Math 332 Winter 2023, Lecture 20: Modules

website: `https://www.cip.ifi.lmu.de/~grinberg/t/23wa`

2. Modules

2.4. Module morphisms (cont'd)

2.4.4. General properties of linearity

Fix a ring R.

Recall that R-module morphisms (aka R-linear maps) are maps between R-modules that respect all the relevant structure (addition, scaling and zero).

The following facts about R-module morphisms are analogues of known facts about ring morphisms:

Proposition 2.4.2. Let M and N be two left R-modules. Let $f : M \to N$ be an invertible left R-module morphism. Then, f is a left R-module isomorphism.

Proposition 2.4.3. A composition of two left R-module morphisms is again a left R-module morphism.

Proposition 2.4.4. A composition of two left R-module isomorphisms is again a left R-module isomorphism.

Proposition 2.4.5. The inverse of a left R-module isomorphism is again a left R-module isomorphism.

Proposition 2.4.6. The relation $\cong$ between left R-modules is an equivalence relation.

The proofs of all these propositions are similar to the analogous proofs for ring morphisms.

Again, there is an isomorphism principle: Any intrinsic property of an R-module M (that is, any property that does not depend on what the elements of M "are") automatically holds for any R-module isomorphic to M.

All of what we said about left R-modules holds equally well for right R-modules (because right R-modules are equivalent to left R^{op}-modules; see §3.1.4 in the text for details). In the future, this will go without saying. This gives us a good excuse to ignore right R-modules (at least until the point where we consider "bimodules": hybrid modules with a ring acting on the left and another ring acting on the right).

Let me also recall that if R is commutative, then we treat left R-modules and right R-modules as being the same thing (up to the notational issue of whether the scalar is written on the left or on the right of the vector), and we just call them "R-modules".

2.4.5. Kernels and images

Just like ring morphisms, module morphisms have kernels and images. These behave slightly better than those of ring morphisms.

We fix a ring R.

> **Definition 2.4.7.** Let R be a ring. Let M and N be two left R-modules. Let $f : M \to N$ be a left R-module morphism. Then, the **kernel** (aka **nullspace**) of f (denoted $\operatorname{Ker} f$ or $\ker f$) is defined to be the subset
>
> $$\operatorname{Ker} f := \{a \in M \mid f(a) = 0_N\}$$
>
> of M.

Some examples:

- Let R be a commutative ring. Let $b \in R$. Then, the map

  $$R \to R,$$
  $$r \mapsto br$$

  is an R-module morphism (called "multiplication by b"). Its kernel is

  $$\{r \in R \mid br = 0\}.$$

  If b is not zero and not a zero-divisor, then this kernel is $\{0\}$.

- Both $\mathbb{Z}^3$ and $\mathbb{Z} \times (\mathbb{Z}/2)$ are abelian groups, and thus are $\mathbb{Z}$-modules (since we have seen in §2.3 (in Lecture 19) that every abelian group is a $\mathbb{Z}$-module). The map

  $$\mathbb{Z}^3 \to \mathbb{Z} \times (\mathbb{Z}/2),$$
  $$(a,b,c) \mapsto \left(a-b,\ \overline{b-c}\right)$$

  is a $\mathbb{Z}$-module morphism (one among many). Its kernel is

  $$\begin{aligned} &\left\{(a,b,c) \in \mathbb{Z}^3 \mid \left(a-b,\ \overline{b-c}\right) = 0_{\mathbb{Z}\times(\mathbb{Z}/2)}\right\} \\ &= \left\{(a,b,c) \in \mathbb{Z}^3 \mid a-b = 0 \text{ and } \overline{b-c} = 0\right\} \\ &= \left\{(a,b,c) \in \mathbb{Z}^3 \mid a = b \text{ and } b \equiv c \bmod 2\right\}. \end{aligned}$$

From linear algebra, you should be familiar with some properties of kernels (aka nullspaces). They still hold in our more general context:

Theorem 2.4.8. Let M and N be two left R-modules. Let $f : M \to N$ be a left R-module morphism. Then:

(a) The kernel $\operatorname{Ker} f$ is an R-submodule of M.

(b) The image $\operatorname{Im} f = f(M)$ is a R-submodule of N.

Lemma 2.4.9. Let M and N be two left R-modules. Let $f : M \to N$ be a left R-module morphism. Then, f is injective if and only if $\operatorname{Ker} f = \{0_M\}$.

