Tensor Products are Projective Geometries

C.R. Leedham-Green   E.A. O'Brien

Abstract

Given a finite dimensional vector space V, we construct a family of projective geometries whose flats are certain subspaces of V, and show that there is a one-to-one correspondence between this family of projective geometries and the set of equivalence classes of tensor decompositions of V. This provides a practical method for finding a tensor decomposition of a finite dimensional KG-module, or proving that no non-trivial tensor decomposition exists.

1 Introduction

The object of this paper is to give an internal description of a tensor decomposition of a finite dimensional vector space, and to use this description to determine whether or not a KG-module V of finite dimension over the ground field K can be decomposed as the tensor product of two modules of smaller dimension.

A tensor decomposition of a KG-module V consists of a KE-isomorphism between V and $U \otimes W$, where E is a covering group of G, and U and W are KE-modules. The kernel C of the homomorphism $E \to G$ is a central cyclic subgroup of E that acts as scalars on both U and W. If $g \in C$ acts as α on U, then g acts as α^{-1} on W. We shall generally assume that $G \leq GL(V)$ and hence that G acts faithfully on V. We shall also assume, when searching for a tensor decomposition of V, that the projective action of G on at least one of the tensor factors is irreducible.

In Section 2 we discuss what is meant by a tensor decomposition of a vector space. A direct decomposition of a module can be defined externally or internally; the internal definition, as a pair of complementary modules, is used in deciding whether or not a module has a direct decomposition. A tensor decomposition of a module, on the other hand, is usually given in terms of an external description. This is not useful in deciding whether or not a given module has a tensor decomposition. In the case that particularly concerns us, when we are considering KG-modules for some field K and group G, one would in principle have to consider the tensor product of all pairs of modules of suitable dimensions over covering groups of G.

Instead, we provide here, apparently for the first time, a description of a tensor decomposition of a vector space V in terms of a projective geometry whose flats are certain subspaces of V. Such a projective geometry on V is an *internal* description of a tensor decomposition of V in that it consists of a set of subspaces of V that are required to satisfy axioms that depend only on the structure of V as a vector space, and on the factorisation of the dimension d of V.

Note that if we regard a point in our projective geometry as a subset of V, then the union of all the points in the projective geometry is a Segre variety, first studied by Segre [8]. (We thank Tim Penttila for this reference.)

Our geometrical approach presented here forms a central part of an algorithm that, given a KG-module V as input, decides whether or not V has a tensor decomposition. In Sections 3 and 4, we outline the theory of two components of this algorithm. We have developed an implementation of our algorithm and find that it generally performs well for matrix groups of moderate degree. The algorithm is presented and its performance analysed in Leedham-Green & O'Brien [5]. Finally, in Section 5, we illustrate these ideas by considering a simple example.

Apart from its intrinsic interest, another motivation for our work is its application to the on-going matrix group "recognition" project. Aschbacher [1] classified the maximal subgroups of $GL(d, q)$ into nine categories, one of which is that a subgroup preserves a tensor decomposition. A first step in "recognising" a matrix group is to determine (at least one of) its categories in the Aschbacher classification. Algorithms have already been developed for some of the other categories: for example, Neumann & Praeger [6], in a seminal paper, propose an algorithm for recognising the special linear group in its natural representation; more recently, Niemeyer & Praeger [7] present a recognition algorithm for classical groups in their natural representations; Holt, Leedham-Green, O'Brien & Rees [2] present an algorithm for primitivity testing.

2 Tensor Products as Projective Geometries

The usual construction of a projective space is to take the subspaces of a linear space. We generalise this to take certain subspaces of dimension a multiple of a fixed divisor u of the dimension d of a vector space V over a field K. These subspaces correspond to the subspaces of the form $U \otimes X$ of $U \otimes W$, where U and W are K-spaces of dimension u and w respectively, and X varies over the subspaces of W, under an isomorphism taking V onto $U \otimes W$. All tensor products are, of course, taken over K.

We start by stating a number of simple axioms that our projective geometry must satisfy.

Definition 2.1 *A set of subspaces $P_0, \ldots, P_w$ of a vector space V is said to be in general position if, for all i, one has $V = \oplus_{j \neq i} P_j$.*

This of course implies that the P_i are all of the same dimension. Given $P_0, P_1, \ldots, P_w$ in general position, and vectors $v_i \in P_i$ such that $\sum_i v_i = 0$, each v_i determines the others. So $v_0 \mapsto v_i$ defines a linear isomorphism θ_i of P_0 onto P_i; in other words, $-\theta_i$ is the restriction to P_0 of the projection of V onto P_i with respect to the decomposition $V = \oplus_{j=1}^{w} P_j$. We sometimes regard θ_i as a linear map of P_0 into V.

