

Relative Linear Word Observables and Representation Complexity of Semidirect Products

Luca Blanchi

June 2026

Abstract

We introduce and study relative linear word-observable profiles for marked semidirect products

$$G_M = M \rtimes Q,$$

where Q is a finite group and M is a finite-dimensional module over the group algebra kQ . A relative linear word observable is obtained by evaluating expressions in variables ranging over the abelian normal subgroup M , with conjugations by elements of Q and with scalar operations from k explicitly included in the marked observable language. This distinction is essential: when $k = \mathbb{F}_p$, the scalar operations are ordinary powers in the elementary abelian p -group M , while over a general finite field $\mathbb{F}_q$ they should be regarded as part of a k -linear relative observable language rather than as ordinary group words.

The first main result identifies these observables exactly with matrix-rank observables on the kQ -module M : every matrix over kQ is realized by a finite tuple of relative linear word observables, and the corresponding value distribution determines the rank of the induced linear map on M . This gives a bridge between word-statistical observation of semidirect products and modular representation theory.

We prove that the full relative matrix word profile is a complete invariant of marked semidirect products: if all matrix-rank observables agree on two finite-dimensional kQ -modules M, N , then $M \cong N$, hence $M \rtimes Q \cong N \rtimes Q$ as marked semidirect products. The proof uses the standard module-theoretic fact that pp-pairs, equivalently finitely presented functors, detect indecomposable summands of finite-dimensional modules.

We then analyze finite-budget observables. In the semisimple fixed-quotient case, primitive central idempotents give bounded complete observables. For cyclic modular quotients, difference observables recover ranks of radical powers and hence the Jordan partition. In wild modular regimes, the full profile remains complete, but every finite set of matrix-rank observables is generically blind on positive-dimensional algebraic families. Thus the paper exhibits a trichotomy:

semisimple fixed quotient $\Rightarrow$ bounded complete observables,

cyclic modular quotient $\Rightarrow$ radical-depth observables,

wild modular quotient $\Rightarrow$ complete infinite profile but finite-budget residual fibers.

Declaration of generative AI and AI-assisted technologies in the writing process

This article belongs to the Presentation theory section. Generative AI, including ChatGPT by OpenAI, was used to assist with mathematical drafting, formalization, review, and editing. The note may be preliminary, may have been only partially reviewed, and in some cases may be only AI-reviewed.

1 Introduction

A finite group may be represented in many ways: by generators and relations, by matrices, by a polycyclic presentation, by a short straight-line description, by a normal coordinate system, or by an extension

$$1 \rightarrow N \rightarrow G \rightarrow Q \rightarrow 1.$$

Such representations do more than describe the abstract group. They determine which operations are easy, which observables are available, what data can be recovered, and how much information is visible at bounded complexity.

This paper studies a concrete representation system: marked semidirect products

$$G_M = M \rtimes Q,$$

where Q is a fixed finite group and M is a finite-dimensional kQ -module. The marking means that the normal abelian subgroup M , the quotient Q , and the action of Q on M are part of the represented object. This avoids ambiguities caused by automorphisms of Q or by noncanonical choices of complements.

The central question is: *how much of the kQ -module M is visible through bounded-complexity relative observables in $M \rtimes Q$?* The answer is sharp.

First, relative linear word observables are exactly matrix-rank observables. Given

$$a = \sum_{q \in Q} c_q q \in kQ,$$

we form the relative observable

$$w_a(y) = \prod_{q \in Q} q y^{c_q} q^{-1},$$

where y ranges over the abelian normal subgroup M and y^{c_q} denotes the scalar multiple $c_q y$. If $k = \mathbb{F}_p$, this scalar operation is an ordinary group power in M . If $k = \mathbb{F}_q$ with $q > p$, it is part of the marked k -linear observable language. Since M is abelian,

$$w_a(v) = a \cdot v.$$

More generally, a matrix $A \in M_{s \times t}(kQ)$ gives a tuple of relative observables whose evaluation is exactly the linear map

$$A_M : M^t \rightarrow M^s.$$

When the input is uniformly distributed over M^t , the output is uniformly distributed on $\text{im } A_M$. Therefore the value distribution determines $\text{rank } A_M$.

