Computing direct sum decompositions and applications in algebraic geometry

Devlin Mallory
Basque Center for Applied Mathematics

Mahrud Sayrafi
McMaster University

June 27, 2025

The plan

My goal today is to present a computational algorithm for determining indecomposable decompositions of coherent sheaves or modules; this algorithm is already implemented in Macaulay2 and available publicly.

I will try to motivate the question for the audience today: why would an algebraic (birational?) geometer care about decomposing coherent sheaves or modules?

The plan is as follows:

1. Introduce the question of determining direct sum decompositions of coherent sheaves or modules.
2. Discuss some example applications of determining direct sum decompositions in algebraic geometry.
3. Present our algorithm.
4. Give some applications and related questions.

Indecomposable decompositions

Definition

- A sheaf E on a variety X is *indecomposable* if it cannot be written as a nontrivial direct sum $E = E_1 \oplus E_2$.
- Likewise, a module M over a ring R is indecomposable if it cannot be written as a nontrivial direct sum $M = M_1 \oplus M_2$.
- An *indecomposable decomposition* is an expression $M = M_1 \oplus M_2 \oplus \cdots \oplus M_r$ with each M_i indecomposable.

Coherent sheaves (or finitely generated modules) always have indecomposable decompositions.

If X is projective (or R is graded), then coherent sheaves over X (or f.g. modules over R) satisfy the Krull–Schmidt property: the summands are unique up to isomorphism.

(The same holds if R is *complete* local, but not all local rings satisfy the Krull–Schmidt property!)

A toy example

Consider the subvarieties of $\mathbb{P}^5$ defined by

$$X = V(x_4^2 - x_3x_5, x_2x_4 - x_1x_5, x_1x_4 - x_0x_5, x_0x_5 - x_2x_3, x_0x_4 - x_1x_3, x_1^2 - x_0x_2)$$

and

$$Y = V(x_4^2 - x_3x_5, x_2x_4 - x_1x_5, x_1x_4 - x_2x_3, x_0x_5 - x_2^2, x_0x_4 - x_1x_2, x_1^2 - x_0x_3)$$

Computer algebra systems can tell you many things about these varieties: their (co)tangent sheaves, that they're smooth, etc.

They can even tell you that both are Fano of dimension 2 with $(-K_X)^2 = 8$, so X, Y must be either $\mathrm{Bl}_p \mathbb{P}^2$ or $\mathbb{P}^1 \times \mathbb{P}^1$.

But how can you tell which is which, or if $X \cong Y$?

A toy example, continued

Say that a computer algebra system could tell you that

$$T_X = M_1 \oplus M_2,$$

while

$$T_Y \text{ is indecomposable.}$$

Then we would know that $X \cong \mathbb{P}^1 \times \mathbb{P}^1$ and $Y \cong \mathrm{Bl}_p \mathbb{P}^2$.

(In fact, Beauville showed that if X is a Fano surface and T_X is a direct sum, then $X = X_1 \times X_2$. Note also that T_X a direct sum doesn't always imply X is a product; consider abelian varieties!)

Our results

With Mahrud Sayrafi, we provide algorithms answering the following questions:

- If a coherent sheaf or f.g. module is decomposable, how to find its summands?
- How to verify if a coherent sheaf is indecomposable?

These questions are mentioned as outstanding in the literature, e.g., in Chapter 15 of Eisenbud's *Commutative Algebra*, and to our knowledge were not yet answered.

In fact, we'll see that the answer requires little more than linear algebra!

Our algorithms are implemented already in Macaulay2, and are available for anyone interested in decomposing modules or sheaves.

First, however, a bit more motivation and context.

Related notions

Some related notions, which we mention but won't use:

Definition

A coherent sheaf E on an irreducible variety X is:

- μ-stable if for any proper subsheaf $F \subset E$ we have $\mu(F) < \mu(E)$.
- simple if $\operatorname{End} E \cong k$.