We identify W with K^w, and for $x = (x_1, \ldots, x_w) \in W$ define $\theta_x \in \mathrm{Hom}(P_0, V)$ by $\theta_x = \sum_j x_j \theta_j$. Now for $X \leq W$, define $\mathcal{G}(X) = \cup_{x \in X} \theta_x(P_0)$, where $\mathcal{G} = (P_0, \ldots, P_w)$

is an ordered $(w+1)$-tuple in general position. Note that θ_x depends on $\mathcal{G}$. If $x = (x_1, \ldots, x_w)$, we will sometimes write $[x_1, \ldots, x_w]$ for $\mathcal{G}(\langle x \rangle)$.

We can now state our main definition.

Definition 2.2 *Let V be a K-space of dimension $d = uw$ and let $\mathcal{G}$ be a set of $w+1$ subspaces of V of dimension u in general position. Let $\mathcal{F} = \mathcal{F}(\mathcal{G})$ be the collection of subspaces $\mathcal{G}(X)$, for all subspaces X of W. The u-projective geometry defined by $\mathcal{G}$ is the collection of subspaces $\mathcal{F}(\mathcal{G})$.*

A subspace $\mathcal{G}(X)$ is called a *flat*; if X has dimension one, then $\mathcal{G}(X)$ is a *point*.

Observe that the map $X \mapsto \mathcal{G}(X)$ is an isomorphism between the projective geometry $PG(W)$ and the subgeometry $\mathcal{F}$ of $PG(V)$. In particular, the flats in a u-projective geometry on V form a sublattice of the lattice of subspaces of V.

It would be useful to have a characterisation of a u-projective geometry purely in terms of incidence relations. The following example shows that such a task would be difficult.

Let V be the 4-dimensional space of homogeneous polynomials of degree three in two variables over an arbitrary field K. For f a homogeneous polynomial of degree one in these variables, let P_f be the 2-dimensional space of homogeneous polynomials of degree three that are multiples of f^2. Then take $\mathcal{F}$ to consist of the set of these subspaces, together with $\langle 0 \rangle$ and V. It is easy to see that any two distinct subspaces P_f and P_g intersect trivially, so any three are in general position. However, if we define P_0, P_1 and P_2 to be P_{x+y}, P_x and P_y respectively, then

$$[\alpha, \beta] = \langle x(\alpha x^2 + 2\alpha xy + \beta y^2), y(\alpha x^2 + 2\beta xy + \beta y^2) \rangle;$$

this cannot be of the form P_h for any linear homogeneous polynomial h, unless $\alpha = 0$ or $\beta = 0$ or $\alpha = \beta$. Hence, $\mathcal{F}$ is not a 2-projective geometry if K has more than two elements. However, if K is a perfect field of characteristic 2, then V has a tensor decomposition, as $GL(2, K)$-module, as the tensor product of the natural module for $GL(2, K)$ with the module obtained from this one by applying the Frobenius automorphism to the coefficients of the elements of G.

Definition 2.3 *A u-tensor decomposition of V is a linear isomorphism from $U \otimes W$ onto V, where U and W are fixed vector spaces, with U of dimension u. If α and β are u-tensor decompositions of V, they are equivalent if there are linear automorphisms ϕ and ψ of U and W respectively such that $\alpha = \beta(\phi \otimes \psi)$.*

We now justify our definition of a u-projective geometry by showing that it gives an internal definition of a tensor decomposition.

Theorem 2.4 *Let V be a vector space of dimension uw. For each u-tensor decomposition $\alpha : U \otimes W \mapsto V$, define $\mathcal{F}(\alpha)$ to be $\{\alpha(U \otimes X) : X \leq W\}$. Then the map $[\alpha] \longmapsto \mathcal{F}(\alpha)$ is a bijection between the set of equivalence classes $[\alpha]$ of u-tensor decompositions of V and the set of u-projective geometries on V.*

Proof. First note that $\mathcal{F}(\alpha) = \mathcal{F}(\mathcal{G}(\alpha))$ where

$$\mathcal{G}(\alpha) = (\alpha(Kw_0), \alpha(Kw_1), \ldots, \alpha(Kw_w)),$$

w_i is the ith standard basis vector for W for $i > 0$, and $\sum_{i=0}^{w} w_i = 0$. Hence $\mathcal{F}(\alpha)$ is a u-projective geometry. Since $\phi \otimes \psi$ permutes $PG(W)$, it follows that $\mathcal{F}(\alpha) = \mathcal{F}(\alpha \circ (\phi \otimes \psi))$. Hence the map is independent of the representative α of the class $[\alpha]$. Fix a basis $u_1, \ldots, u_u$ of U. Given $\mathcal{G} = (P_0, \ldots, P_w)$ in general position, pick a basis $x_1, \ldots, x_u$ for P_0 and define $\alpha_{\mathcal{F}} : U \otimes W \mapsto V$ by $\alpha_{\mathcal{F}}(u_i \otimes w_j) = \theta_j(x_i)$. Then $\mathcal{F} \mapsto [\alpha_{\mathcal{F}}]$ is an inverse for $[\alpha] \mapsto \mathcal{F}(\alpha)$, so the map is a bijection. $\|$