Second, the full profile is complete. If

$$\text{rank } A_M = \text{rank } A_N$$

for all matrices A over kQ , then $M \cong N$ as kQ -modules. Hence the full relative matrix word profile classifies marked semidirect products. The proof uses standard tools from model theory of modules: ranks of matrices determine dimensions of pp-definable subgroups and pp-pairs; pp-pairs, equivalently finitely presented functors, detect indecomposable summands; Krull-Schmidt then recovers the module.

Third, finite-budget behavior separates the representation-theoretic regimes. If kQ is semisimple, primitive central idempotents give bounded complete observables. If $Q = C_{p^a}$ in characteristic p , iterated difference observables recover the ranks of radical powers and hence the Jordan partition. In wild modular regimes, finite observables cannot classify generic positive-dimensional families, even though the full profile is complete.

2 Marked Semidirect Products and Relative Observables

Let k be a field, Q a finite group, and M a finite-dimensional left kQ -module. We write the associated semidirect product as

$$G_M = M \rtimes Q.$$

The multiplication is

$$(m, q)(n, r) = (m + q \cdot n, qr).$$

We regard G_M as a marked semidirect product: the subgroup M , the quotient Q , and the action of representatives of Q by conjugation are part of the representation.

Definition 2.1 (Relative linear word observables). *A relative linear word observable in variables $y_1, \dots, y_t$ ranging over M is an expression built from:*

- multiplication in the abelian group M ;
- conjugation by marked elements of Q ;
- scalar multiplication $y \mapsto cy$ for $c \in k$.

Equivalently, each output coordinate is a kQ -linear combination of the variables $y_1, \dots, y_t$.

Remark 2.2 (Ordinary group words). *When $k = \mathbb{F}_p$, the additive group of M is elementary abelian, and scalar multiplication by $c \in \mathbb{F}_p$ is the ordinary group power $y \mapsto y^c$. In this prime-field case, the relative linear observables above can be realized by ordinary relative group words using conjugation by Q .*

When $k = \mathbb{F}_q$ with $q = p^f$ and $f > 1$, scalar multiplication by a general $c \in \mathbb{F}_q$ is not an ordinary group word operation on the underlying abelian group. For this reason the present paper uses the more precise term relative k -linear word observable. The ordinary-word interpretation should be read literally only in the prime-field case or after adding the scalar operations to the marked language.

2.1 Unary observables

For

$$a = \sum_{q \in Q} c_q q \in kQ,$$

define the unary relative observable

$$w_a(y) = \prod_{q \in Q} q y^{c_q} q^{-1}.$$

The order of the product is irrelevant because all factors lie in the abelian subgroup M .

Lemma 2.3 (Group-algebra evaluation). *For every $v \in M$,*

$$w_a(v) = a \cdot v.$$

Proof. Each factor satisfies

$$q v^{c_q} q^{-1} = c_q(q \cdot v).$$

Multiplying in the abelian subgroup M gives

$$w_a(v) = \sum_{q \in Q} c_q(q \cdot v) = a \cdot v.$$

□

2.2 Matrix observables

Let

$$A = (a_{ij}) \in M_{s \times t}(kQ).$$

It defines a k -linear map

$$A_M : M^t \rightarrow M^s$$

by

$$A_M(v_1, \dots, v_t) = \left(\sum_{j=1}^t a_{1j}v_j, \dots, \sum_{j=1}^t a_{sj}v_j \right).$$

For each row i , define

$$w_i(y_1, \dots, y_t) = \prod_{j=1}^t w_{a_{ij}}(y_j),$$

and let $W_A = (w_1, \dots, w_s)$.

Proposition 2.4 (Matrix observable evaluation). *For every $\mathbf{v} = (v_1, \dots, v_t) \in M^t$,*

$$W_A(\mathbf{v}) = A_M(\mathbf{v}).$$

Proof. Using the unary evaluation lemma,

$$w_i(\mathbf{v}) = \sum_{j=1}^t a_{ij}v_j.$$

Thus the tuple W_A is exactly the map A_M . □

3 Value Distributions and Rank Profiles

Assume now that $k = \mathbb{F}_q$ is finite. Let M be finite-dimensional over k and sample $\mathbf{v} \in M^t$ uniformly. Since $A_M : M^t \rightarrow M^s$ is linear, the random variable $A_M(\mathbf{v})$ is uniformly distributed on $\text{im } A_M$. Hence the distribution of W_A determines

$$|\text{im } A_M| = q^{\text{rank } A_M}.$$

In particular,

$$\Pr_{\mathbf{v} \in M^t} [W_A(\mathbf{v}) = 0] = q^{-\text{rank } A_M}.$$

Definition 3.1 (Matrix rank profile). *Let A be a finite-dimensional k -algebra and M a finite-dimensional left A -module. The full matrix rank profile of M is*

$$\text{MRP}_A(M) = (\text{rank}_k B_M)_{B \in M_{s \times t}(A), s, t \geq 0}.$$