Again, the proofs are easy (and very similar to the analogous proofs for ring morphisms).

2.4.6. Quotient modules

Again, we fix a ring R.

Quotient modules are an analogue of quotient rings and quotient groups:[1]

Definition 2.4.10. Let M be a left R-module. Let I be a left R-submodule of M. Thus, I is a subgroup of the additive group $(M, +, 0)$, hence a normal subgroup (since any subgroup of an abelian group is normal). Thus, the quotient group M/I itself becomes an abelian group. Its elements are the cosets $a + I$ of I in M. We will denote such a coset $a + I$ by $\overline{a}$, and call it a **residue class**.

Note that the addition of this group M/I is given by

$$(a + I) + (b + I) = (a + b) + I \qquad \text{for all } a, b \in M,$$

i.e., by

$$\overline{a} + \overline{b} = \overline{a + b} \qquad \text{for all } a, b \in M.$$

Now, we define an action of R on M/I by setting

$$r(a + I) = ra + I \qquad \text{for all } r \in R \text{ and } a \in M,$$

i.e., by setting

$$r\overline{a} = \overline{ra} \qquad \text{for all } r \in R \text{ and } a \in M.$$

The set M/I, equipped with the addition and the action we just defined and with the element $0 + I = \overline{0}$ as the zero vector, is a left R-module. This left R-module is called the **quotient module** of M by the submodule I (or, for short, "M **modulo** I"). It is denoted by M/I.

Theorem 2.4.11. This is indeed a left R-module.

Proof. Easy. □

[1] We are using the letter I for a submodule here because submodules are, in a sense, a generalization of ideals (literally when R is commutative, roughly in the general case).

Theorem 2.4.12. Let I be a left R-submodule of a left R-module M. Then, the map

$$\pi : M \to M/I,$$
$$a \mapsto \overline{a} = a + I$$

is a surjective R-module morphism with kernel I. This morphism π is called the **canonical projection** from M to M/I.

Proof. Same as for rings, mutandis mutandis[2]. □

Examples of quotient modules come from various places:

- Quotients of abelian groups are instances of quotient modules, since abelian groups are $\mathbb{Z}$-modules.
- Quotients of vector spaces are instances of quotient modules, since vector spaces are modules over a field.

  For instance, consider the 3-dimensional vector space (i.e., $\mathbb{R}$-module) $\mathbb{R}^3$ over the ring $\mathbb{R}$ of real numbers. This vector space $\mathbb{R}^3$ is typically viewed as a model for three-dimensional space. Define a vector subspace (i.e., $\mathbb{R}$-submodule) I of $\mathbb{R}^3$ by

  $$I = \left\{ (x,y,z) \in \mathbb{R}^3 \mid x+y+z=0 \right\}.$$

  Geometrically, this is a hyperplane through the origin of $\mathbb{R}^3$. Now, consider the quotient $\mathbb{R}$-module (i.e., quotient vector space) $\mathbb{R}^3/I$. Its elements are residue classes of the form $\overline{(x,y,z)}$, where two vectors (x,y,z) and (x',y',z') belong to the same residue class if and only if their entrywise difference $(x-x',\ y-y',\ z-z')$ belongs to I (that is, if we have $(x-x')+(y-y')+(z-z')=0$). For instance, the two residue classes $\overline{(3,0,0)}$ and $\overline{(1,1,1)}$ are identical (since $(3-1)+(0-1)+(0-1)=0$), but the two residue classes $\overline{(1,0,0)}$ and $\overline{(2,0,0)}$ are not. It is not hard to see that each element of $\mathbb{R}^3/I$ can be uniquely written in the form $\overline{(r,0,0)}$ for some $r \in \mathbb{R}$. This shows that the vector space $\mathbb{R}^3/I$ is 1-dimensional.
- If R is any ring, and M is any left R-module, then the two obvious R-submodules $\{0_M\}$ and M of M lead to uninteresting quotient modules: The quotient module $M/\{0_M\}$ is isomorphic to M, whereas the quotient module M/M is trivial (i.e., has only one element).

[2] This incantation means "if you change what needs to be changed". For instance, instead of respecting multiplication, π now needs to respect scaling.