Corollary 2.5 *If $\mathcal{Q} = (Q_1, \ldots, Q_r)$ is an r-tuple of linearly independent points in the u-projective geometry $\mathcal{F}$, then $\mathcal{Q}$ extends in $\mathcal{F}$ to $\mathcal{P} = (Q_0, \ldots, Q_w)$ in general position and $\mathcal{F} = \mathcal{F}(\mathcal{P})$.*

We now discuss linear maps between flats. The most general class of linear maps we are interested in are those that map subflats of one flat to subflats of a second. These will be called *geometric transformations*. A geometric endomorphism of V is conjugated by a u-projective decomposition of V into an endomorphism of $U \otimes W$ of the form $A \otimes C$, where $A \in \mathrm{End}(U)$ and $C \in \mathrm{End}(W)$.

We shall make considerable use of the more special class of geometric endomorphisms of V to itself that are conjugated into endomorphisms of the form $A \otimes C$ where A is scalar. We call such endomorphisms of V to itself *projectivities*; these will play a crucial role in Section 3 and will be useful for the remainder of this section.

More generally we define projectivities between arbitrary flats of V.

Definition 2.6 *Let V be a K-space of dimension uw and let $\mathcal{F}$ be a u-projective geometry defined on V. A projectivity (or u-projectivity) is a linear map between two flats in $\mathcal{F}$ that is constructed according to the following rules.*

(1) *If $Q_0, \ldots, Q_w$ is a set of points in general position, and if $0 \leq i, j \leq w$, then there is a linear map θ_{ij} that maps Q_i to Q_j defined as follows. Let $v_0 + v_1 + \cdots + v_w = 0$; then $v_i\theta_{ij} = v_j$, and θ_{ij} is a projectivity.*

(2) *If $F_1 = \oplus_i P_i$ and $F_2 = \oplus_j Q_j$ are flats, where P_i and Q_j are points, and if $\phi_{ij} : P_i \to Q_j$ is a projectivity, and $\alpha_{ij} \in K$ for all i, j, then $\sum \alpha_{ij}\phi_{ij} : F_1 \to F_2$ is a projectivity.*

(3) *Every projectivity between two flats can be constructed using this definition.*

We use this definition of a projectivity because of its intrinsic nature. The following description of a projectivity between points is not intrinsic, but is in some ways more instructive.

Lemma 2.7 *Let $Q = [\alpha_1, \ldots, \alpha_w]$ and $R = [\beta_1, \ldots, \beta_w]$ be points. Then the projectivities from Q to R are the linear multiples of $\phi_Q^{-1}\phi_R$, where $\phi_Q = \sum \alpha_i\theta_i$ and $\phi_R = \sum \beta_i\theta_i$.*

Proof. We may assume that Q and R are distinct points, so let $Q_i = [\alpha_{i1}, \ldots, \alpha_{iw}]$ for $0 \le i \le w$ be a set of points in general position where $Q = Q_k$ and $R = Q_\ell$, say. Thus the $(w+1) \times w$ matrix α_{ij} has rank w. Let $\sum_i w_i = 0$, where $w_i \in Q_i$ for all i. Now $w_i = \sum_j \alpha_{ij} v_i \theta_j$ where $v_i \in P_0$ for all i. Then $\sum_{ij} \alpha_{ij} v_i \theta_j = 0$. Since $P_1, \ldots, P_w$ span their direct sum, this implies that $\sum_i \alpha_{ij} v_i \theta_j = 0$ for all j, and hence $\sum_i \alpha_{ij} v_i = 0$ for all j. This gives w independent linear relations between $w+1$ vectors, which must thus be scalar multiples of some fixed vector, say $v_i = \lambda_i v$ for all i. So $w_i = \sum_j \lambda_i \alpha_{ij} v \theta_j$. Now $\theta_{k\ell}$ takes w_k to w_ℓ, and hence is $\lambda_k^{-1} \lambda_\ell \phi_Q^{-1} \phi_R$. The result follows. ||

It follows at once that the set of projectivities between two flats is a K-algebra. More precisely, if F_1 and F_2 are of dimensions ru and su respectively, then the algebra of projectivities from F_1 to F_2 is isomorphic to $K_{r \times s}$.