Simple $\implies$ μ-stable $\implies$ indecomposable (and not conversely!)

Simplicity is the most straightforward to check algorithmically. Crucially, not all indecomposable sheaves are simple: for example, if X is an elliptic curve and E the unique nonsplit extension of $\mathcal{O}_X$ by itself, then E is indecomposable but $\operatorname{End} E = k^2$.

I don't know any algorithms for μ-stability (but would love to!).

Indecomposable summands of sheaves on $\mathbb{P}^n$

It is already very interesting to ask about finding line bundle summands of sheaves on $\mathbb{P}^n$.

Theorem (Horrock's)

Let E be a vector bundle on $\mathbb{P}^n$. Then E is the direct sum of line bundles if and only if $H^i(\mathbb{P}^n, E(d)) = 0$ for all $0 < i < n$ and $d \in \mathbb{Z}$.

Conjecture (Hartshorne)

Let E be a vector bundle of rank 2 on $\mathbb{P}^7$. If $n \geq 6$, then E is the direct sum of line bundles.

Remark

A consequence of Hartshorne's conjecture would be that a smooth subvariety of $\mathbb{P}^n$ of codimension 2 is a complete intersection. The statement is false for $n = 4$: there is the famous Horrock–Mumford bundle, which is indecomposable of rank 2, and a general section of which gives an abelian surface in $\mathbb{P}^4$.

Some extensions and limitations of Horrock's criteria

Horrock's-type criteria, which describe indecomposability in terms of some cohomological vanishing, exist for:

- quadrics and Grassmannians (Ottaviani)
- certain toric varieties (Eisenbud–Erman–Schreyer, Sayrafi–Brown, Sayrafi).

However, all these results use special properties of such varieties, e.g., the existence of a nice resolution of the diagonal $X \subset X \times X$, and the information this yields about $D^b_{\mathrm{coh}}(X)$.

In general, it seems quite hard to describe indecomposable summands via purely cohomological information.

For example, on any elliptic curve, there are vector bundles E_1, E_2 such that $H^i(E_1(d)) = H^i(E_2(d))$ for all i, d, but E_1 is indecomposable and E_2 is not.

Frobenius summands

A huge source of motivation for computing indecomposable decompositions comes from the Frobenius morphism.
Let X be a variety over a field of characteristic p, and write $F^e : X \to X$ for the e-th iterated Frobenius morphism.

Definition

The indecomposable summands of $F^e_* \mathcal{O}_X$ for $e \geq 1$ are called *Frobenius summands* of X.

The Frobenius summands of X carry a great deal of information about the global properties of X. For example:

Theorem (Smith)

If $\mathcal{O}_X$ is a Frobenius summand of X (i.e., if X is globally F-split) then X satisfies Kodaira vanishing.

Curves

Let X be a smooth curve of genus g. There is then a trichotomy:

- If $g = 0$, then $F^e_* \mathcal{O}_X = \mathcal{O}_X \oplus \mathcal{O}_X(-1)^{p^e-1}$, so the only Frobenius summands are $\mathcal{O}_X$ and $\mathcal{O}_X(-1)$.
- If $g = 1$, then $F^e_* \mathcal{O}_X$ is the direct sum of the p^e-torsion line bundles if X is ordinary, and is indecomposable otherwise.

  Thus, if X is ordinary then the Frobenius summands are all the p^e-torsion line bundles for $e \geq 1$, and if X is supersingular then the Frobenius summands are certain indecomposable vector bundles of rank $p, p^2, \ldots$.
- If $g \geq 2$, then $F^e_* \mathcal{O}_X$ is indecomposable of rank p^e.

So, the Frobenius summands might reflect the difference between general-type and Fano varieties.

Abelian varieties

More generally, let X be an abelian variety in characteristic p.