- Let R be a ring. As we recall from §2.1.3 (Lecture 18), the left R-module $R^{\mathbb{N}}$ has an R-submodule $R^{(\mathbb{N})}$. How does the quotient module $R^{\mathbb{N}}/R^{(\mathbb{N})}$ look like? Its elements are residue classes of the form $\overline{(a_0, a_1, a_2, \ldots)}$, where two infinite sequences $(a_0, a_1, a_2, \ldots)$ and $(b_0, b_1, b_2, \ldots)$ belong to the same residue class if and only if their entrywise difference $(a_0 - b_0,\ a_1 - b_1,\ a_2 - b_2,\ \ldots)$ belongs to $R^{(\mathbb{N})}$ (that is, if the two sequences $(a_0, a_1, a_2, \ldots)$ and $(b_0, b_1, b_2, \ldots)$ agree at all but finitely many positions). Thus, we can view an element $\overline{(a_0, a_1, a_2, \ldots)}$ of $R^{\mathbb{N}}/R^{(\mathbb{N})}$ as an "infinite sequence determined up to finite change" (where "finite change" means changing finitely many entries). This kind of construction is frequent in analysis: For instance, the limit $\lim\limits_{n\to\infty} a_n$ of a sequence $(a_0, a_1, a_2, \ldots)$ of real numbers does not depend on finite changes (i.e., it does not change if we change finitely many entries of our sequence), and thus (if it exists) can be viewed as a property of the residue class $\overline{(a_0, a_1, a_2, \ldots)} \in R^{\mathbb{N}}/R^{(\mathbb{N})}$.

For quotient rings, we have previously proved a universal property (Theorem 1.9.6 in Lecture 10) and a first isomorphism theorem (Theorem 1.9.10 in Lecture 11). Both of these have analogues for quotient modules. Let me just state the analogue of the universal property:

Theorem 2.4.13 (Universal property of quotient modules, elementwise form)**.** Let M be a left R-module. Let I be a left R-submodule of M.

Let N be a left R-module. Let $f : M \to N$ be a left R-module morphism. Assume that $f(I) = 0$ (that is, $f(i) = 0$ for each $i \in I$). Then, the map

$$\begin{aligned} f' : M/I &\to N, \\ \overline{a} &\mapsto f(a) \qquad \text{(for all } a \in M\text{)} \end{aligned}$$

is well-defined and is a left R-module morphism.

Proof. Analogous to the ring case. □

As in the case of rings, the map f' in Theorem 2.4.13 makes the following diagram commute:

$$\begin{tikzcd} M \arrow[d, "\pi"'] \arrow[dr, "f"] & \\ M/I \arrow[r, dashed, "f'"'] & N \end{tikzcd}$$

(where $\pi : M \to M/I$ is the canonical projection). The commutativity of this diagram is saying precisely that $f'(\pi(a)) = f(a)$ for all $a \in M$, that is, that $f'(\overline{a}) = f(a)$ for all $a \in M$; it thus is just the definition of f' in picture form.

2.5. Spanning, linear independence, bases and free modules

Again, we fix a ring R.

2.5.1. Definitions

We shall now define some classical notions of linear algebra (spanning, linear independence and bases) but in the generality of arbitrary R-modules.

> **Definition 2.5.1.** Let M be a left R-module. Let $m_1, m_2, \ldots, m_n$ be finitely many vectors in M. Then:
>
> **(a)** A **linear combination** of $m_1, m_2, \ldots, m_n$ means a vector of the form
>
> $$r_1 m_1 + r_2 m_2 + \cdots + r_n m_n \qquad \text{with } r_1, r_2, \ldots, r_n \in R.$$
>
> **(b)** The set of all linear combinations of $m_1, m_2, \ldots, m_n$ is called the **span** of $(m_1, m_2, \ldots, m_n)$, and is denoted by $\operatorname{span}(m_1, m_2, \ldots, m_n)$.
>
> **(c)** If the span of $(m_1, m_2, \ldots, m_n)$ is M, then we say that the vectors $m_1, m_2, \ldots, m_n$ **span** M (or **generate** M).
>
> **(d)** We say that the vectors $m_1, m_2, \ldots, m_n$ are **linearly independent** if the following holds: If $r_1, r_2, \ldots, r_n \in R$ satisfy
>
> $$r_1 m_1 + r_2 m_2 + \cdots + r_n m_n = 0,$$
>
> then $r_1 = r_2 = \cdots = r_n = 0$.
>
> **(e)** We say that the n-tuple $(m_1, m_2, \ldots, m_n)$ is a **basis** of the R-module M if $m_1, m_2, \ldots, m_n$ are linearly independent and span M.
>
> **(f)** All of this terminology depends on R. If R is not clear from the context, then we will say "R-linear combination", "R-span", etc.