We now turn to the representation of projectivities with respect to suitable bases for V. Suppose that we have fixed flats $P_0, \ldots, P_w$ as in Definition 2.2. Take an ordered basis B_0 for P_0, and let B_i be the image of B_0 under θ_i for $1 \le i \le w$. Now the concatenation of $B_1, B_2, \ldots, B_w$ is an ordered basis for V. Call this the *standard* basis for V with respect to $P_0, P_1, \ldots, P_w$ and B_0. Now every point Q in the projective geometry is the image of P_0 under some linear map $\sum_i \alpha_i \theta_i$, and so Q has an ordered basis B that is the image of B_0 under this map. Clearly B is unique up to scalar multiple. Call B a *geometric* basis for Q. Now every flat is a direct sum of points, and by concatenating geometric bases for these points we obtain a basis for the flat. Such bases we again call *geometric*. We shall always take matrices of linear transformations between flats with respect to geometric bases.

Let A be an $r \times s$ matrix over K, where r and s are multiples of u. We call A a *u-projective matrix* if A is built up of $u \times u$ blocks, each of which is scalar: that is, A is of the form $I_u \otimes C$ where C is an $r/u \times s/u$ matrix. We can now prove a key result.

Theorem 2.8 *Let F_1 and F_2 be flats in a u-projective geometry on the uw-dimensional K-space V. Let f be a linear transformation from F_1 to F_2, and A be the matrix of f with respect to geometric bases for F_1 and F_2. Then f is a projectivity if and only if A is u-projective.*

Proof. It is clear that the set of u-projective matrices is closed under combining and extracting $u \times u$ blocks, and under forming linear combinations. Hence every projectivity corresponds to a u-projective matrix, since by Lemma 2.7, a projectivity between points maps a geometric basis of the first point to a scalar multiple of a geometric basis of the second point.

Conversely, a u-projective matrix defining a map f between points is, by definition, scalar; since there exists a projectivity between any two points, and such a projectivity must, by Lemma 2.7, be defined by an (arbitrary) scalar matrix, it follows that f is a projectivity. As an arbitrary u-projective matrix is built up from blocks of $u \times u$ scalar matrices, it follows that every u-projective matrix defines a projectivity. ||

We use the following corollaries of this result repeatedly.

Corollary 2.9 *The group of invertible u-projectivities from the uw-dimensional K-space V to V acts transitively on the set of ordered $(w+1)$-tuples of points in general position, and preserves $\mathcal{F}$.*

Corollary 2.10 *An invertible projectivity between two points is unique up to scalar multiplication.*

Corollary 2.11 *A projectivity that fixes each element of a set of $w+1$ points in general position is a scalar.*

Corollary 2.12 *The kernel of a projectivity is a flat.*

Recall that a geometric transformation is a linear map between flats that maps subflats of the first flat to subflats of the second.

Lemma 2.13 *A geometric automorphism of V that fixes each of a set of $w+1$ points in general position, and acts as a scalar on one of them, is a scalar.*

Proof. Let $P_0, P_1, \ldots, P_w$ be the set of points that are fixed and let $v_0+v_1+\cdots+v_w = 0$, where $v_i \in P_i$ for $i = 0, \ldots, w$. Assume that one, and hence all, of the v_i are non-zero. Applying the geometric automorphism, g say, gives

$$v_0^g + v_1^g + \cdots + v_w^g = 0 \qquad (*)$$

where, for some i, we have $v_i^g = \alpha v_i$ for some fixed non-zero $\alpha \in K$. However any one summand in $(*)$ determines the others, hence $v_j^g = \alpha v_j$ for all j. The result follows. $\|$

It is now easy to prove the main results of this section.

Theorem 2.14 *Let g be a linear automorphism of the uw-dimensional K-space V. Let $\mathcal{F}$ be a u-projective geometry on V. Then g is geometric with respect to $\mathcal{F}$ if and only if a u-tensor decomposition α of V corresponding to $\mathcal{F}$ conjugates g to an endomorphism of the form $A \otimes C$, where A and C lie in $GL(u, K)$ and $GL(w, K)$ respectively.*

Proof. Let g be conjugated by α to an automorphism of $U \otimes W$ of the required form. If A is the identity matrix, we have seen that g is a projectivity, and hence is geometric. If C is the identity and A is arbitrary, g maps every flat to itself, and hence is geometric. Since the composite of geometric transformations is geometric, it follows that g is geometric.

Conversely, assume that g is geometric. Multiplying g by some projective automorphism of V, we may assume, using Corollary 2.9, that g fixes every point in the set of points in general position that was used to co-ordinatise V. Multiplying g by a geometric transformation of the form $A \otimes I_w$ we may also assume that g acts as the identity on one of these points. Then, by Lemma 2.13, we conclude that g is now the identity. The result follows. $\|$

The motivation for this work is to produce an algorithm for determining whether or not a KG-module has a non-trivial tensor decomposition. The next result gives the first step along this path.

