README: Extension calculations

Normal reduction, Ext^1, Ext^1_0 and the embedding

The calculator takes the two t-action matrices Φ_t and Ψ_t and constructs the action of t on extension classes. It also computes the subgroup $\mathrm{Ext}^1_0(\Phi,\Psi)$, its inclusion into $\mathrm{Ext}^1(\Phi,\Psi)$, and the quotient projection. The six output panels are described below. The matrices R and S are auxiliary matrices for coordinate extraction and constant-direction diagnostics, respectively.

1. Quick start

Symbolic mode. Select the `None` checkbox to work in arbitrary positive characteristic without fixing a numerical prime. Computations may be performed over the rational function field $\mathbb{F}_q(\theta)$, where q is an unspecified prime power, θ is an indeterminate, and $\tau a = a^q\tau$. Frobenius fixes $\mathbb{F}_q$ and sends θ to θ^q; its iterates are retained symbolically.

Concrete mode. Choose p and $q = p^{\mathtt{n}}$ to perform arithmetic over $\mathbb{F}_q(\theta)$ in the selected characteristic. The following steps describe a concrete run.

1. For a concrete computation, leave the `None` checkbox unselected. Set `prime` to a prime p and `n` to a positive integer. The Frobenius is $q = p^{\mathtt{n}}$; the coefficient characteristic is p.
2. Set the matrix sizes: Φ_t is $d \times d$ and Ψ_t is $e \times e$. These sizes are independent of the exponent `n` above.
3. Enter the two matrices in the mathematical editors or the corresponding `LaTeX` boxes. For the scalar example on page 5, take $p = 3$, $\mathtt{n} = 1$, $d = e = 1$ and enter:

```
Phi_t:  \begin{pmatrix}\theta+\tau^2\end{pmatrix}
Psi_t:  \begin{pmatrix}\theta+\tau\end{pmatrix}
```

4. Click `Run` and wait for `COMPLETE`. After an input change, `RESULTS STALE` means the displayed matrices belong to an earlier calculation. Run again before using or copying them.

2. The output panels at a glance

Write $m = \dim \mathrm{Ext}^1(\Phi,\Psi)$, $m_0 = \dim \mathrm{Ext}^1_0(\Phi,\Psi)$ and $h = m - m_0$. These are dimensions of the underlying additive groups.

Output label	Size	Meaning
Output 1: Ext1 Matrix	$m \times m$	Π_t: the t-action on full normal coordinates of $\mathrm{Ext}^1(\Phi,\Psi)$.
Output 2: Ext0 Matrix	$m_0 \times m_0$	$\Pi_{0,t}$: the induced t-action on $\mathrm{Ext}^1_0(\Phi,\Psi)$.
Output 3: K Matrix	$m \times m_0$	K: inclusion; y in Ext0 coordinates becomes Ky in full coordinates.
Output 4: R Matrix	$m_0 \times m$	R: auxiliary matrix for extracting Ext0 coordinates; $RK = I_{m_0}$.
Output 5: P Matrix	$h \times m$	P: quotient projection; a full coordinate vector x lies in Ext0 exactly when $Px = 0$.
Output 6: S Matrix	$e \times d$	$S = \varepsilon$: auxiliary matrix recording zero/nonzero constant inner directions.

Use `Copy As LaTeX` next to an output to copy its matrix. The `Compatible Basis` panel displays U and U^{-1}, used to construct the reduction. The `Leaf Count` and `Byte Count` metrics measure expression size and storage, not dimensions or ranks.

3. Algebra, assumptions and normal coordinates

The coefficient convention

Let L be a coefficient field of characteristic p containing $\mathbb{F}_q$, with characteristic parameter $\theta \in L$. The symbol t is the generator of $\mathbb{F}_q[t]$; θ is its image under the characteristic map. The operator τ acts by $c \mapsto c^q$.

Calculations take place in the Ore ring $L\{\tau\}$, with coefficients on the left:

$$\tau a = a^q \tau, \qquad (a\tau^i)(b\tau^j) = a\, b^{q^i} \tau^{i+j}.$$

Thus $\sum_j a_j \tau^j$ represents the additive polynomial $c \mapsto \sum_j a_j c^{q^j}$. Products of output action matrices mean composition of additive maps, using this Ore rule.

