing matrices for the g_i expressed with respect to this basis contain only coefficients from F. As $SB(V, w, (g_1, \ldots, g_m))$ is K-linearly independent, it is a K-basis of V, and we have found our base change matrix t explicitly.

Finally, we determine the smallest subfield F of K containing all coefficients of $t^{-1}g_i t$ for $1 \leq i \leq m$.

Step 1 requires $O((d^3 + R_{K_0[G]}) \log \delta^{-1})$, and all other steps are $O(md^3)$, proving our claims. If the search for c is repeated indefinitely, the probability of success tends to 1 and the expected number of tries is $1/(1 - b)$. □

In summary, our Las Vegas algorithm to write G over a subfield proceeds as follows. We assume that we have already tested V for absolute irreducibility, and hence know the degree $e = \dim_K(\mathrm{End}_{KG}(V))$ of the splitting field.

1. Choose a uniformly distributed random element $c \in K_0G$ in its action on V and compute $\ker_V(c)$. Repeat this until $\dim_K(\ker_V(c)) = e$ or fail after $O(\log \delta^{-1})$ tries.

2. Take $0 \neq w \in \ker_V(c)$ and compute $B := SB(V, w, (g_1, \ldots, g_m))$ using K-linear independence.

3. Let $t^{-1} \in \mathrm{GL}(d, K)$ have the vectors in B as rows, and find the smallest subfield of K containing all entries of all $t^{-1}g_it$.

By Theorem 5.9 this algorithm either fails in step 1 with bounded probability or finds the smallest possible subfield F of K together with an explicit base change matrix t to write G over F. If G cannot be written over a smaller field then $F = K$ in step 3.

6 Restriction to a subgroup of the derived group

We now consider the case of a matrix group $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ acting absolutely irreducibly on the natural module $V = \mathbb{F}_q^d$, that cannot be written over a smaller field with trivial scalars. Our algorithm finds a reduction provided that G lies in $\mathcal{C}_3$ or $\mathcal{C}_5$, or the derived group of G is not absolutely irreducible. If none of these is the case, it might still find a reduction but might also report that G does not lie in $\mathcal{C}_3$ or $\mathcal{C}_5$ and that G' is absolutely irreducible.

In Sections 6.2 and following we refer to a normal subgroup N of G that is contained in the derived group G'. In Section 6.1 we describe a method of computing a subgroup H of such an N which can be used instead. However, note that H is produced via a Monte Carlo algorithm, so if H does not act on V in the same way as some normal subgroup N, it is essential that no incorrect answer is returned. In each of the following sections, we analyse the complexity of the algorithms used in terms of number of field operations.

Note that some of these complexity results involve a prescribed bound δ for the failure probability. If we do several such steps consecutively, we have to adjust the individual bounds because the complete procedure fails if any of the intermediate steps fails. We analyse the overall picture in Section 7.

6.1 Computing a normal subgroup of the derived group

For the first nonscalar generator g_i of G we test in $O(md^3)$ whether $[g_i, g_j]$ is scalar for $j > i$. If this holds for all j then Proposition 6.9 applies. Otherwise, we now have a nonempty set S of nonscalar commutators. We compute a subgroup H of a normal subgroup N of G that is contained in G' by the methods of Section 4. Namely, we produce a set T of $s_{\delta,d,q}$ elements of $N = \langle S^G \rangle$ in $O(R_{\langle H^G \rangle} s_{\delta,d,q})$ field operations.

By Corollary 4.4 the group $H = \langle T \rangle \leq N$ has the same submodule structure and (if N is irreducible) centraliser algebra as N with probability at least $1-\delta$. That is, we can use H instead of N in the methods described in subsequent sections. In each case we discuss the possibility that H is a proper subgroup of N and show that we do not return a wrong result. This ensures that our overall algorithm is Las Vegas rather than Monte Carlo.

6.2 A case analysis for $N \trianglelefteq G$ with $N \leq G'$

From now on we assume that H is given by s generators and is a subgroup of a normal subgroup N of G that is contained in G'. Note that $s = s_{\delta,d,q} = O(\log \delta^{-1} + d \log q)$ if we use the method from Section 6.1, but our algorithms in Sections 6.3 to 6.6 can be applied to any normal subgroup. The group H might be smaller than N, but with probability $1 - \delta$ the structure of the natural module is the same for both groups.

Since $N \trianglelefteq G$ there are only five possibilities, by Clifford's Theorem.

1. N is absolutely irreducible on V.
2. N is irreducible but not absolutely irreducible.
3. N is reducible, and V is a direct sum of more than one homogeneous components.
4. N is reducible, and V splits into a direct sum of isomorphic irreducible N-submodules of dimension greater than 1, so that in particular N is nonscalar.
5. N is reducible, and V splits into a direct sum of isomorphic 1-dimensional submodules, so that N is scalar.

We proceed differently in each of these five cases, but, assuming G is in $\mathcal{C}_3$ or $\mathcal{C}_5$ or N is not absolutely irreducible, in each case we find a reduction with probability δ of failure. By a reduction we mean a nontrivial homomorphism onto a smaller group or an isomorphism to a situation with smaller input size. To distinguish these cases, we first run the MeatAxe on H in place of N. This uses $O((R_{\mathbb{F}_q[H]} + sd^3) \log \delta^{-1})$ field operations since H is given by s generators, where δ is the upper bound for the failure probability for this step. This MeatAxe run decides whether we run the algorithms for case 1, case 2, or one of cases 3 and 4. In case 2, it returns a field generator of the endomorphism algebra. In cases 3 and 4, it returns a proper H-submodule. Note that if we use the methods in Section 6.1 to compute H, then case 5 is never found here because it is detected earlier on (see Section 6.7).

6.3 Absolutely irreducible normal subgroup

We continue to assume that N is a normal subgroup of G that is contained in G', and add the assumption that the MeatAxe has shown that H and hence N act absolutely irreducibly.

We first note the following lemma, which rules out case 1 for $\mathcal{C}_3$.

Lemma 6.1
If G lies in class $\mathcal{C}_3$ then G', and hence H and N, are not absolutely irreducible.

PROOF: Assume that there is an $\mathbb{F}_{q^e}$-vector space structure on V, such that G acts semilinearly. Then G' acts $\mathbb{F}_{q^e}$-linearly and thus $\mathrm{End}_{\mathbb{F}_q G'}(V) \neq \mathbb{F}_q$. Thus G' is not absolutely irreducible. □

We can therefore assume in this section that G lies in class $\mathcal{C}_5$.