If G does preserve a u-projective geometry on V, then this induces a projective representation of G on the corresponding tensor factors, which need not be linear. If a tensor decomposition of V as a KG-module exists, this may give rise to a linear action of some covering group E of G on each factor. Explicitly, an element of E may act as a non-identity scalar α say on U and as α^{-1} on W, and hence act trivially on V.

Theorem 2.15 *Let G be a subgroup of $GL(V)$ where V is a uw-dimensional K-space. There is a u-projective geometry on V that is preserved by G if and only if there is an isomorphism of $U \otimes W$ onto V such that G lies in the induced image of the central product $GL(U) \circ GL(W)$, where U is a u-dimensional K-space and W is a w-dimensional K-space.*

This is just a rewording of Theorem 2.14. It follows that G preserves a u-projective geometry if and only if it preserves a w-projective geometry. If $u = 1$, a case that does not interest us, this is the classical dual.

Theorem 2.16 *Let V be a uw-dimensional K-space with a u-projective geometry and corresponding isomorphism to $U \otimes W$. The invertible projectivities of V form a normal subgroup P of the group G of invertible geometric transformations of V, with P isomorphic to $GL(W)$, and G/P isomorphic to $PGL(U)$.*

This follows at once from the above results. We shall need the following slight generalisation:

Theorem 2.17 *Let g be a geometric automorphism of V, and let F_1 and F_2 be flats. If f is a projectivity from F_1 to F_2 then f^g is a projectivity of F_1^g to F_2^g.*

Proof. There is a flat that is a complement to F_2 in V, and the projection of V onto F_2 defined by this complement is a projectivity. Hence f can be extended to a projectivity of V to itself. Then its g-conjugate is a projectivity, and hence restricts to a projectivity of F_1^g to F_2^g. $\|$

3 Finding a projective geometry from a flat

We now outline two central components of an algorithm that, given a KG-module V as input, constructs a G-invariant projective geometry on V or proves that no such geometry exists. In this section, we describe how a point in the geometry may be obtained from a flat of higher dimension; in Section 4, we discuss how we find such a flat.

Suppose that we have a vector space V of dimension d over a field K, a proper subspace F of V, and a set of generators for a group G acting K-linearly on V. How do we decide whether or not V has, for some chosen $u > 1$, a u-projective geometry that is preserved by G and which has F as a flat?

The motivation for asking this question is that, given G, we expect to find one or more subspaces of V such that, if G does preserve a u-projective geometry of V, then it preserves a u-projective geometry that has one of these subspaces as a flat.

We assume that d has a proper factorisation as $u \times w$ and look for a tensor decomposition of V as $U \otimes W$, where U has dimension u and W has dimension w. We assume that the projective action of G on W is irreducible but we do not require that G acts irreducibly on V. We also assume that G acts faithfully on V, though this is just a formality.

Clearly, if F is a flat in a u-projective geometry, then u must divide the dimension of F. If this is so, we try to find a set of G-images of F in general position. This may fail in one of three ways.

1. It may be that the G-images of F do not span V. In this case no tensor decomposition as required can exist since we assume that the projective action of G on W is irreducible.

2. We may find that some G-image F_1 of F intersects some direct sum of G-images of F in some proper non-zero subspace F_2 of F_1. Then F_2 must be a flat in our putative geometry, and we start again with F_2 in place of F provided that u divides the dimension of F_2.

3. It may be that the images of F under G form a system of imprimitivity, in which case again no such geometry exists.

We have now reduced to the situation in which we have found a set of G-images of F in general position. Let $\mathcal{F}$ be the projective geometry they define, and take a geometric basis for V with respect to $\mathcal{F}$.

If the geometry is preserved by G, this can be read off from the matrices of the given generators of G written with respect to this basis; each of these will be a block matrix in which any two non-zero $f \times f$ blocks will differ by a scalar multiple, where f is the dimension of F.

If this is not the case, but F is a flat (not now a point) in some non-trivial projective geometry preserved by G, then the $f \times f$ blocks of which the matrix of an element is composed define geometric transformations between flats. In fact they are a composite of projectivities and a geometric automorphism. Since, by Corollary 2.12, the kernel of a projectivity is a flat, and a geometric automorphism takes flats to flats, it follows that the kernel of each of these blocks is again a flat. Hence if one of these blocks is non-zero but singular, then its kernel defines a smaller dimensional flat, and we start again with this flat.