These features can be defined not just for a finite list $(m_1, m_2, \ldots, m_n)$ of vectors, but for any family $(m_i)_{i \in I}$ of vectors. There is just one extra complication: We don't allow "truly infinite" linear combinations like $1m_0 + 1m_1 + 1m_2 + \cdots$ (because infinite sums are not defined in a general R-module). Thus, a linear combination of a family $(m_i)_{i \in I}$ has to be defined as a vector of the form $\sum\limits_{i \in I} r_i m_i$, where all but finitely many $i \in I$ satisfy $r_i = 0$. This latter condition ("all but finitely many $i \in I$ satisfy $r_i = 0$") ensures that the sum $\sum\limits_{i \in I} r_i m_i$ has only finitely many nonzero addends, and thus is well-defined (because it is just a finite sum inflated with a possibly infinite supply of zeroes). Of course, if the set I itself is finite, then this extra condition is automatically satisfied.

For example, for an infinite family $(m_i)_{i \in \mathbb{N}} = (m_0, m_1, m_2, \ldots)$ of vectors, the infinite sum $m_0 + m_1 + m_2 + \cdots$ does not count as a linear combination (even if this sum is defined to begin with[3]), but only finite sums like $m_1 + m_3$ or $m_0 + 17m_2 - 19m_3 + m_8$ do.

Likewise, linear independence for a family $(m_i)_{i \in I}$ is also defined in terms of finite sums only. Thus, the extension of Definition 2.5.1 to arbitrary families $(m_i)_{i \in I}$ looks as follows:

[3] In **some** modules, infinite sums like $m_0 + m_1 + m_2 + \cdots$ are **sometimes** defined, but we cannot count on this (this is not part of the structure of a module).

Definition 2.5.2. Let M be a left R-module. Let $(m_i)_{i\in I}$ be a family of vectors in M (with I being any set).

(a) A **linear combination** of $(m_i)_{i\in I}$ means a vector of the form

$$\sum_{i\in I} r_i m_i$$

for some family $(r_i)_{i\in I}$ of scalars (i.e., for some choice of $r_i \in R$ for each $i \in I$) with the property that

$$\text{all but finitely many } i \in I \text{ satisfy } r_i = 0. \qquad (1)$$

Here, the sum $\sum\limits_{i\in I} r_i m_i$ is an infinite sum, but all but finitely many of its addends are zero (thanks to the condition (1)). Such a sum is simply defined to be the sum of the nonzero addends. For example, $3+2+0+0+0+\cdots = 3+2=5$.

(b) The set of all linear combinations of $(m_i)_{i\in I}$ is called the **span** of $(m_i)_{i\in I}$, and is denoted by $\operatorname{span}(m_i)_{i\in I}$.

(c) If the span of $(m_i)_{i\in I}$ is M, then we say that the family $(m_i)_{i\in I}$ **spans** M (or **generates** M).

(d) We say that the family $(m_i)_{i\in I}$ is **linearly independent** if the following holds: If some family $(r_i)_{i\in I}$ of scalars $r_i \in R$ has the properties that

$$\text{all but finitely many } i \in I \text{ satisfy } r_i = 0 \qquad (2)$$

and that

$$\sum_{i\in I} r_i m_i = 0,$$

then $r_i = 0$ for all $i \in I$.

(e) We say that the family $(m_i)_{i\in I}$ is a **basis** of the R-module M if $(m_i)_{i\in I}$ is linearly independent and spans M.

(f) All of this terminology depends on R. Thus, if R is not clear from the context, then we will say "R-linear combination", "R-span", etc.

2.5.2. Spans are submodules

As in linear algebra, we can generate submodules of a given module M by taking spans of (families of) vectors in M:

Proposition 2.5.3. Let M be a left R-module. Let $(m_i)_{i\in I}$ be a family of vectors in M. Then, the span of this family is an R-submodule of M.