If $A = kQ$ and $G_M = M \rtimes Q$, the full relative matrix word profile is

$$\text{RelMRWP}(G_M) = \text{MRP}_{kQ}(M).$$

A finite-budget version is obtained by restricting to matrices B whose entries are described by observables of complexity at most L .

Theorem 3.2 (Relative observables equal matrix-rank observables). *Let $G_M = M \rtimes Q$ be a marked semidirect product over a finite field k . Then relative k -linear multi-word value distributions determine exactly the matrix rank profile of the kQ -module M :*

$$\text{RelMRWP}(G_M) = \text{MRP}_{kQ}(M).$$

Proof. The matrix observable construction realizes every matrix A over kQ by a tuple of relative observables W_A , and the value distribution determines $\text{rank } A_M$.

Conversely, every relative k -linear word expression in the abelian variables $y_1, \dots, y_t$, using scalar operations and conjugations by elements of Q , has each output coordinate of the form

$$\sum_j a_{ij} y_j$$

with $a_{ij} \in kQ$. Thus every such observable is represented by a matrix over kQ . $\square$

4 Completeness of the Full Matrix Rank Profile

Theorem 4.1 (Full matrix rank profile completeness). *Let A be a finite-dimensional k -algebra, and let M, N be finite-dimensional left A -modules. If*

$$\text{MRP}_A(M) = \text{MRP}_A(N),$$

then $M \cong N$ as A -modules. Consequently, for marked semidirect products $G_M = M \rtimes Q$ and $G_N = N \rtimes Q$,

$$\text{RelMRWP}(G_M) = \text{RelMRWP}(G_N)$$

implies $G_M \cong G_N$ as marked semidirect products.

4.1 From matrix ranks to pp-pairs

A pp-formula for left A -modules is a system of homogeneous linear equations with existentially quantified variables. It has the form

$$\phi(x) \equiv \exists y \quad Cx + Dy = 0,$$

where C, D are matrices over A . For a module M , the solution set $\phi(M)$ is a k -subspace of M^n .

Let

$$[C \ D]_M : M^{n+m} \rightarrow M^r$$

be the induced linear map, and let $D_M : M^m \rightarrow M^r$ be the map induced by D . Then

$$\dim_k \phi(M) = \dim_k \ker [C \ D]_M - \dim_k \ker D_M.$$

Equivalently,

$$\dim_k \phi(M) = n \dim_k M - \text{rank} [C \ D]_M + \text{rank } D_M.$$

Therefore the matrix rank profile determines the dimension of $\phi(M)$ for every pp-formula ϕ .

If $\psi \leq \phi$ are pp-formulas, the pp-pair ϕ/ψ has value $\phi(M)/\psi(M)$, and

$$\dim_k(\phi(M)/\psi(M)) = \dim_k \phi(M) - \dim_k \psi(M).$$

Therefore the matrix rank profile determines the dimensions of all pp-pairs.

4.2 pp-pairs detect indecomposable summands

For a finite-dimensional algebra A , the category of finite-dimensional A -modules is Krull-Schmidt. Every finite-dimensional module decomposes uniquely, up to isomorphism and permutation, as a finite direct sum of indecomposable modules.

We use the following standard consequence of the Auslander category of finitely presented functors, equivalently the model theory of modules.

Lemma 4.2 (Indecomposable-detecting pp-pairs). *For every finite-dimensional indecomposable A -module U , there exists a pp-pair*

$$\phi_U/\psi_U$$

such that, for every finite-dimensional A -module M ,

$$\dim_k(\phi_U(M)/\psi_U(M)) = m_U(M) \cdot \dim_k D_U,$$

where $m_U(M)$ is the multiplicity of U as a direct summand of M , and

$$D_U = \text{End}_A(U)/\text{rad End}_A(U).$$

Equivalently, the simple finitely presented functor

$$S_U = \text{Hom}_A(U, -)/\text{rad}(U, -)$$

is represented by a pp-pair, and its value on M detects the number of direct summands isomorphic to U .