Now we are reduced to the case in which the $f \times f$ blocks of the matrices representing elements of G are either zero or non-singular. In this case, for $g \in G$, we can find two such blocks C_{ij} and C_{kl} that are both non-singular. Then $C = C_{ij}C_{kl}^{-1}$ is the matrix of the transformation

$$B_i \to V \xrightarrow{g} V \to B_j \to B_l \to V \xrightarrow{g^{-1}} V \to B_k \to B_i.$$

Here B_t is the t-th f-dimensional flat from which the standard basis for V has been constructed, and unnamed maps are projectivities. Thus, by Theorem 2.17, C is the matrix with respect to a geometric basis of a projectivity, provided that the supposed geometry exists. We may assume that C is not scalar since we have not found a tensor decomposition. Since the set of projectivities is closed under addition and scalar multiplication, the eigenspaces of C are flats. So if C has a non-zero eigenspace, we can find a smaller flat. More generally, every polynomial in C over K represents a projectivity, so if the minimal polynomial $m(x)$ of C is reducible we can find a smaller flat as the null-space of $f(C)$ where $f(x)$ is an irreducible factor of $m(x)$.

Finally we consider the case in which the minimal polynomial $m(x)$ of C is irreducible. More precisely, we can proceed as above unless every element in the algebra generated by matrices of the form $C = C_{ij}C_{kl}^{-1}$ is either zero or non-singular. This can only be the case if this algebra is a field $\overline{K}$, which is a proper extension of K, since at least one of the matrices C is not scalar. Suppose that $\overline{K} = GF(q^e)$; of course, e must divide f. If we allow $GF(q^e)$ to act on V by left multiplication by matrices of the form $\mathrm{diag}(C, C, \ldots, C)$, where C is as above, and we use the geometrical basis constructed for V, then clearly G acts semi-linearly on V as $\overline{K}$-module. This gives a tensor decomposition $V = F \otimes W_0 \hat{\otimes} \overline{K}$, where the tensor products are over K, and the expression $W_0 \hat{\otimes} \overline{K}$ denotes the fact that G acts on this tensor product not as a subgroup of the central product of the linear groups on the tensor factors, but rather as a subgroup of the group that acts semi-linearly on the tensor product as $\overline{K}$-space and linearly as K-space.

4 Finding a flat

Assume, as before, that we are looking for a tensor decomposition of the KG-module V of dimension d as $U \otimes W$, where U and W have dimensions u and w respectively.

We now present two approaches to find a flat in a suitable G-invariant projective geometry, or to prove that no such geometry exists.

4.1 Using a projectivity

It may be that G does not act faithfully modulo scalars on one of the factors in the putative tensor decomposition. If a non-scalar element g of G acts as a scalar on U, then g is a u-projectivity mapping V to itself.

Then the characteristic polynomial $f(x)$ of g is a u-th power. If $f(x)$ is not a power of an irreducible polynomial, then we choose an irreducible factor $h(x)$ of $f(x)$; now, by Corollary 2.12, the kernel of $h(A)$ is a flat.

If $f(x)$ is a power of an irreducible polynomial, we can search in the K-algebra generated by g and its conjugates under G for an element whose characteristic polynomial has more than one irreducible factor. This search will fail if the algebra is a field $\overline{K}$. In this case, V may be regarded as a $\overline{K}$-space on which G acts semi-linearly, and we terminate our investigation at this point. Of course, we can easily find generators for the subgroup G_0 of G that acts linearly over $\overline{K}$, and look for a tensor decomposition of V under the action of G_0, working over $\overline{K}$. If V does have a proper tensor decomposition as $\overline{K}G_0$-module, this gives *a fortiori* a tensor decomposition of V as KG-module in which g acts as a projectivity. If V does not have a proper tensor decomposition as $\overline{K}G_0$-module, it may still have a proper decomposition as KG-module with g acting as a projectivity, since G could be a central product of a group of $\overline{K}$-automorphisms of V defined over K with a subgroup of the group $\overline{K}^*$ acting on itself.

Thus we have a procedure IsProjectivity that takes as input generators of $G \le GL(V)$ and $g \in G$, and returns one of the following:

(a) A non-trivial tensor decomposition of V on which g acts as a projectivity.

(b) "False", if a proof has been constructed that no such tensor decomposition exists.

(c) A KG-isomorphism of V with some semi-linear $\overline{K}G$-space, where $\overline{K}$ is a proper extension of K.

IsProjectivity performs a task that could also be carried out using the more general algorithm Smash. That algorithm, described in Holt, Leedham-Green, O'Brien & Rees [3], investigates whether G has certain decompositions with respect to a normal subgroup; however it requires that G acts irreducibly on V.