Definition

The p-rank of X is $r_X = \dim_{\mathbb{F}_p} \operatorname{Hom}(\mu_p, A[p])$, where μ_p is the group scheme of p-th roots of unity,

Thus, p^{r_X} is the number of p-torsion points of A (over $\bar{k}$), so $r_X = \dim X$ if X is ordinary, and $r_X = 0$ if X is supersingular.

Theorem (Sannai–Tanaka)

$F_*^e \mathcal{O}_X$ is the direct sum of indecomposable vector bundles of rank $p^{e(\dim X - r_X)}$.

So, the Frobenius summands capture some delicate arithmetic information about X. In particular, the Frobenius summands are all line bundles if and only if X is ordinary.

Toric varieties

The decomposition of F_*L for other line bundles L can also help us understand X. For example:

Theorem (Achinger)

Let X be a variety over a field of characteristic p such that F_*L is the direct sum of divisorial sheaves for any invertible sheaf L. Then X is a toric variety.

(The converse is also true, and much easier: if X is toric, then F_*L is the direct sum of divisorial sheaves whenever L is.)
Note also:

Theorem (Sannai–Tanaka)

Let X be a variety of nonnegative Kodaira dimension. If $F_*^e\mathcal{O}_X$ is a direct sum of line bundles for all e, then X is an abelian variety.

So, hopefully this gives some motivation for the question of finding indecomposable summands of coherent sheaves or modules.

Finding indecomposable decompositions

So, if you have a presentation of a module M over a ring R, how do you find indecomposable summands of M?

Our algorithms take in either:

- A homogeneous, finitely generated module M over a (multi)graded k-algebra R, or
- A finitely generated module M over a local k-algebra $(R, \mathfrak{m})$ with $k \subset \overline{\mathbb{F}}_p$.

and returns a decomposition

$$M = M_1 \oplus M_2 \oplus \cdots \oplus M_r,$$

with each M_i (probabilistically) indecomposable.

Applications to coherent sheaves

This yields an algorithm for finding indecomposable summands of a coherent sheaf F on an embedded projective variety $X \subset \mathbb{P}^N$:

If $R = \bigoplus_{d \geq 0} H^0(X, \mathcal{O}_X(d))$ and $\Gamma_*(F) := \bigoplus_{\mathbb{Z}} H^0(X, F(d))$, sheafifying an indecomposable decomposition of $\Gamma_*(F)$ yields a direct sum decomposition of F into indecomposable sheaves.

(Warning: it is possible to have M a graded module with $\widetilde{M}$ decomposable but M indecomposable. This is fixed by replacing M by $\Gamma_*(\widetilde{M})$.)

Our algorithm works equally well over multigraded rings, including the Cox rings of toric varieties, which allows for computing indecomposable summands of coherent sheaves on (non-toric!) subvarieties of toric varieties.

The local case: via idempotents

Both algorithms essentially reduce to linear algebra. Let me first sketch the local case, because it's a bit simpler in some ways.

If M is a module over a ring R, a direct sum decomposition $M = M_1 \oplus M_2$ is equivalent to an idempotent endomorphism, i.e., $\varphi \in \mathrm{End}_R(M)$ such that $\varphi^2 = \varphi$: if φ is idempotent, then the short exact sequence

$$0 \to \ker \varphi \to M \xrightarrow{\varphi} \operatorname{im} \varphi \to 0$$

is split by the inclusion $\operatorname{im} \varphi \hookrightarrow M$.

Reduction to linear algebra

Let $(R, \mathfrak{m})$ be a local ring. The following lemma is an easy consequence of Artin–Rees and the fact that surjectivity can be checked after completion:

Lemma

Let $\varphi \in \operatorname{End}_R(M)$. If the induced map $\bar{\varphi} : M/\mathfrak{m}M \to M/\mathfrak{m}M$ is idempotent, then $M = \ker \varphi \oplus \operatorname{im} \varphi$.

So, it suffices to produce $\varphi \in \operatorname{End}_R(M)$ such that $\bar{\varphi}$ is idempotent. (Note: φ may not be an idempotent!)