Proof of Theorem 4.1. Assume $\text{MRP}_A(M) = \text{MRP}_A(N)$. By the preceding subsection, M and N have the same dimensions for all pp-pairs. For each indecomposable U , apply the lemma:

$$\dim_k(\phi_U(M)/\psi_U(M)) = \dim_k(\phi_U(N)/\psi_U(N)).$$

Thus

$$m_U(M) \dim_k D_U = m_U(N) \dim_k D_U,$$

so $m_U(M) = m_U(N)$. Therefore M and N have the same multiplicities of all indecomposable direct summands. By Krull-Schmidt, $M \cong N$.

If $A = kQ$, an isomorphism of kQ -modules $M \cong N$ induces an isomorphism of marked semidirect products

$$M \rtimes Q \cong N \rtimes Q$$

by $(m, q) \mapsto (Sm, q)$, where $S : M \rightarrow N$ is the module isomorphism. □

Remark 4.3. *The completeness theorem is not a claim that a finite list of observables classifies arbitrary modules. Rather, the full infinite profile is complete. Finite-budget profiles form a hierarchy of approximations to isomorphism. In semisimple fixed-quotient cases this hierarchy collapses at bounded budget; in wild cases it typically does not.*

5 Semisimple Fixed Quotients

Let $k = \mathbb{F}_q$ be finite and let Q be a finite group with $\text{char } k \nmid |Q|$. By Maschke's theorem, kQ is semisimple. Let

$$1 = e_1 + \cdots + e_s$$

be the decomposition of the identity into primitive central idempotents of kQ . Let $S_1, \dots, S_s$ be the corresponding simple kQ -modules, and set $d_i = \dim_k S_i$. Every finite-dimensional kQ -module decomposes as

$$M \cong \bigoplus_{i=1}^s S_i^{\oplus m_i}.$$

The operator e_i acts as the projection onto the S_i -isotypic component. Hence

$$\text{rank}(e_i)_M = m_i d_i.$$

Theorem 5.1 (Bounded observable classification for fixed semisimple quotients). *Fix $k = \mathbb{F}_q$ and a finite group Q with $\text{char } k \nmid |Q|$. There exists a constant $L = L(k, Q)$ such that, for every finite-dimensional kQ -module M , the marked semidirect product $G_M = M \rtimes Q$ is classified by the relative identity profiles of unary k -linear observables of complexity at most L .*

Proof. For each primitive central idempotent

$$e_i = \sum_{q \in Q} c_{i,q} q,$$

form the unary observable w_{e_i} . By the group-algebra evaluation lemma, $w_{e_i}(v) = e_i v$. Therefore

$$\Pr_{v \in M} [w_{e_i}(v) = 0] = q^{-\text{rank}(e_i)_M}.$$

The profile of w_{e_i} determines $\text{rank}(e_i)_M = m_i d_i$, hence the multiplicity m_i . The idempotents e_i depend only on k and Q , so the corresponding observables have uniformly bounded description complexity. The multiplicities m_i determine the semisimple module M , hence the marked semidirect product. $\square$

Remark 5.2. *If one insists on ordinary group words rather than k -linear relative observables, the prime-field case $k = \mathbb{F}_p$ is literal. Over larger finite fields, the theorem should be read in the marked k -linear observable language, or after replacing scalar operations by a fixed marked implementation when such an implementation is part of the representation.*

6 Cyclic Modular Quotients and Radical Depth

Let $k = \mathbb{F}_p$ and

$$Q = C_{p^a} = \langle t \mid t^{p^a} = 1 \rangle.$$

A finite-dimensional kQ -module is equivalent to an operator T satisfying $T^{p^a} = I$. Set

$$N = T - I.$$

In characteristic p ,

$$T^{p^a} - I = (T - I)^{p^a} = N^{p^a},$$

so N is nilpotent. Let $G_T = V \rtimes_T C_{p^a}$.