Finding a (possible) projectivity of V to itself is considered in Leedham-Green & O'Brien [5]. In summary, we may find an element of G whose projective order dictates that some power of it would have to be a projectivity. The *projective order* of $g \in GL(V)$ is the order of the image of g in $PGL(V)$.

4.2 Using reducible subgroups

We now consider the case where G acts faithfully modulo scalars on each of the factors in every tensor decomposition of V.

When considering direct decompositions, one naturally turns to the submodule structure of V. This suggests a second approach to finding a flat – that is, we consider the H-submodule structure of V for "suitable" subgroups H of G. A subgroup is suitable if it is guaranteed to act reducibly on at least one of the tensor factors in any putative tensor decomposition.

Suppose that we also assume that the projective action of G on W is irreducible. Let H be a subgroup of G that acts reducibly on W. Then at least one of the H-invariant subspaces of V is a non-trivial flat in the corresponding u-projective geometry. Hence, we seek to find a subgroup of G that normalises sufficiently few subspaces of V that we can process these subspaces, but which also acts reducibly on W if the required tensor factorisation exists.

One natural class of such subgroups, which has proved useful in practice, can be found as follows. We assume that the ground field is finite, and take H to be p-local for some prime p: namely, H is contained in the normaliser of some non-trivial p-group. The problem that we wish to address here is that of proving, when possible, that H acts reducibly on W if the tensor decomposition exists.

The first and simplest criterion is as follows. If p is the characteristic of the ground field, then H cannot act irreducibly in any dimension greater than one: the subspace of V centralised by a p-group must be non-trivial, and this space is normalised by H.

Now suppose that H normalises a cyclic p-subgroup $P = \langle g \rangle$ of G, where p is not the characteristic of K. We give a criterion, in terms of g, which guarantees that H will act reducibly on one of the tensor factors if the tensor factorisation exists.

Theorem 4.1 *If H is a subgroup of G that acts irreducibly on W, and normalises the cyclic group generated by g, where g has projective order p, then the characteristic polynomial of g acting on W is of the form $\prod_i f_i(x)^t$, where f_1 is an irreducible factor of $x^p - \lambda$ for some scalar λ, and the f_i are the conjugates of f_1 under the action of H. That is to say, if $f_1(x) = \prod_j (x - \phi_j)$, and $h \in H$ conjugates g to g^s for some integer s, then h conjugates $f_1(x)$ to $\prod_j (x - \phi_j^s)$.*

Proof. Since H acts irreducibly on W, by Clifford's theorem (see, for example, Huppert, [4], p. 565) W is a direct sum of homogeneous P-submodules where $P = \langle g \rangle$. If a homogeneous subspace T is the direct sum of t isomorphic irreducible P-submodules, then the characteristic polynomial of g restricted to T is $f(x)^t$, where $f(x)$ is an irreducible factor of $x^p - \lambda$ and λ is some scalar in K. Now the action of g on the other homogeneous factors of W is obtained by conjugating the action on T by each element of a transversal of the inertia group of the action of H on T. The result follows. $\|$

Similar criteria can be developed when P is not cyclic. The construction of p-local subgroups is described in Leedham-Green & O'Brien [5].

If $\prod_i (x - \lambda_i)$ and $\prod_j (x - \mu_j)$ are monic polynomials, their *tensor product* is defined to be $\prod_{ij} (x - \lambda_i \mu_j)$. Clearly the tensor product of two polynomials with coefficients in a field K has its coefficients also in K. The problem of finding a tensor factorisation of a polynomial is discussed in Leedham-Green & O'Brien [5].

We use these ideas as follows. Suppose that we can find $g \in G$ of projective order p and generators for the subgroup H of G that normalises $\langle g \rangle$. A necessary condition for H to act irreducibly on W is that the characteristic polynomial $f(x)$ of g has a tensor factorisation as $u(x) \otimes w(x)$, where $w(x)$ is of degree w and satisfies the condition of Theorem 4.1. So if $f(x)$ does not have such a factorisation, then H cannot act irreducibly on W, and some invariant H-submodule of V must be a proper flat in a u-projective geometry on V.

5 A sample calculation

We illustrate the algorithm embodied in the above description by applying it to the following simple example. Let G be the subgroup of $GL(4, 5)$ generated by the following matrices:

$$A = \begin{pmatrix} 3 & 1 & 1 & 1 \\ 3 & 0 & 4 & 0 \\ 3 & 4 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}; B = \begin{pmatrix} 3 & 4 & 4 & 4 \\ 4 & 2 & 1 & 2 \\ 4 & 1 & 2 & 2 \\ 3 & 1 & 1 & 2 \end{pmatrix}; C = \begin{pmatrix} 4 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 \\ 2 & 0 & 1 & 0 \\ 0 & 0 & 0 & 4 \end{pmatrix}.$$

The problem is to determine whether or not G preserves a tensor decomposition of the natural module V. One first observes that C and AC are of order 2, and hence generate a dihedral group. Since the characteristic polynomial of A is $(x-1)^4$, it follows that A has order 5. So the group H generated by A and C is dihedral of order 10, and is 5-local.