Lemma 6.2 (Compare [11, Lemma 4.1])
Assume that G can be written over $\mathbb{F}_{q_0}$ modulo scalars in $\mathbb{F}_q$. Then $H \leq G'$ can be written over $\mathbb{F}_{q_0}$. If furthermore H acts absolutely irreducibly on the natural module and $H^t \leq \mathrm{GL}(d, q')$ for some $t \in \mathrm{GL}(d, q)$ such that $\mathbb{F}_{q'}$ is a proper subfield of $\mathbb{F}_{q_0}$, then $G^t \leq \mathrm{GL}(d, q_0) \cdot \mathbb{F}_q$.

PROOF: Multiplying each of $g, h \in G$ by a fixed scalar does not change the value of $[g, h]$, so the first claim follows.

Assume now that $G = \langle g_1, \ldots, g_m \rangle$ and that there are $\lambda_1, \ldots, \lambda_m \in \mathbb{F}_q$ and $s \in \mathrm{GL}(d, q)$ such that $\lambda_i g_i^s \in \mathrm{GL}(d, q_0)$. Let $\tilde{G} := \langle \lambda_1 g_1, \ldots, \lambda_m g_m \rangle$. Then $\tilde{G}^s \leq \mathrm{GL}(d, q_0)$. Suppose furthermore that $H^t \leq \mathrm{GL}(d, q')$ with $\mathbb{F}_{q'}$ being a subfield of $\mathbb{F}_{q_0}$. Then $H^s \leq \mathrm{GL}(d, q_0)$ since $H \leq G'$ and G' is equal to the derived group of $\tilde{G}$. But then H^t and H^s are two representations of the group H over $\mathbb{F}_{q_0}$ which are equivalent over the extension field $\mathbb{F}_q$. Thus by [7, (29.7)] they are equivalent over $\mathbb{F}_{q_0}$ and there is an element $r \in \mathrm{GL}(d, q_0)$ with $n^t = n^{sr}$ for all $n \in H$. Since H acts absolutely irreducibly, the matrix $srt^{-1} \in \mathrm{GL}(d, q)$ is scalar. Thus $\tilde{G}^t = \tilde{G}^{sr} \leq \mathrm{GL}(d, q_0)$ proving our claim. □

Theorem 6.3 (Recognition of $\mathcal{C}_5$)
Consider $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ or its projective version $\overline{G}$ and let $1 > \delta > 0$ be given. Let $H = \langle n_1, \ldots, n_s \rangle \leq N \trianglelefteq G$ with $N \leq G'$ and let H be known (by a MeatAxe run) to be absolutely irreducible. Then in $O((d^3 + R_{\mathbb{F}_p[H]}) \log \delta^{-1} + (s + m)d^3)$ field operations we can construct a homomorphism from G to $\mathrm{PGL}(d, q_0)$ for minimal q_0 or prove that G and $\overline{G}$ are not in $\mathcal{C}_5$. The algorithm returns `fail` *with probability at most δ.*

PROOF: Since H is absolutely irreducible we use the methods of Section 5 to find a matrix t such that $t^{-1}Ht \leq \mathrm{GL}(d, q')$ with δ as an upper bound on the failure probability. This automatically finds the smallest prime power q' with this property. Notice that the vector chosen in the kernel of c by the standard basis method is unique up to multiplication by elements of $\mathrm{End}_{\mathbb{F}_q H}(V) = \mathbb{F}_q$. By the second statement in Lemma 6.2 the matrix t conjugates G modulo scalars into the smallest possible field. Therefore, from this point on the algorithm is guaranteed to determine whether G lies in $\mathcal{C}_5$.

We examine $h_i := t^{-1} g_i t$ and check whether it can be written as a product of a scalar $\lambda_i \in \mathbb{F}_q$ and an element of $\mathrm{GL}(d, q_0)$ for $q' < q_0 < q$. For this, notice that if $h_i \in \lambda_i \mathrm{GL}(d, q_0)$, then the quotient between any two nonzero entries in h_i lies in $\mathbb{F}_{q_0}$. Therefore we may take λ_i to be any nonzero entry of h_i and then find the minimal field $\mathbb{F}_{q_0}$ containing all entries of h_i/λ_i. This enables us to set up a homomorphism from G to $\mathrm{PGL}(d, q_0)$ with kernel $G \cap Z(\mathrm{GL}(d, q))$, and so a reduction has been completed. For the projective group $\overline{G}$, we get a homomorphism into $\mathrm{PGL}(d, q_0)$, which could be an isomorphism. Even if this is the case, we have reduced to a smaller field.

If no smaller field is found, the procedure reports that G does not lie in $\mathcal{C}_3$ or $\mathcal{C}_5$ and that G' is absolutely irreducible. □

Note that although we work with H instead of N, since H is absolutely irreducible so is N.

6.4 Irreducible but not absolutely irreducible

We continue to assume that N is a normal subgroup of G that is contained in G'. As described in Section 6.2, we assume that the MeatAxe has proved that H acts irreducibly but not absolutely irreducibly, and so N is guaranteed to act irreducibly, but with probability at most δ the endomorphism ring $\mathrm{End}_{\mathbb{F}_q N}(V)$ may be smaller than $\mathrm{End}_{\mathbb{F}_q H}(V)$. We will deal with this possibility at the end of this section, and in general talk about N rather than H.

Proposition 6.4
If G is absolutely irreducible and $N \trianglelefteq G$ is irreducible but not absolutely irreducible then G is semilinear.

PROOF: Since N is irreducible but not absolutely irreducible, $E = \mathrm{End}_{\mathbb{F}_q N}(V) = \mathbb{F}_{q^e}$ for some $e > 1$. Let $C \in \mathrm{GL}(d, q)$ generate the multiplicative group of E.

For all $h \in N$, $g \in G$, by definition $hC = Ch$, thus $(hC)^g = (Ch)^g = h^g C^g = C^g h^g = h_1 C^g = C^g h_1$, for some $h_1 \in N$. As h varies over N the element h_1 takes every value in N, therefore $\langle C \rangle^g = \langle C \rangle$, and so $C^g = C^k$ for some k. Suppose that $C^i + C^j = C^l$, then $(C^i)^g + (C^j)^g = (C^l)^g$ so g acts as field automorphisms on $\mathbb{F}_{q^e}$ and thus G is semilinear. □

Theorem 6.5 (Recognition of $\mathcal{C}_3$)
Let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ or its projective version. Let $N = \langle n_1, \ldots, n_s \rangle \trianglelefteq G$ be known to be irreducible but not absolutely irreducible. In deterministic $O(d^4 \log q + md^3)$ field operations we can construct two homomorphisms, one to the cyclic group of order e for some divisor e of d and a second from the kernel of the first to $\mathrm{GL}(d/e, q^e)$ or $\mathrm{PGL}(d/e, q^e)$.