Define the difference observable

$$\Delta(y) = tyt^{-1}y^{-1}.$$

Then

$$\Delta(v) = Tv - v = Nv.$$

Define iterates by

$$\Delta^{(0)}(y) = y, \quad \Delta^{(e+1)}(y) = \Delta(\Delta^{(e)}(y)).$$

Then

$$\Delta^{(e)}(v) = N^e v.$$

Theorem 6.1 (Modular cyclic radical-depth classification). *The relative identity profiles of*

$$\Delta^{(e)}(y), \quad 1 \leq e \leq p^a,$$

determine the Jordan partition of $N = T - I$, hence classify the marked semidirect product $V \rtimes_T C_{p^a}$.

Proof. For $v \in V$ uniformly distributed,

$$\Delta^{(e)}(v) = N^e v.$$

Thus

$$\Pr[\Delta^{(e)}(v) = 0] = \frac{|\ker N^e|}{|V|} = p^{\dim \ker N^e - \dim V}.$$

So the identity profile determines $\dim \ker N^e$ for each e . If the Jordan block sizes of N are

$$\lambda = (\lambda_1 \geq \lambda_2 \geq \dots),$$

then

$$\dim \ker N^e = \sum_i \min(e, \lambda_i).$$

Therefore

$$\dim \ker N^e - \dim \ker N^{e-1} = |\{i : \lambda_i \geq e\}|.$$

These numbers determine the partition λ , hence the similarity class of N and of $T = I + N$. $\square$

Proposition 6.2 (Pure-difference lower bound). *For every $D \geq 1$, there exist two nonisomorphic kC_{p^a} -modules, for a large enough, that are indistinguishable by the observables*

$$\Delta^{(e)}, \quad 1 \leq e \leq D.$$

Proof. Consider nilpotent operators with Jordan partitions

$$\lambda = (D+1, D+1), \quad \mu = (D+2, D).$$

Both have total dimension $2D+2$. For every $1 \leq e \leq D$,

$$\sum_i \min(e, \lambda_i) = 2e = \sum_i \min(e, \mu_i).$$

Thus $\dim \ker N_\lambda^e = \dim \ker N_\mu^e$ for all $e \leq D$, while the partitions are distinct. $\square$

Remark 6.3. *The lower bound concerns the pure-difference hierarchy $\Delta^{(e)}$. In a straight-line program model, certain high radical powers can be probed more efficiently. For example,*

$$t^{p^s} y t^{-p^s} y^{-1}$$

evaluates to

$$(T^{p^s} - I)y = (T - I)^{p^s} y = N^{p^s} y.$$

Thus the observable complexity of probing N^e depends on the operator-polynomial complexity of $(z - 1)^e$ modulo $(z - 1)^{p^a}$, not simply on e .

7 Elementary Abelian Quotients and Rank Varieties

Let

$$Q = E_r = C_p^r$$

and let k have characteristic p . Then

$$kE_r \cong k[\epsilon_1, \dots, \epsilon_r] / (\epsilon_1^p, \dots, \epsilon_r^p).$$

A finite-dimensional kE_r -module is a tuple of commuting nilpotent operators

$$N_1, \dots, N_r, \quad N_i^p = 0.$$

For every $\lambda = (\lambda_1, \dots, \lambda_r) \in k^r$, define

$$N_\lambda = \lambda_1 N_1 + \dots + \lambda_r N_r.$$

The Jordan type of N_λ is determined by the ranks of N_λ^e for $1 \leq e \leq p$.

Theorem 7.1 (Rank-variety observables). *The relative matrix-rank profile of $M \rtimes E_r$ recovers the rank and Jordan-type data of the restrictions of M to all rank-one shifted cyclic subgroups. In particular, it recovers the usual rank-variety and constant-Jordan-type loci associated to M .*

Proof. For each λ and e , the operator N_λ^e is an element of kE_r . The unary observable associated to this group-algebra element evaluates on $v \in M$ as $N_\lambda^e v$. The identity probability determines rank N_λ^e , and these ranks determine the Jordan type of N_λ . $\square$

Remark 7.2. *Rank varieties and constant Jordan type are therefore contained in the relative matrix-rank observable profile. The full matrix word profile is broader: it includes ranks of arbitrary matrices over kE_r , not only ranks of powers of one-parameter nilpotent restrictions.*

8 Wild Regimes

Let A be a finite-dimensional algebra over an algebraically closed field k . Let X be an irreducible algebraic variety parametrizing a family of A -modules

$$(M_x)_{x \in X}$$

of fixed dimension. Assume the family is generically non-isomorphic: there is a dense open subset of X whose distinct generic points represent nonisomorphic modules.