Proof. We must prove that this span is closed under addition, closed under scaling, and contains 0.

Let's only check "closed under addition": This means proving that any sum of two linear combinations of $(m_i)_{i\in I}$ is again a linear combination of $(m_i)_{i\in I}$.

Let's prove this: Consider two arbitrary linear combinations $\sum_{i\in I} a_i m_i$ and $\sum_{i\in I} b_i m_i$ of $(m_i)_{i\in I}$. We have to prove that $\sum_{i\in I} a_i m_i + \sum_{i\in I} b_i m_i$ is again a linear combination of $(m_i)_{i\in I}$.

This is almost trivial:

$$\sum_{i\in I} a_i m_i + \sum_{i\in I} b_i m_i = \sum_{i\in I} (a_i m_i + b_i m_i) = \sum_{i\in I} (a_i + b_i) m_i.$$

But wait: We still need to check that the right hand side here is a legit linear combination, i.e., that all but finitely many $i \in I$ satisfy $a_i + b_i = 0$. In other words, we need to check that only finitely many $i \in I$ satisfy $a_i + b_i \neq 0$.

However, by the definition of a linear combination, we know that

- only finitely many $i \in I$ satisfy $a_i \neq 0$ (since $\sum_{i\in I} a_i m_i$ is a linear combination);
- only finitely many $i \in I$ satisfy $b_i \neq 0$ (since $\sum_{i\in I} b_i m_i$ is a linear combination).

Since the union of two finite sets is always finite, we thus conclude that only finitely many $i \in I$ have the property that at least one of a_i and b_i is nonzero. Therefore, only finitely many $i \in I$ satisfy $a_i + b_i \neq 0$ (because for $a_i + b_i$ to be $\neq 0$, it must hold that at least one of a_i and b_i is nonzero). Therefore, $\sum_{i\in I} (a_i + b_i) m_i$ is a legit linear combination.

This completes the proof of "closed under addition" for the span of $(m_i)_{i\in I}$. As we said, the other axioms are similar or easier. Thus, Proposition 2.5.3 is proved. □

2.5.3. Coordinates

The notions of linear independence and spanning can be described in a slightly different (but equivalent) form:

Proposition 2.5.4. Let M be a left R-module. Let $(m_i)_{i\in I}$ be a family of vectors in M (with I being any set). Then:

(a) The family $(m_i)_{i\in I}$ spans M if and only if each vector $v \in M$ can be written as an R-linear combination of $(m_i)_{i\in I}$ in **at least one** way.

(b) The family $(m_i)_{i\in I}$ is linearly independent if and only if each vector $v \in M$ can be written as an R-linear combination of $(m_i)_{i\in I}$ in **at most one** way (i.e., there is **at most one** family $(r_i)_{i\in I}$ of scalars such that $v = \sum_{i\in I} r_i m_i$ and such that all but finitely many $i \in I$ satisfy $r_i = 0$).

(c) The family $(m_i)_{i\in I}$ is a basis of M if and only if each vector $v \in M$ can be written as an R-linear combination of $(m_i)_{i\in I}$ in **exactly one** way (i.e., there is **exactly one** family $(r_i)_{i\in I}$ of scalars such that $v = \sum\limits_{i\in I} r_i m_i$ and such that all but finitely many $i \in I$ satisfy $r_i = 0$).

Proposition 2.5.4 **(c)** shows that a basis of an R-module M can be used as a "coordinate system" on M, allowing to identify each vector $v \in M$ by a family of scalars $(r_i)_{i\in I}$ (which are the "coordinates" of v with respect to this basis).

Proof of Proposition 2.5.4. **(a)** This is a trivial consequence of the definitions of span and spanning.

(b) $\Longrightarrow$: Assume that the family $(m_i)_{i\in I}$ is linearly independent. We must prove that each vector $v \in M$ can be written as an R-linear combination of $(m_i)_{i\in I}$ in **at most one** way.

So let $v \in M$ be a vector. Let $v = \sum\limits_{i\in I} a_i m_i$ and $v = \sum\limits_{i\in I} b_i m_i$ be two ways to write v as an R-linear combination of $(m_i)_{i\in I}$ (where a_i and b_i are scalars such that all but finitely many $i \in I$ satisfy $a_i = 0$, and such that all but finitely many $i \in I$ satisfy $b_i = 0$). We must prove that these two ways are actually identical, i.e., that we have $(a_i)_{i\in I} = (b_i)_{i\in I}$.