Warning: there are many idempotent elements of $\operatorname{End}_{R/\mathfrak{m}}(M/\mathfrak{m}M)$, since $M/\mathfrak{m}M$ is a vector space, but most will not lift to $\operatorname{End}_R(M)$.

Thus, to produce φ, we want to only do operations that "lift to $\operatorname{End}_R(M)$".

A trick

Proposition

Let $k \subset \bar{\mathbb{F}}_p$ and let A be an endomorphism of a k-vector space. For e sufficiently large, $A^{p^e(p^e-1)}$ is idempotent.
Moreover, if λ is not the only eigenvalue of A over $\bar{k}$, then $(A - \lambda \cdot \mathrm{id})^{p^e(p^e-1)}$ is a nontrivial idempotent.

Combining with the previous lemma, we have

Corollary

Let $(R, \mathfrak{m})$ be a local ring with $k := R/\mathfrak{m} \subset \bar{\mathbb{F}}_p$ and let M be a finitely generated R-module. If there is an endomorphism $\varphi \in \operatorname{End}_R M$ such that $A = \bar{\varphi}$ has multiple eigenvalues over $\bar{k}$, then M is a direct sum $\ker \varphi^{p^e(p^e-1)} \oplus \operatorname{im} \varphi^{p^e(p^e-1)}$.

Proof

We may assume $k = \bar{k}$, since if A is idempotent over $\bar{k}$, it is idempotent over k. Thus, we can assume A is in Jordan canonical form. The n-th power of an $r \times r$ Jordan matrix is easy to compute: if $\lambda \neq 0$, then

$$\begin{pmatrix} \lambda & 1 & 0 & \dots & 0 \\ 0 & \lambda & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \lambda \end{pmatrix}^n = \begin{pmatrix} \lambda^n & \binom{n}{1}\lambda^{n-1} & \binom{n}{2}\lambda^{n-2} & \dots & \binom{n}{r}\lambda^{n-r} \\ 0 & \lambda^n & \binom{n}{1}\lambda^{n-1} & \dots & \binom{n}{r-1}\lambda^{n-r+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \lambda^n \end{pmatrix}$$

If $n_0 = p^e > r$, then $p \mid \binom{p^e}{r}$, and so all off-diagonal terms vanish, leaving a diagonal matrix with λ^{n_0} on the diagonal.

If e is taken large enough that $\lambda \in \mathbb{F}_{p^e}$, then $(\lambda^n)^{p^e-1} = 1$, so this block is the identity.

Taking e large enough to do this for every Jordan block of A, we obtain a power of A that is idempotent.

The local algorithm, summarized

Thus, we have our algorithm:

1. Pick a general endomorphism $\varphi \in \mathrm{End}_R(M)$.
2. Let $\lambda \in k$ be an eigenvalue of $\bar{\varphi}$; replace φ by $\varphi - \lambda \cdot \mathrm{id}$.
3. Set $A = \bar{\varphi} : M/\mathfrak{m}M \to M/\mathfrak{m}M$, and let e be such that $A^{p^e(p^e-1)}$ is idempotent.
4. If $A^{p^e(p^e-1)}$ is not zero, then set $M_1 = \ker \varphi^{p^e(p^e-1)}$ and $M_2 = \mathrm{im}\, \varphi^{p^e(p^e-1)}$; then $M = M_1 \oplus M_2$.
5. Repeat for M_1 and M_2.

This algorithm relies on picking an endomorphism φ with multiple eigenvalues. I will describe later why, if M is decomposable, a random φ satisfies this with sufficiently high probability.

(In practice, we split M into the eigenspaces of powers of $\varphi - \lambda_i \cdot \mathrm{id}$ for each eigenvalue λ_i of $\bar{\varphi}$; this requires fewer recursions.)

Note: you may need to extend the field k in order to decompose M.