Admissible input for this reduction

Write

$$\Phi_t = \theta I_d + N_\Phi + \sum_{r \geq 1} A_r \tau^r, \qquad \Psi_t = \theta I_e + N_\Psi + \sum_{r \geq 1} B_r \tau^r.$$

The matrices N_Φ and N_Ψ must be nilpotent. The positive coefficient blocks of Φ_t must admit the compatible jump basis required by the reduction algorithm. It selects columns

$$u_s = A_{\nu_s} \mathbf{e}_s, \qquad U = (u_1 \ \cdots \ u_d), \qquad \nu_s \geq 1,$$

with U invertible and a compatible column ordering σ. The selected columns must respect the successive rank jumps of the positive coefficient blocks; an arbitrary invertible choice is insufficient. The degree hypothesis is

$$\deg_\tau \Psi_t < \min_{1 \leq s \leq d} \nu_s.$$

For a scalar $\Phi_t = \theta + \tau^a$, this reads $\deg_\tau \Psi_t < a$.

What is reduced

A biderivation is determined by its value $V = \delta(t) \in \mathrm{Mat}_{e \times d}(L\{\tau\})$. Its inner changes have the form

$$\delta_H(t) = H\Phi_t - \Psi_t H, \qquad H \in \mathrm{Mat}_{e \times d}(L\{\tau\}).$$

Accordingly, extension classes are represented by these matrices modulo inner changes. Under the reduction hypotheses, full normal representatives satisfy

$$\deg_\tau V_{is} < \nu_s.$$

The full normal coordinate family is

$$\mathcal{B} = \{E_{is}\tau^r : 1 \leq i \leq e,\ 1 \leq s \leq d,\ 0 \leq r < \nu_s\}, \qquad m = e \sum_{s=1}^{d} \nu_s.$$

The triples (i, s, r) are ordered lexicographically, with the degree varying fastest. Here E_{is} is an $e \times d$ elementary matrix.

How Output 1 is computed

The t-action on a class is $t \cdot [V] = [\Psi_t V]$. For each normal coordinate direction, the algorithm reduces $\Psi_t(cE_{is}\tau^r)$ with a formal coefficient c. If an output coordinate is $\sum_j a_j c^{q^j}$, the corresponding entry of Π_t is $\sum_j a_j \tau^j$. Keeping c formal is essential: substituting $c = 1$ would lose the Frobenius exponents.

4. Ext0 and the exact sequence

The constant-term obstruction

Define the L-linear map on constant matrices

$$D : \mathrm{Mat}_{e\times d}(L) \longrightarrow \mathrm{Mat}_{e\times d}(L), \qquad D(H) = HN_\Phi - N_\Psi H.$$

The constant coefficient of an inner change is

$$(H\Phi_t - \Psi_t H)_0 = D(H_0).$$

Consequently, the class of V_0 modulo $\operatorname{im} D$ depends only on the extension class. The constant-term map gives the exact sequence

$$0 \longrightarrow \mathrm{Ext}^1_0(\Phi,\Psi) \xrightarrow{K} \mathrm{Ext}^1(\Phi,\Psi) \xrightarrow{P} \operatorname{coker} D \longrightarrow 0.$$

Here $\mathrm{Ext}^1_0(\Phi,\Psi)$ consists of the classes admitting a representative with zero constant coefficient. For a *full normal representative*, the correct condition is

$$[V] \in \mathrm{Ext}^1_0(\Phi,\Psi) \quad\Longleftrightarrow\quad V_0 \in \operatorname{im} D.$$

Its normal constant coefficient need not itself be zero: subtracting a constant inner change removes it but can introduce higher powers of τ.

Dimensions and the basis of Ext0

Put $\rho = \operatorname{rank}_L D$ and $h = ed - \rho$. Since D is an endomorphism of an ed-dimensional space, h also equals $\dim_L \ker D$, although the quotient in the sequence is $\operatorname{coker} D$.

$$\boxed{m_0 = e\sum_s(\nu_s - 1) + \rho = m - h.}$$

The Ext0 coordinate basis consists of a basis of $\operatorname{im} D$ in degree zero, followed by all positive full normal monomials $E_{is}\tau^r$ with $1 \le r < \nu_s$. These positive monomials retain their order from $\mathcal{B}$. In particular, the inclusion K may mix constant coordinates.

The action on the quotient

The quotient inherits the action

$$[C] \longmapsto [(\theta I_e + N_\Psi)C], \qquad [C] \in \mathrm{Mat}_{e\times d}(L)/\operatorname{im} D.$$

In the selected quotient coordinates, write this matrix as Q_t. It is constant in τ and has the form $\theta I_h +$ a nilpotent matrix. It equals θI_h when $N_\Phi = N_\Psi = 0$, but it can be nonscalar for nilpotent input parts.