PROOF: When N is irreducible but not absolutely irreducible, the MeatAxe returns a generator C of the field $\mathbb{F}_{q^e} = \mathrm{End}_{\mathbb{F}_q N}(V)$ realised as a matrix in $\mathrm{GL}(d, q)$ together with e. Note that $e \leq d$. The matrix C need not generate the multiplicative group of $\mathbb{F}_{q^e}$, but its powers $C^0, C^1, \ldots, C^{e-1}$ are $\mathbb{F}_q$-linearly independent.

As shown in Proposition 6.4, the group G acts by conjugation as field automorphisms on $\mathbb{F}_{q^e} = \mathrm{End}_{\mathbb{F}_q N}(V)$ and thus on the group $\langle C \rangle$. We can immediately read off this action using $O(md^3)$ field operations by computing the matrices $C, C^q, C^{q^2}, \ldots, C^{q^{e-1}}$, conjugating C with the generators of G and looking up the result. Computing these matrices requires at most $O(d^4 \log q)$ field operations and space for $O(d)$ matrices, since $e \leq d$. Computing this action provides a homomorphism from G to the cyclic group of order e, because the above mentioned matrices are the possible images of C under automorphisms of $\mathbb{F}_{q^e}$. Note that if every C^{g_i} is equal to one of the C^{q^j}, it follows that $\{C^{q^j} \mid 0 \leq j < e\}$ is a union of orbits of the conjugation action of G on $\mathrm{GL}(d, q)$. Thus in this case we have computationally proved that we have found a homomorphism of G into the cyclic group of order e.

In addition, C gives an explicit $\mathbb{F}_{q^e}$-vector space structure on V. To get the $\mathbb{F}_{q^e}$-span of a vector $v \in V$ we compute $v, vC, vC^2, \ldots, vC^{e-1}$. In this way we can perform a spinning algorithm for V as an $\mathbb{F}_{q^e}$-vector space. All computations are with vectors over $\mathbb{F}_q$ but whenever we produce a new vector v that does not lie in the $\mathbb{F}_q$-span of what we already have, we not only add v but also $vC, vC^2, \ldots, vC^{e-1}$ by repeatedly multiplying with C. This spinning algorithm gives us a base change to an $\mathbb{F}_{q^e}$-adapted basis. It needs at most $O(md^3)$ field operations.

The kernel of the action as field automorphisms acts $\mathbb{F}_{q^e}$-linearly on the original space and we read off this action using the above base change to the $\mathbb{F}_{q^e}$-adapted basis. This therefore also leads to a reduction for the kernel by reducing the input size to $(d/e) \times (d/e)$-matrices over $\mathbb{F}_{q^e}$.

Altogether, we have found a significant reduction using $O(d^3(d \log q + m))$ field operations and memory for $O(d)$ matrices.

Note that since scalars from $\mathbb{F}_q$ do not alter the action of elements of G as field automorphisms on $\mathbb{F}_{q^e}$, the same procedure works for the projective case $\overline{G}$. The homomorphism is

the same as in the matrix case, the kernel is a subgroup of $\mathrm{PGL}(d,q)$ and we construct a map from the kernel into $\mathrm{PGL}(d/e, q^e)$ by writing the matrices over the bigger field. This map in turn has a kernel, since we divide out more scalars. However, this second kernel only contains $\mathbb{F}_{q^e}$-scalars modulo $\mathbb{F}_q$-scalars, which can be handled easily. Thus, this case can be handled in the projective situation. □

We finish with a discussion of the possibility that $\mathbb{F}_{q^e} = \mathrm{End}_{\mathbb{F}_q H}(V) \neq \mathrm{End}_{\mathbb{F}_q N}(V)$, which happens with probability bounded by δ. If $\mathrm{End}_{\mathbb{F}_q H}(V) \neq \mathrm{End}_{\mathbb{F}_q N}(V)$ then the elements of G need not act as field automorphisms on $\mathbb{F}_{q^e}$, and indeed G need not even be semilinear, which we notice during the above computation. For the claim in the Main Theorem, it suffices to return `fail` if this occurs.

However, we can do better than this. If G is contained in $\mathcal{C}_3$ then the generators of G *will* act as field automorphisms on some subfield of $\mathbb{F}_{q^e}$ that properly contains $\mathbb{F}_q$. Thus, since $e \leq d$ is a small number, we test for each divisor i of e whether G acts as field automorphisms on $\mathbb{F}_{q^i}$.

To find a generating element C' for the field $\mathbb{F}_{q^i}$ we proceed as follows. A high percentage of elements in $\mathbb{F}_{q^e}$ have order $q^e - 1$, so picking a random linear combination of $1, C, C^2, ..., C^{e-1}$ has good chances to find an element $\tilde{C}$ of order $q^e - 1$. This can be done using $O(ed^2)$ elementary field operations using e random integers in the range $0, \ldots, q-1$. The element $C' := \tilde{C}^{(q^e-1)/(q^i-1)}$ is then contained in $\mathbb{F}_{q^i}$ and generates this field with even higher probability. We can now compute

$$C', C'^q, C'^{q^2}, \ldots, C'^{q^{i-1}}$$

using $O(d^4 \log q)$ elementary field operations. If C' is in fact contained in a proper subfield of $\mathbb{F}_{q^i}$ we notice this now since the above list will have repetitions. We can then either give up and try the next divisor of e or try another random element $\tilde{C}$. If C' is a field generator of $\mathbb{F}_{q^i}$, we check whether C'^{g_j} is contained in this list for all generators g_j with $1 \leq j \leq m$. If so, we have found an action of G as field automorphisms of $\mathbb{F}_{q^i}$. If not, we try the next divisor i of e.

If no divisor works then the algorithm reports `fail`, that G is not in $\mathcal{C}_3$ and that N and the derived group are absolutely irreducible. Note that if N is not absolutely irreducible then the algorithm is guaranteed to find that G is in $\mathcal{C}_3$ at this point, therefore if failure is reported the algorithm adds generators to H until it is absolutely irreducible, and then returns to the test of Section 6.3.

6.5 More than one homogeneous component

We continue to assume that N is a normal subgroup of G that is contained in G'. As described in Section 6.2 we assume that the MeatAxe has proved that H acts reducibly by finding an explicit proper nontrivial submodule V' of the natural module.

First we prove a lemma which will eventually be used to find an irreducible H-submodule that with probability $1 - \delta$ is an N-submodule.

Lemma 6.6 (Finding an irreducible module)
Let $N = \langle n_1, \ldots, n_s \rangle \leq \mathrm{GL}(d,q)$ act reducibly on the natural module V and let $1 > \delta > 0$ be given. Given a submodule $V' < V$, an irreducible N-subfactor can be found in Las Vegas $O((R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d))$ field operations, with probability of failure at most δ.