The graded algorithm

For a graded ring R and a homogeneous R-module M, write $\mathrm{End}_R^0(M)$ for the vector space of degree-0 endomorphisms of M.

We pick a random element $\varphi \in \mathrm{End}_R^0(M)$.

The key idea is to show if M has r generators and $\varphi \in \mathrm{End}_R^0(M)$ then $\ker \varphi \subset \ker \varphi^2 \subset \dots$ stabilizes in at most r steps.

We claim that if $\ker \varphi^i = \ker \varphi^{i+1}$, then $M = \ker \varphi^i \oplus \operatorname{im} \varphi^i$.

By replacing φ by $\varphi - \lambda \cdot \mathrm{id}$ for some eigenvalue λ of φ, we may assume that $\ker \varphi^i \neq 0$. If φ has multiple eigenvalues then this yields a nontrivial decomposition.

We will describe later why, if M is decomposable, a random φ satisfies this with sufficiently high probability, and explain what we mean by "eigenvalue" of φ.

(Again, one may need to extend the base field so that $\lambda \in k$.)

Kernel stabilization $\Longrightarrow$ splitting

Key point: if if $\ker \varphi^i = \ker \varphi^{i+1}$, then φ is injective on $\operatorname{im} \varphi^i$, and thus $\varphi^i|_{\operatorname{im} \varphi^i} : \operatorname{im} \varphi^i \to \operatorname{im} \varphi^i$ is injective as well.

Crucially, a degree-0 injective endomorphism of a finitely generated homogeneous module N is an isomorphism (it is injective, hence surjective, on each finite-dimensional N_d for each degree d).

Since $\varphi^i|_{\operatorname{im} \varphi^i} : \operatorname{im} \varphi^i \to \operatorname{im} \varphi^i$ is an isomorphism, we can split the short exact sequence

$$0 \to \ker \varphi^i \to M \xrightarrow{\varphi^i} \operatorname{im} \varphi^i \to 0.$$

(This step is the only place that we need the homogeneity assumption.)

Stabilization of kernels

Linear algebra says that if A is an endomorphism of a vector space V of dimension r, then $\ker A \subset \ker A^2 \subset \dots$ stabilizes in at most r steps. However, $\ker \psi/\mathfrak{m} \cdot \ker \psi \neq \ker \bar{\psi}$, so this doesn't immediately apply to our case.

Instead, we use a trick. Cayley–Hamilton over the ring R says that

$$0 = \chi_\varphi(\varphi) = \prod (\varphi - \lambda_i \cdot \mathrm{id})^{\mu_i},$$

where the $\lambda_i \in k$ are the eigenvalues of φ and μ_i their multiplicities.

(One has to be slightly careful: *a priori* the eigenvalues, i.e., the zeroes of $\det(\varphi - t \cdot \mathrm{id}) \in R[t]$, live in $\overline{R}$; however, the degree-0 condition forces them to lie in k.)

Stabilization of kernels, continued

Replacing φ by $\varphi - \lambda_i \cdot \mathrm{id}$, we may assume that 0 is an eigenvalue, so

$$0 = \chi_\varphi(\varphi) = \prod(\varphi - \lambda_i \cdot \mathrm{id})^{\mu_i} \cdot \varphi^{\mu_0},$$

We write $\prod(\varphi - \lambda_i \cdot \mathrm{id})^{\mu_i} = \varphi^{r-\mu_0} + \cdots + c_0$, with $c_0 \in k$ nonzero, and thus

$$\begin{aligned}\varphi^{\mu_0} &= c_0^{-1}\left(\varphi^{\mu(M)-\mu_0} - \cdots + c_1\varphi\right) \circ \varphi^{\mu_0} \\ &= c_0^{-1}\left(\varphi^{\mu(M)-\mu_0-1} - \cdots + c_1\right) \circ \varphi \circ \varphi^{\mu_0}.\end{aligned}$$

Thus, in particular, φ must be injective on the image of φ^{μ_0}, and so the kernel stabilizes after $\mu_0 \leq r$ steps.