Thus, in additive coordinates, the sequence is

$$0 \longrightarrow (\mathbb{G}_a^{m_0}, \Pi_{0,t}) \xrightarrow{K} (\mathbb{G}_a^{m}, \Pi_t) \xrightarrow{P} (\mathbb{G}_a^{h}, Q_t) \longrightarrow 0.$$

The six output panels do not include a separate Q_t panel. The auxiliary matrix S records constant-direction diagnostics.

The auxiliary matrix R

For $x \in \ker P$, Rx gives the unique Ext0 coordinate vector with $K(Rx) = x$. For a general x, it extracts the Ext0 component of a chosen additive decomposition. The matrix R need not commute with the t-actions, so its existence does not assert a splitting of the exact sequence as t-modules.

5. Coordinate maps and auxiliary matrices

How K, R and P fit together

For the constant calculation, $\mathrm{vec}(C)$ stacks the columns of an $e \times d$ matrix. This order differs from the row/column/degree order of the full normal family $\mathcal{B}$; constant coordinates are inserted at their corresponding positions.

Let W have columns forming a basis of $\operatorname{im} D$ in column-vectorized coordinates, and choose a complementary matrix Z so that $(W\ Z)$ is invertible. Define

$$\begin{pmatrix} C_{\mathrm{im}} \\ P_{\mathrm{c}} \end{pmatrix} = (W\ Z)^{-1}.$$

If $c = \mathrm{vec}(V_0)$ and x_+ lists the positive normal coordinates, the full coordinate maps are characterized by

$$\begin{aligned} K(y_{\mathrm{im}}, y_+): \quad (c, x_+) &= (W y_{\mathrm{im}}, y_+), \\ Rx = (C_{\mathrm{im}} c, x_+), &\qquad Px = P_{\mathrm{c}} c. \end{aligned}$$

Let Z_{full} insert the columns of Z into the constant positions of the full coordinates. It is an additive section of P. The following identities express the coordinate decomposition and exactness:

$$\begin{gathered} PK = 0, \qquad RK = I_{m_0}, \qquad PZ_{\mathrm{full}} = I_h, \\ KR + Z_{\mathrm{full}} P = I_m. \end{gathered}$$

Neither Z_{full} nor R is required to be a t-module morphism.

Actions and Frobenius twists

The inclusion and quotient projection intertwine the t-actions:

$$\Pi_t K = K \Pi_{0,t}, \qquad P \Pi_t = Q_t P.$$

All products involving τ use Ore multiplication. In particular, if $\Pi_t = \sum_j \Pi_j \tau^j$, then

$$\boxed{\Pi_{0,t} = R \Pi_t K = \sum_j R \Pi_j K^{(q^j)} \tau^j.}$$

The superscript (q^j) means entrywise Frobenius. Ordinary coefficientwise multiplication would omit these twists when K has coefficients outside the Frobenius-fixed field. Similarly, $Q_t = P \Pi_t Z_{\mathrm{full}}$ as an Ore identity.

Output 6: the auxiliary matrix S

For the compatible basis U, set $H_{is} = E_{is} U^{-1}$. Then

$$S_{is} = \varepsilon_{is} = \begin{cases} 0, & H_{is} N_\Phi - N_\Psi H_{is} = 0, \\ 1, & H_{is} N_\Phi - N_\Psi H_{is} \neq 0. \end{cases}$$

The directions with $S_{is} = 0$ give an independent partial basis of $\ker D$. Linear combinations of directions with nonzero images can supply additional kernel vectors. Therefore counting zeros of S need not give h, and S alone cannot determine m_0. Use the rank of the complete map D. The J_2 example on page 6 exhibits this distinction.

6. A complete scalar example

Choose $p = 3$, $\mathtt{n} = 1$, hence $q = 3$, and take

$$\Phi_t = \theta + \tau^2, \qquad \Psi_t = \theta + \tau.$$

The compatible basis is $U = (1)$, the jump degree is $\nu_1 = 2$, and the full normal family is $(1, \tau)$. A full coordinate vector $(x_0, x_1)^T$ represents the class of $x_0 + x_1\tau$.

Since $N_\Phi = N_\Psi = 0$, the map D vanishes. Thus

$$m = 2, \qquad \rho = 0, \qquad h = 1, \qquad m_0 = 1,$$

and the Ext0 coordinate family is (τ).