PROOF: We repeatedly use the MeatAxe to find an irreducible subfactor of $V|_{\mathbb{F}_q N}$. Initially, we have a submodule $V' < V$. We run the MeatAxe either on V/V' or on V', whichever has the smaller dimension. If we find a proper submodule, we repeat the same technique. Since we halve the dimension in each step, this terminates after at most $\log d$ runs of the MeatAxe using at most $O((R_{\mathbb{F}_q[N]} + sd^3)2^{-3i} \log \delta'^{-1})$ field operations in step i, where δ' is an upper bound for the failure probability in each step. To bound the overall failure probability of this whole procedure by δ, we define $\delta' := \delta / \log d$. Since $\sum_{i=1}^{\infty} 2^{-3i} < 1$, the overall cost for finding an irreducible subfactor is $O((R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d))$. □

Theorem 6.7 (Construction of a block action)
Let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ or its projective version, and let $1 > \delta > 0$ be given. Let $N = \langle n_1, \ldots, n_s \rangle \trianglelefteq G$ be known (by a MeatAxe run) to be reducible. In Las Vegas $O((R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d))$ field operations we can either construct a homomorphism from G to a permutation group with kernel the pointwise stabiliser of the set of homogeneous components of $V|_{\mathbb{F}_q N}$ or prove that $V|_{\mathbb{F}_q N}$ has a single homogeneous component. The probability of failure is bounded from above by δ.

PROOF: By assumption $V|_{\mathbb{F}_q N}$, as an $\mathbb{F}_q N$-module, is a direct sum of homogeneous components $C_1, \ldots, C_k$ with $k \geq 1$. The C_i form a block system exhibiting an imprimitive action of G and N is a normal subgroup of the kernel of the action on blocks. We only have to find the action on this block system to find a reduction.

By Lemma 6.6 we can find an irreducible N-subfactor in $O((R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d))$ field operations with probability of failure δ. This subfactor is an irreducible module $\tilde{S}$ and we can now apply the isomorphism testing procedure described in Theorem 5.8 once to give a homomorphism of $\tilde{S}$ into $V|_{\mathbb{F}_q N}$ and thus an irreducible submodule S with $O(sd^3)$ field operations. Note that when we proved the final subfactor in the procedure described in Lemma 6.6 to be irreducible we constructed the algebra word c that is needed for isomorphism testing, namely a word describing an algebra element with nullity the dimension of the centraliser of N.

Such an irreducible module S is all we need to run the MINBLOCKS procedure described in [13] which needs $O(sd^3)$ field operations to compute the block system or reports that there is none. If the latter occurs, then there is a single homogeneous constituent and we apply the algorithms of Section 6.6. Otherwise this provides a nontrivial homomorphism onto a permutation group and thus a reduction. The overall complexity is $O((R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d))$.

Since $\mathbb{F}_q$-scalars act trivially on the set of homogeneous components, the homomorphism onto the permutation group has all scalars in its kernel. Therefore we can use the same homomorphism for the projective situation with $\overline{G}$. Thus, this case can be handled in the projective situation. □

Of course in practice we work with a subgroup H of N. If H has a submodule that is *not* an N-submodule then it is possible that we will not be able to find a homomorphism from the irreducible H-subfactor to V. In this case, all that is required for the Main Theorem is that the algorithm reports `fail`: note that this occurs with probability at most δ.

However, it is possible to rerun the algorithm starting at Section 6.2 with a new version of H that has submodule structure closer to that of N. To see this, note that the subfactor is described

by two H-submodules of V, at least one of which is not preserved by N. Therefore a simple argument shows that at least half of the elements of N must fail to fix at least one of the two H-submodules. Thus we add a new generator to H and return to the MeatAxe run of Section 6.2 to determine whether the new H is (absolutely) irreducible.

6.6 Isomorphic irreducible submodules of dimension at least 2

We continue to assume that N is a normal subgroup of G that is contained in G'. As described in Section 6.2 we assume that the MeatAxe has proved that H acts reducibly by finding an explicit proper nontrivial submodule V' of the natural module.

As described in Theorem 6.7 we first find an irreducible H-submodule S and run the MIN-BLOCKS procedure. If this fails to find a block system, then (assuming H has the same submodule structure as N) there is only one homogeneous component, corresponding to S.

Theorem 6.8 (Reduction for single homogeneous component)
Let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ or its projective version be absolutely irreducible, and let $1 > \delta > 0$ be given. Let $N = \langle n_1, \ldots, n_s \rangle \trianglelefteq G$ be reducible, with a single homogeneous component of dimension $n > 1$, and let an irreducible N-submodule S be given. In deterministic $O((s+m)d^3)$ field operations we can construct a proper nontrivial homomorphism from G into $\mathrm{PGL}(d/n, q^e)$ for e the $\mathbb{F}_q$-dimension of $\mathrm{End}_{\mathbb{F}_q N}(S)$.

PROOF: We first find an explicit decomposition of $V|_{\mathbb{F}_q N}$ as a direct sum of copies of S. This can be done using a variant of the isomorphism testing procedure described in Theorem 5.8 to compute a basis of the space of all homomorphisms of S into $V|_{\mathbb{F}_q N}$. Namely, we compute the action on V of the algebra word $c \in \mathbb{F}_q N$ that proved that S is simple and determine its kernel K. Since $V|_{\mathbb{F}_q N}$ is isomorphic to a direct sum of copies of S, we can choose an arbitrary nonzero vector from K, compute the standard basis with respect to the generators of N starting at that vector and thereby find a summand S_1 of $V|_{\mathbb{F}_q N}$ together with an explicit isomorphism of S to S_1. By choosing further vectors from K that are not contained in the direct sum of previous copies of S and repeating this procedure, we inductively get an explicit direct sum decomposition of $V|_{\mathbb{F}_q N}$ into summands that are all isomorphic to S. This automatically leads to a base change such that every element of N is represented by a block diagonal matrix in which all diagonal blocks are identical of size $n := \dim_{\mathbb{F}_q}(S)$ in $O(sd^3)$ field operations.