For a fixed matrix

$$B \in M_{s \times t}(A),$$

the map $B_{M_x} : M_x^t \rightarrow M_x^s$ depends algebraically on x . Hence the function

$$x \mapsto \text{rank } B_{M_x}$$

is maximal on a dense open subset of X .

Theorem 8.1 (Generic finite-observable blindness). *Let $(M_x)_{x \in X}$ be an irreducible algebraic family of A -modules. For every finite set of matrix-rank observables*

$$\mathcal{S} = \{B_1, \dots, B_N\},$$

there exists a dense open subset $U_{\mathcal{S}} \subseteq X$ such that all ranks

$$\text{rank}(B_i)_{M_x}$$

are constant on $U_{\mathcal{S}}$. If the family is generically non-isomorphic and $\dim X > 0$, then no finite set $\mathcal{S}$ can classify the generic members of the family.

Proof. For each B_i , the locus where $\text{rank}(B_i)_{M_x}$ is maximal is open and dense. Since X is irreducible, the finite intersection of these dense open subsets is dense open and nonempty. On that intersection all ranks are constant.

If the family is generically non-isomorphic and positive-dimensional, the dense open subset contains infinitely many pairwise nonisomorphic generic modules. They have the same $\mathcal{S}$ -profile. $\square$

Corollary 8.2. *Let $G_x = M_x \rtimes Q$ be a marked semidirect family with $A = kQ$. For any finite family of relative matrix-rank observables, there is a dense open subset of X on which all corresponding observable profiles are constant. Thus finite observables detect rank-degeneration strata, not generic parameters, in positive-dimensional wild families.*

9 Fox-Linear Homomorphism Lifting

We include a further application connecting relative observables with homomorphism counts.

Let

$$\Gamma = \langle x_1, \dots, x_k \mid r_1, \dots, r_m \rangle$$

be a finitely presented group. Let $G = M \rtimes Q$ be a split marked semidirect product. Fix a homomorphism $\phi : \Gamma \rightarrow Q$, determined by $q_i = \phi(x_i)$ satisfying $r_j(q_1, \dots, q_k) = 1$.

A lift of ϕ to G is obtained by choosing $a_i \in M$ and mapping

$$x_i \mapsto (a_i, q_i).$$

The relations become linear equations in the a_i . Let $\partial r_j / \partial x_i \in \mathbb{Z}F_k$ be the Fox derivatives. Evaluating at q_i gives elements of $\mathbb{Z}Q$, hence operators on M . Define the Fox-Jacobian

$$J_\phi : M^k \rightarrow M^m$$

by

$$(J_\phi(a_1, \dots, a_k))_j = \sum_{i=1}^k \left(\frac{\partial r_j}{\partial x_i} \right) (q_1, \dots, q_k) \cdot a_i.$$

Theorem 9.1 (Split Fox-Hom lifting formula). *For the split semidirect product $G = M \rtimes Q$,*

$$|\mathrm{Hom}_\phi(\Gamma, G)| = |\ker J_\phi|.$$

Consequently,

$$|\mathrm{Hom}(\Gamma, G)| = \sum_{\phi \in \mathrm{Hom}(\Gamma, Q)} |\ker J_\phi|.$$

Proof. Evaluating a relator r_j on lifts (a_i, q_i) yields, in the M -coordinate, the linear Fox expression

$$\sum_i \left(\frac{\partial r_j}{\partial x_i} \right) (q_1, \dots, q_k) \cdot a_i.$$

The Q -coordinate is already trivial because ϕ is a homomorphism. Thus the lifts are precisely the solutions of $J_\phi(a_1, \dots, a_k) = 0$. $\square$

Remark 9.2 (Nonsplit extensions). *For a nonsplit extension with abelian kernel, the same computation gives an affine system*

$$J_\phi(\mathbf{a}) + B_c(\phi) = 0,$$

where $B_c(\phi)$ is a cocycle-dependent obstruction term. Hence

$$|\mathrm{Hom}_\phi(\Gamma, G)| = |\ker J_\phi| \mathbf{1}_{-B_c(\phi) \in \mathrm{im} J_\phi}.$$

10 Marked and Unmarked Semidirect Products

All main classification theorems above concern marked semidirect products. This is the natural setting for representation complexity, because the action module M and quotient Q are part of the representation.