The six outputs

$$\Pi_t = \begin{pmatrix} \theta & 0 \\ \tau & \theta + \tau^2 \end{pmatrix}, \quad \Pi_{0,t} = \begin{pmatrix} \theta + \tau^2 \end{pmatrix},$$

$$K = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \qquad R = \begin{pmatrix} 0 & 1 \end{pmatrix},$$

$$P = \begin{pmatrix} 1 & 0 \end{pmatrix}, \qquad S = \begin{pmatrix} 0 \end{pmatrix}.$$

In particular,

$$t \cdot \begin{pmatrix} x_0 \\ x_1 \end{pmatrix} = \begin{pmatrix} \theta x_0 \\ x_0^q + \theta x_1 + x_1^{q^2} \end{pmatrix}.$$

The map P retains the constant coordinate, and K inserts an Ext0 coordinate as $(0, y)^T$.

Why the entry τ^2 occurs

For the second coordinate, begin with $V = c\tau$. Its t-multiple is

$$\Psi_t(c\tau) = \theta c\tau + c^q\tau^2.$$

The inner change associated to the constant $H = c^q$ is

$$H\Phi_t - \Psi_t H = c^q\tau^2 - c^{q^2}\tau.$$

Subtracting it leaves $(\theta c + c^{q^2})\tau$. Hence the second diagonal entry of Π_t is $\theta + \tau^2$.

The exact sequence and the auxiliary matrix R

Here $Q_t = (\theta)$ and

$$0 \longrightarrow (\mathbb{G}_a, \theta + \tau^2) \xrightarrow{\binom{0}{1}} (\mathbb{G}_a^2, \Pi_t) \xrightarrow{(1 \;\; 0)} (\mathbb{G}_a, \theta) \longrightarrow 0.$$

Although $RK = 1$, in this example

$$R\Pi_t = (\tau \;\; \theta + \tau^2) \neq (0 \;\; \theta + \tau^2) = \Pi_{0,t}R.$$

Thus the auxiliary matrix R is a coordinate left inverse of the inclusion, but is not itself a t-module retraction.

7. Matrix examples and the partial basis issue

The pair with J_2 in both constant parts

Let $q = 3$ and

$$J_2 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \qquad \Phi_t = (\theta + \tau^2)I_2 + J_2, \qquad \Psi_t = (\theta + \tau)I_2 + J_2.$$

Here $U = I_2$, $\nu_1 = \nu_2 = 2$ and $m = 8$. For $H = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$,

$$D(H) = HJ_2 - J_2H = \begin{pmatrix} -c & a - d \\ 0 & c \end{pmatrix}.$$

Therefore $\rho = 2$ and

$$\ker D = \left\{ \begin{pmatrix} a & b \\ 0 & a \end{pmatrix} : a, b \in L \right\}, \qquad h = 2, \qquad m_0 = 6.$$

The auxiliary matrix S is

$$S = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}.$$

It has only one zero: E_{12} is the only elementary matrix in $\ker D$. The additional kernel direction $I_2 = E_{11} + E_{22}$ is a linear combination of directions whose individual images under D are nonzero. Counting zeros of S would therefore give the wrong quotient dimension.

A basis of $\operatorname{im} D$ can be taken as E_{12} and $-E_{11} + E_{22}$. The Ext0 basis combines these two constant directions with the four positive directions $E_{is}\tau$. In the quotient coordinates used by the checked output,

$$Q_t = \begin{pmatrix} \theta & 0 \\ 1 & \theta \end{pmatrix}.$$

The quotient has underlying additive group $\mathbb{G}_a^2$, with a nonscalar t-action. Different valid basis choices can change the displayed matrices while representing isomorphic coordinate descriptions.

An upper-right entry $1 + \tau$

Another admissible input at $q = 3$ is

$$\Phi_t = \theta + \tau^2, \qquad \Psi_t = \begin{pmatrix} \theta + \tau & 1 + \tau \\ 0 & \theta + \tau \end{pmatrix}.$$

Set the matrix sizes to $d = 1$ and $e = 2$. Here $N_\Phi = 0$, $N_\Psi = J_2$, and

$$D \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} -b \\ 0 \end{pmatrix}.$$

Thus $m = 4$, $\rho = 1$, $h = 1$ and $m_0 = 3$. A normal representative is in Ext0 exactly when its constant column lies in the span of $(1, 0)^T$. The positive off-diagonal term τ affects Π_t and $\Pi_{0,t}$; the constant obstruction D uses only N_Φ and N_Ψ.