As $N \trianglelefteq G$, for all $h \in N$ and $g \in G$, the product $g^{-1}hg \in N$ and thus $g^{-1}hg$ is also a block diagonal matrix in which all $n \times n$-blocks along the diagonal are identical. Fixing g, we conclude that $g \cdot (g^{-1}hg) = hg$ for all $h \in N$. If we now cut g into $n \times n$-blocks, we get:

$$g \cdot h^g = \left[\begin{array}{c|c|c|c} g_{1,1} & g_{1,2} & \cdots & g_{1,d/n} \\ \hline g_{2,1} & g_{2,2} & \cdots & g_{2,d/n} \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline g_{d/n,1} & g_{d/n,2} & \cdots & g_{d/n,d/n} \end{array}\right] \cdot \left[\begin{array}{c|c|c|c} D^{g^{-1}}(h) & 0 & \cdots & 0 \\ \hline 0 & D^{g^{-1}}(h) & \cdots & 0 \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline 0 & 0 & \cdots & D^{g^{-1}}(h) \end{array}\right]$$

$$= \left[\begin{array}{c|c|c|c} D(h) & 0 & \cdots & 0 \\ \hline 0 & D(h) & \cdots & 0 \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline 0 & 0 & \cdots & D(h) \end{array}\right] \cdot \left[\begin{array}{c|c|c|c} g_{1,1} & g_{1,2} & \cdots & g_{1,d/n} \\ \hline g_{2,1} & g_{2,2} & \cdots & g_{2,d/n} \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline g_{d/n,1} & g_{d/n,2} & \cdots & g_{d/n,d/n} \end{array}\right] = hg$$

where the $g_{i,j}$ are $n \times n$-matrices, $D(h)$ is a matrix representing h on the module S and $D^{g^{-1}}(h) = D(g^{-1}hg)$ is the same representation twisted by the element g^{-1}. By the block diagonal structure of the matrices in N we get $g_{i,j} \cdot D^{g^{-1}}(h) = D(h) \cdot g_{i,j}$ for all i and j and all $h \in N$.

But by hypothesis, the matrix representations D and $D^{g^{-1}}$ of N are isomorphic. Thus there is a nonzero matrix $T \in \mathbb{F}_q^{n \times n}$ with $T \cdot D^{g^{-1}}(h) = D(h) \cdot T$ for all $h \in N$. By Schur's lemma and since the representation D is irreducible, the matrix T is invertible and unique up to left multiplication by an element of $C_{\mathrm{GL}(n,q)}(D(N))$, which is isomorphic as a group to the group of units of the extension field $\mathrm{End}_{\mathbb{F}_q N}(S) \cong \mathbb{F}_{q^e}$.

This shows that for every pair $(i,j) \in \{1,\ldots,d/n\} \times \{1,\ldots,d/n\}$ there is a unique element $e_{i,j} \in \mathrm{End}_{\mathbb{F}_q N}(S)$ (possibly 0) with $g_{i,j} = e_{i,j} \cdot T$. Thus we have shown that with respect to the above choice of basis, every element g is equal to a Kronecker product of some matrix in $U \in \mathbb{F}_{q^e}^{d/n \times d/n}$ with a matrix $T \in \mathbb{F}_q^{n \times n}$. Since g is invertible both U and T are invertible.

This provides an explicit embedding of $\mathbb{F}_q^d$ into a tensor product $\mathbb{F}_{q^e}^{d/n} \otimes_{\mathbb{F}_q} \mathbb{F}_q^n$, where one factor can be over an extension field if the $\mathbb{F}_q N$-module S is not absolutely irreducible. This embedding can be computed explicitly because the above base change is constructive. Using another $O(md^3)$ field operations we compute the generators of G after the base change from which we can read off the tensor decomposition.

Thus we get a nontrivial homomorphism of G into $\mathrm{PGL}(d/n, q^e)$ with N lying in the kernel which is a significant reduction. The kernel of this homomorphism can immediately be reduced further since its elements are block diagonal matrices with identical $n \times n$-diagonal blocks.

The projective situation can be handled identically, by viewing the kernel as a projective group. □

If H is a proper subgroup of N, then our algorithm can fail in two ways, both of which must be recognised for the algorithm to be Las Vegas. Firstly, $V|_{\mathbb{F}_q H}$ might not be isomorphic to a direct sum of copies of the irreducible H-module S. In this case there are not enough homomorphisms from S into the socle of $V|_{\mathbb{F}_q H}$ to span the whole of V. Secondly, even if $V|_{\mathbb{F}_q H}$ is a direct sum of copies of S, the generators of G might not be Kronecker products after a corresponding base change, which we detect during the setup of the homomorphism. In both cases, the error is detected and the algorithm reports `fail`. However, by Corollary 4.4 this happens with probability at most δ.

6.7 Normal subgroup is scalar

The remaining case is that the restriction of the natural module to N has only one homogeneous component and all irreducible N-constituents are one-dimensional, so that N consists of scalars. The algorithms in this section are applicable to any group G with a fixed noncentral generator g_i such that $[g_i, g_j]$ is scalar for all generators g_j.

We start with a proposition giving a homomorphism into the multiplicative group of the field that need not necessarily correspond to an imprimitive decomposition of the natural module.

Proposition 6.9 (Scalar homomorphism)
Let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ or its projective version $\overline{G}$ be an absolutely irreducible group such that the commutator of a nonscalar generator g_i with all other generators is known to be scalar. Then we can construct a nontrivial homomorphism from G into the multiplicative group of $\mathbb{F}_q$ at no further cost.

PROOF: We are given a nonscalar generator g_i, such that all commutators of it with all other generators are scalar matrices. Thus g_i is central in G modulo scalars, thus the commutators of g_i with all elements of G are scalar. Therefore, the map $\psi_{g_i} : G \to \mathbb{F}_q, g \mapsto [g, g_i]$ is a group homomorphism into the scalar matrices. This is proved exactly as Lemma 3.2.5.

The kernel of ψ_{g_i} is $C_G(g_i)$. Since g_i is noncentral, ψ_{g_i} is nontrivial. Multiplying generators by scalars does not change commutators, so these algorithms will also work in the projective case. □

Since $g_i \in C_G(g_i)$, the kernel is not an absolutely irreducible group, and may not even be irreducible. If G' is known to be scalar then the derived group of the kernel $C_G(g_i)$ is also central, and hence a hint can be passed to the kernel to return to the techniques of this section once an absolutely irreducible representation has been found.

Finally we give a deterministic decomposition algorithm for groups with scalar derived group that are not r-groups. We can apply this algorithm if G has a very small number of nonscalar generators, so that all commutators of generators can be cheaply calculated — in this case the m^2 vanishes from the complexity. This algorithm can easily be modified to decompose any black box group with order oracle that is known to be nilpotent and not a p-group. The assumption that the prime factors of $q^i - 1$ are known for $i \leq d$ is reasonable in practice, and is relied upon for many other algorithms: see [3] for details of currently maintained lists of such factors.