Indeed, subtracting the equalities $v = \sum\limits_{i\in I} a_i m_i$ and $v = \sum\limits_{i\in I} b_i m_i$ from one another, we find

$$0 = \sum_{i\in I} a_i m_i - \sum_{i\in I} b_i m_i = \sum_{i\in I} (a_i m_i - b_i m_i) = \sum_{i\in I} (a_i - b_i)\, m_i.$$

Moreover, it is easy to see that all but finitely many $i \in I$ satisfy $a_i - b_i = 0$. Thus, from $\sum\limits_{i\in I} (a_i - b_i)\, m_i = 0$, we conclude that $a_i - b_i = 0$ for each $i \in I$ (since $(m_i)_{i\in I}$ is linearly independent). In other words, $a_i = b_i$ for each $i \in I$. In other words, $(a_i)_{i\in I} = (b_i)_{i\in I}$.

This completes our proof of the "$\Longrightarrow$" direction of part **(b)**.

$\Longleftarrow$: Assume that each vector $v \in M$ can be written as an R-linear combination of $(m_i)_{i\in I}$ in **at most one** way. Applying this assumption to $v = 0$, we conclude that 0 can be written as an R-linear combination of $(m_i)_{i\in I}$ in **at most one** way. But clearly, one way to write 0 as an R-linear combination of $(m_i)_{i\in I}$ is $0 = \sum\limits_{i\in I} 0 m_i$. Hence, this must be the **only** way to write 0 as an R-linear combination of $(m_i)_{i\in I}$ (since there is at most one way). In other words, if $(r_i)_{i\in I}$ is a family of scalars satisfying $0 = \sum\limits_{i\in I} r_i m_i$ (and having the property that all but finitely many $i \in I$ satisfy $r_i = 0$), then we must have $r_i = 0$ for all $i \in I$. But this is saying precisely that the family $(m_i)_{i\in I}$ is linearly independent. This proves the "$\Longleftarrow$" direction of part **(b)**.

(c) This follows by combining parts **(a)** and **(b)**. □

2.5.4. Free modules

Thanks partly to Proposition 2.5.4 **(c)**, modules that have bases are the nicest and simplest modules around. They have a name:

Definition 2.5.5. **(a)** A left R-module M is said to be **free** if it has a basis.

(b) Let $n \in \mathbb{N}$. A left R-module M is said to be **free of rank** n if it has a basis of size n (that is, a basis consisting of n vectors).

Note that not every free R-module has a rank in this sense, since its basis could be infinite. (Also, a free R-module can have several ranks at the same time, although this doesn't happen very often.)

For vector spaces (i.e., modules over a field), freeness comes for free:

Theorem 2.5.6. If F is a field, then every F-module (= F-vector space) has a basis.

Proof. I will not prove this. Proofs for the finitely generated case (i.e., for F-modules that are spanned by a finite list of vectors) are easy to find (e.g., Keith Conrad's `https://kconrad.math.uconn.edu/blurbs/linmultialg/dimension.pdf` , or [Treil21, Chapter 1, Proposition 2.8], or all sorts of textbooks). Proofs for the general case involve some set theory (including the Axiom of Choice), and can be found in more advanced literature (e.g., Keith Conrad's note `https://kconrad.math.uconn.edu/blurbs/zorn1.pdf` on Zorn's lemma). □

For example, Theorem 2.5.6 shows that the $\mathbb{Q}$-vector space $\mathbb{R}$ is free. In other words (restated using Proposition 2.5.4 **(c)**), this is saying that there is a family $(b_i)_{i\in I}$ of real numbers such that every real number can be uniquely written as a $\mathbb{Q}$-linear combination of this family (i.e., as a sum $\sum\limits_{i\in I} r_i b_i$ where $r_i \in \mathbb{Q}$ are rational and all but finitely many of them are zero). No one can actually find such a family, because the proof of its existence is not constructive. Such bases are called **Hamel bases**.

To find more interesting examples, we have to consider rings that are not fields. We'll discuss this next time.

References

[Treil21] Serge Treil, *Linear Algebra Done Wrong*, 11 January 2021. `https://sites.google.com/a/brown.edu/sergei-treil-homepage/linear-algebra-done-wrong`