The next step is to find the H-invariant 2-dimensional subspaces of V. Any such subspace must contain an eigenspace for A; that is, a subspace that is centralised by A. It is easy to see that the subspace of V centralised by A has a basis $\{v_1, v_2\}$, where $v_1 = (0, 1, -1, 0)$ and $v_2 = (0, 0, 0, 1)$. One checks that the only one-dimensional subspaces of $\langle v_1, v_2 \rangle$ that are also invariant under C are $\langle v_1 \rangle$ and $\langle v_2 \rangle$. The H-invariant 2-dimensional subspaces of V that contain v_2 are those of the form $X = \langle (x, x+y, -y, 0), v_2 \rangle$, and no other cases arise from subspaces that contain v_1. So we have 6 subspaces to consider.

If we take X and X^B, these subspaces must complement each other, as they cannot be equal, and this implies that $x \neq 0$ and $x + 2y \neq 0$. With these conditions, $P_0 = X$ and $P_1 = X^B$ and $P_2 = X^{BC}$ are in general position. So we take $x = 1$ and $y = 0$.

Now X has a basis $\{x_1, x_2\}$, where $x_1 = (0, 0, 0, 1)$, and $x_2 = (1, 1, 0, 0)$. One finds that $x_1 = f_1 + f_3$ where $f_1 = (-2, 0, -2, -2) \in P_1$ and $f_3 = (2, 0, 2, -2) \in P_2$. Similarly $x_2 = f_2 + f_4$ where $f_2 = (1, -2, 0, -2) \in P_1$ and $f_4 = (0, -2, 0, 2) \in P_2$. So we write A, B and C with respect to the basis f_1, f_2, f_3, f_4, and obtain:

$$A' = \begin{pmatrix} 4 & 3 & 0 & 0 \\ 2 & 3 & 0 & 0 \\ 4 & 3 & 4 & 3 \\ 2 & 3 & 2 & 3 \end{pmatrix}; B' = \begin{pmatrix} 4 & 2 & 4 & 2 \\ 3 & 3 & 3 & 3 \\ 0 & 0 & 4 & 2 \\ 0 & 0 & 3 & 3 \end{pmatrix}; C' = \begin{pmatrix} 0 & 4 & 0 & 0 \\ 4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Thus G acts projectively as $PSL(2, 5)$ on each of the tensor factors, and is in fact isomorphic to $PSL(2, 5)$.

Acknowledgements

O'Brien is supported by the Humboldt Foundation, Bonn, which also partially funded a visit by him to the University of London via a European Research Fellowship. Both authors were supported by the European Union HCM Grant for Computational Group Theory. We thank the School of Mathematical Sciences at the Australian National University and the Department of Mathematics, University of Western Australia, for their hospitality while part of this work was carried out. We are grateful to Stephen Donkin, Stephen Glasby, L.G. Kovács, and I.G. Macdonald for helpful discussions and thank the referee for many helpful suggestions.

References

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

[2] Derek F. Holt, C.R. Leedham-Green, E.A. O'Brien and Sarah Rees, "Testing Matrix Groups for Primitivity", *J. Algebra* **184** (3), 1996.

[3] Derek F. Holt, C.R. Leedham-Green, E.A. O'Brien and Sarah Rees, "Computing matrix group decompositions with respect to a normal subgroup", *J. Algebra* **184** (3), 1996.

[4] B. Huppert, *Endliche Gruppen I*, Grundlehren Math. Wiss., **134**. Springer-Verlag, Berlin, Heidelberg, New York, 1967.

[5] C.R. Leedham-Green and E.A. O'Brien, "Recognising tensor products of matrix groups", preprint.

[6] Peter M. Neumann and Cheryl E. Praeger, "A recognition algorithm for special linear groups", *Proc. London Math. Soc.* (3), **65**, 555–603, 1992.

[7] Alice C. Niemeyer and Cheryl E. Praeger, "A Recognition Algorithm for Classical Groups over Finite Fields", preprint.

[8] C. Segre, "Sulle varietà che rappresentano le coppie di punti di due piani o spazi", *Rend. Circ. Mat. Palermo* **5**, 192–204, 1891.

School of Mathematical Sciences
Queen Mary and Westfield College
University of London
London E1 4NS
United Kingdom
C.R.Leedham-Green@qmw.ac.uk

Lehrstuhl D für Mathematik
RWTH
Templergraben 64
52062 Aachen
Germany
obrien@math.rwth-aachen.de