Lemma 6.10
Let $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ be an absolutely irreducible group whose derived group consists only of scalars. Suppose that the order of G is divisible by k primes for some $k > 1$, and that the prime divisors of $q^i - 1$ are known for $1 \leq i \leq d$. Then $k < \log(q-1)$ and in $O(m^2 d^3 \log q \log(d \log q))$ field operations we can compute a homomorphism from G whose kernel and image have order divisible by $\lfloor k/2 \rfloor$ and $\lceil (k+1)/2 \rceil$ primes respectively. Both the kernel and image have at most m generators.

PROOF: By Proposition 3.3.3 the order of G is a divisor of $o := (q-1)^{m+1}$, which is divisible by less than $\log(q-1)$ distinct primes.

The group G is a direct product of its Sylow subgroups by Lemma 3.2.2. We compute the order o_i of g_i for $1 \leq i \leq m$ in $O(md^3 \log q \log(d \log q))$ field operations (see [4]), and find a set of primes $\{p_1, \ldots, p_k\}$ such that each o_i is a product of powers of these primes. For $1 \leq i \leq k$ we find the highest exponent α_i such that $p_i^{\alpha_i}$ divides o.

Let $a := \lfloor k/2 \rfloor$ and define $r = p_1^{\alpha_1} \cdots p_a^{\alpha_a}$ and $r' = p_{a+1}^{\alpha_{a+1}} \cdots p_k^{\alpha_k}$. First run the extended Euclidean algorithm to find s and s' such that $1 = sr + s'r'$, then let $N = s'r'$ and $M = sr$.

Clearly for all $g \in G$ the order of g^N divides r whilst $|g^M|$ divides r'. Therefore for all $g \in G$ the only way to write g as a product of an element of order dividing r and an element of order dividing r' is $g = g^N g^M$. Since G is nilpotent, the map $x \mapsto x^N$ is a homomorphism from G to $\mathrm{Syl}(G, p_{a+1}) \times \cdots \times \mathrm{Syl}(G, p_k)$ with kernel $\mathrm{Syl}(G, p_1) \times \cdots \times \mathrm{Syl}(G, p_a)$.

Notice that $p_i^{\alpha_i}$ divides o for all i so $N < o$, and hence we can raise each generator to the power N in $O(m^2 d^3 \log q)$ field operations to get generators for the image. Some of them could be trivial, so we get at most m generators. Multiplying each generator by the inverse of its image in $O(md^3)$ field operations will produce at most m generators for the kernel, again ignoring trivial ones. □

7 Complexity summary

In this section we summarise our complexity results, mainly for the sake of a good overview, but also to explicitly give our assumptions.

We begin by describing the complexity of a "MeatAxe run". Although this result is well-known we want to say exactly what results underlie our complexity analysis.

Recall that we let $R_{\langle H^L \rangle}$ denote the cost of producing an independent, uniformly-distributed, random element of the normal closure of a group H in a group L, and we let $R_{K[L]}$, where K is a finite field and L is a group, denote the cost of producing an independent, uniformly-distributed, random element of $K[L]$.

In practice, by using Product Replacement, Rattle, and recent work by Dixon [5, 8, 17, 18] for groups and normal closure, and by taking random linear combinations of random products of generators for algebras, each of these costs is $O(d^3)$, at least after an initialisation phase. However, these methods are not proven to produce independent, uniformly-distributed random elements in general.

Lemma 7.1 (MeatAxe)
Let F be a finite field, $\mathcal{A}$ a finite-dimensional F-algebra, V an $\mathcal{A}$-module of F-dimension d, given by the action of m generators of $\mathcal{A}$ as matrices in $F^{d \times d}$, and let $0 < \delta < 1$ be given. There is a Las Vegas algorithm with failure probability less than δ that determines whether V is irreducible in $O((R_{\mathcal{A}} + md^3) \log \delta^{-1})$ elementary field operations. In the case of success the result is either a proper nontrivial submodule or the answer "irreducible" together with a field generator of the endomorphism ring $\mathrm{End}_{\mathcal{A}}(V)$. Running the algorithm until success gives an algorithm which terminates with probability 1, in which a step needs $O(R_{\mathcal{A}} + md^3)$ field operations and the expected value of the number of such steps is bounded by a constant not depending on $|F|$, d and m.

PROOF: All of this is proved in [14, 15], since it is shown that a certain percentage (not depending on $|F|$ or d or m) of all matrix algebra elements are usable to reach a decision and all operations in one step are $O(R_{\mathcal{A}} + md^3)$. □

We now summarise our complexity results, all given in terms of the number of field operations. The whole procedure contains several subalgorithms of Las Vegas type, namely in steps 2, 5, 6, 7 and 8. However, at most 4 of them are possibly executed sequentially (namely in the

execution path with steps 1, 2, 3, 4, 5, 8, 9). Thus if we prescribe a failure probability of $\delta/4$ in each Las Vegas step, we get a Las Vegas algorithm with overall failure probability bounded from above by δ. Notice that the factor of 4 does not affect the "big O" complexity. We follow the numbering in our summary of the complete procedure in Section 2:

1. Assume $G = \langle g_1, \ldots, g_m \rangle \leq \mathrm{GL}(d, q)$ acting irreducibly, $E = \mathrm{End}_{\mathbb{F}_q G}(\mathbb{F}^d) = \mathbb{F}_{q^e}$ and field generator $C \in \mathrm{GL}(d, q)$ of E^* are known. If $e > 1$ find an explicit base change in $O(md^3)$ field operations.

2. Try to write G over a subfield with $\beta_i = 1$ for $1 \leq i \leq m$ in $O((d^3 + R_{\mathbb{F}_p[G]}) \log \delta^{-1} + md^3)$ field operations (see Theorem 5.9).

3. Immediately get a reduction in either the non-absolutely irreducible or the subfield case.

4. Test in $O(md^3)$ whether all commutators of the first nonscalar generator g_i with other generators are scalar. If so, jump to step 10.

5. Compute $s = s_{\delta,d,q} = O(d \log q + \log \delta^{-1})$ generators for $H \leq N \trianglelefteq G$ with N contained in G' in $O(R_{\langle H^G \rangle} s_{\delta,d,q})$ field operations (see Section 6.1).

   Run the MeatAxe to distinguish cases for N in $O((R_{\mathbb{F}_q[N]} + sd^3) \log \delta^{-1})$ field operations (see Section 6.2).

6. If N is absolutely irreducible, check whether G is in $\mathcal{C}_5$ in $O((d^3 + R_{\mathbb{F}_p[N]}) \log \delta^{-1} + (s + m)d^3)$ field operations (see Section 6.3).

7. If N is irreducible but not absolutely irreducible, check whether G is in $\mathcal{C}_3$ in $O(d^4 \log q + md^3)$ field operations as in Section 6.4.