For unmarked groups, additional equivalences appear. If $Q = C_m$, an automorphism of Q may send

$$t \mapsto t^u, \quad u \in (\mathbb{Z}/m\mathbb{Z})^\times.$$

Thus the action T may be replaced by T^u . Accordingly, unmarked profiles can classify at best the orbit of the action under the relevant automorphism group, unless the subgroup M and quotient Q are intrinsically characterized in the abstract group.

In general, unmarked classification requires quotienting module classification by the natural action of $\mathrm{Aut}(Q)$ and by any further automorphisms of the extension structure. We do not pursue this here. The marked theory is the basic representational object.

11 Summary and Open Problems

The results give the following hierarchy for marked semidirect products $G_M = M \rtimes Q$.

- In the semisimple fixed-quotient case, bounded observables classify arbitrary-dimensional semidirect products.
- In the cyclic modular case, observables $\Delta^{(e)}$ recover radical depth and the Jordan partition.
- In elementary abelian and wild modular cases, relative matrix-rank observables include rank-variety and constant-Jordan-type data; the full profile is complete, but finite profiles are generically blind on positive-dimensional wild families.

Problem 11.1. For a wild algebra A and a finite-budget observable class, estimate the residual fiber size over finite fields, or the geometric residual fiber dimension over algebraically closed fields.

Problem 11.2. For tame algebras, determine whether finite or logarithmically growing observable budgets recover the parameters of one-parameter families.

Problem 11.3. For $kC_{p^a} \cong k[\epsilon]/(\epsilon^{p^a})$, determine the straight-line complexity of constructing the operator ϵ^e as a relative observable.

Problem 11.4. Determine conditions under which the abelian normal subgroup M and quotient Q are intrinsic in the abstract group $M \rtimes Q$, so that marked classification descends to unmarked classification.

Problem 11.5. Describe precisely which finite-budget matrix-rank observables correspond to known rank varieties, constant Jordan type, radical type, and socle type invariants for elementary abelian p -groups.

References

- [1] M. Auslander. Representation theory of Artin algebras I, II. *Communications in Algebra*, 1970s.
- [2] D. J. Benson. *Representations and Cohomology I: Basic Representation Theory of Finite Groups and Associative Algebras*. Cambridge University Press.
- [3] J. F. Carlson, E. M. Friedlander, and J. Pevtsova. Modules of constant Jordan type. *Journal für die reine und angewandte Mathematik* 614 (2008), 191–234.
- [4] Y. A. Drozd. Tame and wild matrix problems. In *Representation Theory II*, Lecture Notes in Mathematics 832, Springer, 1980.
- [5] R. H. Fox. Free differential calculus. I. Derivation in the free group ring. *Annals of Mathematics* 57 (1953), 547–560.
- [6] R. H. Fox. Free differential calculus. II. The isomorphism problem of groups. *Annals of Mathematics* 59 (1954), 196–210.
- [7] P. Gabriel and A. V. Roiter. *Representations of Finite-Dimensional Algebras*. Springer.
- [8] M. Larsen and A. Shalev. Word maps and Waring type problems. *Journal of the American Mathematical Society* 22 (2009), 437–466.
- [9] M. Larsen, A. Shalev, and P. H. Tiep. The Waring problem for finite simple groups. *Annals of Mathematics* 174 (2011), 1885–1950.
- [10] M. Lohrey. *The Compressed Word Problem for Groups*. SpringerBriefs in Mathematics, 2014.
- [11] S. Mac Lane. *Homology*. Springer.
- [12] M. Prest. *Model Theory and Modules*. Cambridge University Press, 1988.
- [13] M. Prest. *Purity, Spectra and Localisation*. Cambridge University Press, 2009.
- [14] J.-P. Serre. *Linear Representations of Finite Groups*. Springer.
- [15] C. A. Weibel. *An Introduction to Homological Algebra*. Cambridge University Press.