The graded algorithm, summarized

Thus, we have our algorithm:

1. Pick a general endomorphism $\varphi \in \mathrm{End}_R^0(M)$.
2. Let $\lambda \in k$ be an eigenvalue of φ; replace φ by $\varphi - \lambda \cdot \mathrm{id}$.
3. Calculate the characteristic polynomial $\chi_\varphi(t)$ and let μ_0 be the multiplicity of the eigenvalue 0 in $\chi_\varphi(t)$.
4. If φ^{μ_0} is not zero, then set $M_1 = \ker \varphi^{\mu_0}$ and $M_2 = \mathrm{im}\, \varphi^{\mu_0}$; then $M = M_1 \oplus M_2$.
5. Repeat for M_1 and M_2.

Just as in the local case, in practice, we split M into the eigenspaces of powers of $\varphi - \lambda_i \cdot \mathrm{id}$ for each eigenvalue λ_i of φ.

Probabilistic indecomposability

In either the graded or local case then, from a randomly chosen $\varphi \in \mathrm{End}_R(M)$ *with distinct eigenvalues* we can produce a nontrivial direct sum decomposition of M.

For this to terminate in an indecomposable decomposition, we need a randomly chosen endomorphism of a decomposable module M to have distinct eigenvalues with high probability.

If M is decomposable, there is *an* endomorphism $\varphi \in \mathrm{End}_R(M)$ with distinct eigenvalues: take φ to be the projection to a summand.

It is also not hard to show that the locus of endomorphisms $\varphi \in \mathrm{End}_R(M)$ with distinct eigenvalues is Zariski-open.

Over an infinite field things are great: a randomly chosen endomorphism $\varphi \in \mathrm{End}_R(M)$ will thus have distinct eigenvalues with probability 1, yielding a nontrivial decomposition of M.

Probability over finite fields

Over a finite field, however, the situation is more delicate, as you can have nonempty Zariski-open sets of $\mathbb{A}^n_{\mathbb{F}_q}$ containing very few points.

All potential splittings are obtained as k-linear combinations of

- minimal generators of $\mathrm{End}_R(M)$ in the local case, or
- a basis of $\mathrm{End}^0_R(M)$ in the graded case.

If N is the number of generators or the dimension of the basis, then we can think of $\mathbb{A}^N_k$ as the parameter space for our choice of endomorphism. The set of endomorphisms with distinct eigenvalues is a Zariski-open subset $U \subset \mathbb{A}^N_k$, with complement Z.

To estimate the probability that a randomly chosen endomorphism $\varphi \in \mathrm{End}_R(M)$ has distinct eigenvalues, we need to understand the degrees of the equations defining Z.

Finding the equations for Z

Say M has m generators, and let let $\varphi_1, \ldots, \varphi_N$ be a basis for $\mathrm{End}_R^0(M)$ in the graded case or minimal generators for $\mathrm{End}_R(M)$ in the local case.
A general endomorphism is of the form

$$\varphi = a_1\varphi_1 + a_2\varphi_2 + \cdots + a_N\varphi_N,$$

and the eigenvalues of φ are the roots of the polynomial

$$\begin{aligned}\det(\varphi - t \cdot \mathrm{id}) &= \det(a_1\varphi_1 + a_2\varphi_2 + \cdots + a_N\varphi_N - t \cdot \mathrm{id}) \\ &= t^m + c_{m-1}t^{m-1} + \cdots + c_0,\end{aligned}$$

which is a univariate polynomial in t with coefficients c_i a polynomial of degree $m - i$ in the a_j.

We want to find polynomial conditions guaranteeing that the polynomial $\det(\varphi - t \cdot \mathrm{id})$ is the m-th power of a linear form.