8. If N is reducible, look for more than one homogeneous component and if so find an imprimitive decomposition of G in $O((R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d))$ field operations as in Section 6.5.

9. If N is reducible with a single homogeneous component with irreducible N-submodules dimension greater than 1, find a tensor decomposition of G in $O((s + m)d^3)$ field operations as in Section 6.6.

10. If one nonscalar generator has only scalar commutators with other generators, we have already constructed a nontrivial homomorphism from G to $\mathbb{F}_q^\times$ as in Section 6.7.

The above algorithm can stop after either of steps 3, 6, 7, 8, 9 or 10. We summarise the complexity statements for each of the possible paths through the above steps in Table 1.

The worst cases are the last two, where the overall complexity is bounded from above by

$$O(md^3 + R_{\mathbb{F}_p[G]} \log \delta^{-1} + (R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d) + sR_{\langle H^G \rangle})$$

where $s = O(\log \delta^{-1} + d \log q)$. Let $R_{\mathcal{A}} = \max\{R_{\mathbb{F}_p[G]}, R_{\mathbb{F}_q[N]}\}$, then this becomes:

$$O(d^3(m + (d \log q + \log \delta^{-1}) \log(\delta^{-1} \log d)) + R_{\mathcal{A}} \log(\delta^{-1} \log d) + (d \log q + \log \delta^{-1}) R_{\langle H^G \rangle}).$$

Path	Cost in finite field operations, where $s = O(\log \delta^{-1} + d \log q)$
1, 2, 3	$O((d^3 + R_{\mathbb{F}_p[G]}) \log \delta^{-1} + md^3)$
1, 2, 3, 4, 10	$O((d^3 + R_{\mathbb{F}_p[G]}) \log \delta^{-1} + md^3)$
1, 2, 3, 4, 5, 6	$O(md^3 + (sd^3 + R_{\mathbb{F}_p[G]} + R_{\mathbb{F}_p[N]} + R_{\mathbb{F}_q[N]}) \log \delta^{-1} + sR_{\langle H^G \rangle})$
1, 2, 3, 4, 5, 7	$O(md^3 + d^4 \log q + (sd^3 + R_{\mathbb{F}_p[G]} + R_{\mathbb{F}_q[N]}) \log \delta^{-1} + sR_{\langle H^G \rangle})$
1, 2, 3, 4, 5, 8	$O(md^3 + R_{\mathbb{F}_p[G]} \log \delta^{-1} + (R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d) + sR_{\langle H^G \rangle})$
1, 2, 3, 4, 5, 8, 9	$O(md^3 + R_{\mathbb{F}_p[G]} \log \delta^{-1} + (R_{\mathbb{F}_q[N]} + sd^3) \log(\delta^{-1} \log d) + sR_{\langle H^G \rangle})$

Table 1: Complexity of algorithm for different cases

Fixing $\delta > 0$ this simplifies to

$$O(d^3(m + d \log(\log d) \log q) + R_{\mathcal{A}} \log(\log d) + R_{\langle H^G \rangle} d \log q).$$

8 Implementation and performance

All of our algorithms have been implemented in the forthcoming GAP package recog for constructive group recognition. In general we make $\lfloor (d \log q)/20 \rfloor$ random elements when producing generators for N (see Section 6.1), but always at least 5 and at most 40, which seems to work well in practice. The division by 20 indicates that our analysis of the generation of a sufficiently large subgroup of N underestimates the probability of success in most cases.

In Table 2 we give timing results. All experiments have been done on a machine with an Intel Core2 Quad CPU Q6600 running at 2.40GHz with 8 GB of main memory using the development version of GAP and the recog package. All times are in milliseconds and were averaged over several runs. Note that due to randomisation the runtimes can vary significantly between runs.

Columns "d" and "q" give the matrix dimension and field size of the input matrix group. Column "m" states the number of generators. Column "Group" contains a structural description of the input group. The notation "$S.G$" indicates that the input group is a central extension of a group G by all scalars of $\mathbb{F}_q$. The notation "$G.A$" indicates that the input group is a group G extended by a group A of field automorphisms of a centralising matrix.

Column "Case" describes the type of reduction found. Here, "Subfield" means that an immediate base change to write the group over a smaller field was found. "NotAbsIrr" means that the input group did not act absolutely irreducibly and was written over a larger field with smaller dimension. "$\mathcal{C}_5$" means that a subgroup $H < N$ was computed which acted irreducibly and then a base change was found to write the (projective) group over a smaller field. "$\mathcal{C}_3$" means that a subgroup $H < N$ was computed which acted irreducibly (but not absolutely irreducibly) and an action as field automorphisms was found. This automatically gives the information required to write the kernel in a smaller dimension over the appropriate extension field. "Components" means that an imprimitive action was found. "Tensor" means that a tensor decomposition was found. "Scalar" means that the group was reduced using commutators with a single noncentral element.

No	Group	d	q	m	Case	d'	q'	Time	Total
1	M_{11}	10	3^5	2	Subfield	10	3	9	70
2	$S.M_{11}$	10	3^5	2	$\mathcal{C}_5$	10	3	25	110
3	J_2	13	3^{10}	2	Subfield	13	3^2	25	2791
4	$S.J_2$	13	3^{10}	2	$\mathcal{C}_5$	13	3^2	184	1687
5	Co_3	22	2^{16}	2	Subfield	22	2	103	3509
6	$S.Co_3$	22	2^{16}	2	$\mathcal{C}_5$	22	2	788	4173
7	$2.\mathrm{A}_8$	24	5^6	2	Subfield	24	5^2	119	1147
8	$S.(2.\mathrm{A}_8)$	24	5^6	2	$\mathcal{C}_5$	24	5^2	954	1711
9	$\mathrm{GL}_{90}(7)$	90	7^5	2	Subfield	90	7	8849	—
10	$S.\mathrm{GL}_{90}(7)$	90	7^5	2	$\mathcal{C}_5$	90	7	51133	—
11	J_2	28	2	2	NotAbsIrr	14	2^2	52	1255
12	$J_2.2$	28	2	2	$\mathcal{C}_3$	14	2^2	27	888
13	M_{22}	90	3	2	NotAbsIrr	45	3^2	622	10262
14	$(S.M_{24}).A$	69	5	2	$\mathcal{C}_3$	23	5^3	594	3096
15	3.3^9	81	19	256	Scalar	81	19	5140	50085
16	$\mathrm{S}_{14} \wr C_5$	320	2	3	Components	320	2	2117	106100
17	$3^{1+4} \otimes HS$	189	25	8	Tensor	9	25	24053	150226

Table 2: Timing results for a few example groups

Columns "d'" and "q'" contain the dimension and field size after the reduction. The "Time" column indicates the time needed for the first reduction. The "Total" column contains the runtime to build a complete composition tree for the given group. This value is occasionally omitted, which indicates that not enough leaf methods are implemented yet to recognise this group fully.

To produce the examples, we first changed the field by embedding or blowing up and then conjugated by a random element in the general linear group over the new field.

To make example 14 we took the Mathieu group M_{24} represented in $\mathrm{GL}_{23}(5)$, multiplied the generators by scalars in $\mathbb{F}_{5^3}$, blew everything up into $\mathrm{GL}_{69}(5)$ and added a new generator that acts as a field automorphism of $\mathbb{F}_{5^3}$ when conjugating the centralising matrix. In example 14 the algorithm then writes the kernel as a subgroup of $\mathrm{GL}_{23}(5^3)$ and finally recognises the $\mathcal{C}_5$ case and goes back to $\mathrm{GL}_{23}(5)$.

Example 15 is an absolutely irreducible 3-group in $\mathrm{GL}_{81}(19)$ such that all commutators are scalar matrices. After one reduction by the commutator action the kernel becomes reducible.

Example 16 is a wreath product of the symmetric group S_{14} with the cyclic group of order 5. We started with an absolutely irreducible 64-dimensional representation of S_{14} over $\mathbb{F}_2$ and made a 320-dimensional absolutely irreducible representation for $\mathrm{S}_{14} \wr C_5$ over $\mathbb{F}_2$. Our algorithm computes the action on the five homogeneous components, then tells the kernel node that the group is reducible and in block form so that a MeatAxe call is not necessary.

Example 17 is a central product of an extraspecial 3-group of order 243 in its irreducible representation of dimension 9 over $\mathbb{F}_{25}$ with the sporadic finite simple group HS in an irreducible representation of dimension 21 over $\mathbb{F}_{25}$. The latter representation came from one over

$\mathbb{F}_5$ where the representing matrices of the group generators were multiplied by elements from $\mathbb{F}_{25}$. The extraspecial factor vanishes when computing a subgroup of the derived group, since it is one example of our characterisation in Section 3, exhibiting the tensor decomposition. In subsequent nodes the extraspecial factor is taken apart using commutators as in Section 6.7.

Acknowledgements: The first author is grateful to the RWTH in Aachen for their kind hospitality and to the Humboldt Foundation for support while parts of this paper were being written. The second and third authors are grateful to Ákos Seress for his generous hospitality while parts of this paper were being planned. We thank the anonymous referee for a very detailed reading and helpful comments to improve the exposition of this paper.

References

[1] M. Aschbacher, On the maximal subgroups of the finite classical groups, *Invent. Math.* **76** (1984) 469–514.

[2] L. Babai, G. Cooperman, L. Finkelstein, E. Luks, A. Seress, Fast Monte Carlo algorithms for permutation groups, *J. Comput. System Sci.* **50** (1995) 296–308.

[3] J. Brillhart, D. H. Lehmer, J. L. Selfridge, B. Tuckerman, S. S. Wagstaff, Jr., *Factorizations of $b^n \pm 1$*, Vol. 22 of *Contemporary Mathematics*, Amer. Math. Soc., Providence, RI, third edition, 2002. `http://www.cerias.purdue.edu/homes/ssw/cun/index.html`

[4] F. Celler, C. R. Leedham-Green, *Calculating the order of an invertible matrix*, in: *Groups and computation, II (New Brunswick, NJ, 1995)*, Amer. Math. Soc., Providence, RI, 1997, pp.55–60.

[5] F. Celler, C. R. Leedham-Green, S. H. Murray, A. C. Niemeyer, E. A. O'Brien, Generating random elements of a finite group, *Comm. Algebra* **23** (1995) 4931–4948.

[6] C. W. Curtis, I. Reiner, *Methods of Representation Theory*, Vol. I, John Wiley & Sons, New York, Chichester, Brisbane, Toronto, Singapore, 1981.

[7] C. W. Curtis, I. Reiner, *Representation Theory of Finite Groups and Associative Algebras*, John Wiley & Sons, New York, Chichester, Brisbane, Toronto, Singapore, 1962.

[8] J. D. Dixon, Generating random elements in finite groups, *Electron. J. Combin.* **15** (2008) R94.

[9] The GAP Group, *GAP – Groups, Algorithms, and Programming, Version 4.4.10*, 2007. `http://www.gap-system.org`

[10] S. P. Glasby, R. B. Howlett, Writing representations over minimal fields, *Comm. Algebra* **25** (1997) 1703–1711.

[11] S. P. Glasby, C. R. Leedham-Green, E. A. O'Brien, Writing projective representations over subfields, *J. Algebra* **295** (2006) 51–61.

[12] D. F. Holt, C. R. Leedham-Green, E. A. O'Brien, S. Rees, Computing matrix group decompositions with respect to a normal subgroup, *J. Algebra* **184** (1996) 818–838.

[13] D. F. Holt, C. R. Leedham-Green, E. A. O'Brien, S. Rees, Testing matrix groups for primitivity, *J. Algebra* **184** (1996) 795–817.

[14] D. F. Holt, S. Rees, Testing modules for irreducibility, *J. Austral. Math. Soc. Ser. A* **57** (1994) 1–16.

[15] G. Ivanyos, K. Lux, Treating the exceptional cases of the MeatAxe, *Experiment. Math.* **9** (2000) 373–381.

[16] C. R. Leedham-Green, The computational matrix group project, in: *Groups and Computation III*, Vol. 8 of *OSU Mathematical Research Institute Publications*, Walter de Gruyter, 2001, pp.229–247.

[17] C. R. Leedham-Green, S. H. Murray, Variants of product replacement, *Contemporary Math.* **298** (2002) 97–104.

[18] C. R. Leedham-Green, E. A. O'Brien, Recognising tensor products of matrix groups, *Internat. J. Algebra Comput.* **7** (1997) 541–559.

[19] P. M. Neumann, C. E. Praeger, A recognition algorithm for special linear groups, *Proc. London Math. Soc.* **65** (1992) 555–603.

[20] M. Neunhöffer, Á. Seress, A data structure for a uniform approach to computations with finite groups, in: *ISSAC '06: Proceedings of the 2006 international symposium on Symbolic and algebraic computation*, 2006, pp.254–261.

[21] R. A. Parker, The computer calculation of modular characters (The Meat-Axe), in: *Computational Group Theory*, ed. M. D. Atkinson, Academic Press, 1984, pp.267–274.

[22] Á. Seress, *Permutation Group Algorithms*, Cambridge Tracts in Mathematics **152**, Cambridge University Press, 2003.

[23] H. J. Zassenhaus, *The Theory of Groups*, Chelsea Publishing Company, New York, 1958.