The m-fold root locus in the c_i coordinates

We can write the conditions defining Z in terms of the c_i: equating

$$t^m + c_{m-1}t^{m-1} + \cdots + c_0 = (t-y)^m,$$

we obtain (weighted-homogeneous) equations

$$c_{m-i} = (-1)^i \binom{m}{i} y^i.$$

If $p \nmid m$, eliminating y gives us the (radical) ideal generated by the m elements

$$\binom{m}{m-i} \cdot c_{m-1}^i \pm m^i \cdot c_{m-i}..$$

If instead $m = p^d m_0$, then the only possible nonzero coefficients are c_{m-ip^d} for $i = 1, \ldots, m_0$, yielding the radical ideal generated by the m_0 elements

$$\binom{m}{m-ip^d} \cdot c_{m-p^d}^i \pm \binom{m}{m-p^d}^i c_{m-ip^d}$$

and the $m - m_0$ elements c_i for $p^d \nmid i$.

The m-fold root locus in the a_j coordinates

Since the c_i are themselves polynomials in the a_j, we have a set of equations in the a_i that defines the locus of endomorphisms with only one eigenvalue.

Note that since c_{m-i} is a degree-i polynomial in the a_j, these equations are homogeneous of degree $\leq m$ in the a_j.

Moreover, if M is decomposable, there is an endomorphism with distinct eigenvalues, the resulting homogeneous equations in the a_j are not all zero, i.e., $Z \neq \mathbb{A}_k^N$.

We thus have that Z is contained inside a homogeneous hypersurface of degree $\leq m$ in $\mathbb{A}_k^N$.

Write $\widetilde{Z} \subset \mathbb{P}_k^{N-1}$ for the corresponding projective hypersurface.

Probability of distinct eigenvalues and Lang–Weil

$\widetilde{Z}$ is contained in an $(N-2)$-dimensional hypersurface of degree $\leq m$. We can thus apply a form of the Lang–Weil estimates (due to Ghorpade and Lachaud) to conclude that

$$|\widetilde{Z}(k)| \leq m \cdot |\mathbb{P}_k^{N-2}(k)| = m\frac{|k|^{N-1}-1}{|k|-1}.$$

Thus, the proportion of endomorphisms with exactly one eigenvalue is bounded above by

$$m \cdot \frac{|\mathbb{P}_k^{N-2}(k)|}{|\mathbb{P}_k^{N-1}(k)|} = m \cdot \frac{|k|^{N-1}-1}{|k|^N-1}$$

and so the proportion of endomorphisms with distinct eigenvalues is

$$\geq 1 - m \cdot \frac{|k|^{N-1}-1}{|k|^N-1} \approx 1 - m/|k|.$$

By enlarging the field k (based on the known value of m), we can ensure that this is as close to 1 as we like.

Applications

- Cuong, Dao, Eisenbud, Kobayashi, Polini and Ulrich used our algorithm to study the summands of syzygy modules over certain Artinian rings $(R, \mathfrak{m})$. They found that the syzygy modules of the residue field are direct sums of only three indecomposable modules: the residue field k, the maximal ideal $\mathfrak{m}$, and an additional module $N = \mathrm{Hom}_R(\mathfrak{m}, R)$.
- Finding Frobenius summands on non-toric Fano varieties (work in progress).
- A better understanding of the behavior of the Frobenius morphism for non-F-split rational double points in low characteristic (work in progress with Ilya Smirnov).
- More to come? If you have questions that would benefit from this algorithm, please let me know!

Remaining questions

- It would be great to have an algorithm for local rings in characteristic 0, or to better understand how to lift a direct sum decomposition from characteristic p to characteristic 0.
- Note that you can test indecomposability after reduction modulo p for a sufficiently large prime p; however, it is not clear how to lift a decomposition back to characteristic 0.
- The bottleneck in our algorithm is the computation of $\mathrm{End}_R(M)$. Given that one needs only random elements rather than a description of all generators and relations, it seems possible to circumvent this calculation.